

P2020 QorIQ Integrated Processor Reference Manual

Supports: P2020 and P2010

Document Number: P2020RM
Rev. 2, 12/2012

Contents

Section number	Title	Page
Chapter 1 Overview		
1.1	Overview.....	67
1.1.1	Block diagram.....	67
1.1.2	Critical performance parameters.....	68
1.1.3	Chip-level features.....	69
1.2	Application examples.....	70
1.2.1	LTE and WiMax baseband application.....	70
1.2.2	Line card control plane application.....	71
1.3	Architecture overview.....	72
1.3.1	e500v2 cores and memory unit.....	72
1.3.2	e500 coherency module (ECM) and address map.....	73
1.3.3	Integrated security engine (SEC).....	74
1.3.4	Enhanced three-speed Ethernet controllers.....	74
1.3.5	Universal serial bus (USB) 2.0.....	76
1.3.6	Enhanced secure digital host controller.....	76
1.3.7	Enhanced serial peripheral interface (eSPI).....	77
1.3.8	DDR SDRAM controller.....	77
1.3.9	High speed I/O interfaces.....	78
1.3.9.1	PCI Express interfaces.....	78
1.3.9.2	Serial RapidIO interfaces.....	79
1.3.9.3	SGMII.....	79
1.3.9.4	High-speed interface multiplexing.....	79
1.3.10	Programmable interrupt controller (PIC).....	80
1.3.11	DMA, I2C, DUART, and eLBC.....	80
1.3.12	Device boot locations.....	81
1.3.13	Boot sequencer.....	82

Section number	Title	Page
1.3.14	Power management.....	82
1.3.15	System performance monitor.....	82
Chapter 2 Memory Map		
2.1	Overview.....	83
2.2	Configuration, control, and status registers.....	84
2.2.1	Accessing CCSR memory from the local processor.....	85
2.2.2	Accessing CCSR memory from external masters.....	85
2.2.3	Organization of CCSR space.....	86
2.2.3.1	General utilities registers.....	86
2.2.3.1.1	General utilities register organization.....	87
2.2.3.2	Programmable interrupt controller registers.....	88
2.2.3.3	Serial RapidIO registers.....	89
2.2.3.4	Device-specific utilities registers.....	90
2.2.4	CCSR address map.....	91
2.3	Local access windows.....	94
2.3.1	Precedence of local access windows.....	95
2.3.2	Configuring local access windows.....	95
2.3.3	Distinguishing local access windows from other mapping functions.....	95
2.3.4	Illegal interaction between local access windows and DDR chip selects.....	96
2.3.5	Local address map example.....	96
2.4	Local access window registers.....	97
2.4.1	Local access window 0 base address register (LAW_LAWBAR _n).....	99
2.4.2	Local access window 0 attribute register (LAW_LAWAR _n).....	99
2.5	Address translation and mapping units.....	101
2.5.1	Address translation	101
2.5.2	Outbound ATMUs.....	102
2.5.3	Inbound ATMUs.....	103
2.5.3.1	Illegal interaction between inbound ATMUs and LAWs.....	103

Section number	Title	Page
Chapter 3		
Signal Descriptions		
3.1	Signals overview.....	105
3.2	Configuration signals sampled at reset.....	116
3.3	Output signal states during reset.....	118
Chapter 4		
Reset, Clocking, and Initialization		
4.1	Overview.....	119
4.2	External signal descriptions.....	119
4.2.1	System control signals.....	120
4.2.2	Clock signals.....	121
4.3	Accessing configuration, control, and status registers.....	122
4.3.1	Updating CCSRBAR.....	122
4.3.2	Accessing alternate configuration space.....	123
4.3.3	Boot page translation.....	124
4.3.4	Boot sequencer.....	124
4.4	Reset Memory Map/Register Definition.....	124
4.4.1	Configuration, control, and status registers base address register (Reset_CCSRBAR).....	125
4.4.2	Alternate configuration base address register (Reset_ALTCBAR).....	126
4.4.3	Alternate configuration attribute register (Reset_ALTCAR).....	126
4.4.4	Boot page translation register (Reset_BPTR).....	127
4.5	Functional description.....	127
4.5.1	Reset operations.....	127
4.5.1.1	Soft reset.....	128
4.5.1.2	Hard reset.....	128
4.5.2	Power-on reset sequence.....	128
4.5.3	Power-on reset configuration.....	130
4.5.3.1	System PLL ratio.....	131
4.5.3.2	DDR PLL Ratio.....	132

Section number	Title	Page
4.5.3.3	e500 core PLL ratios.....	132
4.5.3.4	Boot ROM location.....	133
4.5.3.5	Host/agent configuration.....	134
4.5.3.6	I/O port selection.....	135
4.5.3.7	CPU boot configuration.....	138
4.5.3.8	Boot sequencer configuration.....	139
4.5.3.9	DDR SDRAM type.....	140
4.5.3.10	SerDes reference clock configuration.....	140
4.5.3.11	eTSECn physical interface mode configuration.....	141
4.5.3.11.1	eTSEC2 SGMII mode.....	142
4.5.3.11.2	eTSEC3 SGMII mode.....	142
4.5.3.11.3	eTSEC1 and eTSEC2 width.....	143
4.5.3.11.4	eTSEC1 protocol.....	143
4.5.3.11.5	eTSEC2 protocol.....	144
4.5.3.11.6	eTSEC3 protocol.....	144
4.5.3.12	RapidIO device ID.....	145
4.5.3.13	RapidIO system size.....	146
4.5.3.14	Memory debug configuration.....	146
4.5.3.15	DDR debug configuration.....	147
4.5.3.16	General-purpose POR configuration	147
4.5.3.17	Engineering use POR configuration.....	148
4.5.3.18	eLBC ECC enable configuration.....	148
4.5.3.19	System speed.....	149
4.5.3.20	Platform speed.....	149
4.5.3.21	Core 0 speed.....	149
4.5.3.22	Core 1 speed.....	150
4.5.3.23	DDR controller speed.....	150
4.5.3.24	SerDes PLL time-out enable.....	151
4.5.3.25	eSDHC card-detect polarity select.....	152

Section number	Title	Page
4.5.4	Voltage selection configuration.....	152
4.5.5	Clocking.....	153
4.5.5.1	System clock and DDR controller clock.....	153
4.5.5.2	RapidIO, SGMII and PCI Express clocks.....	154
4.5.5.2.1	Minimum frequency requirements.....	155
4.5.5.3	Ethernet clocks.....	155
4.5.5.4	Real time clock.....	155
4.6	Initialization/applications information.....	156
4.6.1	eSDHC boot.....	157
4.6.1.1	eSDHC boot overview.....	157
4.6.1.2	eSDHC boot features.....	158
4.6.1.3	SD/MMC card data structure.....	159
4.6.1.3.1	Configuration words section.....	160
4.6.1.3.2	Notes on compatibility with FAT12/FAT16/FAT32 filesystems.....	162
4.6.1.4	eSDHC controller initial configuration.....	163
4.6.1.5	eSDHC controller boot sequence.....	164
4.6.1.6	eSDHC boot error handling.....	164
4.6.2	eSPI boot ROM.....	166
4.6.2.1	eSPI boot overview.....	166
4.6.2.2	eSPI boot features.....	167
4.6.2.3	EEPROM data structure.....	168
4.6.2.3.1	EEPROM configuration words section.....	169
4.6.2.4	eSPI controller configuration.....	171
4.6.3	Default e500 addressing during system boot.....	173

Chapter 5 e500 Core Integration Details

5.1	Overview.....	175
5.2	e500 core block diagram.....	176
5.3	e500 core integration and the core complex bus (CCB).....	178

Section number	Title	Page
5.4	Summary of core integration details.....	180
5.4.1	Processor version register (PVR) and system version register (SVR).....	182

Chapter 6 L2 Look-Aside Cache/SRAM

6.1	Introduction.....	183
6.2	L2 cache overview.....	183
6.2.1	L2 cache and SRAM features.....	184
6.3	L2 cache and SRAM organization.....	186
6.3.1	Accessing the on-chip array as an L2 cache.....	187
6.3.2	Accessing the on-chip array as an SRAM.....	188
6.3.3	Connection of the on-chip memory to the system.....	190
6.4	L2 cache external write registers.....	191
6.5	L2 memory-mapped SRAM registers.....	192
6.6	L2 error registers.....	192
6.7	L2_Cache memory map/register definition.....	193
6.7.1	L2 control register (L2_Cache_L2CTL).....	194
6.7.2	L2 cache way allocation for processors (L2_Cache_L2CWAP).....	199
6.7.3	L2 cache external write address register n (L2_Cache_L2CEWAR n).....	201
6.7.4	L2 cache external write address register extended address n (L2_Cache_L2CEWAREA n).....	201
6.7.5	L2 cache external write control register n (L2_Cache_L2CEWCR n).....	202
6.7.6	L2 memory-mapped SRAM base address register n (L2_Cache_L2SRBAR n).....	203
6.7.7	L2 memory-mapped SRAM base address register extended address n (L2_Cache_L2SRBAREA n).....	204
6.7.8	L2 error injection mask high register (L2_Cache_L2ERRINJHI).....	205
6.7.9	L2 error injection mask low register (L2_Cache_L2ERRINJLO).....	206
6.7.10	L2 error injection tag/ECC control register (L2_Cache_L2ERRINJCTL).....	206
6.7.11	L2 error data high capture register (L2_Cache_L2CAPTDATAHI).....	207
6.7.12	L2 error data low capture register (L2_Cache_L2CAPTDATALO).....	208
6.7.13	L2 error syndrome register (L2_Cache_L2CAPTECC).....	208
6.7.14	L2 error detect register (L2_Cache_L2ERRDET).....	210

Section number	Title	Page
6.7.15	L2 error disable register (L2_Cache_L2ERRDIS).....	212
6.7.16	L2 error interrupt enable register (L2_Cache_L2ERRINTEN).....	214
6.7.17	L2 error attributes capture register (L2_Cache_L2ERRATTR).....	216
6.7.18	L2 error address capture register low (L2_Cache_L2ERRADDRL).....	218
6.7.19	L2 error address capture register high (L2_Cache_L2ERRADDRH).....	218
6.7.20	L2 error control register (L2_Cache_L2ERRCTL).....	219
6.8	External writes to the L2 cache (cache stashing).....	219
6.8.1	Stash-only cache regions.....	220
6.9	L2 cache timing.....	220
6.10	L2 cache and SRAM coherency.....	222
6.10.1	L2 cache coherency rules	222
6.10.2	Memory-mapped SRAM coherency rules.....	224
6.11	L2 cache locking.....	224
6.11.1	Locking the entire L2 cache.....	224
6.11.2	Locking programmed memory ranges.....	225
6.11.3	Locking selected lines.....	225
6.11.4	Clearing locks on selected lines.....	226
6.11.5	Flash clearing of instruction and data locks.....	226
6.11.6	Locks with stale data.....	226
6.12	PLRU L2 replacement policy.....	227
6.12.1	PLRU bit update considerations.....	227
6.12.2	Allocation of lines.....	228
6.13	L2 cache operation.....	229
6.13.1	Initialization.....	229
6.13.1.1	L2 cache initialization.....	229
6.13.1.2	Memory-mapped SRAM initialization.....	230
6.13.2	Flash invalidation of the L2 cache.....	230
6.13.3	Managing errors.....	231
6.13.3.1	ECC errors	231

Section number	Title	Page
6.13.3.2	Tag parity errors	231
6.13.4	L2 cache states	231
6.13.5	L2 state transitions.....	232
6.14	Error checking and correcting (ECC).....	237

Chapter 7 e500 Coherency Module

7.1	Introduction.....	239
7.1.1	Overview.....	240
7.1.2	Features.....	240
7.2	ECM memory map/register definition.....	241
7.2.1	ECM CCB address configuration register (ECM_EEBACR).....	242
7.2.2	ECM CCB port configuration register (ECM_EEBPCR).....	243
7.2.3	ECM IP Block Revision Register 1 (ECM_EIPBRR1).....	245
7.2.4	ECM IP Block Revision Register 2 (ECM_EIPBRR2).....	245
7.2.5	ECM error detect register (ECM_EEDR).....	246
7.2.6	ECM error enable register (ECM_EEER).....	247
7.2.7	ECM error attributes capture register (ECM_EEATR).....	248
7.2.8	ECM error low address capture register (ECM_EELADR).....	250
7.2.9	ECM error high address capture register (ECM_EEHADR).....	250
7.3	Functional description.....	251
7.3.1	I/O arbiter.....	251
7.3.2	CCB arbiter.....	251
7.3.3	Transaction queue	252
7.3.4	Global data multiplexor.....	252
7.3.5	CCB interface.....	252
7.4	Initialization/application information.....	252

Chapter 8 DDR Memory Controller

8.1	Introduction.....	255
-----	-------------------	-----

Section number	Title	Page
8.2	Features.....	256
8.2.1	Modes of operation.....	257
8.3	External signal descriptions.....	257
8.3.1	Signals overview.....	257
8.3.2	Detailed signal descriptions.....	259
8.3.2.1	Memory interface signals.....	259
8.3.2.2	Clock interface signals.....	263
8.4	DDR memory map/register definition.....	264
8.4.1	Chip select n memory bounds (DDR_CS _n _BNDS).....	266
8.4.2	Chip select n configuration (DDR_CS _n _CONFIG).....	267
8.4.3	Chip select n configuration 2 (DDR_CS _n _CONFIG_2).....	269
8.4.4	DDR SDRAM timing configuration 3 (DDR_TIMING_CFG_3).....	270
8.4.5	DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0).....	272
8.4.6	DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1).....	276
8.4.7	DDR SDRAM timing configuration 2 (DDR_TIMING_CFG_2).....	280
8.4.8	DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG).....	284
8.4.9	DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2).....	287
8.4.10	DDR SDRAM mode configuration (DDR_DDR_SDRAM_MODE).....	290
8.4.11	DDR SDRAM mode configuration 2 (DDR_DDR_SDRAM_MODE_2).....	291
8.4.12	DDR SDRAM mode control (DDR_DDR_SDRAM_MD_CNTL).....	291
8.4.13	DDR SDRAM interval configuration (DDR_DDR_SDRAM_INTERVAL).....	295
8.4.14	DDR SDRAM data initialization (DDR_DDR_DATA_INIT).....	295
8.4.15	DDR SDRAM clock control (DDR_DDR_SDRAM_CLK_CNTL).....	296
8.4.16	DDR training initialization address (DDR_DDR_INIT_ADDR).....	297
8.4.17	DDR training initialization extended address (DDR_DDR_INIT_EXT_ADDR).....	298
8.4.18	DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4).....	299
8.4.19	DDR SDRAM timing configuration 5 (DDR_TIMING_CFG_5).....	302
8.4.20	DDR ZQ calibration control (DDR_DDR_ZQ_CNTL).....	304
8.4.21	DDR write leveling control (DDR_DDR_WRLVL_CNTL).....	306

Section number	Title	Page
8.4.22	DDR Self Refresh Counter (DDR_DDR_SR_CNTR).....	309
8.4.23	DDR Register Control Words 1 (DDR_DDR_SDRAM_RCW_1).....	310
8.4.24	DDR Register Control Words 2 (DDR_DDR_SDRAM_RCW_2).....	311
8.4.25	DDR write leveling control 2 (DDR_DDR_WRLVL_CNTL_2).....	312
8.4.26	DDR write leveling control 3 (DDR_DDR_WRLVL_CNTL_3).....	314
8.4.27	DDR Debug Status Register 1 (DDR_DDRDSR_1).....	317
8.4.28	DDR Debug Status Register 2 (DDR_DDRDSR_2).....	318
8.4.29	DDR Control Driver Register 1 (DDR_DDRCDR_1).....	318
8.4.30	DDR Control Driver Register 2 (DDR_DDRCDR_2).....	322
8.4.31	DDR IP block revision 1 (DDR_DDR_IP_REV1).....	323
8.4.32	DDR IP block revision 2 (DDR_DDR_IP_REV2).....	323
8.4.33	Memory data path error injection mask high (DDR_DATA_ERR_INJECT_HI).....	324
8.4.34	Memory data path error injection mask low (DDR_DATA_ERR_INJECT_LO).....	324
8.4.35	Memory data path error injection mask ECC (DDR_ERR_INJECT).....	325
8.4.36	Memory data path read capture high (DDR_CAPTURE_DATA_HI).....	326
8.4.37	Memory data path read capture low (DDR_CAPTURE_DATA_LO).....	326
8.4.38	Memory data path read capture ECC (DDR_CAPTURE_ECC).....	327
8.4.39	Memory error detect (DDR_ERR_DETECT).....	328
8.4.40	Memory error disable (DDR_ERR_DISABLE).....	330
8.4.41	Memory error interrupt enable (DDR_ERR_INT_EN).....	332
8.4.42	Memory error attributes capture (DDR_CAPTURE_ATTRIBUTES).....	334
8.4.43	Memory error address capture (DDR_CAPTURE_ADDRESS).....	336
8.4.44	Memory error extended address capture (DDR_CAPTURE_EXT_ADDRESS).....	336
8.4.45	Single-Bit ECC memory error management (DDR_ERR_SBE).....	337
8.5	Functional description.....	337
8.5.1	DDR SDRAM interface operation.....	343
8.5.1.1	Supported DDR SDRAM organizations.....	344
8.5.2	DDR SDRAM address multiplexing.....	345
8.5.3	JEDEC standard DDR SDRAM interface commands.....	360

Section number	Title	Page
8.5.4	DDR SDRAM interface timing.....	362
8.5.5	DDR SDRAM registered DIMM mode.....	365
8.5.6	DDR SDRAM write timing adjustments.....	366
8.5.7	DDR SDRAM refresh.....	367
8.5.7.1	DDR SDRAM refresh timing.....	368
8.5.7.2	DDR SDRAM refresh and power-saving modes.....	369
8.5.7.2.1	Self-refresh in sleep mode.....	370
8.5.8	DDR data beat ordering.....	371
8.5.9	Page mode and logical bank retention.....	372
8.5.10	Error checking and correcting (ECC).....	373
8.5.11	Error management.....	375
8.6	Initialization/application information.....	376
8.6.1	Programming differences between memory types.....	380
8.6.2	DDR SDRAM initialization sequence.....	386

Chapter 9 Programmable Interrupt Controller (PIC)

9.1	Introduction.....	387
9.1.1	Overview.....	387
9.1.2	The PIC in multiple-processor implementations.....	390
9.1.3	Interrupts to the e500 processor core.....	390
9.1.4	Modes of operation.....	391
9.1.4.1	Mixed mode (GCR[M] = 1).....	391
9.1.4.2	Pass-through mode (GCR[M] = 0).....	391
9.1.5	Interrupt sources.....	392
9.1.5.1	Interrupt routing-mixed mode.....	393
9.1.5.2	Interrupt destinations.....	393
9.1.5.3	Internal interrupt sources.....	394
9.2	PIC external signal descriptions.....	395
9.2.1	Signal overview.....	395

Section number	Title	Page
9.2.2	Detailed signal descriptions.....	395
9.3	PIC memory map/register definition.....	396
9.3.1	Block revision register 1 (PIC_BRR1).....	407
9.3.2	Block revision register 2 (PIC_BRR2).....	407
9.3.3	Interprocessor n dispatch register (PIC_IPIDR _n).....	408
9.3.4	Current task priority register (PIC_CTPR).....	409
9.3.5	Who am I register (PIC_WHOAMI).....	410
9.3.6	Interrupt acknowledge register (PIC_IACK).....	410
9.3.7	End of interrupt register (PIC_EOI).....	411
9.3.8	Feature reporting register (PIC_FRR).....	412
9.3.9	Global configuration register (PIC_GCR).....	413
9.3.10	Vendor identification register (PIC_VIR).....	414
9.3.11	Processor core initialization register (PIC_PIR).....	414
9.3.12	Interprocessor interrupt n vector/priority register (PIC_IPIVPR _n).....	415
9.3.13	Spurious vector register (PIC_SVR).....	416
9.3.14	Timer frequency reporting register group X (PIC_TFRR _n).....	416
9.3.15	Global timer n current count register group A (PIC_GTCCR _A _n).....	417
9.3.16	Global timer n base count register group A (PIC_GTBCR _A _n).....	418
9.3.17	Global timer n vector/priority register group A (PIC_GTVPR _A _n).....	419
9.3.18	Global timer n destination register group A (PIC_GTDRA _n).....	420
9.3.19	Timer control register group n (PIC_TCR _n).....	420
9.3.20	External interrupt summary register (PIC_ERQSR).....	423
9.3.21	IRQ_OUT_B summary register 0 (PIC_IRQSR0).....	423
9.3.22	IRQ_OUT_B summary register 1 (PIC_IRQSR1).....	424
9.3.23	IRQ_OUT_B summary register 2 (PIC_IRQSR2).....	424
9.3.24	Critical interrupt summary register 0 (PIC_CISR0).....	425
9.3.25	Critical interrupt summary register 1 (PIC_CISR1).....	425
9.3.26	Critical interrupt summary register 2 (PIC_CISR2).....	426
9.3.27	Performance monitor n mask register 0 (PIC_PM _n MR0).....	426

Section number	Title	Page
9.3.28	Performance monitor n mask register 1 (PIC_PM n MR1).....	427
9.3.29	Performance monitor n mask register 2 (PIC_PM n MR2).....	427
9.3.30	Message register n (PIC_MSGR n).....	428
9.3.31	Message enable register (PIC_MER).....	428
9.3.32	Message status register (PIC_MSR).....	429
9.3.33	Shared message signaled interrupt register n (PIC_MSIR n).....	429
9.3.34	Shared message signaled interrupt status register (PIC_MSISR).....	430
9.3.35	Shared message signaled interrupt index register (PIC_MSIIR).....	431
9.3.36	Global timer n current count register group B (PIC_GTCCRB n).....	432
9.3.37	Global timer n base count register group B (PIC_GTBCRB n).....	433
9.3.38	Global timer n vector/priority register group B (PIC_GTVPRB n).....	434
9.3.39	Global timer n destination register group B (PIC_GTDRB n).....	435
9.3.40	Message register n (PIC_MSGR n).....	435
9.3.41	Message enable register (PIC_MERa).....	436
9.3.42	Message status register (PIC_MSRa).....	437
9.3.43	External interrupt n (IRQn) vector/priority register (PIC_EIVPR n).....	437
9.3.44	External interrupt n (IRQn) destination register (PIC_EIDR n).....	439
9.3.45	Internal interrupt n vector/priority register (PIC_IIVPR n).....	440
9.3.46	Internal interrupt n destination register (PIC_IIDR n).....	441
9.3.47	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR n).....	442
9.3.48	Messaging interrupt n (MSGn) destination register (PIC_MIDR n).....	443
9.3.49	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR n).....	444
9.3.50	Shared message signaled interrupt destination register n (PIC_MSIDR n).....	445
9.3.51	Processor core 0 interprocessor n dispatch register (PIC_IPIDR_CPU0 n).....	446
9.3.52	Processor core current task priority register 0 Processor core (PIC_CTPR_CPU0).....	446
9.3.53	Processor core 0 who am I register (PIC_WHOAMI_CPU0).....	447
9.3.54	Processor core 0 interrupt acknowledge register (PIC_IACK_CPU0).....	448
9.3.55	Processor core 0 end of interrupt register (PIC_EOI_CPU0).....	449
9.3.56	Processor core 1 interprocessor n dispatch register (PIC_IPIDR_CPU1 n).....	450

Section number	Title	Page
9.3.57	Processor core 1 current task priority register (PIC_CTPR_CPU1).....	450
9.3.58	Processor core 1 who am I register (PIC_WHOAMI_CPU1).....	451
9.3.59	Processor core 1 interrupt acknowledge register (PIC_IACK_CPU1).....	452
9.3.60	Processor core 1 end of interrupt register (PIC_EOI_CPU1).....	453
9.4	Functional description.....	453
9.4.1	Programming model considerations.....	453
9.4.1.1	Global registers.....	453
9.4.1.2	Global timer registers.....	454
9.4.1.3	IRQ_OUT_B and critical interrupt summary registers.....	454
9.4.1.4	Performance monitor mask registers (PMMRs).....	455
9.4.1.5	Message registers.....	455
9.4.1.6	Shared message signaled registers.....	455
9.4.1.7	Interrupt source configuration registers.....	455
9.4.1.8	Per-CPU (private access) registers.....	457
9.4.2	Flow of interrupt control.....	459
9.4.2.1	Interrupts routed to cint or IRQ_OUT_B.....	459
9.4.2.2	Interrupts routed to int.....	460
9.4.2.2.1	Interrupt source priority.....	462
9.4.2.2.2	Interrupt acknowledge.....	462
9.4.2.2.3	Spurious vector generation.....	463
9.4.2.2.4	Nesting of interrupts.....	463
9.4.3	Interprocessor interrupts.....	464
9.4.4	Message interrupts.....	464
9.4.5	Shared message signaled interrupts.....	464
9.4.6	PCI Express INTx/IRQn sharing.....	465
9.4.7	Global timers.....	466
9.4.8	Resets.....	467
9.4.9	Resetting the PIC.....	467
9.4.9.1	Processor core initialization.....	467

Section number	Title	Page
9.5	Initialization/application information.....	467
9.5.1	Programming guidelines.....	468
9.5.1.1	PIC registers.....	468
9.5.1.2	Changing interrupt source configuration.....	469

Chapter 10 I2C Interfaces

10.1	Overview.....	471
10.2	Introduction to I2C.....	471
10.2.1	What is the I2C module?.....	471
10.2.2	I2C module block diagram.....	472
10.2.3	Features	472
10.2.4	Advantages of the I2C bus.....	473
10.2.5	Modes of operation.....	473
10.2.6	I2C-specific conditions.....	473
10.3	I2C external signal descriptions.....	474
10.3.1	Signal overview.....	474
10.3.2	Detailed signal descriptions.....	474
10.4	I2C memory map/register definition.....	475
10.4.1	I2C address register (I2Cx_I2CADR).....	476
10.4.2	I2C frequency divider register (I2Cx_I2CFDR).....	477
10.4.3	I2C control register (I2Cx_I2CCR).....	479
10.4.4	I2C status register (I2Cx_I2CSR).....	480
10.4.5	I2C data register (I2Cx_I2CDR).....	481
10.4.6	I2C digital filter sampling rate register (I2Cx_I2CDFSRR).....	482
10.5	Functional description.....	482
10.5.1	Transaction protocol.....	482
10.5.1.1	START condition.....	483
10.5.1.2	Slave address transmission.....	483
10.5.1.3	Repeated START condition.....	484

Section number	Title	Page
10.5.1.4	STOP condition.....	485
10.5.1.5	Protocol implementation details.....	485
10.5.1.5.1	Transaction monitoring-implementation details.....	485
10.5.1.5.2	Control transfer-implementation details.....	485
10.5.1.6	Address compare-implementation details.....	486
10.5.2	Arbitration procedure.....	487
10.5.2.1	Arbitration control.....	487
10.5.3	Handshaking.....	488
10.5.4	Clock control.....	488
10.5.4.1	Clock synchronization.....	488
10.5.4.2	Input synchronization and digital filter.....	489
10.5.4.2.1	Input signal synchronization.....	489
10.5.4.2.2	Filtering of SCL and SDA lines.....	489
10.5.4.3	Clock stretching.....	489
10.5.5	Boot sequencer mode.....	490
10.5.5.1	EEPROM calling address.....	491
10.5.5.2	EEPROM data format.....	491
10.6	Initialization/application information.....	494
10.6.1	Initialization sequence.....	494
10.6.2	Generation of START.....	495
10.6.3	Post-transfer software response.....	495
10.6.4	Generation of STOP.....	496
10.6.5	Generation of repeated START.....	496
10.6.6	Generation of SCL when SDA low.....	496
10.6.7	Slave mode interrupt service routine.....	497
10.6.7.1	Slave transmitter and received acknowledge.....	497
10.6.7.2	Loss of arbitration and forcing of slave mode.....	497
10.6.8	Interrupt service routine flowchart.....	498

Section number	Title	Page
Chapter 11		
DUART		
11.1	Introduction.....	501
11.1.1	Overview.....	501
11.1.1.1	Features.....	502
11.1.1.2	Modes of operation.....	503
11.2	DUART external signal descriptions.....	503
11.3	DUART memory map/register definition.....	504
11.3.1	Receiver Buffer Registers (DUART_URBR _n).....	506
11.3.2	Transmitter Holding Registers (DUART_UTHR _n).....	506
11.3.3	Divisor Least Significant Byte Registers (DUART_UDLB _n).....	507
11.3.4	Divisor Most Significant Byte Registers (DUART_UDMB _n).....	508
11.3.5	Interrupt Enable Register (DUART_UIER _n).....	509
11.3.6	Interrupt ID Registers (DUART_UIIR _n).....	510
11.3.7	FIFO Control Registers (DUART_UFCR _n).....	511
11.3.8	Alternate Function Registers (DUART_UAFR _n).....	513
11.3.9	Line Control Registers (DUART_ULCR _n).....	513
11.3.10	Modem Control Registers (DUART_UMCR _n).....	515
11.3.11	Line Status Registers (DUART_ULSR _n).....	516
11.3.12	Modem Status Registers (DUART_UMSR _n).....	517
11.3.13	Scratch Registers (DUART_USCR _n).....	518
11.3.14	DMA Status Registers (DUART_UDSR _n).....	518
11.4	Functional description.....	520
11.4.1	Serial interface.....	520
11.4.1.1	START bit.....	521
11.4.1.2	Data transfer.....	521
11.4.1.3	Parity bit.....	522
11.4.1.4	STOP bit.....	522
11.4.2	Baud-rate generator logic.....	522

Section number	Title	Page
11.4.3	Local loopback mode.....	523
11.4.4	Errors.....	523
11.4.4.1	Framing error.....	523
11.4.4.2	Parity error.....	524
11.4.4.3	Overrun error.....	524
11.4.5	FIFO mode.....	524
11.4.5.1	FIFO interrupts.....	524
11.4.5.2	DMA mode select.....	525
11.4.5.3	Interrupt control logic.....	525
11.5	DUART initialization/application information.....	526

Chapter 12 Enhanced local bus controller (eLBC)

12.1	eLBC introduction.....	527
12.1.1	Overview.....	528
12.1.2	Features.....	529
12.1.3	Modes of operation.....	530
12.1.3.1	eLBC bus clock and clock ratios.....	530
12.1.3.2	Source ID debug mode.....	531
12.2	eLBC external signal descriptions.....	531
12.3	Enhanced Local Bus Controller (eLBC) Memory Map.....	534
12.3.1	Base register 0 (eLBC_BR0).....	538
12.3.2	Options register 0 layout for GPCM Mode (eLBC_ORg0).....	540
12.3.3	Options register 0 layout for FCM Mode (eLBC_ORf0).....	543
12.3.4	Options register 0 layout for UPM Mode (eLBC_ORu0).....	547
12.3.5	Base register n (eLBC_BRn).....	551
12.3.6	Options register n layout for GPCM Mode (eLBC_ORgn).....	552
12.3.7	Options register n layout for FCM Mode (eLBC_ORfn).....	556
12.3.8	Options register n layout for UPM Mode (eLBC_ORun).....	560
12.3.9	UPM address register (eLBC_MAR).....	562

Section number	Title	Page
12.3.10	UPMn mode register (eLBC_MnMR).....	563
12.3.11	Memory refresh timer prescaler register (eLBC_MRTPR).....	566
12.3.12	UPM data register (eLBC_MDRu).....	566
12.3.13	FCM data register (eLBC_MDRf).....	567
12.3.14	Special operation initiation register (eLBC_LSOR).....	567
12.3.15	UPM refresh timer (eLBC_LURT).....	568
12.3.16	Transfer error status register (eLBC_LTESR).....	569
12.3.17	Transfer error disable register (eLBC_LTEDR).....	571
12.3.18	Transfer error interrupt register (eLBC_LTEIR).....	572
12.3.19	Transfer error attributes register (eLBC_LTEATR).....	573
12.3.20	Transfer error address register (eLBC_LTEAR).....	574
12.3.21	Transfer error ECC register (eLBC_LTECCR).....	575
12.3.22	Configuration register (eLBC_LBCR).....	576
12.3.23	Clock ratio register (eLBC_LCRR).....	577
12.3.24	Flash mode register (eLBC_FMR).....	579
12.3.25	Flash instruction register (eLBC_FIR).....	581
12.3.26	Flash command register (eLBC_FCR).....	582
12.3.27	Flash block address register (eLBC_FBAR).....	583
12.3.28	Flash page address register [Large Page Device (ORx[PGS] = 1)] (eLBC_FPARI).....	583
12.3.29	Flash page address register [Small Page Device (ORx[PGS] = 0)] (eLBC_FPARs).....	584
12.3.30	Flash byte count register (eLBC_FBCR).....	585
12.3.31	Flash ECC block n registers (eLBC_FECCn).....	586
12.4	eLBC functional description.....	586
12.4.1	Basic architecture.....	588
12.4.1.1	Address and address space checking.....	588
12.4.1.2	External address latch enable signal (LALE).....	588
12.4.1.3	Data transfer acknowledge (TA).....	590
12.4.1.4	Data buffer control (LBCTL).....	591
12.4.1.5	Parity generation and checking (LDP).....	592

Section number	Title	Page
12.4.1.6	Bus monitor.....	592
12.4.1.7	PLL Bypass mode.....	593
12.4.2	General-purpose chip-select machine (GPCM).....	594
12.4.2.1	GPCM read signal timing.....	596
12.4.2.2	GPCM write signal timing.....	598
12.4.2.3	Chip-select assertion timing.....	600
12.4.2.3.1	Programmable wait state configuration.....	601
12.4.2.3.2	Chip-select and write enable negation timing.....	601
12.4.2.3.3	Relaxed timing.....	603
12.4.2.3.4	Output enable (LOE_B) timing.....	609
12.4.2.3.5	Extended hold time on read accesses-GPCM.....	609
12.4.2.4	External access termination (LGTA_B).....	611
12.4.2.5	GPCM boot chip-select operation.....	612
12.4.3	Flash control machine (FCM).....	613
12.4.3.1	FCM buffer RAM	616
12.4.3.1.1	Buffer layout and page mapping for small-page NAND flash devices	617
12.4.3.1.2	Buffer layout and page mapping for large-page NAND flash devices	618
12.4.3.1.3	Error correcting codes and the spare region	619
12.4.3.2	Programming FCM.....	621
12.4.3.2.1	FCM command instructions	622
12.4.3.2.2	FCM no-operation instruction	623
12.4.3.2.3	FCM address instructions	623
12.4.3.2.4	FCM data read instructions	624
12.4.3.2.5	FCM data write instructions	624
12.4.3.3	FCM signal timing.....	625
12.4.3.3.1	FCM chip-select timing	625
12.4.3.3.2	FCM command, address, and write data timing	625
12.4.3.3.3	FCM ready/busy timing	628
12.4.3.3.4	FCM read data timing	629

Section number	Title	Page
	12.4.3.3.5 FCM extended read hold timing	631
12.4.3.4	FCM boot chip-select operation	632
	12.4.3.4.1 FCM bank 0 reset initialization	632
	12.4.3.4.2 Boot block loading into the FCM buffer RAM	633
12.4.4	User-programmable machines (UPMs).....	635
12.4.4.1	UPM requests.....	636
	12.4.4.1.1 Memory access requests.....	637
	12.4.4.1.2 UPM refresh timer requests.....	638
	12.4.4.1.3 Software requests-RUN command.....	638
	12.4.4.1.4 Exception requests.....	639
12.4.4.2	Programming the UPMs.....	639
	12.4.4.2.1 UPM programming example (two sequential writes to the RAM array).....	640
	12.4.4.2.2 UPM programming example (two sequential reads from the RAM array)....	641
12.4.4.3	UPM signal timing.....	642
12.4.4.4	RAM array.....	643
	12.4.4.4.1 RAM words.....	643
	12.4.4.4.2 Chip-select signal timing (CSTn).....	647
	12.4.4.4.3 Byte select signal timing (BSTn).....	648
	12.4.4.4.4 General-purpose signals (GnTn, GOn).....	649
	12.4.4.4.5 Loop control (LOOP).....	649
	12.4.4.4.6 Repeat execution of current RAM word (REDO).....	650
	12.4.4.4.7 Address multiplexing (AMX).....	650
	12.4.4.4.8 Data valid and data sample control (UTA).....	652
	12.4.4.4.9 LGPL[0:5] signal negation (LAST).....	653
	12.4.4.4.10 Wait mechanism (WAEN).....	653
12.4.4.5	Synchronous sampling of LUPWAIT for early transfer acknowledge.....	655
12.4.4.6	Extended hold time on read accesses-UPM.....	656

Section number	Title	Page
12.5	eLBC initialization/application information.....	656
12.5.1	Interfacing to peripherals in different address modes.....	656
12.5.1.1	Multiplexed address/data bus for 32-bit addressing.....	656
12.5.1.2	Non-multiplexed address and data buses.....	657
12.5.1.3	Peripheral hierarchy on the local bus for high bus speeds.....	658
12.5.1.4	GPCM timings.....	659
12.5.2	Bus turnaround.....	660
12.5.2.1	Address phase after previous read.....	661
12.5.2.2	Read data phase after address phase.....	661
12.5.2.3	Read-modify-write cycle for parity protected memory banks.....	661
12.5.2.4	UPM cycles with additional address phases.....	662
12.5.3	Interface to different port-size devices.....	662
12.5.4	Command sequence examples for NAND flash EEPROM.....	663
12.5.4.1	NAND flash soft reset command sequence example.....	663
12.5.4.2	NAND flash read status command sequence example.....	664
12.5.4.3	NAND flash read identification command sequence example.....	665
12.5.4.4	NAND flash page read command sequence example.....	665
12.5.4.5	NAND flash block erase command sequence example.....	666
12.5.4.6	NAND flash program command sequence example.....	667
12.5.5	Interfacing to fast-page mode DRAM using UPM.....	668
12.5.6	Interfacing to ZBT SRAM using UPM.....	679

Chapter 13 DMA Controller

13.1	DMA overview.....	681
13.1.1	DMA features summary.....	682
13.1.2	DMA modes of operation.....	682
13.2	DMA external signal description.....	685
13.2.1	Signal overview.....	685
13.2.2	DMA signal descriptions.....	686

Section number	Title	Page
13.3	DMA controller memory map.....	687
13.3.1	DMA n mode register (DMAx_MRn).....	693
13.3.2	DMA n status register (DMAx_SRn).....	697
13.3.3	DMA n current link descriptor extended address register (DMAx_ECLNDARn).....	698
13.3.4	DMA n current link descriptor address register (DMAx_CLNDARn).....	699
13.3.5	DMA n source attributes register (DMAx_SATRn).....	701
13.3.6	DMA n source address register (DMAx_SARn).....	702
13.3.7	DMA n destination attributes register (DMAx_DATRn).....	703
13.3.8	DMA n destination address register (DMAx_DARn).....	704
13.3.9	DMA n byte count register (DMAx_BCRn).....	705
13.3.10	DMA n extended next link descriptor address register (DMAx_ENLNDARn).....	705
13.3.11	DMA n next link descriptor address register (DMAx_NLNDARn).....	706
13.3.12	DMA n extended current list descriptor address register (DMAx_ECLSDARn).....	707
13.3.13	DMA n current list descriptor address register (DMAx_CLSDARn).....	708
13.3.14	DMA n extended next list descriptor address register (DMAx_ENLSDARn).....	709
13.3.15	DMA n next list descriptor address register (DMAx_NLSDARn).....	709
13.3.16	DMA n source stride register (DMAx_SSRn).....	710
13.3.17	DMA n destination stride register (DMAx_DSRn).....	711
13.3.18	DMA general status register (DMAx_DGSR).....	712
13.4	DMA functional description.....	715
13.4.1	DMA channel operation.....	715
13.4.1.1	Source/destination transaction size calculations.....	716
13.4.1.2	Basic DMA mode transfer.....	718
13.4.1.2.1	Basic direct mode.....	718
13.4.1.2.2	Basic, direct, single-write start mode.....	718
13.4.1.2.3	Basic chaining mode.....	719
13.4.1.2.4	Basic chaining, single-write start mode.....	720
13.4.1.3	Extended DMA mode transfer.....	721
13.4.1.3.1	Extended direct mode	721

Section number	Title	Page
13.4.1.3.2	Extended direct, single-write start mode.....	721
13.4.1.3.3	Extended chaining mode.....	721
13.4.1.3.4	Extended chaining, single-write start mode.....	722
13.4.1.4	External control mode transfer.....	723
13.4.1.5	Channel continue mode for cascading transfer chains.....	724
13.4.1.5.1	Basic mode.....	725
13.4.1.5.2	Extended mode.....	725
13.4.1.6	Channel abort.....	725
13.4.1.7	Bandwidth control.....	725
13.4.1.8	Channel state.....	726
13.4.1.9	Illustration of stride size and stride distance.....	726
13.4.2	DMA transfer interfaces.....	727
13.4.3	DMA errors.....	727
13.4.4	DMA descriptors.....	727
13.4.5	DMA controller limitations and restrictions.....	731
13.5	DMA system considerations.....	732
13.5.1	Unusual DMA scenarios.....	733
13.5.1.1	DMA to core.....	733
13.5.1.2	DMA to configuration, control, and status registers (CCSR).....	733
13.5.1.3	DMA to I2C.....	734
13.5.1.4	DMA to DUART.....	734

Chapter 14

Enhanced three-speed Ethernet controllers (eTSEC)

14.1	eTSEC overview.....	735
14.2	eTSEC features.....	736
14.3	eTSEC modes of operation.....	738
14.4	eTSEC external signals description.....	740
14.4.1	eTSEC detailed signal descriptions.....	743

Section number	Title	Page
14.5	eTSEC memory map/register definition.....	748
14.5.1	eTSEC top-level module memory map.....	748
14.5.2	eTSEC Memory Map/Register Definition.....	749
14.5.2.1	Controller ID register (ETSECx_TSEC_ID).....	776
14.5.2.2	Controller ID register (ETSECx_TSEC_ID2).....	776
14.5.2.3	Interrupt event register (ETSECx_IEVENT).....	777
14.5.2.4	Interrupt mask register (ETSECx_IMASK).....	783
14.5.2.5	Error disabled register (ETSECx_EDIS).....	785
14.5.2.6	Ethernet control register (ETSECx_ECNTL).....	787
14.5.2.7	Pause time value register (ETSECx_PTV).....	791
14.5.2.8	DMA control register (ETSECx_DMACTRL).....	791
14.5.2.9	TBI PHY address register (ETSECx_TBIPA).....	793
14.5.2.10	Transmit control register (ETSECx_TCTRL).....	795
14.5.2.11	Transmit status register (ETSECx_TSTAT).....	798
14.5.2.12	Default VLAN control word (ETSECx_DFVLAN).....	802
14.5.2.13	Transmit interrupt coalescing register (ETSECx_TXIC).....	803
14.5.2.14	Transmit queue control register (ETSECx_TQUEUE).....	804
14.5.2.15	TxBD Rings 0-3 round-robin weightings (ETSECx_TR03WT).....	805
14.5.2.16	TxBD Rings 4-7 round-robin weightings (ETSECx_TR47WT).....	806
14.5.2.17	Tx data buffer pointer high bits (ETSECx_TBDBPH).....	807
14.5.2.18	TxBD pointer for ring n (ETSECx_TBPTRn).....	807
14.5.2.19	TxBD base address high bits (ETSECx_TBASEH).....	808
14.5.2.20	TxBD base address of ring n (ETSECx_TBASEn).....	809
14.5.2.21	Tx time stamp identification tag (set n) register (ETSECx_TMR_TXTSn_ID).....	809
14.5.2.22	Tx time stamp high (ETSECx_TMR_TXTSn_H).....	810
14.5.2.23	Tx time stamp low (ETSECx_TMR_TXTSn_L).....	810
14.5.2.24	Receive control register (ETSECx_RCTRL).....	811
14.5.2.25	Receive status register (ETSECx_RSTAT).....	814
14.5.2.26	Receive interrupt coalescing register (ETSECx_RXIC).....	816

Section number	Title	Page
14.52.27	Receive queue control register. (ETSECx_RQUEUE).....	817
14.52.28	Receive bit field extract control register (ETSECx_RBIFX).....	819
14.52.29	Receive queue filing table address register (ETSECx_RQFAR).....	821
14.52.30	Receive queue filing table control register (ETSECx_RQFCR).....	822
14.52.31	Receive queue filing table property register (ETSECx_RQFPR).....	824
14.52.32	Maximum receive buffer length register (ETSECx_MRBLR).....	830
14.52.33	Rx data buffer pointer high bits (ETSECx_RBDBPH).....	831
14.52.34	RxBD pointer for ring n (ETSECx_RBPTRn).....	831
14.52.35	RxBD base address high bits (ETSECx_RBASEH).....	832
14.52.36	RxBD base address of ring n (ETSECx_RBASEn).....	833
14.52.37	Rx timer time stamp register high (ETSECx_TMR_RXTS_H).....	833
14.52.38	Rx timer time stamp register low (ETSECx_TMR_RXTS_L).....	834
14.52.39	MAC configuration register 1 (ETSECx_MACCFG1).....	834
14.52.40	MAC configuration register 2 (ETSECx_MACCFG2).....	837
14.52.41	Inter-packet/inter-frame gap register (ETSECx_IPGIFG).....	839
14.52.42	Half-duplex control (ETSECx_HAFDUP).....	840
14.52.43	Maximum frame length (ETSECx_MAXFRM).....	841
14.52.44	MII management configuration (ETSECx_MIIMCFG).....	842
14.52.45	MII management command (ETSECx_MIIMCOM).....	844
14.52.46	MII management address (ETSECx_MIIMADD).....	845
14.52.47	MII management control (ETSECx_MIIMCON).....	845
14.52.48	MII management status (ETSECx_MIIMSTAT).....	846
14.52.49	MII management indicator (ETSECx_MIIMIND).....	846
14.52.50	Interface status (ETSECx_IFSTAT).....	847
14.52.51	MAC station address register 1 (ETSECx_MACSTNADDR1).....	849
14.52.52	MAC station address register 2 (ETSECx_MACSTNADDR2).....	849
14.52.53	MAC exact match address n, part 1 (ETSECx_MACnADDR1).....	850
14.52.54	MAC exact match address n, part 2 (ETSECx_MACnADDR2).....	851
14.52.55	Transmit and receive 64-byte frame counter (ETSECx_TR64).....	851

Section number	Title	Page
14.52.56	Transmit and receive 65- to 127-byte frame counter (ETSECx_TR127).....	852
14.52.57	Transmit and receive 128- to 255-byte frame counter (ETSECx_TR255).....	852
14.52.58	Transmit and receive 256- to 511-byte frame counter (ETSECx_TR511).....	853
14.52.59	Transmit and receive 512- to 1023-byte frame counter (ETSECx_TR1K).....	853
14.52.60	Transmit and receive 1024- to 1518-byte frame counter (ETSECx_TRMAX).....	853
14.52.61	Transmit and receive 1519- to 1522-byte good VLAN frame count (ETSECx_TRMGV)...	854
14.52.62	Receive byte counter (ETSECx_RBYT).....	854
14.52.63	Receive packet counter (ETSECx_RPKT).....	855
14.52.64	Receive FCS error counter (ETSECx_RFCFS).....	855
14.52.65	Receive multicast packet counter (ETSECx_RMCA).....	855
14.52.66	Receive broadcast packet counter (ETSECx_RBCA).....	856
14.52.67	Receive control frame packet counter (ETSECx_RXCF).....	856
14.52.68	Receive PAUSE frame packet counter (ETSECx_RXPF).....	857
14.52.69	Receive unknown OP code counter (ETSECx_RXUO).....	857
14.52.70	Receive alignment error counter (ETSECx_RALN).....	857
14.52.71	Receive frame length error counter (ETSECx_RFLR).....	858
14.52.72	Receive code error counter (ETSECx_RCDE).....	858
14.52.73	Receive carrier sense error counter (ETSECx_RCSE).....	859
14.52.74	Receive undersize packet counter (ETSECx_RUND).....	859
14.52.75	Receive oversize packet counter (ETSECx_ROVR).....	859
14.52.76	Receive fragments counter (ETSECx_RFRG).....	860
14.52.77	Receive jabber counter (ETSECx_RJBR).....	860
14.52.78	Receive drop counter (ETSECx_RDRP).....	861
14.52.79	Transmit byte counter (ETSECx_TBYT).....	861
14.52.80	Transmit packet counter (ETSECx_TPKT).....	861
14.52.81	Transmit multicast packet counter (ETSECx_TMCA).....	862
14.52.82	Transmit broadcast packet counter (ETSECx_TBCA).....	862
14.52.83	Transmit PAUSE control frame counter (ETSECx_TXPF).....	863
14.52.84	Transmit deferral packet counter (ETSECx_TDFR).....	863

Section number	Title	Page
14.52.85	Transmit excessive deferral packet counter (ETSECx_TEDF).....	863
14.52.86	Transmit single collision packet counter (ETSECx_TSCL).....	864
14.52.87	Transmit multiple collision packet counter (ETSECx_TMCL).....	864
14.52.88	Transmit late collision packet counter (ETSECx_TLCL).....	865
14.52.89	Transmit excessive collision packet counter (ETSECx_TXCL).....	865
14.52.90	Transmit total collision counter (ETSECx_TNCL).....	865
14.52.91	Transmit drop frame counter (ETSECx_TDRP).....	866
14.52.92	Transmit jabber frame counter (ETSECx_TJBR).....	866
14.52.93	Transmit FCS error counter (ETSECx_TFCS).....	867
14.52.94	Transmit control frame counter (ETSECx_TXCF).....	867
14.52.95	Transmit oversize frame counter (ETSECx_TOVR).....	867
14.52.96	Transmit undersize frame counter (ETSECx_TUND).....	868
14.52.97	Transmit fragments frame counter (ETSECx_TFRG).....	868
14.52.98	Carry register one registerCleared on read. (ETSECx_CAR1).....	869
14.52.99	Carry register two register 3 (ETSECx_CAR2).....	871
14.52.100	Carry register one mask register (ETSECx_CAM1).....	873
14.52.101	Carry register two mask register (ETSECx_CAM2).....	875
14.52.102	Receive filer rejected packet counter (ETSECx_RREJ).....	876
14.52.103	Individual/group address register n (ETSECx_IGADDRn).....	877
14.52.104	Group address register n (ETSECx_GADDRn).....	878
14.52.105	Attribute register (ETSECx_ATTR).....	879
14.52.106	Attribute extract length and extract index register (ETSECx_ATTRELI).....	881
14.52.107	Receive Queue Parameters register n (ETSECx_RQPRMn).....	882
14.52.108	Last Free RxBD pointer for ring n (ETSECx_RFBPTRn).....	883
14.5.3	ETSEC_TBI Memory Map/Register Definition.....	883
14.53.1	Control (ETSEC_TBI_CR).....	885
14.53.2	Status (ETSEC_TBI_SR).....	886
14.53.3	AN advertisement (ETSEC_TBI_ANA).....	887
14.53.4	AN link partner base page ability (ETSEC_TBI_ANLPBPA).....	889

Section number	Title	Page
14.53.5	AN expansion (ETSEC_TBI_ANEX).....	891
14.53.6	AN next page transmit (ETSEC_TBI_ANNPT).....	891
14.53.7	AN link partner ability next page (ETSEC_TBI_ANLPANP).....	892
14.53.8	Extended status (ETSEC_TBI_EXST).....	893
14.53.9	Jitter diagnostics (ETSEC_TBI_JD).....	894
14.53.10	TBI control (ETSEC_TBI_TBICON).....	896
14.5.4	ETSEC_1588 Memory Map/Register Definition.....	897
14.54.1	Timer control register (ETSEC_1588_TMR_CTRL).....	898
14.54.2	Time stamp event register (ETSEC_1588_TMR_TEVENT).....	901
14.54.3	Timer event mask register (ETSEC_1588_TMR_TEMASK).....	903
14.54.4	Time stamp event register (ETSEC_1588_TMR_PEVENT).....	904
14.54.5	Timer event mask register (ETSEC_1588_TMR_PEMASK).....	905
14.54.6	Time stamp status register (ETSEC_1588_TMR_STAT).....	906
14.54.7	Timer counter high register (ETSEC_1588_TMR_CNT_H).....	906
14.54.8	Timer counter low register (ETSEC_1588_TMR_CNT_L).....	907
14.54.9	Timer drift compensation addend register (ETSEC_1588_TMR_ADD).....	908
14.54.10	Timer accumulator register (ETSEC_1588_TMR_ACC).....	909
14.54.11	Timer prescale (ETSEC_1588_TMR_PRSC).....	909
14.54.12	Timer offset high (ETSEC_1588_TMROFF_H).....	910
14.54.13	Timer offset low (ETSEC_1588_TMROFF_L).....	910
14.54.14	Timer alarm n high register (ETSEC_1588_TMR_ALARMn_H).....	911
14.54.15	Timer alarm n low register (ETSEC_1588_TMR_ALARMn_L).....	911
14.54.16	Timer fixed period interval (ETSEC_1588_TMR_FIPERn).....	912
14.54.17	Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTSn_H).....	913
14.54.18	Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTSn_L).....	914
14.6	Functional description.....	914
14.6.1	Programming Model Considerations.....	914
14.6.1.1	MAC Functionality.....	915
14.6.1.1.1	Configuring the MAC	915

Section number	Title	Page
14.6.1.1.2	Controlling CSMA/CD	915
14.6.1.1.3	Handling Packet Collisions	916
14.6.1.1.4	Controlling Packet Flow	917
14.6.1.1.5	Controlling PHY Links	918
14.6.1.2	MIB Registers	918
14.6.1.3	Hash Function Registers	919
14.6.1.4	Lossless Flow Control Configuration Registers.....	919
14.6.1.5	IEEE 1588-Compatible Timestamping Registers.....	920
14.6.1.6	Ten-Bit Interface (TBI).....	920
14.6.1.6.1	TBI Transmit Process.....	920
14.6.1.6.1.1	Packet Encapsulation.....	920
14.6.1.6.1.2	8B10B Encoding.....	921
14.6.1.6.1.3	Preamble Shortening.....	921
14.6.1.6.2	TBI Receive Process.....	921
14.6.1.6.2.1	Synchronization.....	921
14.6.1.6.2.2	Auto-Negotiation for 1000BASE-X.....	921
14.6.2	Connecting to physical interfaces on Ethernet.....	922
14.6.2.1	Media-independent interface (MII).....	922
14.6.2.2	Reduced media-independent interface (RMII).....	923
14.6.2.3	Gigabit media-independent interface (GMII).....	924
14.6.2.4	Reduced gigabit media-independent interface (RGMII).....	925
14.6.2.5	Ten-bit interface (TBI).....	926
14.6.2.6	Reduced ten-bit interface (RTBI).....	927
14.6.2.7	Serial gigabit media-independent interface (SGMII).....	928
14.6.2.8	Ethernet physical interfaces signal summary.....	929
14.6.3	Gigabit Ethernet controller channel operation.....	933
14.6.3.1	Initialization sequence.....	933
14.6.3.1.1	Hardware controlled initialization.....	933
14.6.3.1.2	User initialization.....	933

Section number	Title	Page
14.6.3.2	Soft reset and reconfiguring procedure.....	934
14.6.3.3	Gigabit Ethernet frame transmission.....	935
14.6.3.4	Gigabit Ethernet frame reception.....	937
14.6.3.5	Ethernet preamble customization.....	939
14.6.3.5.1	User-defined preamble transmission.....	939
14.6.3.5.2	User-visible preamble reception.....	940
14.6.3.6	RMON support.....	941
14.6.3.7	Frame recognition.....	941
14.6.3.7.1	Destination address recognition and frame filtering.....	941
14.6.3.7.2	Hash table algorithm.....	943
14.6.3.8	Magic Packet mode.....	945
14.6.3.9	Flow control.....	945
14.6.3.10	Interrupt handling.....	946
14.6.3.10.1	Interrupt coalescing.....	947
14.6.3.10.2	Interrupt coalescing by frame count threshold.....	948
14.6.3.10.3	Interrupt coalescing by timer threshold.....	948
14.6.3.11	Inter-frame gap time.....	949
14.6.3.12	Internal and external loop back.....	950
14.6.3.13	Error-handling procedure.....	950
14.6.4	TCP/IP off-load.....	952
14.6.4.1	Frame control blocks.....	953
14.6.4.2	Transmit path off-load and Tx PTP packet parsing.....	954
14.6.4.3	Receive path off-load.....	956
14.6.5	Quality of service (QoS) provision.....	958
14.6.5.1	Receive parser.....	958
14.6.5.2	Receive queue filter.....	960
14.6.5.2.1	Filtering rules	961
14.6.5.2.2	Comparing properties with bit masks.....	963
14.6.5.2.3	Special-case rules.....	963

Section number	Title	Page
14.6.5.2.4	Filer interrupt events.....	964
14.6.5.2.5	Setting up the receive queue filer table.....	965
14.6.5.2.6	Filer example-802.1p priority filing.....	965
14.6.5.2.7	Filer example-IP diff-serv code points filing.....	966
14.6.5.2.8	Filer example-TCP and UDP port filing	967
14.6.5.2.9	Filer example-interrupt from deep sleep.....	968
14.6.5.3	Transmission scheduling.....	970
14.6.5.3.1	Priority-based queuing (PBQ).....	971
14.6.5.3.2	Modified weighted round-robin queuing (MWRR).....	971
14.6.6	Lossless flow control.....	972
14.6.6.1	Back pressure determination through free buffers.....	972
14.6.6.2	Software use of hardware-initiated back pressure.....	975
14.6.6.2.1	Initialization	975
14.6.6.2.2	Operation.....	975
14.6.7	Hardware assist for IEEE 1588-compatible timestamping.....	976
14.6.7.1	Features.....	976
14.6.7.2	Timer logic overview.....	977
14.6.7.3	Time-stamp insertion on the received packets.....	979
14.6.7.3.1	Timestamp point.....	979
14.6.7.4	PTP packet parsing.....	979
14.6.7.4.1	General purpose filer rule.....	980
14.6.7.5	Time-stamp insertion on transmit packets.....	981
14.6.7.5.1	Interrupts.....	981
14.6.7.5.2	Error condition.....	982
14.6.7.6	Tx PTP packet parsing.....	983
14.6.8	Buffer descriptors.....	985
14.6.8.1	Data buffer descriptors.....	985
14.6.8.2	Transmit data buffer descriptors (TxBD).....	987
14.6.8.3	Receive buffer descriptors (RxBD).....	990

Section number	Title	Page
14.7	Initialization/application information.....	992
14.7.1	Interface mode configuration.....	992
14.7.1.1	MII interface mode.....	993
14.7.1.2	GMII interface mode.....	996
14.7.1.3	TBI interface mode	1000
14.7.1.4	RGMII interface mode.....	1005
14.7.1.5	RMII interface mode.....	1009
14.7.1.6	RTBI interface mode.....	1014
14.7.1.7	SGMII interface support.....	1018

Chapter 15 serial RapidIO

15.1	Serial RapidIO overview.....	1023
15.2	Serial RapidIO features summary.....	1024
15.3	Serial RapidIO modes of operation.....	1025
15.3.1	RapidIO port.....	1025
15.3.2	Message unit.....	1026
15.4	LP-serial signal descriptions.....	1026
15.4.1	Serial Rapid I/O overview.....	1026
15.4.2	Serial Rapid I/O detailed signal descriptions.....	1026
15.4.2.1	SD_TXn/SD_TX n_B-outputs.....	1026
15.4.2.2	SD_RXn/SD_RX n_B-inputs.....	1027
15.5	Serial RapidIO Memory Map/Register Definition.....	1027
15.5.1	Device identity capability register (SRIO_DIDCAR).....	1040
15.5.2	Device information capability register (SRIO_DICAR).....	1041
15.5.3	Assembly identity capability register (SRIO_AIDCAR).....	1041
15.5.4	Assembly information capability register (SRIO_AICAR).....	1042
15.5.5	Processing element features capability register (SRIO_PEF CAR).....	1043
15.5.6	Source operations capability register (SRIO_SOCAR).....	1045
15.5.7	Destination operations capability register (SRIO_DOCAR).....	1048

Section number	Title	Page
15.5.8	Mailbox command and status register (SRIO_MCSR).....	1051
15.5.9	Port -Write and doorbell command and status register (SRIO_PWDCSR).....	1053
15.5.10	Processing element logical layer control command and status register (SRIO_PELLCSR).....	1055
15.5.11	Local configuration space base address 1 command and status register (SRIO_LCSBA1CSR).....	1056
15.5.12	Base device ID command and status register (SRIO_BDIDCSR).....	1057
15.5.13	Host base device ID lock command and status register (SRIO_HBDIDLCSR).....	1058
15.5.14	Component tag command and status register (SRIO_CTCSR).....	1058
15.5.15	Port maintenance block header 0 (SRIO_PMBH0).....	1059
15.5.16	Port link time-out control command and status register (SRIO_PLTOCCSR).....	1059
15.5.17	Port response time-out control command and status register (SRIO_PRTOCCSR).....	1060
15.5.18	Port General control command and status register (SRIO_PGCCSR).....	1061
15.5.19	Port 1 Link maintenance request command and status register (SRIO_P1LMREQCSR).....	1062
15.5.20	Port 1 Link maintenance response command and status register (SRIO_P1LMRESPCSR).....	1062
15.5.21	Port 1 Local ackID status command and status register (SRIO_P1LASCsr).....	1063
15.5.22	Port 1 Error and status command and status register (SRIO_P1ESCSR).....	1065
15.5.23	Port 1 Control command and status register (SRIO_P1CCSR).....	1068
15.5.24	Port 2 Link maintenance request command and status register (SRIO_P2LMREQCSR).....	1070
15.5.25	Port 2 Link maintenance response command and status register (SRIO_P2LMRESPCSR).....	1070
15.5.26	Port 2 Local ackID status command and status register (SRIO_P2LASCsr).....	1071
15.5.27	Port 2 Error and status command and status register (SRIO_P2ESCSR).....	1073
15.5.28	Port 2 Control command and status register (SRIO_P2CCSR).....	1076
15.5.29	Error reporting block header (SRIO_ERBH).....	1078
15.5.30	Logical/Transport layer error detect command and status register (SRIO_LTLEDCSR).....	1078
15.5.31	Logical/Transport layer error enable command and status register (SRIO_LTLEECSR).....	1081
15.5.32	Logical/Transport layer address capture command and status register (SRIO_LTLACCSR).....	1082
15.5.33	Logical/Transport layer device ID capture command and status register (SRIO_LTLIDCCSR).....	1083
15.5.34	Logical/Transport layer control capture command and status register (SRIO_LTLCCCSR).....	1084
15.5.35	Port 1 Error detect command and status register (SRIO_P1EDCSR).....	1085
15.5.36	Port 1 Error rate enable command and status register (SRIO_P1ERECSR).....	1086

Section number	Title	Page
15.5.37	Port 1 Error capture attributes command and status register (SRIO_P1ECACSR).....	1088
15.5.38	Port 1 Packet/control symbol error capture command and status register 0 (SRIO_P1PCSECCSR0).....	1089
15.5.39	Port 1 Packet error capture command and status register 1 (SRIO_P1PECCSR1).....	1090
15.5.40	Port 1 Packet error capture command and status register 2 (SRIO_P1PECCSR2).....	1090
15.5.41	Port 1 Packet error capture command and status register 3 (SRIO_P1PECCSR3).....	1091
15.5.42	Port 1 Error rate command and status register (SRIO_P1ERCSR).....	1091
15.5.43	Port 1 Error rate threshold command and status register (SRIO_P1ERTCSR).....	1092
15.5.44	Port 2 Error detect command and status register (SRIO_P2EDCSR).....	1093
15.5.45	Port 2 Error rate enable command and status register (SRIO_P2ERECSR).....	1095
15.5.46	Port 2 Error capture attributes command and status register (SRIO_P2ECACSR).....	1096
15.5.47	Port 2 Packet/control symbol error capture command and status register 0 (SRIO_P2PCSECCSR0).....	1097
15.5.48	Port 2 Packet error capture command and status register 1 (SRIO_P2PECCSR1).....	1098
15.5.49	Port 2 Packet error capture command and status register 2 (SRIO_P2PECCSR2).....	1098
15.5.50	Port 2 Packet error capture command and status register 3 (SRIO_P2PECCSR3).....	1099
15.5.51	Port 2 Error rate command and status register (SRIO_P2ERCSR).....	1099
15.5.52	Port 2 Error rate threshold command and status register (SRIO_P2ERTCSR).....	1100
15.5.53	Logical layer configuration register (SRIO_LLCR).....	1101
15.5.54	Error / port-write interrupt status register (SRIO_EPWISR).....	1103
15.5.55	Logical retry error threshold configuration register (SRIO_LRETCR).....	1104
15.5.56	Physical retry error threshold configuration register (SRIO_PRETCR).....	1105
15.5.57	Port 1 Alternate device ID command and status register (SRIO_P1ADIDCSR).....	1105
15.5.58	Port 1 accept-all configuration register (SRIO_P1AACR).....	1106
15.5.59	Port 1 Logical Outbound Packet time-to-live configuration register (SRIO_P1LOPTTLCR).....	1107
15.5.60	Port 1 Implementation error command and status register (SRIO_P1IECSR).....	1108
15.5.61	Port 1 Physical configuration register (SRIO_P1PCR).....	1109
15.5.62	Port 1 Serial link command and status register (SRIO_P1SLCSR).....	1110
15.5.63	Port 1 Serial link error injection configuration register (SRIO_P1SLEICR).....	1111
15.5.64	Port 2 Alternate device ID command and status register (SRIO_P2ADIDCSR).....	1112
15.5.65	Port 2 accept-all configuration register (SRIO_P2AACR).....	1113

Section number	Title	Page
15.5.66	Port 2 Logical Outbound Packet time-to-live configuration register (SRIO_P2LOPTTLCR).....	1113
15.5.67	Port 2 Implementation error command and status register (SRIO_P2IECSR).....	1114
15.5.68	Port 2 Physical configuration register (SRIO_P2PCR).....	1115
15.5.69	Port 2 Serial link command and status register (SRIO_P2SLCSR).....	1116
15.5.70	Port 2 Serial link error injection configuration register (SRIO_P2SLEICR).....	1117
15.5.71	IP Block Revision Register 1 (SRIO_IPBRR1).....	1118
15.5.72	IP Block Revision Register 2 (SRIO_IPBRR2).....	1119
15.5.73	Port 1 RapidIO outbound window translation address register n (SRIO_PIROWTAR0).....	1119
15.5.74	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_PIROWTEARn).....	1121
15.5.75	Port 1 RapidIO outbound window attributes register n (SRIO_PIROWARn).....	1122
15.5.76	Port 1 RapidIO outbound window translation address register n (SRIO_PIROWTARn).....	1125
15.5.77	Port 1 RapidIO outbound window base address register n (SRIO_PIROWBARn).....	1126
15.5.78	Port 1 RapidIO outbound window segment 1 register n (SRIO_PIROWS1Rn).....	1127
15.5.79	Port 1 RapidIO outbound window segment 2 register n (SRIO_PIROWS2Rn).....	1129
15.5.80	Port 1 RapidIO outbound window segment 3 register n (SRIO_PIROWS3Rn).....	1131
15.5.81	Port 1 RapidIO Inbound window translation address register n (SRIO_P1RIWTARn).....	1133
15.5.82	Port 1 RapidIO Inbound window base address register n (SRIO_P1RIWBARn).....	1134
15.5.83	Port 1 RapidIO inbound window attributes register n (SRIO_P1RIWARn).....	1135
15.5.84	Port 1 RapidIO inbound window attributes register 0 (SRIO_P1RIWAR0).....	1137
15.5.85	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR0).....	1140
15.5.86	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEARn).....	1141
15.5.87	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWARn).....	1142
15.5.88	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTARn).....	1145
15.5.89	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBARn).....	1146
15.5.90	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1Rn).....	1147
15.5.91	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2Rn).....	1149
15.5.92	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3Rn).....	1151
15.5.93	Port 2 RapidIO Inbound window translation address register n (SRIO_P2RIWTARn).....	1153
15.5.94	Port 2 RapidIO Inbound window base address register n (SRIO_P2RIWBARn).....	1154

Section number	Title	Page
15.5.95	Port 2 RapidIO inbound window attributes register n (SRIO_P2RIWAR n).....	1155
15.5.96	Port 2 RapidIO inbound window attributes register 0 (SRIO_P2RIWAR0).....	1157
15.5.97	Outbound message n mode register (SRIO_OM n MR).....	1159
15.5.98	Outbound message n status register (SRIO_OM n SR).....	1163
15.5.99	Extended outbound message n descriptor queue dequeue pointer address register (SRIO_EOM n DQDPAR).....	1165
15.5.100	Outbound message n descriptor queue dequeue pointer address register (SRIO_OM n DQDPAR).....	1166
15.5.101	Extended outbound message n source address register (SRIO_EOM n SAR).....	1167
15.5.102	Outbound message n source address register (SRIO_OM n SAR).....	1167
15.5.103	Outbound message n destination port register (SRIO_OM n DPR).....	1168
15.5.104	Outbound message n destination attributes Register (SRIO_OM n DATR).....	1169
15.5.105	Outbound message n double-word count register (SRIO_OM n DCR).....	1170
15.5.106	Extended outbound message n descriptor queue enqueue pointer address register (SRIO_EOM n DQEPAR).....	1171
15.5.107	Outbound message n descriptor queue enqueue pointer address register (SRIO_OM n DQEPAR).....	1171
15.5.108	Outbound message n retry error threshold configuration register (SRIO_OM n RETCR).....	1173
15.5.109	Outbound message n multicast group register (SRIO_OM n MGR).....	1174
15.5.110	Outbound message n multicast list register (SRIO_OM n MLR).....	1175
15.5.111	Inbound message n mode register (SRIO_IM n MR).....	1176
15.5.112	Inbound message n status register (SRIO_IM n SR).....	1179
15.5.113	Extended inbound message n frame queue dequeue pointer address register (SRIO_EIM n FQDPAR)....	1181
15.5.114	Inbound message n frame queue dequeue pointer address register (SRIO_IM n FQDPAR).....	1181
15.5.115	Extended inbound message n frame queue enqueue pointer address register (SRIO_EIM n FQEPAR)....	1182
15.5.116	Inbound message n frame queue enqueue pointer address register (SRIO_IM n FQEPAR).....	1183
15.5.117	Inbound message n maximum interrupt report interval register (SRIO_IM n MIRIR).....	1184
15.5.118	Outbound doorbell mode register (SRIO_ODMR).....	1184
15.5.119	Outbound doorbell status register (SRIO_ODSR).....	1186
15.5.120	Outbound doorbell n destination port register (SRIO_ODDPR).....	1188
15.5.121	Outbound doorbell n destination attributes register (SRIO_ODDATR).....	1189

Section number	Title	Page
15.5.122	Outbound doorbell n retry error threshold configuration register (SRIO_ODRETCR).....	1190
15.5.123	Inbound doorbell n mode register (SRIO_IDMR).....	1191
15.5.124	Inbound doorbell n status register (SRIO_IDSR).....	1194
15.5.125	Extended inbound doorbell n queue dequeue pointer address register (SRIO_EIDQDPAR).....	1195
15.5.126	Inbound doorbell n queue dequeue Pointer address register (SRIO_IDQDPAR).....	1196
15.5.127	Extended inbound doorbell n queue enqueue pointer address register (SRIO_EIDQEPAR).....	1197
15.5.128	Inbound doorbell n Queue enqueue pointer address register (SRIO_IDQEPAR).....	1197
15.5.129	Inbound doorbell n maximum interrupt report interval register (SRIO_IDMIRIR).....	1198
15.5.130	Inbound port-write n mode register (SRIO_IPWMR).....	1199
15.5.131	Inbound port-write n status register (SRIO_IPWSR).....	1200
15.5.132	Extended inbound port-write n queue base address register (SRIO_EIPWQBAR).....	1201
15.5.133	Inbound port-write n queue base address register (SRIO_IPWQBAR).....	1201
15.6	Serial RapidIO functional description.....	1202
15.6.1	Segmented Outbound Window Description.....	1202
15.6.2	RapidIO transaction summary.....	1204
15.6.3	RapidIO packet format summary.....	1206
15.6.4	RapidIO control symbol summary.....	1207
15.6.5	Accessing configuration registers using rapidIO packets.....	1209
15.6.5.1	Guidelines.....	1210
15.6.5.2	Outbound maintenance accesses.....	1210
15.6.6	RapidIO outbound ATMU.....	1211
15.6.6.1	Valid hits to multiple ATMU windows.....	1211
15.6.6.2	Outbound window boundary crossing errors.....	1212
15.6.7	RapidIO inbound ATMU.....	1213
15.6.7.1	Hits to multiple ATMU windows.....	1213
15.6.7.2	Inbound window boundary crossing errors.....	1214
15.6.8	Generating link-request/reset-device.....	1214
15.6.9	Outbound drain mode.....	1215
15.6.10	Input port disable mode.....	1215

Section number	Title	Page
15.6.11	Software assisted error recovery register support.....	1216
15.6.12	Hot-swap support.....	1217
15.6.12.1	Method 1.....	1217
15.6.12.1.1	Extraction-method 1.....	1217
15.6.12.1.2	Insertion-method 1.....	1218
15.6.12.2	Method 2 with RapidIO port hot-swapped.....	1219
15.6.12.2.1	Extraction-method 2.....	1219
15.6.12.2.2	Insertion-method 2.....	1219
15.6.13	Software re-training.....	1220
15.6.14	Errors and error handling.....	1221
15.6.14.1	RapidIO error description.....	1221
15.6.14.2	Physical layer RapidIO errors detected.....	1222
15.6.14.3	Logical layer errors and error handling.....	1226
15.6.14.3.1	Logical layer RapidIO errors detected.....	1226
15.7	RapidIO message unit	1259
15.7.1	Message unit overview.....	1259
15.7.2	Message unit features.....	1260
15.7.3	Outbound modes of operation.....	1261
15.7.4	Outbound message controller operation.....	1261
15.7.4.1	Direct mode operation.....	1262
15.7.4.1.1	Direct mode interrupts.....	1263
15.7.4.1.2	Direct mode message error response errors.....	1263
15.7.4.1.3	Direct mode packet response time-out errors.....	1264
15.7.4.1.4	Direct mode retry error threshold exceeded errors.....	1264
15.7.4.1.5	Direct mode transaction errors.....	1264
15.7.4.1.6	Direct mode error handling.....	1264
15.7.4.1.7	Disabling and enabling the message controller.....	1265
15.7.4.1.8	Direct mode hardware error handling.....	1265
15.7.4.1.9	Direct mode programming errors.....	1269

Section number	Title	Page
15.7.4.2	Chaining mode.....	1269
15.7.4.2.1	Message controller initialization.....	1270
15.7.4.2.2	Chaining mode operation.....	1271
15.7.4.2.3	Changing descriptor queues in chaining mode.....	1272
15.7.4.2.4	Preventing queue overflow in chaining mode.....	1273
15.7.4.2.5	Switching between direct and chaining modes.....	1273
15.7.4.2.6	Chaining mode descriptor format.....	1273
15.7.4.2.7	Chaining mode controller interrupts.....	1275
15.7.4.2.8	Chaining mode message error response errors.....	1275
15.7.4.2.9	Chaining mode packet response time-out errors.....	1276
15.7.4.2.10	Chaining mode retry error threshold exceeded errors.....	1276
15.7.4.2.11	Chaining mode transaction errors.....	1276
15.7.4.2.12	Chaining mode error handling.....	1277
15.7.4.2.13	Chaining mode hardware error handling.....	1277
15.7.4.2.14	Chaining mode programming errors.....	1278
15.7.4.3	Message controller arbitration.....	1278
15.7.5	Inbound message controller operation.....	1279
15.7.5.1	Inbound message controller initialization.....	1280
15.7.5.2	Inbound controller operation.....	1281
15.7.5.3	Message steering.....	1282
15.7.5.4	Inbound message controller retry response conditions.....	1282
15.7.5.5	Inbound message controller interrupts.....	1283
15.7.5.5.1	Message request time-out errors.....	1283
15.7.5.5.2	Inbound message controller transaction errors.....	1284
15.7.5.5.3	Inbound message controller error handling.....	1284
15.7.5.5.4	Inbound message controller hardware error handling.....	1284
15.7.5.5.5	Programming errors.....	1290
15.7.5.5.6	Disabling and enabling the inbound message controller.....	1290
15.7.6	Message unit message passing logical specification registers.....	1291

Section number	Title	Page
15.8	RapidIO doorbell and port-write unit.....	1291
15.8.1	Doorbell and port-write unit features.....	1292
15.8.2	Doorbell controller.....	1292
15.8.2.1	Outbound doorbell controller.....	1293
15.8.2.1.1	Outbound doorbell controller interrupts.....	1294
15.8.2.1.2	Error response errors.....	1294
15.8.2.1.3	Outbound doorbell controller packet response time-out errors.....	1294
15.8.2.1.4	Outbound doorbell controller retry error threshold exceeded errors.....	1295
15.8.2.1.5	Error handling.....	1295
15.8.2.1.6	Disabling and enabling the doorbell controller-outbound.....	1295
15.8.2.1.7	Outbound doorbell controller hardware error handling.....	1296
15.8.2.1.8	Outbound doorbell controller programming errors.....	1298
15.8.2.2	Inbound doorbell controller.....	1299
15.8.2.2.1	Inbound doorbell controller initialization.....	1299
15.8.2.2.2	Inbound doorbell controller operation.....	1299
15.8.2.2.3	Inbound doorbell queue entry format.....	1300
15.8.2.2.4	Inbound doorbell controller retry response conditions.....	1301
15.8.2.2.5	Doorbell controller interrupts.....	1301
15.8.2.2.6	Doorbell controller transaction errors.....	1302
15.8.2.2.7	Doorbell controller error handling.....	1302
15.8.2.2.8	Inbound doorbell controller hardware error handling.....	1302
15.8.2.2.9	Inbound doorbell controller programming errors.....	1305
15.8.2.2.10	Disabling and enabling the doorbell controller-inbound.....	1306
15.8.2.3	Doorbell message passing logical specification registers.....	1306
15.8.3	Port-write controller.....	1307
15.8.3.1	Port-write controller initialization.....	1308
15.8.3.2	Port-write controller operation.....	1308
15.8.3.3	Port-write controller interrupt.....	1309
15.8.3.4	Discarding port-writes.....	1309

Section number	Title	Page
15.8.3.5	Transaction errors.....	1310
15.8.3.5.1	Port-write controller error handling.....	1310
15.8.3.6	Port-write controller hardware error handling.....	1310
15.8.3.6.1	Port-write controller programming errors.....	1313
15.8.3.7	Disabling and enabling the port-write controller.....	1314
15.8.3.8	Port-write controller message passing logical specification registers.....	1314

Chapter 16

PCI Express Interface Controller

16.1	Introduction.....	1317
16.1.1	Outbound transactions.....	1319
16.1.2	Inbound transactions.....	1320
16.2	PCI Express features summary.....	1321
16.3	PCI Express modes of operation.....	1321
16.3.1	Root complex/endpoint modes.....	1322
16.3.2	Link width.....	1322
16.4	PCI Express signal descriptions.....	1322
16.5	Memory map/register overview.....	1323
16.6	PCI Express memory-mapped registers.....	1324
16.6.1	PCI Express configuration address register (PEX _x _PEX_CONFIG_ADDR).....	1331
16.6.2	PCI Express configuration data register (PEX _x _PEX_CONFIG_DATA).....	1332
16.6.3	PCI Express outbound completion timeout register (PEX _x _PEX_OTB_CPL_TOR).....	1333
16.6.4	PCI Express configuration retry timeout register (PEX _x _PEX_CONF_RTU_TOR).....	1334
16.6.5	PCI Express configuration register (PEX _x _PEX_CONFIG).....	1335
16.6.6	PCI Express PME & message detect register (PEX _x _PEX_PME_MES_DR).....	1336
16.6.7	PCI Express PME & message disable register (PEX _x _PEX_PME_MES_DISR).....	1339
16.6.8	PCI Express PME & message interrupt enable register (PEX _x _PEX_PME_MES_IER).....	1341
16.6.9	PCI Express power management command register (PEX _x _PEX_PMCR).....	1343
16.6.10	IP block revision register 1 (PEX _x _PEX_IP_BLK_REV1).....	1344
16.6.11	IP block revision register 2 (PEX _x _PEX_IP_BLK_REV2).....	1344

Section number	Title	Page
16.6.12	PCI Express outbound translation address register n (PEX _x _PEXOTAR _n).....	1345
16.6.13	PCI Express outbound translation extended address register n (PEX _x _PEXOTEAR _n).....	1345
16.6.14	PCI Express outbound window attributes register n (PEX _x _PEXOWAR0).....	1347
16.6.15	PCI Express outbound window base address register n (PEX _x _PEXOWBAR _n).....	1349
16.6.16	PCI Express outbound window attributes register n (PEX _x _PEXOWAR _n).....	1350
16.6.17	PCI Express outbound window attributes register 3 (PEX _x _PEXOWAR3).....	1353
16.6.18	PCI Express outbound window attributes register 4 (PEX _x _PEXOWAR4).....	1355
16.6.19	PCI Express inbound translation address register n (PEX _x _PEXITAR _n).....	1357
16.6.20	PCI Express inbound window base address register n (PEX _x _PEXIWBAR _n).....	1358
16.6.21	PCI Express inbound window base extended address register n (PEX _x _PEXIWBEAR _n).....	1359
16.6.22	PCI Express inbound window attributes register n (PEX _x _PEXIWAR _n).....	1360
16.6.23	PCI Express error detect register (PEX _x _PEX_ERR_DR).....	1363
16.6.24	PCI Express error interrupt enable register (PEX _x _PEX_ERR_EN).....	1366
16.6.25	PCI Express error disable register (PEX _x _PEX_ERR_DISR).....	1369
16.6.26	PCI Express error capture status register (PEX _x _PEX_ERR_CAP_STAT).....	1371
16.6.27	PCI Express error capture register n (PEX _x _PEX_ERR_CAP_R _n).....	1372
16.7	PCI Express configuration-space registers.....	1373
16.7.1	PCI compatible configuration headers.....	1373
16.8	Type 0 configuration header registers.....	1374
16.8.1	PCI Express Vendor ID Register (Vendor_ID_Register).....	1376
16.8.2	PCI Express Device ID Register (Device_ID_Register).....	1376
16.8.3	PCI Express Command Register (Command_Register).....	1377
16.8.4	PCI Express Status Register (Status_Register).....	1379
16.8.5	PCI Express Revision ID Register (Revision_ID_Register).....	1380
16.8.6	PCI Express Class Code Register (Class_Code_Register).....	1380
16.8.7	PCI Express Cache Line Size Register (Cache_Line_Size_Register).....	1381
16.8.8	PCI Express Latency Timer Register (Latency_Timer_Register).....	1382
16.8.9	PCI Express Header Type Register (Header_Type_Register).....	1382
16.8.10	PCI Express Base Address Register 0 (PEXCSRBAR).....	1383

Section number	Title	Page
16.8.11	PCI Express Base Address Register 1 (BAR1).....	1384
16.8.12	PCI Express Base Address Register 2,4 (BAR n).....	1385
16.8.13	PCI Express Base Address Register 3,5 (BAR n).....	1387
16.8.14	PCI Express Subsystem Vendor ID Register (Subsystem_Vendor_ID_Register).....	1387
16.8.15	PCI Express Subsystem ID Register (Subsystem_ID_Register).....	1388
16.8.16	Capabilities Pointer Register (Capabilities_Pointer_Register).....	1389
16.8.17	PCI Express Interrupt Line Register (Interrupt_Line_Register).....	1389
16.8.18	PCI Express Interrupt Pin Register (Interrupt_Pin_Register).....	1390
16.8.19	PCI Express Minimum Grant Register (Minimum_Grant_Register).....	1390
16.8.20	PCI Express Maximum Latency Register (Maximum_Latency_Register).....	1391
16.9	Type 1 configuration header registers.....	1391
16.9.1	PCI Express Base Address Register 0 (PEXCSRBAR).....	1392
16.9.2	PCI Express Primary Bus Number Register (Primary_Bus_Number_Register).....	1394
16.9.3	PCI Express Secondary Bus Number Register (Secondary_Bus_Number_Register).....	1394
16.9.4	PCI Express Subordinate Bus Number Register (Subordinate_Bus_Number_Register).....	1394
16.9.5	PCI Express I/O Base Register (IO_Base_Register).....	1395
16.9.6	PCI Express I/O Limit Register (IO_Limit_Register).....	1395
16.9.7	PCI Express Secondary Status Register (Secondary_Status_Register).....	1396
16.9.8	PCI Express Memory Base Register (Memory_Base_Register).....	1397
16.9.9	PCI Express Memory Limit Register (Memory_Limit_Register).....	1397
16.9.10	PCI Express Prefetchable Memory Base Register (Prefetchable_Memory_Base_Register).....	1398
16.9.11	PCI Express Prefetchable Memory Limit Register (Prefetchable_Memory_Limit_Register).....	1398
16.9.12	PCI Express Prefetchable Base Upper 32 Bits Register (Prefetchable_Base_Upper_32_Bits_Register).....	1399
16.9.13	PCI Express Prefetchable Limit Upper 32 Bits Register (Prefetchable_Limit_Upper_32_Bits_Register).....	1399
16.9.14	PCI Express I/O Base Upper 16 Bits Register (IO_Base_Upper_16_Bits_Register).....	1400
16.9.15	PCI Express I/O Limit Upper 16 Bits Register (IO_Limit_Upper_16_Bits_Register).....	1400
16.9.16	Capabilities Pointer Register (Capabilities_Pointer_Register).....	1401
16.9.17	PCI Express Interrupt Line Register (Interrupt_Line_Register).....	1401

Section number	Title	Page
16.9.18	PCI Express Interrupt Pin Register (Interrupt_Pin_Register).....	1402
16.9.19	PCI Express Bridge Control Register (Bridge_Control_Register).....	1402
16.10	PCI compatible device-specific configuration space.....	1403
16.10.1	PCI Express Power Management Capability ID Register (Power_Management_Capability_ID_Register).....	1405
16.10.2	PCI Express Power Management Capabilities Register (Power_Management_Capabilities_Register)...	1405
16.10.3	PCI Express Power Management Status and Control Register (Power_Management_Status_and_Control_Register).....	1406
16.10.4	PCI Express Power Management Data Register (Power_Management_Data_Register).....	1407
16.10.5	PCI Express Capability ID Register (Capability_ID_Register).....	1407
16.10.6	PCI Express Capabilities Register (Capabilities_Register).....	1408
16.10.7	PCI Express Device Capabilities Register (Device_Capabilities_Register).....	1409
16.10.8	PCI Express Device Control Register (Device_Control_Register).....	1410
16.10.9	PCI Express Device Status Register (Device_Status_Register).....	1411
16.10.10	PCI Express Link Capabilities Register (Link_Capabilities_Register).....	1412
16.10.11	PCI Express Link Control Register (Link_Control_Register).....	1413
16.10.12	PCI Express Link Status Register (Link_Status_Register).....	1414
16.10.13	PCI Express Slot Capabilities Register (Slot_Capabilities_Register).....	1415
16.10.14	PCI Express Slot Control Register (Slot_Control_Register).....	1416
16.10.15	PCI Express Slot Status Register (Slot_Status_Register).....	1418
16.10.16	PCI Express Root Control Register (Root_Control_Register).....	1419
16.10.17	PCI Express Root Status Register (Root_Status_Register).....	1420
16.10.18	PCI Express MSI Message Capability ID Register (MSI_Message_Capability_ID_Register).....	1421
16.10.19	PCI Express MSI Message Control Register (MSI_Message_Control_Register).....	1421
16.10.20	PCI Express MSI Message Address Register (MSI_Message_Address_Register).....	1422
16.10.21	PCI Express MSI Message Upper Address Register (MSI_Message_Upper_Address_Register).....	1422
16.10.22	PCI Express MSI Message Data Register (MSI_Message_Data_Register).....	1423
16.11	PCI Express extended configuration space.....	1423
16.11.1	PCI Express Advanced Error Reporting Capability ID Register (Advanced_Error_Reporting_Capability_ID_Register).....	1425

Section number	Title	Page
16.11.2	PCI Express Uncorrectable Error Status Register (Uncorrectable_Error_Status_Register).....	1425
16.11.3	PCI Express Uncorrectable Error Mask Register (Uncorrectable_Error_Mask_Register).....	1427
16.11.4	PCI Express Uncorrectable Error Severity Register (Uncorrectable_Error_Severity_Register).....	1428
16.11.5	PCI Express Correctable Error Status Register (Correctable_Error_Status_Register).....	1429
16.11.6	PCI Express Correctable Error Mask Register (Correctable_Error_Mask_Register).....	1430
16.11.7	PCI Express Advanced Error Capabilities and Control Register (Advanced_Error_Capabilities_and_Control_Register).....	1432
16.11.8	PCI Express Header Log Register 1 (Header_Log_Register_DWORD1).....	1433
16.11.9	PCI Express Header Log Register 2 (Header_Log_Register_DWORD2).....	1433
16.11.10	PCI Express Header Log Register 3 (Header_Log_Register_DWORD3).....	1434
16.11.11	PCI Express Header Log Register 4 (Header_Log_Register_DWORD4).....	1435
16.11.12	PCI Express Root Error Command Register (Root_Error_Command_Register).....	1436
16.11.13	PCI Express Root Error Status Register (Root_Error_Status_Register).....	1437
16.11.14	PCI Express Correctable Error Source ID Register (Correctable_Error_Source_ID_Register).....	1438
16.11.15	PCI Express Error Source ID Register (Error_Source_ID_Register).....	1438
16.11.16	LTSSM State Status Register (LTSSM_State_Status_Register).....	1439
16.11.17	PCI Express Controller Core Clock Ratio Register (Controller_Core_Clock_Ratio_Register).....	1440
16.11.18	PCI Express Power Management Timer Register (Power_Management_Timer_Register).....	1441
16.11.19	PCI Express PME Time-Out Register (PME_Time_Out_Register).....	1442
16.11.20	PCI Express Subsystem Vendor ID Update Register (Subsystem_Vendor_ID_Update_Register).....	1443
16.11.21	Configuration Ready Register (Configuration_Ready_Register).....	1443
16.11.22	Flow Control Update Timeout Register (Flow_Control_Update_Timeout_Register).....	1445
16.11.23	Secondary Status Interrupt Mask Register (Secondary_Status_Interrupt_Mask_Register).....	1446
16.12	Functional description.....	1447
16.12.1	Architecture.....	1448
16.12.1.1	PCI Express transactions.....	1448
16.12.1.2	Byte ordering.....	1449
16.12.1.2.1	Address invariance.....	1450
16.12.1.2.2	Byte order for configuration transactions.....	1451

Section number	Title	Page
16.12.1.3	Transaction ordering rules.....	1452
16.12.1.4	PCI Express outbound ATMUs.....	1454
16.12.1.5	PCI Express inbound ATMUs.....	1454
16.12.1.5.1	EP inbound ATMU implementation.....	1454
16.12.1.5.2	RC inbound ATMU implementation.....	1455
16.12.1.6	Memory space addressing.....	1456
16.12.1.7	I/O space addressing.....	1456
16.12.1.8	Configuration space addressing.....	1457
16.12.1.9	PCI Express configuration space access.....	1457
16.12.1.9.1	RC configuration register access.....	1457
16.12.1.9.1.1	PCI Express configuration access register mechanism.....	1458
16.12.1.9.1.2	Outbound ATMU configuration mechanism (RC-only).....	1459
16.12.1.9.2	EP configuration register access.....	1459
16.12.1.10	Serialization of configuration and I/O writes.....	1460
16.12.1.11	Messages.....	1460
16.12.1.11.1	Outbound ATMU message generation.....	1460
16.12.1.11.2	Inbound messages.....	1462
16.12.1.12	Error handling.....	1463
16.12.1.12.1	PCI Express error logging and signaling.....	1464
16.12.1.12.2	PCI Express controller internal interrupt sources.....	1466
16.12.1.12.3	Error conditions	1467
16.12.1.12.4	Error capture registers.....	1470
16.12.1.12.4.1	Error capture registers (outbound error).....	1470
16.12.1.12.4.2	Error capture registers (inbound error).....	1471
16.12.2	Interrupts.....	1474
16.12.2.1	EP interrupt generation.....	1474
16.12.2.1.1	Hardware INTx message generation.....	1474
16.12.2.1.2	Hardware MSI generation.....	1475
16.12.2.1.3	Software INTx message generation.....	1475

Section number	Title	Page
16.12.2.1.4	Software MSI generation.....	1475
16.12.2.2	RC handling of INTx message and MSI interrupts.....	1476
16.12.2.2.1	INTx message handling.....	1476
16.12.2.2.2	MSI handling.....	1476
16.12.3	Initial credit advertisement.....	1476
16.12.4	Power management.....	1477
16.12.4.1	L2/L3 ready link state.....	1478
16.12.5	Hot reset.....	1478
16.12.6	Link down.....	1478
16.13	Initialization/application information.....	1479
16.13.1	EP Boot mode and inbound configuration transactions.....	1479
16.13.2	Automatic link retraining during initialization.....	1479
16.13.3	Configuration accesses and inbound writes to CCSR space.....	1480

Chapter 17

Enhanced secure digital host controller (eSDHC)

17.1	eSDHC overview.....	1481
17.2	eSDHC features summary.....	1483
17.2.1	Data transfer modes.....	1484
17.3	eSDHC external signal description.....	1484
17.4	Enhanced Secure Digital Host Controller (eSDHC) Memory Map.....	1485
17.4.1	DMA system address (eSDHC_DSADDR).....	1487
17.4.2	Block attributes (eSDHC_BLKATTR).....	1487
17.4.3	Command argument (eSDHC_CMDARG).....	1488
17.4.4	Command transfer type (eSDHC_XFERTYP).....	1489
17.4.5	Command response n (eSDHC_CMDRSP _n).....	1492
17.4.6	Data buffer access port (eSDHC_DATPORT).....	1493
17.4.7	Present state (eSDHC_PRSSTAT).....	1494
17.4.8	Protocol control (eSDHC_PROCTL).....	1499
17.4.9	System control (eSDHC_SYSCTL).....	1502

Section number	Title	Page
17.4.10	Interrupt status (eSDHC_IRQSTAT).....	1505
17.4.11	Interrupt status enable (eSDHC_IRQSTATEN).....	1510
17.4.12	Interrupt signal enable (eSDHC_IRQSIGEN).....	1513
17.4.13	Auto CMD12 status (eSDHC_AUTOC12ERR).....	1515
17.4.14	Host controller capabilities (eSDHC_HOSTCAPBLT).....	1518
17.4.15	Watermark level (eSDHC_WML).....	1520
17.4.16	Force event (eSDHC_FEVT).....	1521
17.4.17	Host controller version (eSDHC_HOSTVER).....	1523
17.4.18	DMA control register (eSDHC_DCR).....	1524
17.5	eSDHC functional description.....	1525
17.5.1	Data buffer.....	1525
17.5.1.1	Write operation sequence.....	1526
17.5.1.2	Read operation sequence.....	1527
17.5.1.3	Data buffer size.....	1528
17.5.2	DMA CCB interface.....	1528
17.5.2.1	Internal DMA request.....	1529
17.5.2.2	DMA burst length.....	1529
17.5.2.3	CCB Master interface.....	1529
17.5.3	SD protocol unit.....	1530
17.5.3.1	SD transceiver.....	1530
17.5.3.2	SD clock and monitor.....	1530
17.5.3.3	Command agent.....	1531
17.5.3.4	Data agent.....	1531
17.5.4	Clock and reset manager.....	1532
17.5.5	Clock generator.....	1532
17.5.6	Card insertion and removal detection.....	1533
17.5.7	Power management and wake-up events.....	1533
17.5.7.1	Setting wake-up events.....	1534

Section number	Title	Page
17.6	Initialization/application information.....	1534
17.6.1	Command send and response receive basic operation.....	1535
17.6.2	Card identification mode.....	1535
17.6.2.1	Card detect.....	1536
17.6.2.2	Reset.....	1537
17.6.2.3	Voltage validation.....	1538
17.6.2.4	Card registry.....	1538
17.6.3	Card access.....	1540
17.6.3.1	Block write.....	1540
17.6.3.1.1	Normal write.....	1540
17.6.3.1.2	Write with pause.....	1541
17.6.3.2	Block read.....	1542
17.6.3.2.1	Normal read.....	1542
17.6.3.2.2	Read with pause.....	1543
17.6.3.3	Transfer error.....	1544
17.6.3.3.1	CRC error.....	1544
17.6.3.3.2	Internal DMA error.....	1544
17.6.3.3.3	Auto CMD12 error.....	1545
17.6.3.4	Card interrupt.....	1545
17.6.4	Switch function.....	1545
17.6.4.1	Query, enable and disable SD high speed mode.....	1546
17.6.4.2	Query, enable and disable MMC high speed mode.....	1546
17.6.4.3	Set MMC bus width.....	1547
17.6.5	Commands for MMC/SD.....	1547
17.6.6	Software restrictions.....	1554

Chapter 18 Universal Serial Bus Interface

18.1	Introduction.....	1555
18.1.1	Overview.....	1556

Section number	Title	Page
18.1.2	Features.....	1556
18.1.3	Modes of operation.....	1557
18.2	USB external signals.....	1557
18.2.1	ULPI interface.....	1557
18.2.2	PHY clocks.....	1558
18.3	USB memory map/register definition.....	1559
18.3.1	Identification register (USB_ID).....	1561
18.3.2	Capability register length (USB_CAPLENGTH).....	1562
18.3.3	Host controller interface version number (USB_HCIVERSION).....	1562
18.3.4	Host controller structural parameters (USB_HCSPARAMS).....	1563
18.3.5	Host controller capability parameters (USB_HCCPARAMS).....	1565
18.3.6	Device controller interface version number (USB_DCIVERSION).....	1567
18.3.7	Device controller capability parameters (USB_DCCPARAMS).....	1568
18.3.8	USB command (USB_USBCMD).....	1568
18.3.9	USB status (USB_USBSTS).....	1572
18.3.10	USB interrupt enable (USB_USBINTR).....	1576
18.3.11	USB frame index (USB_FRINDEX).....	1578
18.3.12	Periodic frame list base address [host mode] (USB_PERIODICLISTBASE).....	1579
18.3.13	USB device address [device mode] (USB_DEVICEADDR).....	1580
18.3.14	Next asynchronous list addr [host mode] (USB_ASYNCLISTADDR).....	1580
18.3.15	Address at endpoint list [device mode] (USB_ENDPOINTLISTADDR).....	1581
18.3.16	Master interface data burst size (USB_BURSTSIZE).....	1582
18.3.17	Transmit FIFO tuning controls (USB_TXFILLTUNING).....	1582
18.3.18	ULPI register access (USB_ULPI_VIEWPORT).....	1584
18.3.19	Configured flag register (USB_CONFIGFLAG).....	1586
18.3.20	Port status/control (USB_PORTSC).....	1587
18.3.21	USB device mode (USB_USBMODE).....	1594
18.3.22	Endpoint setup status (USB_ENDPTSETUPSTAT).....	1595
18.3.23	Endpoint initialization (USB_ENDPOINTPRIME).....	1596

Section number	Title	Page
18.3.24	Endpoint flush (USB_ENDPTFLUSH).....	1597
18.3.25	Endpoint status (USB_ENDPTSTATUS).....	1597
18.3.26	Endpoint complete (USB_ENDPTCOMPLETE).....	1598
18.3.27	Endpoint control 0 (USB_ENDPTCTRL0).....	1600
18.3.28	Endpoint control n (USB_ENDPTCTRLn).....	1602
18.3.29	Snoop n (USB_SNOOPn).....	1604
18.3.30	Age count threshold (USB_AGE_CNT_THRESH).....	1605
18.3.31	Priority control (USB_PRI_CTRL).....	1607
18.3.32	System interface control (USB_SI_CTRL).....	1608
18.3.33	Control (USB_CONTROL).....	1609
18.4	Functional description.....	1610
18.4.1	System interface.....	1610
18.4.2	DMA engine.....	1610
18.4.3	FIFO RAM controller.....	1611
18.4.4	PHY interface.....	1611
18.5	Host data structures.....	1612
18.5.1	Periodic frame list.....	1612
18.5.2	Asynchronous list queue head pointer.....	1614
18.5.3	Isochronous (high-speed) transfer descriptor (iTDD).....	1615
18.5.3.1	Next link pointer-iTDD.....	1615
18.5.3.2	iTDD transaction status and control list.....	1616
18.5.3.3	iTDD buffer page pointer list (plus).....	1617
18.5.4	Split transaction isochronous transfer descriptor (siTDD).....	1619
18.5.4.1	Next link pointer-siTDD.....	1619
18.5.4.2	siTDD endpoint capabilities/characteristics.....	1620
18.5.4.3	siTDD transfer state.....	1621
18.5.4.4	siTDD buffer pointer list (plus).....	1622
18.5.4.5	siTDD back link pointer.....	1623

Section number	Title	Page
18.5.5	Queue element transfer descriptor (qTD).....	1623
18.5.5.1	Next qTD pointer.....	1624
18.5.5.2	Alternate next qTD pointer.....	1625
18.5.5.3	qTD token.....	1625
18.5.5.4	qTD buffer page pointer list.....	1628
18.5.6	Queue head.....	1629
18.5.6.1	Queue head horizontal link pointer.....	1629
18.5.6.2	Endpoint capabilities/characteristics	1630
18.5.6.3	Transfer overlay.....	1632
18.5.7	Periodic frame span traversal node (FSTN).....	1633
18.5.7.1	FSTN normal path pointer.....	1634
18.5.7.2	FSTN back path link pointer.....	1635
18.6	Host operations.....	1635
18.6.1	Host controller initialization.....	1635
18.6.2	Power port.....	1637
18.6.3	Reporting over-current.....	1637
18.6.4	Suspend/resume	1637
18.6.4.1	Port suspend/resume.....	1638
18.6.5	Schedule traversal rules.....	1640
18.6.6	Periodic schedule frame boundaries vs. bus frame boundaries.....	1642
18.6.7	Periodic schedule.....	1644
18.6.8	Managing isochronous transfers using iTDs.....	1645
18.6.8.1	Host controller operational model for iTDs.....	1646
18.6.8.2	Software operational model for iTDs.....	1648
18.6.8.2.1	Periodic scheduling threshold.....	1650
18.6.9	Asynchronous schedule.....	1651
18.6.9.1	Adding queue heads to asynchronous schedule.....	1653
18.6.9.2	Removing queue heads from asynchronous schedule.....	1653
18.6.9.3	Empty asynchronous schedule detection	1656

Section number	Title	Page
18.6.9.4	Asynchronous schedule traversal: Start event.....	1657
18.6.9.5	Reclamation status bit (USBSTS Register).....	1657
18.6.10	Managing control/bulk/interrupt transfers via queue heads.....	1658
18.6.10.1	Buffer pointer list use for data streaming with qTDs	1659
18.6.10.2	Adding interrupt queue heads to the periodic schedule	1661
18.6.10.3	Managing transfer complete interrupts from queue heads	1661
18.6.11	Ping control.....	1662
18.6.12	Split transactions.....	1663
18.6.12.1	Split transactions for asynchronous transfers	1663
18.6.12.1.1	Asynchronous-do-start-split.....	1664
18.6.12.1.2	Asynchronous-do-complete-split	1664
18.6.12.2	Split transaction interrupt	1666
18.6.12.2.1	Split transaction scheduling mechanisms for interrupt	1666
18.6.12.2.2	Host controller operational model for FSTNs	1669
18.6.12.2.3	Software operational model for FSTNs.....	1672
18.6.12.2.4	Tracking split transaction progress for interrupt transfers	1673
18.6.12.2.5	Split transaction execution state machine for interrupt	1674
18.6.12.2.6	Periodic interrupt-do-start-split	1675
18.6.12.2.7	Periodic interrupt-do-complete-split	1676
18.6.12.2.8	Managing the QH[FrameTag] field	1680
18.6.12.2.9	Rebalancing the periodic schedule	1681
18.6.12.3	Split transaction isochronous	1682
18.6.12.3.1	Split transaction scheduling mechanisms for isochronous	1682
18.6.12.3.2	Tracking split transaction progress for isochronous transfers	1687
18.6.12.3.3	Split transaction execution state machine for isochronous	1689
18.6.12.3.4	Periodic isochronous-do-start-split	1690
18.6.12.3.5	Periodic isochronous-do complete split.....	1692
18.6.12.3.6	Complete-split for scheduling boundary cases 2a, 2b.....	1695
18.6.12.3.7	Split transaction for isochronous-processing example.....	1697

Section number	Title	Page
18.6.13	Port test modes.....	1699
18.6.14	Interrupts.....	1699
18.6.14.1	Transfer/transaction based interrupts.....	1701
18.6.14.1.1	Transaction error.....	1701
18.6.14.1.2	Serial bus babble.....	1701
18.6.14.1.3	Data buffer error	1702
18.6.14.1.4	USB interrupt (interrupt on completion (IOC)).....	1703
18.6.14.1.5	Short packet.....	1703
18.6.14.2	Host controller event interrupts.....	1703
18.6.14.2.1	Port change events.....	1704
18.6.14.2.2	Frame list rollover.....	1704
18.6.14.2.3	Interrupt on async advance.....	1704
18.6.14.2.4	Host system error.....	1704
18.7	Device data structures.....	1705
18.7.1	Endpoint queue head.....	1706
18.7.1.1	Endpoint capabilities/characteristics	1707
18.7.1.2	Transfer overlay	1708
18.7.1.3	Current dTD pointer.....	1708
18.7.1.4	Setup buffer.....	1709
18.7.2	Endpoint transfer descriptor (dTD).....	1709
18.8	Device operational model.....	1711
18.8.1	Device controller initialization.....	1712
18.8.2	Port state and control.....	1713
18.8.2.1	Bus reset.....	1715
18.8.2.2	Suspend/resume	1716
18.8.2.2.1	Suspend description.....	1716
18.8.2.2.2	Suspend operational model.....	1717
18.8.2.2.3	Resume.....	1717

Section number	Title	Page
18.8.3	Managing endpoints.....	1717
18.8.3.1	Endpoint initialization.....	1718
18.8.3.1.1	Stalling.....	1719
18.8.3.2	Data toggle.....	1719
18.8.3.2.1	Data toggle reset.....	1720
18.8.3.2.2	Data toggle inhibit.....	1720
18.8.3.3	Device operational model for packet transfers.....	1720
18.8.3.3.1	Priming transmit endpoints.....	1721
18.8.3.3.2	Priming receive endpoints.....	1721
18.8.3.4	Interrupt/bulk endpoint operational model.....	1721
18.8.3.4.1	Interrupt/bulk endpoint bus response matrix.....	1723
18.8.3.5	Control endpoint operation model.....	1724
18.8.3.5.1	Setup phase.....	1724
18.8.3.5.2	Data phase.....	1725
18.8.3.5.3	Status phase.....	1725
18.8.3.5.4	Control endpoint bus response matrix.....	1726
18.8.3.6	Isochronous endpoint operational model.....	1726
18.8.3.6.1	Isochronous pipe synchronization.....	1728
18.8.3.6.2	Isochronous endpoint bus response matrix.....	1728
18.8.4	Managing queue heads.....	1729
18.8.4.1	Queue head initialization.....	1729
18.8.4.2	Operational model for setup transfers.....	1730
18.8.5	Managing transfers with transfer descriptors.....	1731
18.8.5.1	Software link pointers.....	1731
18.8.5.2	Building a transfer descriptor.....	1732
18.8.5.3	Executing a transfer descriptor.....	1732
18.8.5.4	Transfer completion.....	1733
18.8.5.5	Flushing/depriming an endpoint.....	1734
18.8.5.6	Device error matrix.....	1734

Section number	Title	Page
18.8.6	Servicing interrupts.....	1735
18.8.6.1	High-frequency interrupts.....	1735
18.8.6.2	Low-frequency interrupts.....	1735
18.8.6.3	Error interrupts.....	1736
18.9	Deviations from the EHCI specifications.....	1736
18.9.1	Embedded transaction translator function.....	1737
18.9.1.1	Capability registers.....	1737
18.9.1.2	Operational registers.....	1737
18.9.1.3	Discovery	1737
18.9.1.4	Data structures.....	1738
18.9.1.5	Operational model.....	1739
18.9.1.5.1	Microframe pipeline.....	1739
18.9.1.5.2	Split state machines.....	1740
18.9.1.5.3	Asynchronous transaction scheduling and buffer management.....	1740
18.9.1.5.4	Periodic transaction scheduling and buffer management.....	1741
18.9.1.5.5	Multiple transaction translators.....	1741
18.9.2	Device operation.....	1741
18.9.3	Non-zero fields the register file.....	1741
18.9.4	SOF interrupt.....	1742
18.9.5	Embedded design.....	1742
18.9.5.1	Frame adjust register.....	1742
18.9.6	Miscellaneous variations from EHCI.....	1742
18.9.6.1	Discovery.....	1743
18.9.6.1.1	Port reset.....	1743
18.9.6.1.2	Port speed detection.....	1743

Chapter 19 Enhanced Serial Peripheral Interface

19.1	Introduction.....	1745
19.1.1	Features.....	1746

Section number	Title	Page
19.1.2	eSPI transmission and reception process.....	1747
19.1.3	Modes of operation.....	1747
19.2	External signal descriptions.....	1748
19.2.1	Overview.....	1748
19.2.2	ESPI detailed signal descriptions	1749
19.3	Enhanced serial peripheral interface (eSPI) memory map.....	1750
19.3.1	eSPI mode register (ESPI_SPMODE).....	1751
19.3.2	eSPI event register (ESPI_SPIE).....	1752
19.3.3	eSPI mask register (ESPI_SPIM).....	1754
19.3.4	eSPI command register (ESPI_SPCOM).....	1756
19.3.5	eSPI transmit FIFO access register (ESPI_SPITF).....	1758
19.3.6	eSPI receive FIFO access register (ESPI_SPIRF).....	1760
19.3.7	eSPI CS0 mode register (ESPI_SPMODE0).....	1761
19.3.8	eSPI CS1 mode register (ESPI_SPMODE1).....	1763
19.3.9	eSPI CS2 mode register (ESPI_SPMODE2).....	1765
19.3.10	eSPI CS3 mode register (ESPI_SPMODE3).....	1767
19.4	eSPI transfer formats.....	1768
19.5	CI and CP values for various eSPI devices.....	1769
19.6	eSPI programming examples.....	1770
19.6.1	24-bit address example.....	1770
19.6.2	16-bit address example.....	1770

Chapter 20 General Purpose I/O (GPIO)

20.1	Introduction.....	1771
20.1.1	Overview.....	1771
20.1.2	Features.....	1772
20.2	External signal description.....	1772
20.2.1	Signals overview.....	1772

Section number	Title	Page
20.3	GPIO memory map/register definition.....	1772
20.3.1	GPIO direction register (GPIO_GPDIR).....	1773
20.3.2	GPIO open drain register (GPIO_GPODR).....	1774
20.3.3	GPIO data register (GPIO_GPDAT).....	1774
20.3.4	GPIO interrupt event register (GPIO_GPIER).....	1775
20.3.5	GPIO interrupt mask register (GPIO_GPIMR).....	1775
20.3.6	GPIO external interrupt control register (GPIO_GPICR).....	1776

Chapter 21 Global Utilities

21.1	Overview.....	1777
21.2	Global utilities features.....	1777
21.2.1	Power management and block disables.....	1777
21.2.2	Accessing current POR configuration settings.....	1777
21.2.3	Clock control.....	1778
21.3	External signal description.....	1778
21.3.1	Signals overview.....	1778
21.3.2	Detailed signal descriptions.....	1778
21.4	GUTS Memory Map/Register Definition.....	1780
21.4.1	POR PLL ratio status register (GUTS_PORPLLSR).....	1782
21.4.2	POR boot mode status register (GUTS_PORBMSR).....	1784
21.4.3	POR device status register (GUTS_PORDEVSR).....	1786
21.4.4	POR debug mode status register (GUTS_PORDBGMSR).....	1789
21.4.5	POR device status register 2 (GUTS_PORDEVSR2).....	1791
21.4.6	General-purpose POR configuration register (GUTS_GPPORCR).....	1794
21.4.7	Alternate function signal multiplex control (GUTS_PMUXCR).....	1794
21.4.8	Device disable control register (GUTS_DEVDISR).....	1796
21.4.9	Power management control and status register (GUTS_POWMGTCSR).....	1800
21.4.10	Power management clock disable register (GUTS_PMCDR).....	1803
21.4.11	Machine check summary register (GUTS_MCPSUMR).....	1804

Section number	Title	Page
21.4.12	Reset request status and control register (GUTS_RSTRSCR).....	1806
21.4.13	Exception reset control register (GUTS_ECTRSTCR).....	1808
21.4.14	Automatic reset status register (GUTS_AUTORSTSR).....	1810
21.4.15	Processor version register (GUTS_PVR).....	1812
21.4.16	System version register (GUTS_SVR).....	1812
21.4.17	Reset control register (GUTS_RSTCR).....	1813
21.4.18	IO voltage select status register (GUTS_IOVSELSR).....	1814
21.4.19	DDR clock disable register (GUTS_DDRCLKDR).....	1815
21.4.20	Clock out control register (GUTS_CLKOCR).....	1816
21.4.21	ECM control register (GUTS_ECMCR).....	1817
21.4.22	SDHC debug control register (GUTS_SDHCDCR).....	1818
21.4.23	SRDS Control Register 0 (GUTS_SRDSCR0).....	1819
21.4.24	SRDS Control Register 1 (GUTS_SRDSCR1).....	1821
21.4.25	SRDS Control Register 2 (GUTS_SRDSCR2).....	1823
21.4.26	SRDS Control Register 4 (GUTS_SRDSCR4).....	1824
21.4.27	SRDS Control Register 5 (GUTS_SRDSCR5).....	1826
21.4.28	SRDS Control Register 6 (GUTS_SRDSCR6).....	1827
21.5	Functional description.....	1828
21.5.1	Power management.....	1828
21.5.1.1	Relationship between both cores and device power management states.....	1828
21.5.1.2	CKSTP_IN0_B/CKSTP_IN1_B is not power management.....	1829
21.5.1.3	Dynamic power management	1830
21.5.1.4	Shutting down unused blocks.....	1830
21.5.1.5	Software-controlled power-down states.....	1830
21.5.1.5.1	Doze mode.....	1831
21.5.1.5.2	Nap mode.....	1831
21.5.1.5.3	Sleep mode.....	1831
21.5.1.6	Power management control fields.....	1832
21.5.1.7	Power-down sequence coordination.....	1832

Section number	Title	Page
21.5.1.8	Interrupts and power management.....	1834
21.5.1.8.1	Interrupts and power management controlled by MSR[WE].....	1835
21.5.1.8.2	Interrupts and power management controlled by POWMGTCR.....	1835
21.5.1.9	Snooping in power-down modes.....	1836
21.5.1.10	Software considerations for power management.....	1836
21.5.1.11	Requirements for reaching and recovering from sleep state.....	1836

Chapter 22 Device Performance Monitor

22.1	Introduction.....	1839
22.1.1	Overview.....	1840
22.1.2	Features.....	1841
22.2	Signal descriptions.....	1842
22.3	PERFMON Memory Map/Register Definition.....	1842
22.3.1	Performance monitor global control register (PERFMON_PMGC0).....	1845
22.3.2	Performance monitor local control register A0 (PERFMON_PMLCA0).....	1846
22.3.3	Performance monitor local control register B0 (PERFMON_PMLCB0).....	1847
22.3.4	Performance monitor counter 0 lower (PERFMON_PMC0_lower).....	1848
22.3.5	Performance monitor counter 0 upper (PERFMON_PMC0_upper).....	1849
22.3.6	Performance monitor local control register A _n (PERFMON_PMLCA _n).....	1849
22.3.7	Performance monitor local control register B _n (PERFMON_PMLCB _n).....	1850
22.3.8	Performance monitor counter <i>n</i> (PERFMON_PMC _n).....	1852
22.4	Functional description.....	1852
22.4.1	Performance monitor interrupt.....	1852
22.4.2	Event counting.....	1853
22.4.3	Threshold events.....	1853
22.4.4	Chaining.....	1854
22.4.5	Triggering.....	1855
22.4.6	Burstiness counting.....	1855
22.4.7	Performance monitor events.....	1858

Section number	Title	Page
22.4.8	Performance monitor examples.....	1870

Chapter 23 Debug Features and Watchpoint Facility

23.1	Introduction.....	1873
23.1.1	Features.....	1875
23.1.2	Modes of operation.....	1875
23.1.2.1	Memory debug mode (eLBC and DDR).....	1876
23.1.2.2	DDR SDRAM interface debug mode.....	1876
23.1.2.3	Watchpoint monitor modes.....	1877
23.1.2.4	Trace buffer modes.....	1877
23.2	External signal description.....	1877
23.2.1	Overview.....	1878
23.2.2	Detailed signal descriptions.....	1879
23.2.2.1	Debug signals-details.....	1879
23.2.2.2	Watchpoint monitor trigger signals-details.....	1880
23.2.2.3	JTAG test signals-details.....	1881
23.3	Debug Memory Map/Register Definition.....	1882
23.3.1	Watchpoint monitor control register 0 (Debug_WMCR0).....	1883
23.3.2	Watchpoint monitor control register 1 (Debug_WMCR1).....	1885
23.3.3	Watchpoint monitor address register (Debug_WMAR).....	1886
23.3.4	Watchpoint monitor address mask register (Debug_WMAMR).....	1886
23.3.5	Watchpoint monitor transaction mask register (Debug_WMTMR).....	1887
23.3.6	Watchpoint monitor status register (Debug_WMSR).....	1888
23.3.7	Trace buffer control register 0 (Debug_TBCR0).....	1889
23.3.8	Trace buffer control register 1 (Debug_TBCR1).....	1891
23.3.9	Trace buffer address register (Debug_TBAR).....	1892
23.3.10	Trace buffer address mask register (Debug_TBAMR).....	1892
23.3.11	Trace buffer transaction mask register (Debug_TBTMR).....	1893
23.3.12	Trace buffer status register (Debug_TBSR).....	1894

Section number	Title	Page
23.3.13	Trace buffer access control register (Debug_TBACR).....	1895
23.3.14	Trace buffer access data high register (Debug_TBADHR).....	1895
23.3.15	Trace buffer access data register (Debug_TBADR).....	1896
23.3.16	Programmed context ID register (Debug_PCIDR).....	1896
23.3.17	Current context ID register (Debug_CCIDR).....	1897
23.3.18	Trigger output source register (Debug_TOSR).....	1897
23.4	Functional description.....	1898
23.4.1	Source and target ID.....	1898
23.4.2	DDR SDRAM interface debug.....	1900
23.4.2.1	Debug information on debug pins.....	1900
23.4.2.2	Debug information on ECC pins.....	1900
23.4.3	Local bus interface debug.....	1901
23.4.4	Watchpoint monitor.....	1901
23.4.4.1	Watchpoint monitor performance monitor events.....	1901
23.4.5	Trace buffer.....	1902
23.4.5.1	Traced data formats (as a function of TBCR1[IFSEL]).....	1902
23.5	Initialization.....	1904

Chapter 1

Overview

The chip combines dual Power Architecture® e500v2 processor cores with system logic required for networking, wireless infrastructure, and telecommunications applications.

1.1 Overview

This section provides an overview of features and functionality of the P2020 QorIQ integrated processor.

The chip combines dual Power Architecture® e500v2 processor cores with system logic required for networking, wireless infrastructure, and telecommunications applications.

The chip offers an excellent combination of protocol and interface support including dual high-performance CPU cores, a large L2 cache, a DDR2/DDR3 memory controller, three enhanced three-speed Ethernet controllers with SGMII support, two serial RapidIO interfaces with a messaging unit, a secure digital interface, a USB 2.0 interface, and three PCI Express controllers. The device also supports the IEEE 1588™ precision time protocol for network synchronization over Ethernet.

This manual is written from the perspective of the P2020; the information also applies to the P2010, with the exception that the P2010 is a single-core chip and refer to *P2010 QorIQ Integrated Processor Hardware Specifications*.

1.1.1 Block diagram

This figure shows the major functional units within the chip.

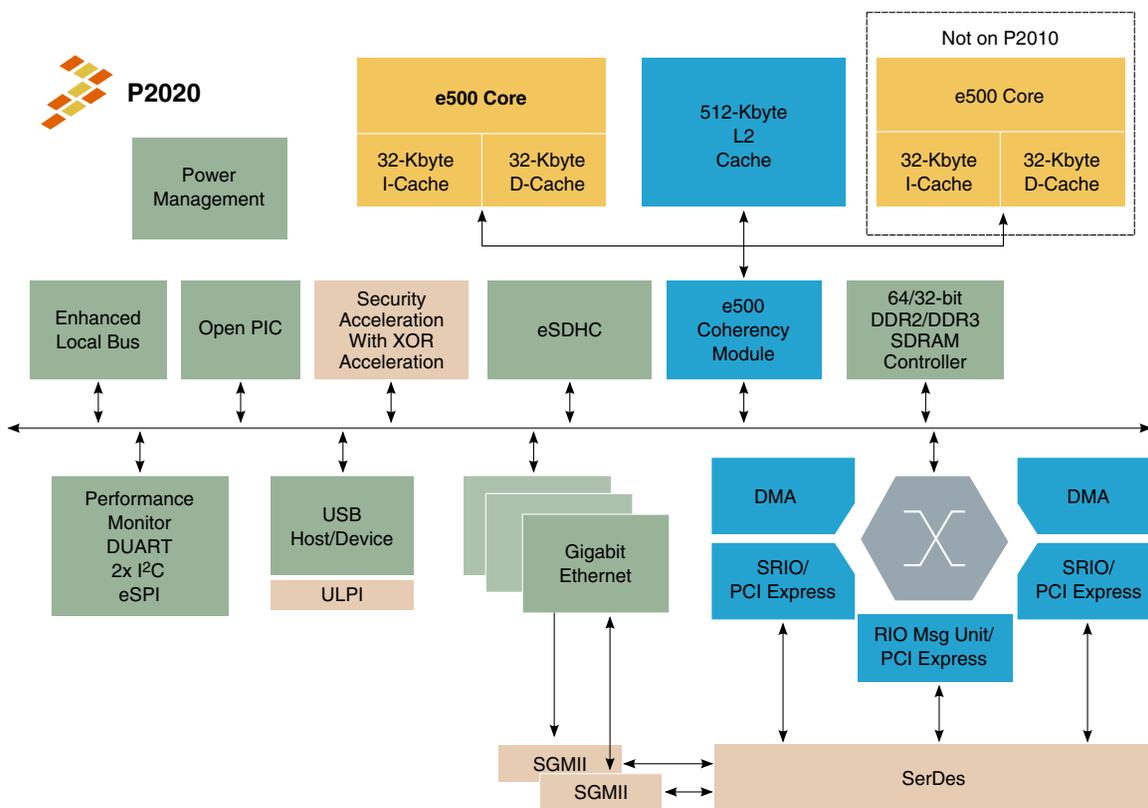


Figure 1-1. P2020/P2010 block diagram

1.1.2 Critical performance parameters

The critical performance parameters are as follows:

- e500v2 core frequency of up to 1.2 GHz
- 45-nm SOI process technology
- Data rate of up to 800 Mbps
- Support for DDR2 and DDR3 SDRAM
- Supply voltages:
 - Core: 1.05 V
 - PCI Express, serial RapidIO: 1.05 V
 - Ethernet: 2.5 or 3.3 V (subject to protocol)
 - Local bus, 8 bit of GPIO: 3.3, 2.5, or 1.8 V
 - DDR: 1.8 V for DDR2, 1.5 V for DDR3 (conforms to JEDEC standard)

- SGMII: 1.05 V
- SPI, eSDHC, USB: 3.3, 2.5, or 1.8 V
- Package: 689-pin WB-TEPBGA II (wire bond temperature-enhanced plastic BGA)

1.1.3 Chip-level features

Key features of the chip include:

- Dual (P2020) or single (P2010) high-performance Power Architecture e500v2 cores
 - 36-bit physical addressing
 - Double-precision floating-point support
 - 32-Kbyte L1 instruction cache and 32-Kbyte L1 data cache for each core
- 512-Kbyte L2 cache with ECC
- Three 10/100/1000 Mbps enhanced three-speed Ethernet controllers (eTSECs)
 - TCP/IP acceleration and classification capabilities
 - IEEE 1588 support (not supported in SGMII 10/100 mode)
- High-speed interfaces:
 - Two x1 or one x4 serial RapidIO interfaces with message unit
 - Three PCI Express interfaces:
 - Two x1, one x2
 - Two x2
 - One x4
 - Two SGMII interfaces
- High-speed USB interface (USB 2.0)
- Enhanced secure digital host controller (SD/MMC)
- Enhanced serial peripheral interface (eSPI)
- Integrated security engine
 - XOR acceleration
 - Protocol support includes SNOW, ARC4, 3DES, AES, RSA/ECC, RNG, single-pass SSL/TLS, Kasumi
- DDR2/DDR3 SDRAM memory controller with ECC support
- Power management controller
- Programmable interrupt controller (PIC) compliant with Open-PIC standard
- Two four-channel DMA controllers
- DUART, timers, and two I²C controllers
- Enhanced local bus controller (eLBC)
- 16 general-purpose I/O signals

These features are described in greater detail in subsequent sections.

NOTE

The chip is also available without a security engine. All specifications other than those relating to security apply to the chip exactly as described in this document.

1.2 Application examples

The following section provides block diagrams of sample applications.

The chip is very flexible and can be configured to meet many system application needs.

Target applications include:

- Networking (switches and routers) - Line card controller, mid-range line card control plane, low-end line card combined control and data plane, shelf controller, business gateway, multiservice router, wireless access points
- Telecommunications - AMC card, controller on ATCA carrier card, channel and control card for NodeB/BTS/WCDMA/4G LTE/WiMax, general-purpose compute blade
- Industrial - Robotics, multifunction printer, single board computers, industrial applications, test and measurement for networking/telecommunications

Both cores can operate in a symmetric multiprocessing mode to achieve higher performance, or they can perform separate tasks, optionally running independent operating systems. This flexibility enables application developers to assign distinct processing resources to distinct tasks that need guaranteed performance. For example, one core can manage a data plane and the other a control plane.

The LTE and WiMax baseband and line card control plane applications are described in the following sections.

1.2.1 LTE and WiMax baseband application

As an integrated dual-core device, the chip is well-suited for long term evolution (LTE) and WiMax channel card applications.

The chip offers dual-core performance without breaking out of a single core power budget. Additionally, in order to maximize subscriber bandwidth, the chip also offers support for the increased performance needed in Layer 2 baseband processing and implementing network interfaces.

This figure shows an example of a channel card application using the chip. The advanced quality of service (QoS) features of the eTSEC ports assist in resource scheduling and medium access control (MAC), implemented by Layer 2 baseband processing. For network interfacing, the chip supports both dual gigabit Ethernet on SGMII (for redundancy) and serial RapidIO. The serial RapidIO interface also allows direct connection to the digital signal processors (DSPs) that implement the layer 1 processing. The security module handles the secure network termination requirement.

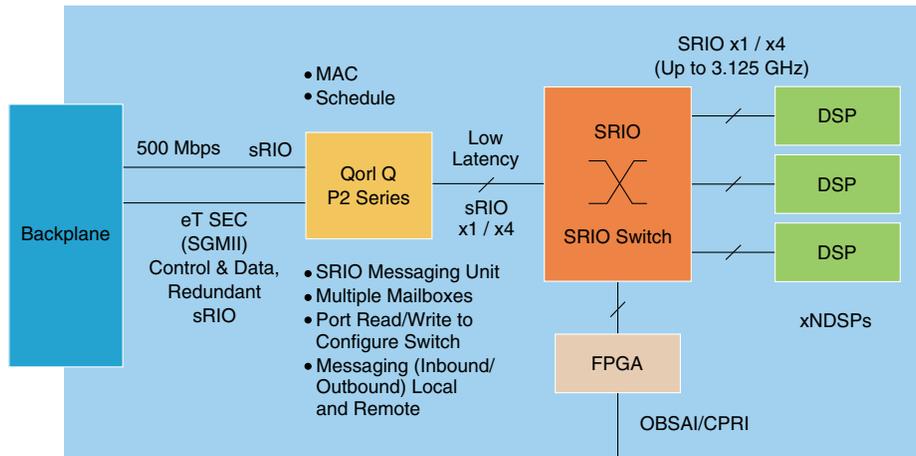


Figure 1-2. Channel card application

1.2.2 Line card control plane application

One of the primary applications for the chip is as a line card control plane processor, as shown in this figure.

Typically the chip sits to the side of the main data path, which has I/O processors on the front panel connected to a backplane interface device. The chip communicates with these components using standard interfaces, such as PCI Express or the local bus. The chip often has front panel interfaces as well. One eTSEC may be used for high-performance debug; the DUART can be used for a low-level command line interface; and the USB port may be used for front-panel code uploads when connecting directly to a PC for maintenance. Redundant gigabit Ethernet ports to the backplane are used as a management interface, which can be used to communicate with a centralized resource for receiving table updates, for instance. The SD/MMC interface enables modularity by supporting booting from Flash cards.

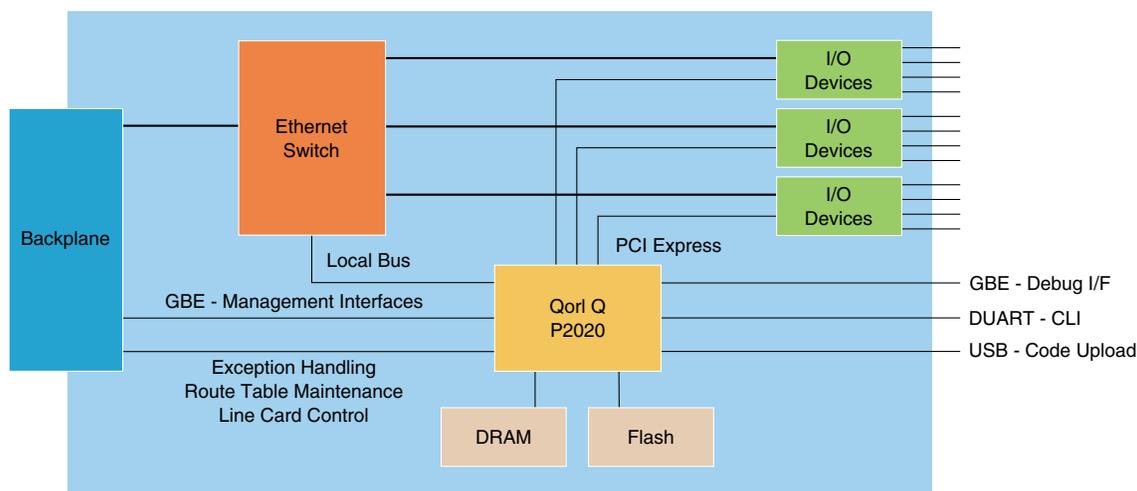


Figure 1-3. Line card control plane application

1.3 Architecture overview

This section contains a high-level view of the device architecture.

1.3.1 e500v2 cores and memory unit

The chip contains two high-performance 32-bit e500v2 cores that implement Power Architecture technology.

In addition to 36-bit physical addressing, this version of the e500 core includes:

- SPE double-precision floating-point instruction set using 64-bit operands
- SPE embedded vector and scalar single-precision floating-point instruction set using 32- or 64-bit operands
- 32-Kbyte L1 instruction cache and 32-Kbyte L1 data cache with parity protection

The device also contains 512 Kbytes of L2 cache/SRAM, as follows:

- 8-way set-associative cache organization with 32-byte cache lines
- Flexible configuration (can be configured as part cache, part SRAM)
- External masters can force data to be allocated into the cache through programmed memory ranges or special transaction types (stashing).
- SRAM features include the following:
 - I/O devices access SRAM regions by marking transactions as snoopable (global).
 - Regions can reside at any aligned location in the memory map.
 - Byte-accessible ECC uses read-modify-write transaction accesses for smaller-than-cache-line accesses.

1.3.2 e500 coherency module (ECM) and address map

The e500 coherency module (ECM) provides a mechanism for I/O-initiated transactions to snoop the bus between the e500v2 cores and the integrated L2 cache in order to maintain coherency across local cacheable memory.

It also provides a flexible switch-type structure for core- and I/O-initiated transactions to be routed or dispatched to target modules on the device.

The chip supports a flexible 36-bit physical address map. Conceptually, the address map consists of local space and external address space. The local address map is supported by twelve local access windows that define mapping within the local 36-bit (64-Gbyte) address space.

The chip can be made part of a larger system address space through the mapping of translation windows. This functionality is included in the address translation and mapping units (ATMUs). Both inbound and outbound translation windows are provided. The ATMUs allows the chip to be part of larger address maps such as those of PCI Express or RapidIO.

1.3.3 Integrated security engine (SEC)

The SEC is a modular and scalable security core optimized to process all the algorithms associated with IPsec, IKE, SSL/TLS, iSCSI, SRTP, IEEE 802.11i™, IEEE 802.16™ (WiMAX), IEEE 802.1AE™ Std. (MACSec), 3GPP, A5/3 for GSM and EDGE, and GEA3 for GPRS.

Although it is not a protocol processor, the SEC is designed to perform multi-algorithmic operations (for example, 3DES-HMAC-SHA-1) in a single pass of the data. The version of the SEC used in the chip is specifically capable of performing single-pass security cryptographic processing for SSL 3.0, SSL 3.1/TLS 1.0, IPsec, SRTP, and 802.11i.

SEC features include the following:

- Compatible with code written for the Freescale MPC8548E, MPC8555E, and MPC8541E devices
- XOR engine for parity checking in RAID storage applications
- Four crypto-channels, each supporting multi-command descriptor chains
- Cryptographic execution units:
 - PKEU-public key execution unit
 - DEU-data encryption standard execution unit
 - AESU-advanced encryption standard unit
 - AFEU-ARC four execution unit
 - MDEU-message digest execution unit
 - KEU-kasumi execution unit
 - CRCU-cyclical redundancy check unit
 - RNGU-random number generator unit
 - STEU-SNOW3G execution unit

See *P2020 Security (SEC 3.1) Reference Manual*, for more information.

1.3.4 Enhanced three-speed Ethernet controllers

The chip has three on-chip enhanced three-speed Ethernet controllers (eTSECs).

The eTSECs incorporate a media access control (MAC) sublayer that supports 10- and 100-Mbps and 1-Gbps Ethernet/802.3 networks with SGMII, GMII, RGMII, MII, RMII, TBI, and RTBI physical interfaces. The eTSECs include 2-Kbyte receive and 10-Kbyte transmit DMA functions.

The chip eTSECs support programmable CRC generation and checking, RMON statistics, and jumbo frames of up to 9.6 Kbytes. Frame headers and buffer descriptors can be forced into the L2 cache to speed classification or other frame processing. They are designed to comply with IEEE Std. 802.3™, 802.3u, 802.3x, 802.3z, 802.3ac, and 802.3ab.

The buffer descriptors are based on the MPC8540 three-speed Ethernet controller programming model. Each eTSEC can emulate a PowerQUICC III TSEC, allowing existing driver software to be re-used with minimal change.

Some of the key features of these controllers include:

- This table lists the flexible configuration for multiple PHY interface configurations.

Table 1-1. eTSEC configuration options¹

eTSEC1	eTSEC2	eTSEC3
<ul style="list-style-type: none"> • Standard interface² • None 	<ul style="list-style-type: none"> • Standard interface • SGMII • None 	<ul style="list-style-type: none"> • SGMII • None
<ul style="list-style-type: none"> • Reduced interface³ • None 	<ul style="list-style-type: none"> • Reduced interface • SGMII • None 	<ul style="list-style-type: none"> • SGMII • None
<ul style="list-style-type: none"> • None 	<ul style="list-style-type: none"> • SGMII • None 	<ul style="list-style-type: none"> • SGMII • None
<ul style="list-style-type: none"> • RGMII • RTBI • RMII • None 	<ul style="list-style-type: none"> • RGMII • RTBI • RMII • SGMII • None 	<ul style="list-style-type: none"> • RGMII • RTBI • RMII • SGMII • None

1. The parallel interfaces must use the same voltage.

2. Standard interfaces are GMII, TBI, MII

3. Reduced interfaces are RGMII, RTBI, RMII

- TCP/IP acceleration and QoS features:
 - IP v4 and IP v6 header recognition on receive
 - IP v4 header checksum verification and generation
 - TCP and UDP checksum verification and generation
 - Per-packet configurable acceleration
 - Recognition of VLAN, stacked (queue in queue) VLAN, 802.2, PPPoE session, MPLS stacks, and ESP/AH IP-security headers
 - Transmission from up to eight physical queues
 - Reception to up to eight physical queues
- Full- and half-duplex Ethernet support (1000 Mbps supports only full duplex):
 - IEEE Std. 802.3 full-duplex flow control (automatic PAUSE frame generation or software-programmed PAUSE frame generation and recognition)
- IEEE Std. 802.1 virtual local area network (VLAN) tags and priority

- VLAN insertion and deletion
 - Per-frame VLAN control word or default VLAN for each eTSEC
 - Extracted VLAN control word passed to software separately
- Programmable Ethernet preamble insertion and extraction of up to 7 bytes
- MAC address recognition
- Ability to force allocation of header information and buffer descriptors into L2 cache

1.3.5 Universal serial bus (USB) 2.0

The chip USB 2.0 controller provides point-to-point connectivity complying with the USB specification, Rev. 2.0.

The USB controller can be configured to operate as a stand-alone host or stand-alone device. The host and device functions are both configured to support the following types of USB transfers:

- Bulk
- Control
- Interrupt
- Isochronous

Other controller features are as follows:

- Supports USB dual-role operation and can be configured as host or device
- Complies with USB specification, Rev. 2.0
- Supports operation as a stand-alone USB device
 - Supports one upstream facing port
 - Supports six programmable USB endpoints
- Supports operation as a stand-alone USB host controller
 - Supports USB root hub with one downstream-facing port
 - Enhanced host controller interface (EHCI) compatible
- Supports high-speed (480 Mbps), full-speed (12 Mbps), and low-speed (1.5 Mbps) operations
- Supports external PHY with UTMI+ low-pin interface (ULPI)

1.3.6 Enhanced secure digital host controller

The enhanced secure digital host controller (eSDHC) provides an interface between the host system and SD/MMC cards.

The eSDHC acts as a bridge, passing host bus transactions to SD/MMC cards by sending commands and performing data accesses to or from the cards. It handles SD/MMC protocols at the transmission level. For the chip, booting from on-chip ROM is supported through the eSDHC, as described in [eSDHC boot](#).

The secure digital (SD) card is an evolution of old multimedia card (MMC) technology. The MMC is a universal low-cost data storage and communication media designed to cover a wide area of applications including mobile video and gaming, which are available from either pre-loaded MMC cards or downloadable from cellular phones, WLAN, or other wireless networks. Old MMC cards are based on a 7-pin serial bus with a single data pin, while the new high-speed MMC communication is based on an advanced 11-pin serial bus designed to operate in a low voltage range.

The SD card is specifically designed to meet the security, capacity, performance, and environment requirements inherent in emerging audio and video consumer electronic devices. The physical form factor, pin assignments, and data transfer protocol are forward compatible with the old MMC.

1.3.7 Enhanced serial peripheral interface (eSPI)

The enhanced serial peripheral interface (eSPI) allows the device to exchange data between other QorIQ family devices, Ethernet PHYs for configuration, and peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.

The SPI is a full-duplex, synchronous, character-oriented channel that supports a four-wire interface (receive, transmit, clock, and slave select). The SPI module consists of transmitter and receiver sections, an independent baud-rate generator, and a control unit. The chip also has the ability to boot from an SPI serial Flash device.

The SPI receiver and transmitter each have a FIFO of 32 bytes to support more efficient transfers to and from SPI devices. The SPI interface supports RapidS for Atmel devices as well as Winbond devices dual read commands; in this mode the SPI uses two bits in parallel for reads.

1.3.8 DDR SDRAM controller

The chip supports DDR2 and DDR3 SDRAM.

The memory interface controls main memory accesses and provides for a maximum of 32 Gbytes of main memory.

The device supports a variety of SDRAM configurations. SDRAM banks can be built using DIMMs or directly-attached memory devices. Sixteen multiplexed address signals provide for device densities from 64 Mbits to 8 Gbits. Four chip select signals support up to four banks of memory. The device supports bank sizes from 16 Mbytes to 8 Gbytes. Nine data mask (MDM[0:8]) are used to provide byte selection for memory bank writes.

The chip can be configured to retain the currently active SDRAM page for pipelined burst accesses. Page mode support of up to 32 simultaneously open pages can dramatically reduce access latencies for page hits. Depending on the memory system design and timing parameters, using page mode can save 3 to 4 clock cycles from subsequent burst accesses that hit in an active page.

Using ECC, the device detects and corrects all single-bit errors and detects all double-bit errors and all errors within a nibble.

The device can invoke a level of system power management by asserting the MCKE SDRAM signal on-the-fly to put the memory into a low-power sleep mode.

The chip offers both hardware and software options to support battery-backed main memory. In addition, the DDR controller offers an initialization bypass feature which system designers may use to prevent re-initialization of main memory during system power-on following abnormal shutdown.

1.3.9 High speed I/O interfaces

The chip supports the SGMII, serial RapidIO, and PCI Express high-speed I/O interface standards.

1.3.9.1 PCI Express interfaces

The chip supports three PCI Express interfaces that are compliant with the *PCI Express Base Specification Revision 1.0a*.

They are configurable at boot time to act as either root complex or endpoint.

The physical layer of the PCI Express interface operates at a transmission rate of 2.5 Gbaud (data rate of 2.0 Gbps) per lane. The theoretical unidirectional peak bandwidth is 2 Gbps per lane. Receive and transmit ports operate independently, resulting in an aggregate theoretical bandwidth of 4 Gbps per lane.

Other features of the PCI Express interface include:

- x4, x2, and x1 link widths supported

- Both 32- and 64-bit addressing and 256-byte maximum payload size
- Full 64-bit decode with 36-bit wide windows

1.3.9.2 Serial RapidIO interfaces

The two serial RapidIO interfaces are based on the *RapidIO Interconnect Specification, Revision 1.2*.

RapidIO is a high-performance, point-to-point, low-pin-count, packet-switched system-level interconnect that can be used in a variety of applications as an open standard. The RapidIO architecture has a rich variety of features including high data bandwidth, low-latency capability, and support for high-performance I/O devices, as well as support for message-passing and software-managed programming models. Key features of the serial RapidIO interface unit include:

- Support for *RapidIO Interconnect Specification, Revision 1.2* (all transaction flows and priorities)
- Both x1 and x4 LP-serial link interfaces, with transmission rates of 1.25, 2.5, or 3.125 Gbaud (data rates of 1.0, 2.0, or 2.5 Gbps) per lane
- Auto detection of x1 or x4 mode operation during port initialization
- 34-bit addressing and up to 256-byte data payload
- Receiver-controlled flow control
- Support for RapidIO error injection

The RapidIO messaging unit supports two inbox/outbox mailboxes (queues) for data and one doorbell message structure. Both chaining and direct modes are provided for the outbox, and messages can hold up to 16 packets of 256 bytes, or a total of 4 Kbytes.

The DMA engines do not support bypass of SRIO outbound ATMUs. This feature was supported on some legacy parts but was made obsolete by SRIO ATMU definition improvements, which now utilize segments and sub-segments.

1.3.9.3 SGMII

The serial gigabit media independent interface (SGMII) is a high-speed interface linking the Ethernet controller with an Ethernet PHY.

SGMII uses differential signaling for electrical robustness. Only four signals are required: receive data and its inverse, and send data and its inverse; no clock signals are required.

1.3.9.4 High-speed interface multiplexing

This table shows the supported high-speed interface configurations. The desired configuration must be selected at power-on reset.

Table 1-2. Supported high-speed interface combinations

cfg_IO_ports[0:3]	Lanes				Gbaud	
	0	1	2	3	0&1	2&3
0000	PCI Express1: x1	off	off	off	2.5	-
0001	off	off	off	off	-	-
0010	PCI Express1: x1	PCI Express 2: x1	PCI Express 3: x2		2.5	2.5
0011	Reserved				-	-
0100	PCI Express 1: x2		PCI Express 3: x2		2.5	2.5
0101	Reserved				-	-
0110	PCI Express 1: x4				2.5	-
0111	SRIO2: x1	SRIO1: x1	off	off	3.125	-
1000	SRIO2: x4				1.25	-
1001	SRIO2: x4				2.5	-
1010	SRIO2: x4				3.125	-
1011	SRIO2: x1	SRIO1: x1	SGMII2	SGMII3	1.25	1.25
1100	SRIO2: x1	SRIO1: x1	SGMII2	SGMII3	2.5	1.25
1101	PCI Express 1: x1	SRIO1: x1	SGMII2	SGMII3	2.5	1.25
1110	PCI Express 1: x1	PCI Express 2: x1	SGMII2	SGMII3	2.5	1.25
1111	PCI Express 1: x2		SGMII2	SGMII3	2.5	1.25

1.3.10 Programmable interrupt controller (PIC)

The chip PIC implements the logic and programming structures of the OpenPIC architecture, providing for external interrupts (with fully nested interrupt delivery), message interrupts, internal-logic driven interrupts, and global high-resolution timers.

Up to 16 programmable interrupt priority levels are supported. The PIC can be bypassed to allow use of an external interrupt controller.

1.3.11 DMA, I²C, DUART, and eLBC

The chip provides two integrated four-channel DMA controllers, which can transfer data between any of its I/O or memory ports or between two devices or locations on the same port.

The DMA controllers can be used as follows:

- To chain (both extended and direct) through local memory-mapped chain descriptors.
- To handle misaligned transfers as well as stride transfers and complex transaction chaining.
- To specify local attributes such as snoop and L2 write stashing.

There are two I²C controllers. These synchronous, multimaster buses can be connected to additional devices for expansion and system development.

The DUART supports full-duplex operation and is compatible with the PC16450 and PC16550D programming models. Both the transmitter and receiver support 16-byte FIFOs.

The enhanced local bus controller (eLBC) port allows connection with a wide variety of external memories, DSPs, and ASICs. Three separate state machines share the same external pins and can be programmed separately to access different types of devices. The general-purpose chip select machine (GPCM) controls accesses to asynchronous devices using a simple handshake protocol. The user programmable machine (UPM) can be programmed to interface to synchronous devices or custom ASIC interfaces. The NAND Flash control machine (FCM) further extends interface options. Each chip select can be configured so that the associated chip interface can be controlled by the GPCM, UPM, or FCM controller. All controllers can exist in the same system. The local bus controller supports the following features:

- Multiplexed 16-bit address and data bus operating at up to 150 MHz
- 32-bit address support
- Eight chip selects support eight external slaves
- Up to eight-beat burst transfers
- 16- and 8-bit port sizes controlled by on-chip memory controller
- Three protocol engines available on a per-chip-select basis
- Parity support
- Default boot ROM chip select with configurable bus width (8 or 16 bits)
- Zero-bus-turnaround (ZBT) RAM Support
- FCM supports NAND Flash, GPCM supports NOR Flash

1.3.12 Device boot locations

The chip may be configured to boot using one of the following interfaces:

- DDR2/DDR3 memory controller
- Any PCI Express interface
- Any Serial RapidIO interface

- Enhanced local bus interface (using the GPCM or FCM)
- SPI Flash
- SD/MMC Flash

1.3.13 Boot sequencer

The chip provides a boot sequencer that uses the I²C1 interface to access an external serial ROM and loads the data into the device's configuration registers.

The boot sequencer is enabled by a configuration pin sampled at the negation of the device's hardware reset signal. If enabled, the boot sequencer holds the processor cores in reset until the boot sequence is complete. If the boot sequencer is not enabled, the processor cores exit reset and fetches boot code in default configurations.

1.3.14 Power management

In addition to low-voltage operation and dynamic power management, which automatically minimizes power consumption of modules when they are idle, four power consumption modes are supported: full on, doze, nap, and sleep.

1.3.15 System performance monitor

The performance monitor facility supports eight 32-bit counters that can count up to 512 counter-specific events.

It supports duration and quantity threshold counting and a burstiness feature that permits counting of burst events with a programmable time between bursts.

Chapter 2

Memory Map

This chapter describes the mechanisms that define the device memory map—the local access windows (LAWs), the address translation and mapping units (ATMUs), and the configuration, control, and status registers (CCSRs).

2.1 Overview

This chapter describes the mechanisms that define the device memory map—the local access windows (LAWs), the address translation and mapping units (ATMUs), and the configuration, control, and status registers (CCSRs).

There are several address domains within the chip, including the following:

- Logical, virtual, and physical (real) address spaces within the e500 core
- Internal local address space
- Internal configuration, control, and status register (CCSR) address space
- External memory, I/O, and configuration address spaces of the PCI Express link
- External memory address space of serial RapidIO

The MMU in the e500 core handles translation of logical (effective) addresses, into virtual addresses, and ultimately to the physical addresses for the local address space. The MMU is described in the *PowerPC e500 Core Family Reference Manual* (E500CORERM).

The local address map refers to the physical 36-bit address space seen by the e500 core as it accesses memory and I/O space. The DMA engines also see this same local address map. All memory controlled by the DDR and enhanced local bus controllers exist in this address map, as do all memory-mapped configuration, control, and status registers (CCSRs). The local address map is defined by a set of twelve local access windows (LAWs). Each of these windows maps a region of the local address space to a specified target interface, such as the DDR controller, enhanced local bus controller, serial

RapidIO controller, PCI Express controller, or other targets. The internal configuration, control, and status registers (CCSRs) for all the functional blocks are located in the local memory space at a specific CCSR window.

If the target mapping performed by the local access windows directs the transaction to one of the external interfaces (as an outbound read or write), the transaction is then mapped into that interface's external address space by address translation and mapping unit (ATMU) windows associated with the external interface. Outbound ATMUs perform the mapping from the local 36-bit address space to the address space of external interface; inbound ATMU windows perform the address translation from the external address space to the local address space.

2.2 Configuration, control, and status registers

All of the memory-mapped configuration, control, and status registers (CCSRs) in this device are contained within a 1-Mbyte address region.

To allow for flexibility, the CCSR block is relocatable in the local address space.

The local address map location of the CCSR block is controlled by the configuration, control, and status base address register (CCSRBAR); see [Configuration, control, and status registers base address register \(Reset_CCSRBAR\)](#). The default value for CCSRBAR is 0x0_FFE0_0000 (or 4 Gbytes-9 Mbytes). No address translation is performed for CCSR space, so there is no associated translation address register. The CCSR window is always enabled with a fixed size of 1 Mbyte; no other attributes are attached, so there is no associated window attribute register.

NOTE

The CCSR window must not overlap a LAW that maps to the DDR controller. Otherwise, undefined behavior occurs.

An example of a top-level memory map with the default location of the configuration, control, and status registers is shown in this figure.

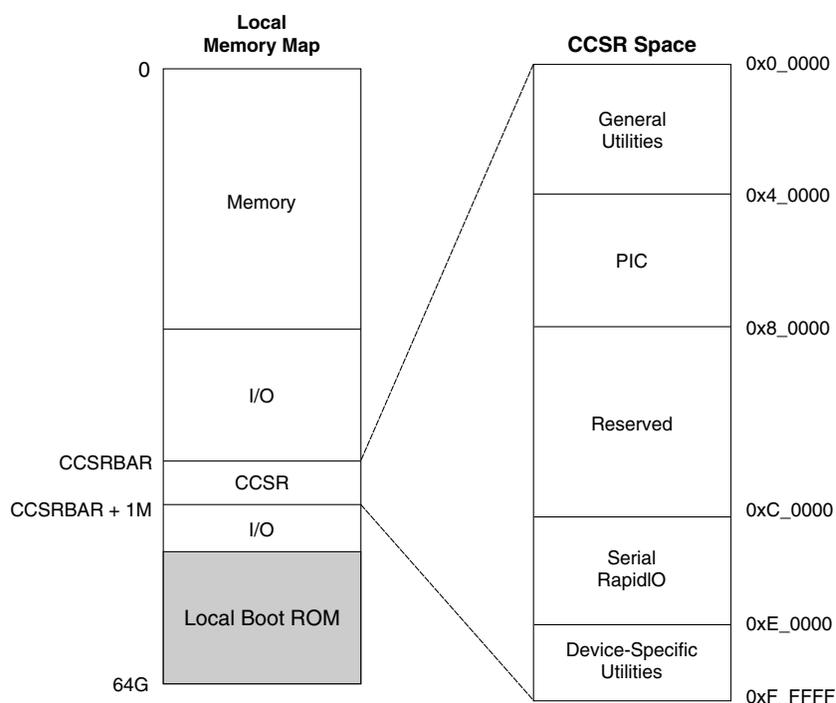


Figure 2-1. CCSR space

2.2.1 Accessing CCSR memory from the local processor

When the local e500 core is used to configure CCSR space, the CCSR memory space should typically be marked as cache-inhibited and guarded.

In addition, many configuration registers affect accesses to other memory regions; therefore writes to these registers must be guaranteed to have taken effect before accesses are made to the associated memory regions.

To guarantee that the results of any sequence of writes to configuration registers are in effect, the final configuration register write should be chased by a read of the same register, and that should be followed by a SYNC instruction. Then accesses can safely be made to memory regions affected by the configuration register write.

2.2.2 Accessing CCSR memory from external masters

In addition to being accessible by the e500 processor, the CCSRs are accessible from the external PCI Express interface.

This allows external masters on the I/O ports to configure the device.

External masters do not need to know the location of the CCSR memory in the local address map. Rather, they access the CCSR region of the local memory map through a window defined by a register in the interface's programming model that is accessible to the external master from its external memory map.

The PCI Express base address for accessing the local CCSR memory is selectable through the PCI Express configuration and status register base address register (PEXCSRBAR), at offset 0x10. An external PCI Express master sets this register by performing a PCI Express configuration cycle to this device. Subsequent memory accesses by a PCI Express master to the PCI Express address range indicated by PEXCSRBAR are translated to the local CCSR address indicated by the current setting of CCSRBAR.

2.2.3 Organization of CCSR space

As shown in above figure the CCSR space is divided into the following groups-general utilities, programmable interrupt controller (PIC), serial RapidIO, and device-specific utilities registers.

Starting at the block base address, the first 3 Kbytes are available for general registers. If the functional block has associated ATMUs, the next 512 bytes are dedicated to address translation and mapping registers. If a functional block has error management registers, they are typically placed starting at offset 0xE00 from the block base address, and any debug registers are typically placed in the final 256 bytes of the block's register space starting at offset 0xF00.

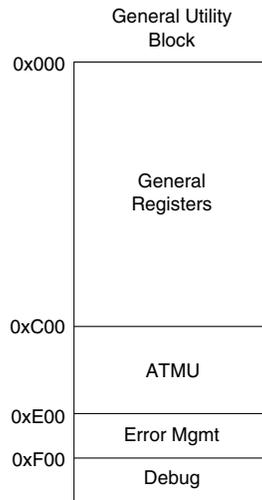


Figure 2-3. General utility register block

NOTE

Refer to detailed register descriptions for each functional block for exact locations, sizes, and access requirements.

2.2.3.2 Programmable interrupt controller registers

The programmable interrupt controller (PIC) follows the OpenPIC programming model which requires a larger register address space than the 4 Kbytes allocated to other blocks within the general utilities space.

For this reason, the PIC is allocated the second 256 Kbytes of CCSR space, beginning at offset 0x4_0000 from CCSRBAR.

The layout of the PIC register space is shown in this figure. Note that the PIC registers should only be accessed with 32-bit accesses.

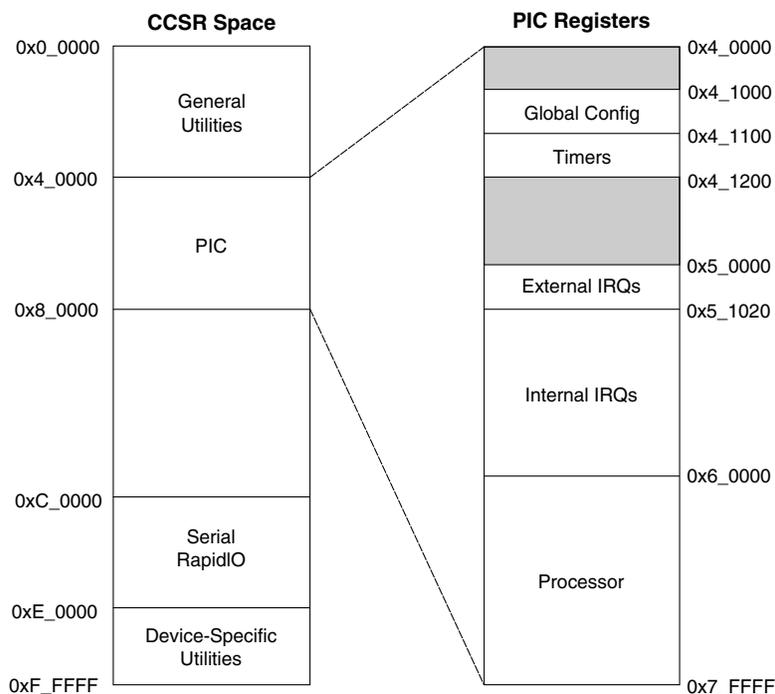


Figure 2-4. PIC register map within CCSR space

2.2.3.3 Serial RapidIO registers

The serial RapidIO module uses 128 Kbytes of CCSR space.

All registers are 32-bits wide and should only be accessed with 32-bit accesses. The 4-kbyte serial RapidIO implementation block has the same internal organization as that defined for the general utilities described above. The serial RapidIO memory space is shown in this figure.

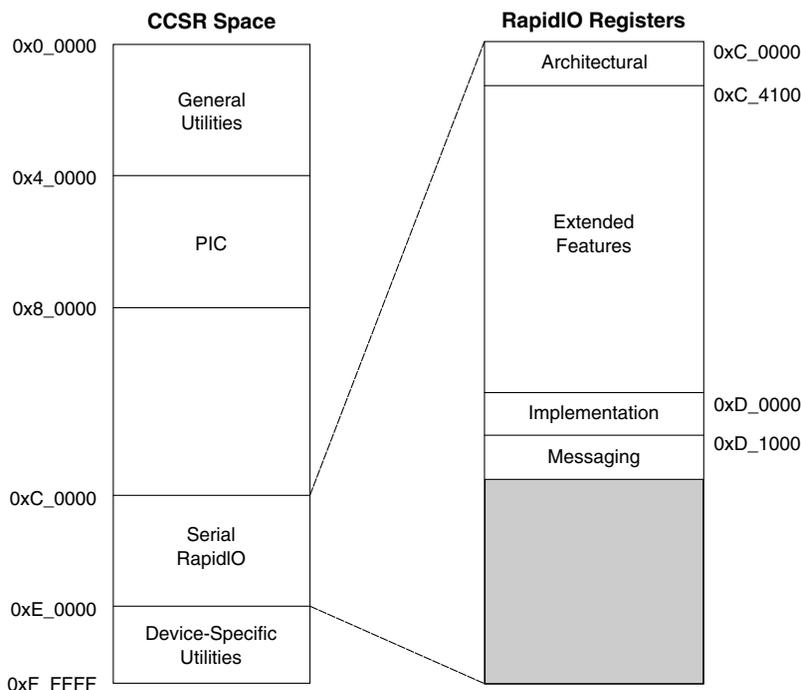


Figure 2-5. Serial RapidIO register map within CCSR space

2.2.3.4 Device-specific utilities registers

The device-specific utilities registers control functions that are not particular to a functional unit but to the device as a whole; they occupy the highest 256 Kbytes of CCSR space.

The device-specific utilities registers consist of power management, performance monitors, and device-wide debug utilities.

This figure shows the layout of the device-specific utilities registers. Note that the device-specific registers are accessible with 32-bit accesses only. Transactions of a size other than 32-bits are considered programming errors and the operation is undefined.

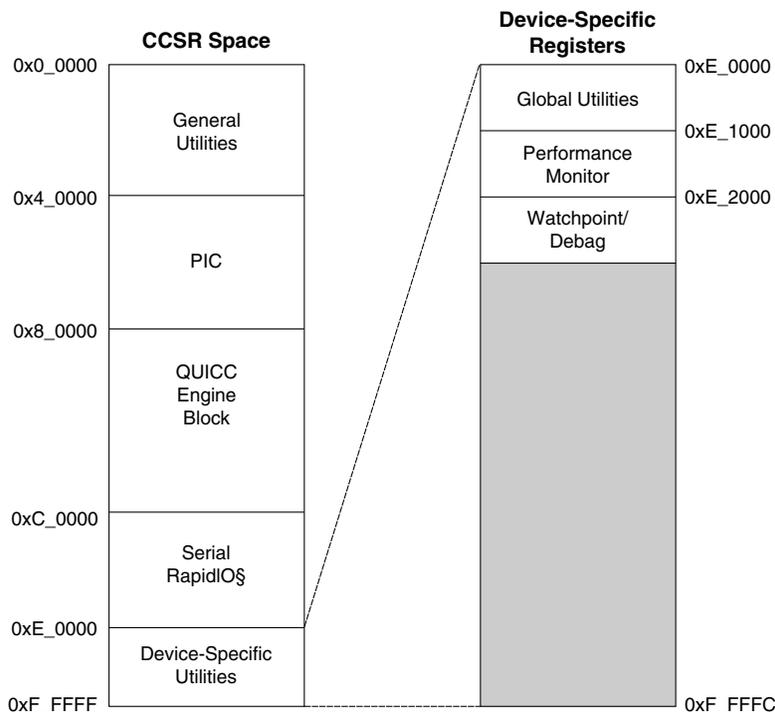


Figure 2-6. Device-specific register map within CCSR space

2.2.4 CCSR address map

The full register address of any CCSR is comprised of the CCSR window base address, specified in CCSRBAR (default address 0x0_FFE0_0000), plus the functional block base address, plus the specific register's offset within that block.

This table shows the location of the functional block base addresses for the entire CCSR space. Cross-references are provided to the CCSR maps for each individual block.

Table 2-1. CCSR block base address map

Block base address (Hex)	Block	Section/page	Comments
General utilities (0x0_0000-0x3_FFFF)			
0x0_0000	Local configuration control	Reset Memory Map/Register Definition	-
	Local access	Local access window registers	-

Table continues on the next page...

Table 2-1. CCSR block base address map (continued)

Block base address (Hex)	Block	Section/page	Comments
0x0_1000	ECM (e500 coherency module)	ECM memory map/register definition	-
0x0_2000	DDR memory controller	DDR memory map/register definition	-
0x0_3000	I ² C controllers (2)	I ² C memory map/register definition	I ² C controller 1: 0x0_3000 I ² C controller 2: 0x0_3100
0x0_4000	DUART	DUART memory map/register definition	UART0: 0x0_4500 UART1: 0x0_4600
0x0_5000	Enhanced local bus controller (eLBC)	Enhanced Local Bus Controller (eLBC) Memory Map	-
0x0_6000- 0x0_6FFF	Reserved		
0x0_7000	Enhanced serial peripheral interface (eSPI)	Enhanced serial peripheral interface (eSPI) memory map	-
0x0_8000	PCI Express controller 3	PCI Express memory-mapped registers	-
0x0_9000	PCI Express controller 2	PCI Express memory-mapped registers	-
0x0_A000	PCI Express controller 1	PCI Express memory-mapped registers	-
0x0_B000- 0x0_BFFF	Reserved		
0x0_C000	DMA controller 2	DMA controller memory map	DMA0: 0x0_C100 DMA1: 0x0_C180 DMA2: 0x0_C200 DMA3: 0x0_C280 General status: 0x0_C300
0x0_D000- 0x0_EFFF	Reserved		
0x0_F000	GPIO	GPIO memory map/register definition	-
0x1_0000- 0x1_FFFF	Reserved		
0x2_0000	L2 Cache	L2_Cache memory map/register definition	-

Table continues on the next page...

Table 2-1. CCSR block base address map (continued)

Block base address (Hex)	Block	Section/page	Comments
0x2_1000	DMA controller 1	DMA controller memory map	DMA0: 0x2_1100 DMA1: 0x2_1180 DMA2: 0x2_1200 DMA3: 0x2_1280 General Status: 0x2_1300
0x2_2000	USB DR controller	USB memory map/register definition	-
0x2_3000- 0x2_3FFF	Reserved		
0x2_4000	eTSEC 1	eTSEC Memory Map/Register Definition	-
0x2_5000	eTSEC 2	eTSEC Memory Map/Register Definition	-
0x2_6000	eTSEC 3	eTSEC Memory Map/Register Definition	-
0x2_7000- 0x2_DFFF	Reserved		
0x2_E000	eSDHC (SD/MMC)	Enhanced Secure Digital Host Controller (eSDHC) Memory Map	-
0x2_F000- 0x2_FFFF	Reserved		
0x3_0000	Security engine	See <i>P2020 Security (SEC 3.1) Reference Manual</i>	-
Programmable interrupt controller (PIC) (0x4_0000-0x6_FFFF)			
0x4_0000-0x6_FFFF	PIC	PIC memory map/register definition	Global registers: 0x4_0000 Interrupt source registers: 0x5_0000 Processor (per-CPU) registers: 0x6_0000
0x6_0000	PIC-processor (per-CPU) registers	PIC memory map/register definition	-
0x7_0000- 0xB_FFFF	Reserved		
Serial RapidIO (0xC_0000-0xD_FFFF)			
0xC_0000	RapidIO architectural registers	Serial RapidIO Memory Map/Register Definition	-

Table continues on the next page...

Table 2-1. CCSR block base address map (continued)

Block base address (Hex)	Block	Section/page	Comments
0xD_0000	RapidIO implementation registers	Serial RapidIO Memory Map/ Register Definition	-
Device-specific utilities (0xE_0000-0xF_FFFF)			
0xE_0000	Global utilities	GUTS Memory Map/Register Definition	-
0xE_1000	Performance monitor	PERFMON Memory Map/ Register Definition	-
0xE_2000	Debug/watchpoint monitor and trace buffer	Debug Memory Map/Register Definition	-
0xE_3000	SerDes control	GUTS Memory Map/Register Definition	-
0xE_3100- 0xE_FFFF	Reserved		
0xF_E000	Internal boot ROM	-	-

2.3 Local access windows

The local address map is defined by a set of twelve local access windows (LAWs).

Each of these windows maps a programmable 4-Kbyte to 32-Gbyte region of the local 36-bit address space to a specified target interface, such as the DDR controller, eLBC, serial RapidIO controller, PCI Express controllers, or other targets. This allows the internal interconnections of the device to route a transaction from its source to the proper target.

Each LAW is defined by a base address register which specifies the starting address for the window, and an attribute register which specifies whether the mapping is enabled, the size of the window, and the target interface for that window. Note that the LAWs do not perform any address translation, and therefore there are no corresponding translation address registers. The local access window registers exist as part of the local access block in the general utilities registers in CCSR space.

With the exception of configuration space (mapped by CCSRBAR) and the default boot ROM space, all addresses used by the system must be mapped by an LAW. This includes addresses that are mapped by inbound ATMU windows. Thus, target mappings of the local access windows and the inbound ATMU windows must be consistent.

2.3.1 Precedence of local access windows

If two local access windows overlap, the lower numbered window takes precedence.

For instance, if two windows are set up as shown in this table, local access window 1 governs the mapping of the 1-Mbyte region from 0x0_7FF0_0000 to 0x0_7FFF_FFF, even though the window described in local access window 2 also encompasses that memory region.

Table 2-2. Overlapping local access windows

Window	Base address	Size	Target interface
1	0x0_7FF0_0000	1 Mbyte	0b0100 (Local bus controller-LBC)
2	0x0_0000_0000	2 Gbytes	0b1111 (DDR controller)

2.3.2 Configuring local access windows

Once a local access window is enabled, it should not be modified while any device in the system may be using the window.

Neither should a new window be used until the effect of the write to the window is visible to all blocks that use the window. This can be guaranteed by completing a read of the last LAW configuration register before enabling any other devices to use the window. For example, if LAWs 0-3 are being configured in order during the initialization process, the last write (to LAWAR3) should be followed by a read of LAWAR3 before any devices try to use any of these windows. If the configuration is being performed by the e500 core, the read of LAWAR3 should be followed by an **isync** instruction.

2.3.3 Distinguishing local access windows from other mapping functions

It is important to distinguish between the mapping function performed by the LAWs and the additional mapping functions that occur at the target interfaces.

The LAWs define how a transaction is routed through the device's internal interconnects from the transaction's source to its target. After the transaction has arrived at its target interface, that interface controller may perform additional mapping. For instance, the DDR controller has chip select registers that map a memory request to a particular external device. Similarly, the enhanced local bus controller has base registers that perform a similar function. The serial RapidIO and PCI Express interfaces have outbound address translation and mapping units (ATMUs) that map the local address into an external address space.

These other mapping functions are configured by programming the CCSRs of the individual interfaces. Note that there is no need to have a one-to-one correspondence between LAWs and chip select regions or outbound ATMU windows. A single LAW can be further decoded to any number of chip selects or to any number of outbound ATMU windows at the target interface.

2.3.4 Illegal interaction between local access windows and DDR chip selects

If a local access window maps an address to an interface other than the DDR controller, there should not be a valid chip select configured for the same address in the DDR controller.

Because DDR chip select boundaries are defined by a beginning and ending address, it is easy to define them so that they do not overlap with LAWs that map to other interfaces.

2.3.5 Local address map example

This figure shows what a typical local address map might look like.

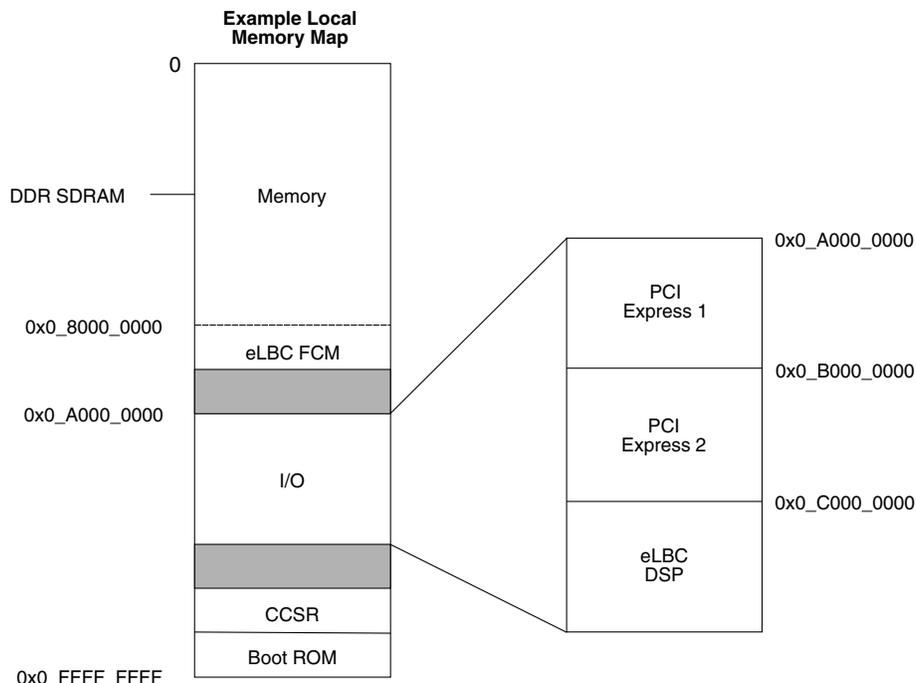


Figure 2-7. Local address map example

This table shows the corresponding set of LAW settings for the example shown in the figure above.

Table 2-3. Local access window settings example

Window	Base address	SIZE	Target interface	TRGT
0	0x0_0000_0000	2 Gbytes	DDR controller	0b01111
1	0x0_8000_0000	1 Mbyte	Enhanced local bus controller (eLBC)-FCM	0b00100
2	0x0_A000_0000	256 Mbytes	PCI Express 1	0b00010
3	0x0_B000_0000	256 Mbytes	PCI Express 2	0b00011
4	0x0_C000_0000	256 Mbytes	Enhanced local bus controller (eLBC)-DSP	0b00100
5-9	Unused			

In this example, it is not necessary to use a LAW to specify the location of the boot ROM. See [Boot page translation](#) for more information.

2.4 Local access window registers

The local access window registers are accessed by reading and writing to an address comprised of the base address (specified in the CCSRBAR), plus the block base address, plus the offset of the specific register to be accessed. For the LAWs, the block base address is 0x0_0000.

Note that all LAW registers should only be accessed a word (4-bytes) at a time. This table shows the memory map for the LAW registers.

LAW memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C08	Local access window 0 base address register (LAW_LAWBAR0)	32	R/W	0000_0000h	2.4.1/99
C10	Local access window 0 attribute register (LAW_LAWAR0)	32	R/W	0000_0000h	2.4.2/99
C28	Local access window 0 base address register (LAW_LAWBAR1)	32	R/W	0000_0000h	2.4.1/99
C30	Local access window 0 attribute register (LAW_LAWAR1)	32	R/W	0000_0000h	2.4.2/99
C48	Local access window 0 base address register (LAW_LAWBAR2)	32	R/W	0000_0000h	2.4.1/99
C50	Local access window 0 attribute register (LAW_LAWAR2)	32	R/W	0000_0000h	2.4.2/99
C68	Local access window 0 base address register (LAW_LAWBAR3)	32	R/W	0000_0000h	2.4.1/99
C70	Local access window 0 attribute register (LAW_LAWAR3)	32	R/W	0000_0000h	2.4.2/99
C88	Local access window 0 base address register (LAW_LAWBAR4)	32	R/W	0000_0000h	2.4.1/99
C90	Local access window 0 attribute register (LAW_LAWAR4)	32	R/W	0000_0000h	2.4.2/99
CA8	Local access window 0 base address register (LAW_LAWBAR5)	32	R/W	0000_0000h	2.4.1/99
CB0	Local access window 0 attribute register (LAW_LAWAR5)	32	R/W	0000_0000h	2.4.2/99
CC8	Local access window 0 base address register (LAW_LAWBAR6)	32	R/W	0000_0000h	2.4.1/99
CD0	Local access window 0 attribute register (LAW_LAWAR6)	32	R/W	0000_0000h	2.4.2/99
CE8	Local access window 0 base address register (LAW_LAWBAR7)	32	R/W	0000_0000h	2.4.1/99
CF0	Local access window 0 attribute register (LAW_LAWAR7)	32	R/W	0000_0000h	2.4.2/99
D08	Local access window 0 base address register (LAW_LAWBAR8)	32	R/W	0000_0000h	2.4.1/99
D10	Local access window 0 attribute register (LAW_LAWAR8)	32	R/W	0000_0000h	2.4.2/99
D28	Local access window 0 base address register (LAW_LAWBAR9)	32	R/W	0000_0000h	2.4.1/99
D30	Local access window 0 attribute register (LAW_LAWAR9)	32	R/W	0000_0000h	2.4.2/99

Table continues on the next page...

LAW memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D48	Local access window 0 base address register (LAW_LAWBAR10)	32	R/W	0000_0000h	2.4.1/99
D50	Local access window 0 attribute register (LAW_LAWAR10)	32	R/W	0000_0000h	2.4.2/99
D68	Local access window 0 base address register (LAW_LAWBAR11)	32	R/W	0000_0000h	2.4.1/99
D70	Local access window 0 attribute register (LAW_LAWAR11)	32	R/W	0000_0000h	2.4.2/99

2.4.1 Local access window 0 base address register (LAW_LAWBAR_n)

The LAWBAR *n* registers define the 24 high-order address bits that fixes the location of each window in the local address space. Note that the minimum size of any LAW is 4 Kbytes, so the 12 lowest-order bits of the base address cannot be specified.

Address: 0h base + C08h offset + (32d × *i*), where *i*=0d to 11d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

LAW_LAWBAR_n field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–31 BASE_ADDR	Identifies the 24 most significant address bits of the base of local access window <i>n</i> . The specified base address must be aligned to the window size, as defined by LAWAR <i>n</i> [SIZE].

2.4.2 Local access window 0 attribute register (LAW_LAWAR_n)

The LAWAR *n* registers are used to enable specific local access windows, define their size and specify the target interface.

The target interface for each LAW is specified using the encodings shown in [Table 2-7](#). Note that configuration registers are mapped by the windows defined by CCSRBAR. The CCSR mapping supersedes local access window mappings, so configuration registers do not appear as a target for local access windows.

Table 2-7. Target Interface Encodings

TRGT	Target Interface
00000	PCI Express 3
00001	PCI Express 2
00010	PCI Express 1
00011	Reserved
00100	Enhanced Local Bus (eLBC)
00101	Reserved
00110	Reserved
00111	Reserved
01000	Reserved
01001	Reserved
01010	Reserved
01011	Reserved
01100	SRIO 1
01101	SRIO 2
01110	Reserved
01111	DDR memory controller

Address: 0h base + C10h offset + (32d × i), where i=0d to 11d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	EN	Reserved						TRGT					Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved											SIZE				
W	Reserved											SIZE				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

LAW_LAWAR_n field descriptions

Field	Description
0 EN	Enable 0 The local access window <i>n</i> (and all other LAWAR <i>n</i> and LAWBAR <i>n</i> fields) are disabled. 1 The local access window <i>n</i> is enabled and other LAWAR <i>n</i> and LAWBAR <i>n</i> fields combine to identify an address range for this window.
1–6 -	This field is reserved. Reserved
7–11 TRGT	Identifies the target interface when a transaction hits in the address range defined by this window. The encodings for TRGT are defined in Table 2-7 .
12–25 -	This field is reserved. Reserved
26–31 SIZE	Identifies the size of the window from the starting address. Window size is 2 ^(SIZE+1) bytes. Patterns not shown are reserved.

Table continues on the next page...

LAW_LAWAR_n field descriptions (continued)

Field	Description
001011	4 Kbytes
001100	8 Kbytes
001101	16 Kbytes
...	$2^{(\text{size}+1)}$ bytes
100010	32 Gbytes

2.5 Address translation and mapping units

To facilitate flexibility in defining the address maps for external interfaces such as serial RapidIO and PCI Express, the device provides address translation and mapping units (ATMUs).

The following types of translation and mapping operations are performed by the ATMUs:

- Translating the local 36-bit address to an external address space
- Translating external addresses to the local 36-bit address space
- Assigning attributes to transactions
- Mapping a local address to a target interface

Outbound address translation and mapping refers to the translation of addresses from the local 36-bit address space to the external address space and attributes of a particular I/O interface.

Inbound address translation and mapping refers to the translation of an address from the external address space of an I/O interface to the local address space understood by the internal interfaces of this device. It also refers to the mapping of transactions to a particular target interface and the assignment of transaction attributes. Note that in mapping the transaction to the target interface, an inbound ATMU window performs a function similar to the local access windows. The target mappings created by an inbound ATMU must be consistent with those of the LAWs. That is, if an inbound ATMU maps a transaction to a given local address and a given target, a LAW must also map that same local address to the same target.

2.5.1 Address translation

All of the configuration registers that define an ATMU window's translation and mapping functions follow the same general register format, summarized in this table.

Table 2-43. Format of ATMU window definitions

Register	Function
Translation address (TAR)	High-order address bits defining location of the window in the target address space
Base address (BAR)	High-order address bits defining location of the window in the initiator address space
Window attributes (WAR)	Window enable, window size, target interface, and transaction attributes

The size of the windows must be a power-of-two. To perform a translation or mapping function, the address of the transaction is compared with the base address register of each window. The number of bits used in the comparison is dictated by each window's size attribute. When an address hits a window, if address translation is being performed, the new translated address is created by concatenating the window offset to the translation address. Again, the windows size attribute dictates how many bits are translated.

2.5.2 Outbound ATMUs

If the target mapping performed by the local access windows directs the transaction to one of the external interfaces (as an outbound read or write), the transaction is then mapped into that interface's external address space by outbound ATMUs associated with the external interface. The outbound ATMUs perform the mapping from the local 36-bit address space to the address spaces of PCI Express and serial RapidIO, which may be much larger than the local space. The outbound ATMUs also map attributes such as transaction type and priority level.

Each serial RapidIO controller has eight outbound ATMU windows plus a default window. If a transaction's address does not hit any of the eight outbound ATMU windows, the translation attributes defined by the default window are used. The default window is always enabled. See [Port 1 RapidIO outbound window translation address register n \(SRIO_P1ROWTAR_n\)](#) for a detailed description of the serial RapidIO outbound ATMU windows.

The PCI Express controller has four outbound ATMU windows plus a default window. If a transaction's address does not hit any of the four outbound ATMU windows, the translation attributes defined by the default window are used. The default window is always enabled. The PCI Express outbound ATMUs include extended translation address

registers so that up to 64 bits of external address space can be supported. See [PCI Express outbound ATMUs](#) for a detailed description of the PCI Express outbound ATMU windows.

2.5.3 Inbound ATMUs

The inbound ATMUs perform the address translation from external address spaces to the local address space, attach attributes and transaction types to the transaction, and also map the transaction to its target interface.

Each serial RapidIO controller has four inbound ATMU windows plus a default. If the inbound transaction's address does not hit any of the four inbound ATMU windows, the translation attributes defined by the default window are used. The default window is always enabled. See [Port 1 RapidIO outbound window translation address register n \(SRIO_PIROWTAR_n\)](#) for a detailed description of the serial RapidIO inbound ATMU windows.

The PCI Express controller has three general inbound ATMU windows plus a default. See [PCI Express inbound ATMUs](#) for a detailed description of the PCI-Express inbound ATMU windows.

2.5.3.1 Illegal interaction between inbound ATMUs and LAWs

Since both local access windows and inbound ATMUs map transactions to a target interface, it is essential that they not contradict one another.

For example, it is considered a programming error to have an inbound ATMU map a transaction to the DDR memory controller (target interface 0b0_1111) if the resulting translated local address is mapped to the PCI Express interface (target interface 0b0_0000) by a local access window. Such programming errors may result in unpredictable system deadlocks.



Chapter 3

Signal Descriptions

This chapter describes the external signals.

3.1 Signals overview

This chapter describes the external signals.

It is organized into the following sections:

- Overview of signals and cross-references for signals that serve multiple functions
- List of reset configuration signals
- List of output signal states at reset

NOTE

A bar over a signal name or a signal name followed by '_B' indicates that the signal is active low, such as $\overline{\text{IRQ_OUT}}$ or IRQ_OUT_B (interrupt output). Active-low signals are referred to as asserted (active) when they are low and negated when they are high. Signals that are not active low, such as IRQ (interrupt input), are referred to as asserted when they are high and negated when they are low.

Internal signals are shown throughout this document as lower case and in italics. For example, *sys_logic_clk* is an internal signal. These are discussed only as necessary for understanding the external functionality of the device.

The signals are grouped as follows:

- DDR memory interface signals
- PCI Express/SGMII/serial RapidIO interface signals
- eTSEC interface signals
- Enhanced local bus interface signals

Signals overview

- eSDHC interface signals
- eSPI interface signals
- USB interface signals
- DMA interface signals
- PIC interface signals
- GPIO interface signals
- System control, general-purpose input/output, power management, and debug signals
- DUART interface signals
- I²C interface signals
- Test, JTAG, configuration, and clock signals

The figures below illustrate the external signals of the chip, showing how the signals are grouped. The device hardware specification has a pinout diagram showing pin numbers and a listing of all the electrical and mechanical specifications. Note that these figures show multiplexed signals multiple times.

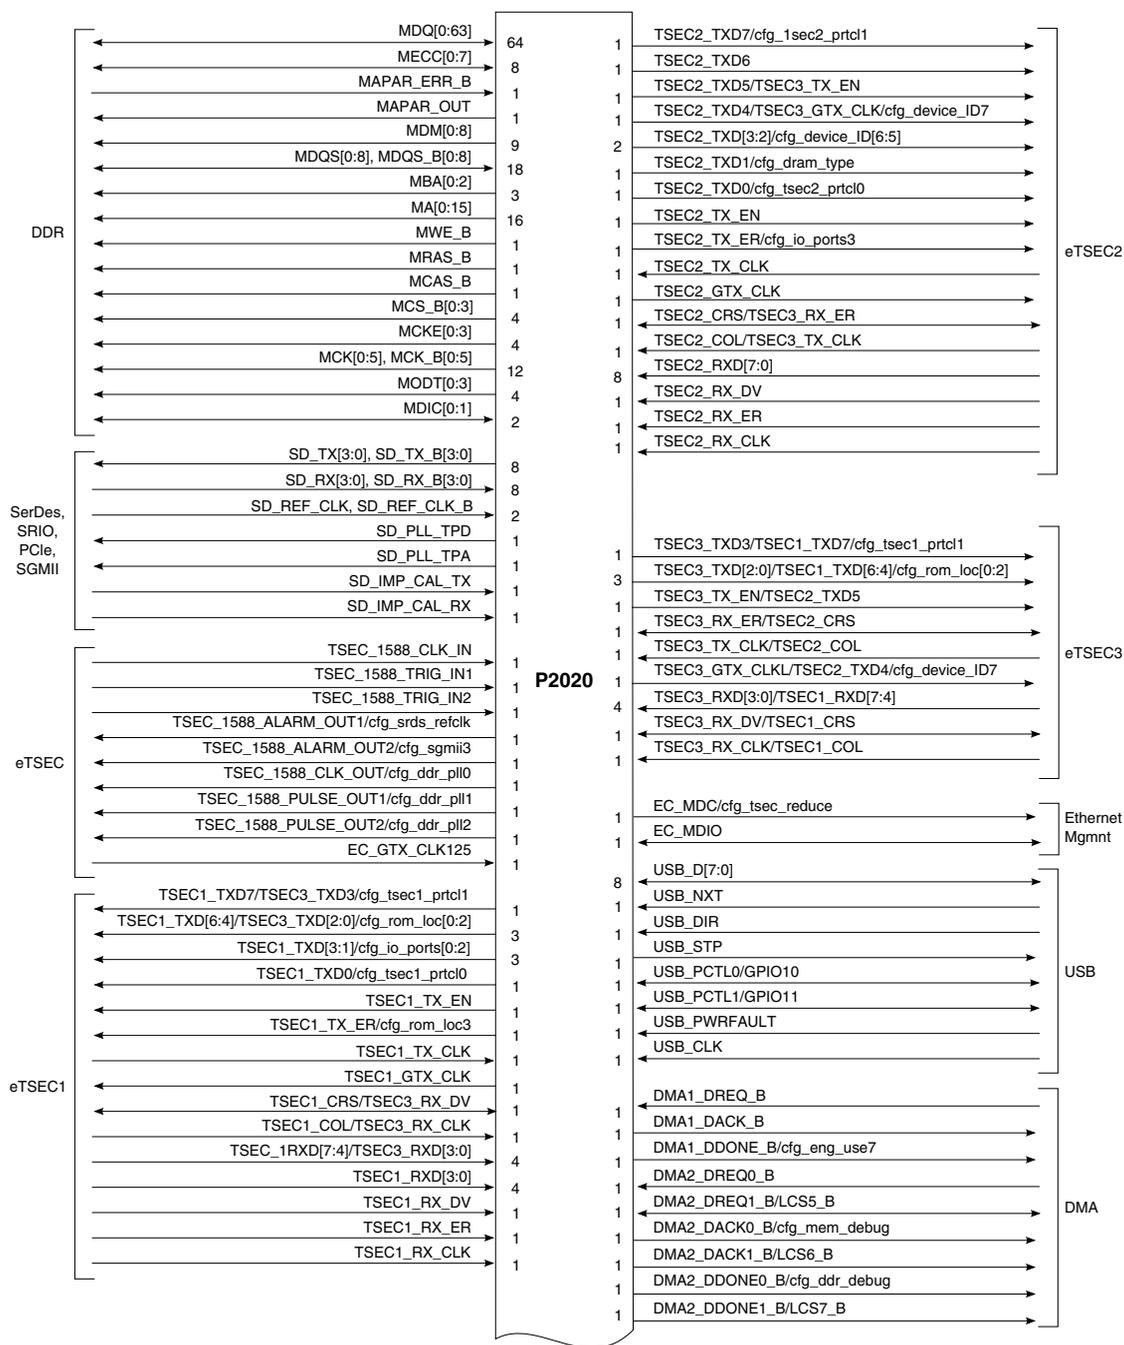


Figure 3-1. P2020 Signal Groupings (1/3)

Figure 3-1. P2020 signal groupings (1/3)

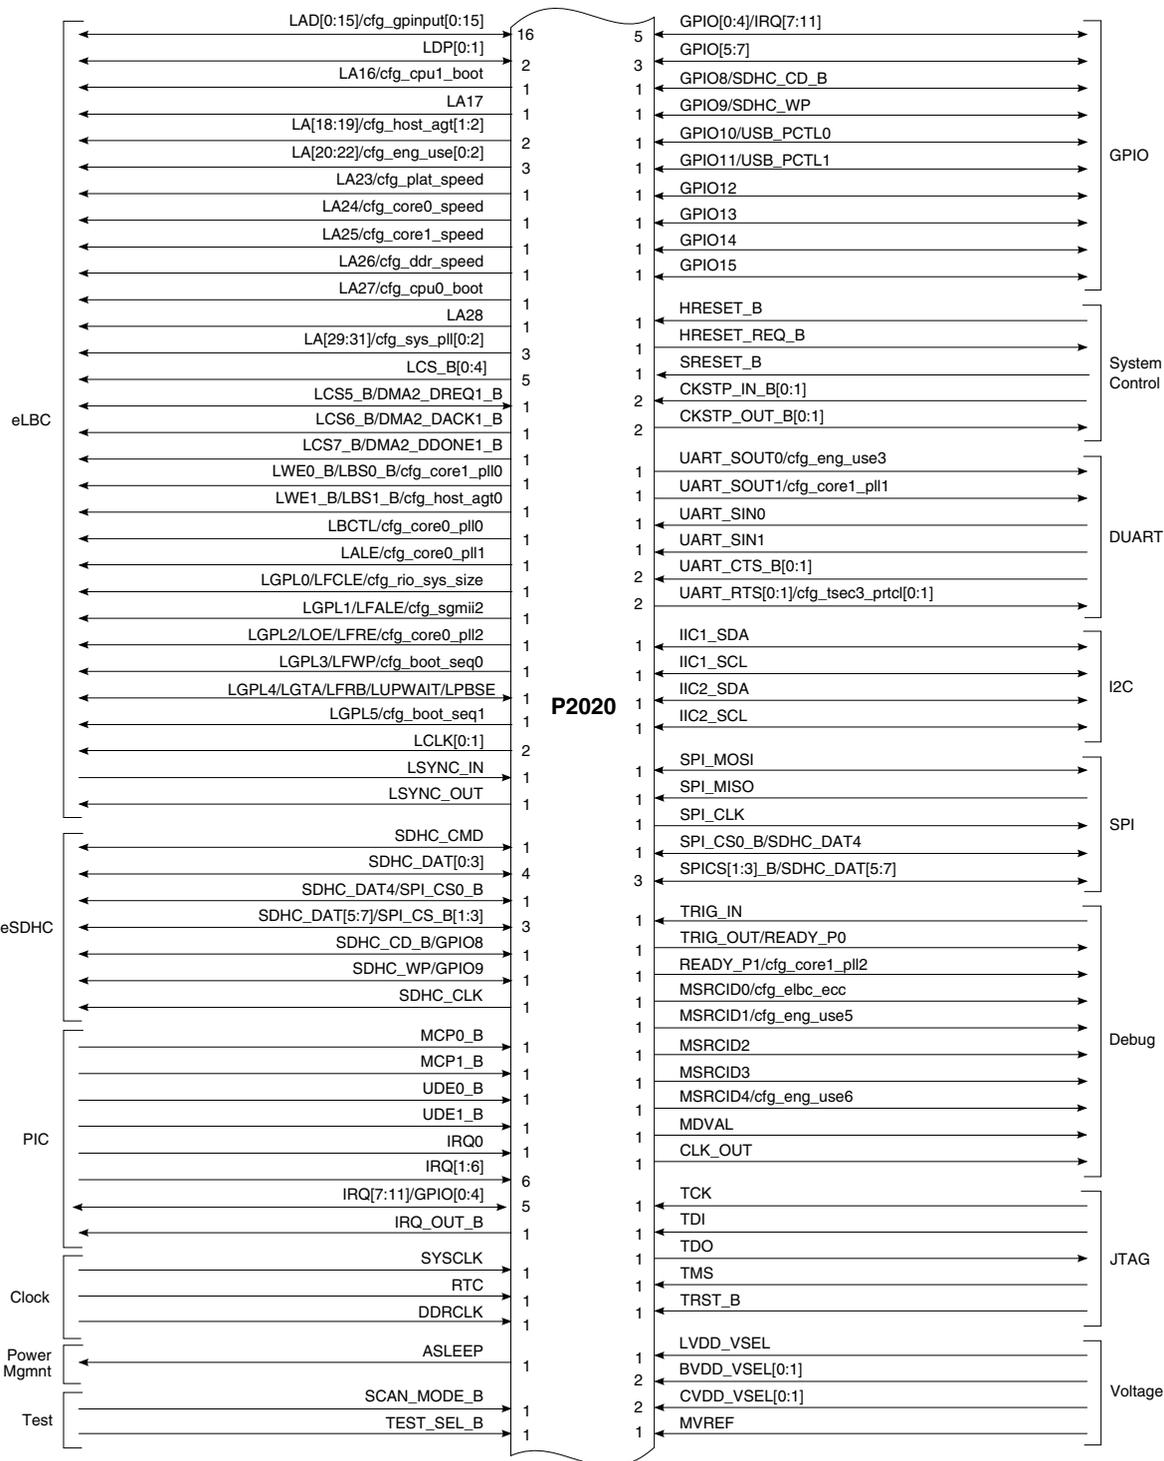


Figure 3-2. P2020 Signal Groupings (2/3) (Continued)

Figure 3-2. P2020 signal groupings (2/3) (Continued)

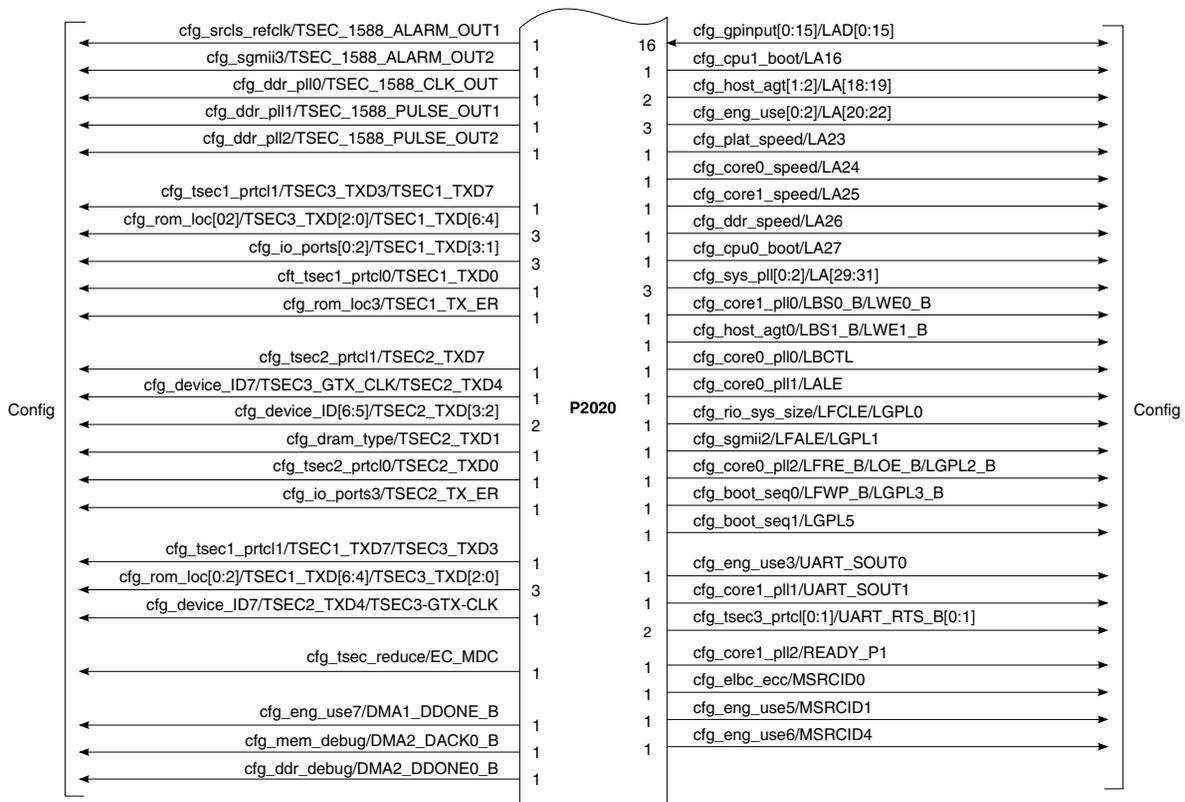


Figure 3-3. P2020 Signal Groupings (3/3) (Continued)

Figure 3-3. P2020 signal groupings (3/3) (Continued)

Note that individual chapters of this document provide details for each signal, describing each signal's behavior when the signal is asserted or negated and when the signal is an input or an output.

The following table provides summaries of signals grouped by function. This table detail the signal name, description, alternate functions, number of signals, and whether the signal is an input, output, or bidirectional. The direction in the right-most column applies to the the multiplexed signal of that row (and does not apply for the state of the reset configuration signals). The table also provides a pointer to the table where the signal function is described.

Table 3-1. Signal reference by functional block

Name	Description	Alternate function(s)	No. of Signals	I/O
DDR Controller				
See DDR external signal descriptions for more details				

Table continues on the next page...

Table 3-1. Signal reference by functional block (continued)

Name	Description	Alternate function(s)	No. of Signals	I/O
MDQ[0:63]	DDR data	-	64	I/O
MECC[0:7]	DDR error correcting code	-	8	I/O
MAPAR_ERR_B	Address parity error	-	1	I
MAPAR_OUT	Address parity out	-	1	O
MDM[0:8]	DDR data mask	-	9	O
MDQS[0:8]	DDR data strobe	-	9	I/O
MDQS_B[0:8]	DDR data strobe (complement)	-	9	I/O
MBA[0:2]	DDR bank select	-	3	O
MA[0:15]	DDR address	-	16	O
MWE_B	DDR write enable	-	1	O
MRAS_B	DDR row address strobe	-	1	O
MCAS_B	DDR column address strobe	-	1	O
MCS_B	DDR chip select (2/DIMM)	-	4	O
MCKE[0:3]	DDR clock enable	-	4	O
MCK[0:5], MCK_B[0:5]	DDR differential clocks (3 pairs/DIMM)	-	12	O
MODT[0:3]	DRAM on-die termination	-	4	O
MDIC[0:1]	Driver impedance calibration	-	2	I/O
SerDes				
See PCI Express signal descriptions for more details				
SD_TX[3:0]	Transmit data	-	4	O
SD_TX_B[3:0]	Transmit data (complement)	-	4	O
SD_RX[3:0]	Receive data	-	4	I
SD_RX_B[3:0]	Receive data (complement)	-	4	I
SD_REF_CLK	PLL reference clock	-	1	I
SD_REF_CLK_B	PLL reference clock (complement)	-	1	I
SD_PLL_TPD	PLL test point digital	-	1	O
SD_PLL_TPA	PLL test point analog	-	1	O
SD_IMP_CAL_TX	Transmitter impedance calibration	-	1	I
SD_IMP_CAL_RX	Receiver impedance calibration	-	1	I
IEEE 1588				
See eTSEC external signals description for more details				
TSEC_1588_CLK_IN	1588 clock-in	-	1	I
TSEC_1588_TRIG_IN1	1588 trigger-in 1	-	1	I
TSEC_1588_TRIG_IN2	1588 trigger-in 2	-	1	I
TSEC_1588_ALARM_OUT1	1588 timer alarm-out 1	cfg_srds_refclk	1	O
TSEC_1588_ALARM_OUT2	1588 timer alarm-out 2	cfg_sgmi3	1	O
TSEC_1588_CLK_OUT	1588 clock-out	cfg_ddr_pll0	1	O

Table continues on the next page...

Table 3-1. Signal reference by functional block (continued)

Name	Description	Alternate function(s)	No. of Signals	I/O
TSEC_1588_PULSE_OUT1	1588 timer pulse-out 1	cfg_ddr_pll1	1	O
TSEC_1588_PULSE_OUT2	1588 timer pulse-out 2	cfg_ddr_pll2	1	O
Ethernet management				
See eTSEC external signals description for more details				
EC_MDC	Ethernet management data clock	cfg_tsec_reduce	1	O
EC_MDIO	Ethernet management data in/out	-	1	I/O
Enhanced three speed ethernet controller 1, 2, 3 (eTSEC1, eTSEC2, eTSEC3)				
See eTSEC external signals description for more details				
EC_GTX_CLK125	Gigabit reference clock	-	1	I
TSEC1_TXD7	eTSEC1 transmit data 7	TSEC3_TXD3/ cfg_tsec1_prctl1	1	O
TSEC1_TXD[6:4]	eTSEC1 transmit data 6-4	TSEC3_TXD[2:0]/ cfg_rom_loc[0:2]	3	O
TSEC1_TXD[3:1]	eTSEC1 transmit data 3-1	cfg_io_ports[0:2]	3	O
TSEC1_TXD0	eTSEC1 transmit data 0	cfg_tsec1_prctl0	1	O
TSEC1_TX_EN	eTSEC1 transmit enable	-	1	O
TSEC1_TX_ER	eTSEC1 transmit error	cfg_rom_loc3	1	O
TSEC1_TX_CLK	eTSEC1 transmit clock in	-	1	I
TSEC1_GTX_CLK	eTSEC1 transmit clock out	-	1	O
TSEC1_CRS	eTSEC1 carrier sense	TSEC3_RX_DV	1	I/O
TSEC1_COL	eTSEC1 collision detect	TSEC3_RX_CLK	1	I
TSEC1_RXD[7:4]	eTSEC1 receive data 7-4	TSEC3_RXD[3:0]	4	I
TSEC1_RXD[3:0]	eTSEC1 receive data 3-0	-	4	I
TSEC1_RX_DV	eTSEC1 receive data valid	-	1	I
TSEC1_RX_ER	eTSEC1 receiver error	-	1	I
TSEC1_RX_CLK	eTSEC1 receive clock	-	1	I
TSEC2_TXD7	eTSEC2 transmit data 7	cfg_tsec2_prctl1	1	O
TSEC2_TXD6	eTSEC2 transmit data 6	-	1	O
TSEC2_TXD5	eTSEC2 transmit data 5	TSEC3_TX_EN	1	O
TSEC2_TXD4	eTSEC2 transmit data 4	TSEC3_GTX_CLK/ cfg_device_ID7	1	O
TSEC2_TXD[3:2]	eTSEC2 transmit data 3-2	cfg_device_ID[6:5]	2	O
TSEC2_TXD1	eTSEC2 transmit data 1	cfg_dram_type	1	O
TSEC2_TXD0	eTSEC2 transmit data 0	cfg_tsec2_prctl0	1	O
TSEC2_TX_EN	eTSEC2 transmit enable	-	1	O
TSEC2_TX_ER	eTSEC2 transmit error	cfg_io_ports3	1	O
TSEC2_TX_CLK	eTSEC2 transmit clock in	-	1	I
TSEC2_GTX_CLK	eTSEC2 transmit clock out	-	1	O

Table continues on the next page...

Table 3-1. Signal reference by functional block (continued)

Name	Description	Alternate function(s)	No. of Signals	I/O
TSEC2_CRS	eTSEC2 carrier sense	TSEC3_RX_ER	1	I/O
TSEC2_COL	eTSEC2 collision detect	TSEC3_TX_CLK	1	I
TSEC2_RXD[7:0]	eTSEC2 receive data	-	8	I
TSEC2_RX_DV	eTSEC2 receive data valid	-	1	I
TSEC2_RX_ER	eTSEC2 receiver error	-	1	I
TSEC2_RX_CLK	eTSEC2 receive clock	-	1	I
TSEC3_TXD3	eTSEC3 transmit data 3	TSEC1_TXD7/ cfg_tsec1_prctl1	1	O
TSEC3_TXD[2:0]	eTSEC3 transmit data 2-0	TSEC1_TXD[6:4]/ cfg_rom_loc[0:2]	3	O
TSEC3_TX_EN	eTSEC3 transmit enable	TSEC2_TXD5	1	O
TSEC3_RX_ER	eTSEC3 receive error	TSEC2_CRS	1	I/O
TSEC3_TX_CLK	eTSEC3 transmit clock in	TSEC2_COL	1	I
TSEC3_GTX_CLK	eTSEC2 transmit clock out	TSEC2_TXD4/ cfg_device_ID7	1	O
TSEC3_RXD[3:0]	eTSEC3 receive data 3-0	TSEC1_RXD[7:4]	4	I
TSEC3_RX_DV	eTSEC3 receive data valid	TSEC1_CRS	1	I/O
TSEC3_RX_CLK	eTSEC3 receive clock	TSEC1_COL	1	I
eLBC				
See eLBC external signal descriptions for more details				
LAD[0:15]	Local bus address/data	cfg_gpinut[0:15]	16	I/O
LDP[0:1]	Local bus data parity	-	2	I/O
LA16	Local bus burst address 16	cfg_cpu1_boot	1	O
LA17	Local bus burst address 17	-	1	O
LA[18:19]	Local bus burst address 18-19	cfg_host_agt[1:2]	2	O
LA[20:22]	Local bus burst address 20	cfg_eng_use[0:2]	3	O
LA23	Local bus burst address 23	cfg_plat_speed	1	O
LA24	Local bus burst address 24	cfg_core0_speed	1	O
LA25	Local bus burst address 25	cfg_core1_speed	1	O
LA26	Local bus burst address 26	cfg_ddr_speed	1	O
LA27	Local bus burst address 27	cfg_cpu0_boot	1	O
LA28	Local bus burst address 28	-	1	O
LA[29:31]	Local bus port address 29-31	cfg_sys_pll[0:2]	3	O
LCS_B[0:4]	Local bus chip select 0-4	-	5	O
LCS5_B	Local bus chip select 5	DMA2_DREQ1_B	1	I/O
LCS6_B	Local bus chip select 6	DMA2_DACK1_B	1	O
LCS7_B	Local bus chip select 7	DMA2_DDONE1_B	1	O
LWE0_B/LBS0_B	Local bus write enable 0	cfg_core1_pll0	1	O
LWE1_B/LBS1_B	Local bus write enable 1	cfg_host_agt0	1	O

Table continues on the next page...

Table 3-1. Signal reference by functional block (continued)

Name	Description	Alternate function(s)	No. of Signals	I/O
LBCTL	Local bus data buffer control	cfg_core0_pll0	1	O
LALE	Local bus address latch enable	cfg_core0_pll1	1	O
LGPL0/ LFCLE	Local bus UPM general purpose line 0 / flash command latch enable	cfg_rio_sys_size	1	O
LGPL1/ LFALE	Local bus GP line 1 / flash address latch enable	cfg_sgmi2	1	O
LGPL2/LOE_B/LFRE_B	Local bus GP line 2 / output enable / flash read enable	cfg_core0_pll2	1	O
LGPL3/ LFWP_B	Local bus GP line 3 / flash write protect	cfg_boot_seq0	1	O
LGPL4/LGTA_B/ LFRB_B/ LUPWAIT/LPBSE	Local bus GP line 4 / transaction termination / flash ready-busy / wait	-	1	I/O
LGPL5	Local bus GP line 5 address	cfg_boot_seq1	1	O
LCLK[0:1]	Local bus clock	-	2	O
LSYNC_IN	Local bus PLL synchronization	-	1	I
LSYNC_OUT	Local bus PLL synchronization	-	1	O
Enhanced secured digital host controller (eSDHC)				
See eSDHC external signal description for more details				
SDHC_CMD	CMD line connect to card	-	1	I/O
SDHC_DAT[0:3]	Dat line 0-3	-	4	I/O
SDHC_DAT4	Dat line 4	SPI_CS0_B	1	I/O
SDHC_DAT[5:7]	Dat line 5-7	SPI_CS[1:3]	3	I/O
SDHC_CD_B	eSDHC card detection	GPIO8	1	I/O
SDHC_WP	eSDHC card write protect	GPIO9	1	I/O
SDHC_CLK	Clock for MMC/SD/SDIO card	-	1	O
Enhanced serial peripheral interface (eSPI)				
See External signal descriptions for more details				
SPI_MOSI	SPI master-out slave-in	-	1	I/O
SPI_MISO	SPI master-in slave-out	-	1	I
SPI_CLK	SPI clock	-	1	O
SPI_CS_B0	SPI slave select 0	SDHC_DAT4	1	I/O
SPI_CS_B[1:3]	SPI slave select 1-3	SDHC_DAT[5:7]	3	I/O
USB ULPI				
See USB external signals for more details				
USB_D[7:0]	Data bus	-	8	I/O
USB_NXT	Next data	-	1	I
USB_DIR	Direction of data bus	-	1	I
USB_STP	End of a transfer on the bus	-	1	O
USB_PCTL0	Port control	GPIO10	1	I/O
USB_PCTL1	Port control	GPIO11	1	I/O

Table continues on the next page...

Table 3-1. Signal reference by functional block (continued)

Name	Description	Alternate function(s)	No. of Signals	I/O
USB_PWRFAULT	USB VBus power fault	-	1	I
USB_CLK	PHY clock	-	1	I
DMA controller				
See DMA external signal description for more details				
DMA1_DREQ_B	DMA1 request	-	1	I
DMA1_DACK_B	DMA1 acknowledge	-	1	O
DMA1_DDONE_B	DMA1 done	cfg_eng_use7	1	O
DMA2_DREQ0_B	DMA2 request 0	-	1	I
DMA2_DREQ1_B	DMA2 request 1	LCS5_B	1	I/O
DMA2_DACK0_B	DMA2 acknowledge 0	cfg_mem_debug	1	O
DMA2_DACK1_B	DMA2 acknowledge 1	LCS6_B	1	O
DMA2_DDONE0_B	DMA2 done 0	cfg_ddr_debug	1	O
DMA2_DDONE1_B	DMA2 done 1	LCS7_B	1	O
Programmable interrupt controller (PIC)				
See PIC external signal descriptions for more details				
MCP0_B	Machine check processor 0	-	1	I
MCP1_B	Machine check processor 1	-	1	I
UDE0_B	Unconditional debug event 0	-	1	I
UDE1_B	Unconditional debug event 1	-	1	I
IRQ0	External interrupt 0	-	1	I
IRQ[1:6]	External interrupt 1-6	-	6	I
IRQ[7:11]	External interrupt 7-11	GPIO[0:4]	5	I/O
IRQ_OUT_B	Interrupt output	-	1	O
GPIO				
See External signal description for more details				
GPIO[0:4]	General-purpose I/O signal 0-4	IRQ[7:11]	5	I/O
GPIO[5:7]	General-purpose I/O signal 5-7	-	3	I/O
GPIO8	General-purpose I/O signal 8	SDHC_CD_B	1	I/O
GPIO9	General-purpose I/O signal 9	SDHC_WP	1	I/O
GPIO10	General-purpose I/O signal 10	USB_PCTL0	1	I/O
GPIO11	General-purpose I/O signal 11	USB_PCTL1	1	I/O
GPIO12	General-purpose I/O signal 12	-	1	I/O
GPIO13	General-purpose I/O signal 13	-	1	I/O
GPIO14	General-purpose I/O signal 14	-	1	I/O
GPIO15	General-purpose I/O signal 15	-	1	I/O
System control				
See System control signals for more details				
HRESET_B	Hard reset	-	1	I

Table continues on the next page...

Table 3-1. Signal reference by functional block (continued)

Name	Description	Alternate function(s)	No. of Signals	I/O
HRESET_REQ_B	Hard reset request	-	1	O
SRESET_B	Soft reset	-	1	I
CKSTP_IN_B[0:1]	Checkstop in	-	2	I
CKSTP_OUT_B[0:1]	Checkstop out	-	2	O
DUART				
See DUART external signal descriptions for more details				
UART_SOUT0	UART serial data out 0	cfg_eng_use3	1	O
UART_SOUT1	UART serial data out 1	cfg_core1_pll1	1	O
UART_SIN0	UART serial data in 0	-	1	I
UART_SIN1	UART serial data in 1	-	1	I
UART_CTS_B[0:1]	UART0-1 clear to send	-	2	I
UART_RTS_B[0:1]	UART0-1 ready to send	cfg_tsec3_ptctl[0:1]	2	O
I2C				
See I2C external signal descriptions for more details				
IIC1_SDA	I2C1 serial data	-	1	I/O
IIC1_SCL	I2C1 serial clock	-	1	I/O
IIC2_SDA	I2C2 serial data	-	1	I/O
IIC2_SCL	I2C2 serial clock	-	1	I/O
Debug				
See External signal description for more details				
TRIG_IN	Watchpoint trigger in	-	1	I
TRIG_OUT/ READY_P0	Watchpoint trigger out / processor 0 ready	-	1	O
READY_P1	Processor 1 ready	cfg_core1_pll2	1	O
MSRCID0	Memory debug source ID 0	cfg_elbc_ecc	1	O
MSRCID1	Memory debug source ID 1	cfg_eng_use5	1	O
MSRCID2	Memory debug source ID 2	-	1	O
MSRCID3	Memory debug source ID 3	-	1	O
MSRCID4	Memory debug source ID 4	cfg_eng_use6	1	O
MDVAL	Memory debug data valid	-	1	O
CLK_OUT	Clock out	-	1	O
Power management				
See External signal description for more details				
ASLEEP	Asleep	-	1	O
Clocking				
See External signal descriptions for more details				
SYSCLK	System clock	-	1	I
RTC	Real time clock	-	1	I
DDRCLK	DDR clock	-	1	I

Table continues on the next page...

Table 3-1. Signal reference by functional block (continued)

Name	Description	Alternate function(s)	No. of Signals	I/O
DFT				
See External signal description for more details				
SCAN_MODE_B	Test select	-	1	I
TEST_SEL_B	Test select	-	1	I
JTAG				
See External signal description for more details				
TCK	Test clock	-	1	I
TDI	Test data in	-	1	I
TDO	Test data out	-	1	O
TMS	Test mode select	-	1	I
TRST_B	Test reset	-	1	I
IO VSEL				
See Voltage selection configuration for more details				
LVDD_VSEL	LVDD voltage select	-	1	I
BVDD_VSEL[0:1]	BVDD voltage select	-	2	I
CVDD_VSEL[0:1]	CVDD voltage select	-	2	I
MVREF	DDR reference voltage	-	1	I

3.2 Configuration signals sampled at reset

The signals that serve alternate functions as configuration input signals during system reset are summarized in the table below.

The detailed interpretation of their voltage levels during reset is described in [Reset, Clocking, and Initialization](#).

Note that throughout this document, the reset configuration signals are described as being sampled at the negation of HRESET_B. However, there is a setup and hold time for these signals relative to the rising edge of HRESET_B, see device hardware specification document. Note that the PLL configuration signals have different setup and hold time requirements than the other reset configuration signals.

The reset configuration signals are multiplexed with other functional signals. The values on these signals during reset are interpreted to be logic one or zero, regardless of whether the functional signal name is defined as active-low. Most of the reset configuration signals have internal pull-up resistors so that if the signals are not driven, the default

value is high (a one), as shown in the table. Some signals do not have pull-up resistors and must be driven high or low during the reset period. For details about all the signals that require external pull-up resistors, see the device hardware specification document.

Note that the multiplexing of various signals on the chip is controlled by the PMUXCR register described in [Global Utilities](#).

Table 3-2. Reset configuration signals

Functional interface	Functional signal name	Reset configuration name	Default
Debug	READY_P1	cfg_core1_pll2	Must be driven
	MSRCID0	cfg_elbc_ecc	1
	MSRCID4	cfg_eng_use6	1
	MSRCID1	cfg_eng_use5	1
	TRIG_OUT/READY_P0	cfg_srds_pll_toe	1
DMA1	DMA1_DDONE_B	cfg_eng_use7	1
DMA2	DMA2_DDONE0	cfg_ddr_debug	1
	DMA2_DACK0_B	cfg_mem_debug	1
DUART	UART_SOUT1	cfg_core1_pll1	Must be driven
	UART_RTS_B[0:1]	cfg_tsec3_ptcl[0:1]	11
	UART_SOUT0	cfg_eng_use3	1
eLBC	LGPL3/LFWP_B	cfg_boot_seq0	1
	LGPL5	cfg_boot_seq1	1
	LBCTL	cfg_core0_pll0	Must be driven
	LALE	cfg_core0_pll1	Must be driven
	LGPL2/LOE_B/LFRE_B	cfg_core0_pll2	Must be driven
	LWE0_B/LBS0_B	cfg_core1_pll0	Must be driven
	LA16	cfg_cpu1_boot	1
	LA[18:19]	cfg_host_agt[1:2]	11
	LA[20:22]	cfg_eng_use[0:2]	1
	LA23	cfg_plat_speed	1
	LA24	cfg_core0_speed	1
	LA25	cfg_core1_speed	1
	LA26	cfg_ddr_speed	1
	LA27	cfg_cpu0_boot	1
	LA[29:31]	cfg_sys_pll[0:2]	Must be driven
	LAD[0:15]	cfg_gpinput[0:15]	
	LWE1_B/LBS1_B	cfg_host_agt0	1
	LGPL0/LFCLE	cfg_rio_sys_size	1
	LGPL1/LFALE	cfg_sgmi2	1
	Ethernet Management	EC_MDC	cfg_tsec_reduce

Table continues on the next page...

Table 3-2. Reset configuration signals (continued)

Functional interface	Functional signal name	Reset configuration name	Default
eTSEC	TSEC_1588_PULSE_OUT1	cfg_ddr_pll1	Must be driven
	TSEC_1588_PULSE_OUT2	cfg_ddr_pll2	Must be driven
	TSEC_1588_CLK_OUT	cfg_ddr_pll0	Must be driven
	TSEC_1588_ALARM_OUT1	cfg_srds_refclk	1
	TSEC_1588_ALARM_OUT2	cfg_sgmi3	1
eTSEC1	TSEC1_TXD[3:1]	cfg_io_ports[0:2]	111
	TSEC1_TXD[6:4]	cfg_rom_loc[0:2]	111
	TSEC1_TX_ER	cfg_rom_loc3	1
	TSEC1_TXD0	cfg_tsec1_ptcl0	1
	TSEC1_TXD7	cfg_tsec1_ptcl1	1
eTSEC2	TSEC2_TX_ER	cfg_io_ports3	1
	TSEC2_TXD0	cfg_tsec2_ptcl0	1
	TSEC2_TXD1	cfg_dram_type	1
	TSEC2_TXD[4:2]	cfg_device_ID[7:5]	111
	TSEC2_TXD7	cfg_tsec2_ptcl1	1

3.3 Output signal states during reset

When a system reset is recognized (HRESET_B is asserted), the chip aborts all current internal and external transactions and releases all bidirectional I/O signals to a high-impedance state.

See [Reset, Clocking, and Initialization](#) for a complete description of the reset functionality.

During reset, the chip ignores most input signals (except for the reset configuration signals) and drives most of the output-only signals to tri-state.

Chapter 4

Reset, Clocking, and Initialization

This chapter describes the reset, clocking, and some overall initialization of the device, including a definition of the reset configuration signals and the options they select.

4.1 Overview

This chapter describes the reset, clocking, and some overall initialization of the chip, including a definition of the reset configuration signals and the options they select.

Additionally, the configuration, control, and status registers are described. Note that other chapters in this book may describe specific aspects of initialization for individual blocks.

The reset, clocking, and control signals provide many options for the operation of the chip. Additionally, many modes are selected with reset configuration signals during a hard reset (assertion of HRESET_B).

4.2 External signal descriptions

[Table 4-1](#) summarizes the external signals described in this chapter. [Table 4-2](#) and [Table 4-3](#) have detailed signal descriptions, but [Table 4-1](#) contains references to additional sections that contain more information.

Table 4-1. Signal summary

Signal	I/O	Description	References (Section/Page)
HRESET_B	I	Hard reset input. Causes a power-on reset (POR) sequence.	System control signals
HRESET_REQ_B	O	Hard reset request output. An internal block requests that HRESET_B be asserted.	System control signals

Table continues on the next page...

Table 4-1. Signal summary (continued)

Signal	I/O	Description	References (Section/Page)
SRESET_B	I	Soft reset input. Causes <i>mcp</i> assertion to the core	System control signals
READY_P0/ TRIG_OUT	O	The device has completed the reset operation, and e500 core 0 is not in a power-down (nap, doze or sleep) or debug state.	Power-on reset sequence
READY_P1	O	The device has completed the reset operation, and e500 core 1 is not in a power-down (nap, doze or sleep) or debug state.	Power-on reset sequence
SYSCLK	I	Primary clock input to the device	System clock and DDR controller clock
RTC	I	Real time clock input	Real time clock
SD_REF_CLK/ SD_REF_CLK_B	I	SerDes high-speed interface reference clock	RapidIO, SGMII and PCI Express clocks
DDRCLK	I	Reference clock for DDR Controller, when running in asynchronous mode	System clock and DDR controller clock

The following sections describe the reset and clock signals in detail.

4.2.1 System control signals

[Table 4-2](#) describes some of the device's system control signals. [Power-on reset configuration](#), describes the signals that also function as reset configuration signals. Note that the CKSTP_IN_B and CKSTP_OUT_B signals are described in [Global Utilities](#).

Table 4-2. System control signals-detailed signal descriptions

Signal	I/O	Description	
HRESET_B	I	Hard reset. Causes the device to abort all current internal and external transactions and set all registers to their default values. HRESET_B may be asserted completely asynchronously with respect to all other signals.	
		State Meaning	Asserted/Negated-See Signals overview and Power-on reset configuration , for more information on the interpretation of other signals during reset.
		Timing	Assertion/Negation-The <i>P2020 Integrated Processor Hardware Specifications</i> gives specific timing information for this signal and the reset configuration signals.

Table continues on the next page...

Table 4-2. System control signals-detailed signal descriptions (continued)

Signal	I/O	Description
HRESET_REQ_B	O	Hard reset request. Indicates to the board (system in which the chip is embedded) that a condition requiring the assertion of HRESET_B has been detected.
		State Meaning Asserted-A watchdog timer expiration sent to the platform (see the <i>PowerPC™ e500 Core Family Reference Manual</i> , for the Timer Control Register(TCR); TCR[WRC] field) a RapidIO command, a boot sequencer failure (see Boot sequencer mode), or an eLBC ECC error (see Reset Request Status and Control Register (RSTRSCR) has triggered a request for hard reset. Negated-Indicates no reset request.
		Timing Assertion/Negation-May occur any time, synchronous to the core complex bus clock. Once asserted, HRESET_REQ_B does not negate until HRESET_B is asserted.
SRESET_B	I	Soft reset. Causes a machine check interrupt to the e500 core. Note that if the e500 core is not configured to process machine check interrupts, the assertion of SRESET_B causes a core checkstop. SRESET_B need not be asserted during a hard reset .
		State Meaning Asserted-Asserting SRESET_B causes a machine check interrupt (edge sensitive) to the e500 core. SRESET_B has no effect while HRESET_B is asserted. However, the POR sequence is paused if SRESET_B is asserted during POR.
		Timing Assertion-May occur at any time, asynchronous to any clock. Negation-Must be asserted for at least two <i>CCB_clk</i> cycles.
READY_P0/ TRIG_OUT	O	Ready processor 0. Multiplexed with TRIG_OUT and QUIESCE_B. See Debug Features and Watchpoint Facility for more information on TOSR and TRIG_OUT.
		State Meaning Asserted-Indicates that the chip has completed the reset operation, and e500 core 0 is not in a power-down state (nap, doze, or sleep) when TOSR[SEL] equals 0b000. See Power-on reset sequence , for more information.
		Timing Assertion/Negation-Initial assertion of READY_P0 after reset is synchronous with SYSCLK. Subsequent assertion/negation due to power down modes occurs asynchronously.
READY_P1	O	Ready processor 1
		State Meaning Asserted-Indicates that the chip has completed the reset operation, and e500 core 1 is not in a power-down state (nap, doze, or sleep).
		Timing Assertion/Negation-Initial assertion of READY_P1 after reset is synchronous with SYSCLK. Subsequent assertion/negation due to power down modes occurs asynchronously.

4.2.2 Clock signals

This table describes the overall clock signals.

Note that some clock signals are specific to blocks within the device, and although some of their functionality is described in [Clocking](#), they are defined in detail in their respective chapters.

Note that there is also a CLK_OUT signal; the signal driven on the CLK_OUT pin is selectable and described in [Clock out control register \(GUTS_CLKOCR\)](#).

Table 4-3. Clock signals-detailed signal descriptions

Signal	I/O	Description
SYSCLK	I	System clock (SYSCLK). SYSCLK is the primary clock input to the device. It is the clock source for the e500 core and for all devices and interfaces that operate synchronously with the core. It is multiplied up with a phased-lock loop (PLL) to create the core complex bus (CCB) clock (also called the platform clock), which is used by virtually all of the synchronous system logic, including the L2 cache, the DDR SDRAM and local bus memory controllers, and other internal blocks such as the DMA and interrupt controllers. The CCB clock, in turn, feeds the PLL in the e500 core and the PLL that creates the local bus memory clocks.
		<table border="1"> <tr> <td>Timing</td> <td>Assertion/Negation-See the <i>P2020 Integrated Processor Hardware Specifications</i> for specific timing information for this signal.</td> </tr> </table>
Timing	Assertion/Negation-See the <i>P2020 Integrated Processor Hardware Specifications</i> for specific timing information for this signal.	
RTC	I	Real time clock. May be used (optionally) to clock the time base of the e500 core. The RTC timing specifications are given in the <i>P2020 Integrated Processor Hardware Specifications</i> , but the maximum frequency should be less than one-quarter of the CCB frequency. See Real time clock . This signal can also be used (optionally) to clock the global timers in the programmable interrupt controller (PIC).
		<table border="1"> <tr> <td>Timing</td> <td>Assertion/Negation-See the <i>P2020 Integrated Processor Hardware Specifications</i> for specific timing information for this signal.</td> </tr> </table>
Timing	Assertion/Negation-See the <i>P2020 Integrated Processor Hardware Specifications</i> for specific timing information for this signal.	
DDRCLK	I	DDR controller clock. DDRCLK is the clock source for the DDR memory controller except in the case where synchronous mode of operation is selected (see DDR PLL Ratio). This clock input is multiplied up with a phased-lock loop (PLL) to create the DDR controller clock. The DDR memory controller clock is the DDR data rate on the external interface.
		<table border="1"> <tr> <td>Timing</td> <td>Assertion/Negation-See the <i>P2020 Integrated Processor Hardware Specifications</i> for specific timing information for this signal.</td> </tr> </table>
Timing	Assertion/Negation-See the <i>P2020 Integrated Processor Hardware Specifications</i> for specific timing information for this signal.	

4.3 Accessing configuration, control, and status registers

The configuration, control, and status registers are memory mapped.

The set of configuration, control, and status registers occupies a 1-Mbyte region of memory. Their location is programmable using the CCSR base address register (CCSRBAR). The default base address for the configuration, control, and status registers is 0xFF70_0000 (CCSRBAR = 0x000F_F700). CCSRBAR itself is part of the local access block of CCSR memory, which begins at offset 0x0 from CCSRBAR. Because CCSRBAR is at offset 0x0 from the beginning of the local access registers, CCSRBAR always points to itself. The contents of CCSRBAR are broadcast internally in the chip to all functional units that need to be able to identify or create configuration transactions.

4.3.1 Updating CCSRBAR

Updates to CCSRBAR that relocate the entire 1-Mbyte region of configuration, control, and status registers require special treatment.

The effect of the update must be guaranteed to be visible by the mapping logic before an access to the new location is seen. To make sure this happens, these guidelines should be followed:

- CCSRBAR should be updated during initial configuration of the device when only one host or controller has access to the device.
 - If the boot sequencer is being used to initialize, it is recommended that the boot sequencer set CCSRBAR to its desired final location.
 - If an external host on PCI Express or RapidIO is configuring the device, it should set CCSRBAR to the desired final location before the e500 core is released to boot.
 - If the e500 core is initializing the device, it should set CCSRBAR to the desired final location before enabling other I/O devices to access the device.
- When the e500 core is writing to CCSRBAR, it should use the following sequence:
 - Read the current value of CCSRBAR using a load word instruction followed by an **isync**. This forces all accesses to configuration space to complete.
 - Write the new value to CCSRBAR.
 - Perform a load of an address that does not access configuration space or the on-chip SRAM, but has an address mapping already in effect (for example, boot ROM). Follow this load with an **isync**.
 - Read the contents of CCSRBAR from its new location, followed by another **isync**.

4.3.2 Accessing alternate configuration space

An alternate configuration space can be accessed by configuring the ALTCCBAR and ALTCCAR registers.

These are intended to be used with the boot sequencer to allow the boot sequencer to access an alternate 1-Mbyte region of configuration space. By loading the proper boot sequencer command in the serial ROM, the base address in the ALTCCBAR can be combined with the 20 bits of address offset supplied from the serial ROM to generate a 36-bit address that is mapped to the target specified in ALTCCAR. Thus, by configuring these registers, the boot sequencer has access to the entire memory map, one 1-Mbyte block at a time. See [Boot sequencer mode](#), for more information.

NOTE

The enable bit in the ALTCCAR register should be cleared either by the boot sequencer or by the boot code that executes after the boot sequencer has completed its configuration operations. This prevents problems with incorrect mappings if subsequent

configuration of the local access windows uses a different target mapping for the address specified in ALTCBAR.

4.3.3 Boot page translation

When each e500 core comes out of reset, its MMU has one 4-Kbyte page defined at `0x0_FFFF_Fnnn`.

Each core begins execution with the instruction at effective address `0x0_FFFF_FFFC`. To get this instruction, the core's first instruction fetch is a burst read of boot code from effective address `0x0_FFFF_FFE0`. For systems in which the boot code resides at a different address, the chip provides boot page translation capability. Boot page translation is controlled by the boot page translation register (BPTR). Note that boot page translation affects transactions initiated by each of the two e500 cores in the same manner.

The boot sequencer can enable boot page translation, or the boot page translation can be set up by an external host when the device is configured to be in boot holdoff mode. If translation is to be performed to a page outside the default boot ROM address range defined in the chip (8 Mbytes at `0x0_FF80_0000` to `0x0_FFFF_FFFF` as defined in [Boot ROM location](#)), the external host or boot sequencer must then also set up a local access window to define the routing of the boot code fetch to the target interface that contains the boot code, because the BPTR defines only the address translation, not the target interface window. [Local address map example](#), and [Boot sequencer mode](#) for more information.

4.3.4 Boot sequencer

The boot sequencer is a DMA engine that accesses a serial ROM on the I²C interface and writes data to CCSR memory or the memory space pointed to by the alternate configuration base address register (ALTCBAR).

See [Accessing alternate configuration space](#). The boot sequencer is enabled by reset configuration pins as described in [Boot sequencer configuration](#). If the boot sequencer is enabled, the e500 core is held in reset until the boot sequencer has completed its operation. For more details, see [Boot sequencer mode](#) in the I²C chapter.

4.4 Reset Memory Map/Register Definition

This section describes the configuration and control registers that control access to the configuration space and to the boot code as well as guidelines for accessing these regions. It also contains a brief description of the boot sequencer which may be used to initialize configuration registers or memory before the CPU is released to boot.

This table shows the memory map for local configuration control registers.

Reset memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
0	Configuration, control, and status registers base address register (Reset_CCSRBAR)	32	R/W	000F_F700h	4.4.1/125
8	Alternate configuration base address register (Reset_ALTCBAR)	32	R/W	0000_0000h	4.4.2/126
10	Alternate configuration attribute register (Reset_ALTCAR)	32	R/W	0000_0000h	4.4.3/126
20	Boot page translation register (Reset_BPTR)	32	R/W	0000_0000h	4.4.4/127

4.4.1 Configuration, control, and status registers base address register (Reset_CCSRBAR)

Address: 0h base + 0h offset = 0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W					-																											
Reset	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	0	1	1	1	0	0	0	0	0	0	0	

Reset_CCSRBAR field descriptions

Field	Description
0–7 -	Write reserved, read = 0.
8–23 BASE_ADDR	Identifies the 16 most-significant address bits of the 36-bit window used for configuration accesses. The base address is aligned on a 1-Mbyte boundary.
24–31 -	Write reserved, read = 0

4.4.2 Alternate configuration base address register (Reset_ALTCBAR)

Address: 0h base + 8h offset = 8h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W					-																											
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Reset_ALTCBAR field descriptions

Field	Description
0–7 -	Write reserved, read = 0
8–23 BASE_ADDR	Identifies the 16 most significant address bits of an alternate window used for configuration accesses.
24–31 -	Write reserved, read = 0

4.4.3 Alternate configuration attribute register (Reset_ALTCAR)

Address: 0h base + 10h offset = 10h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	EN				-						TRGT_ID				-	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Reset_ALTCAR field descriptions

Field	Description
0 EN	Enable for a second configuration window. Like CCSRBAR, it has a fixed size of 1 Mbyte. 0 Second configuration window is disabled. 1 Second configuration window is enabled.
1–7 -	Write reserved, read = 0
8–11 TRGT_ID	Identifies the device ID to target when a transaction hits in the 1-Mbyte address range defined by the second configuration window. Patterns not shown are reserved. 0000 PCI Express interface 3 0001 PCI Express interface 2 0010 PCI Express interface 1

Table continues on the next page...

Reset_ALT CAR field descriptions (continued)

Field	Description
0100	Local bus controller
1000	Configuration, control, status registers
1100	Serial RapidIO interface 1
1101	Serial RapidIO interface 2
1111	Local memory -DDR SDRAM and on-chip SRAM
12–31 -	Write reserved, read = 0

4.4.4 Boot page translation register (Reset_BPTR)

Address: 0h base + 20h offset = 20h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	EN				-				BOOT_PAGE							
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W	BOOT_PAGE															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Reset_BPTR field descriptions

Field	Description
0 EN	Boot page translation enable 0 Boot page is not translated. 1 Boot page is translated as defined in the BPTR[BOOT_PAGE] parameter.
1–7 -	Write reserved, read = 0
8–31 BOOT_PAGE	Translation for boot page. If enabled, the high order 24 bits of accesses to 0x0_FFFF_Fnnn are replaced with this value.

4.5 Functional description

This section describes the various ways to reset the chip, the POR configurations, and the clocking on the device.

4.5.1 Reset operations

The chip has reset input signals for hard and soft reset operation.

4.5.1.1 Soft reset

Assertion of SRESET_B causes a machine check interrupt to both e500 cores.

When this occurs, the soft reset flag is recorded in the machine check summary register in the global utilities block so that software can identify the machine check as a soft reset condition. See the *PowerPC e500 Core Complex Reference Manual* for more information on the machine check interrupt and [Machine check summary register \(GUTS_MCPSUMR\)](#), for more information on the setting of the soft reset flag. Note that if SRESET_B is asserted before a given e500 core is configured to handle a machine check interrupt, a checkstop condition occurs for that particular core, which causes CKSTP_OUT0_B or CKSTP_OUT1_B to assert.

4.5.1.2 Hard reset

The chip can be completely reset by the assertion of the HRESET_B input.

The assertion of this signal by external logic is the equivalent of a POR and causes the sequence of events described in [Power-on reset sequence](#).

Refer to device hardware specifications for the timing requirements for HRESET_B assertion and negation.

The hard reset request output signal (HRESET_REQ_B) indicates to external logic that a hard reset is being requested by hardware or software or a RapidIO device. Hardware causes this signal to assert for a boot sequencer failure (see [Boot sequencer mode](#) and [EEPROM data format](#)), for a local bus uncorrectable ECC error during NAND Flash boot process (see [Boot block loading into the FCM buffer RAM](#)), or when either e500 watchdog timer is configured to cause a hard reset request when it expires (see the *PowerPC™ e500 Core Family Reference Manual*, for the Timer Control Register (TCR); TCR[WRC] field). Software may request a hard reset by setting a bit in a global utilities register; see [Reset control register \(GUTS_RSTCR\)](#). A RapidIO device causes this signal to assert if it sends four consecutive RapidIO link maintenance reset commands without any other intervening packets or control symbols, except idle control symbols.

4.5.2 Power-on reset sequence

The POR sequence for the chip is as follows:

1. Power is applied to meet the specifications in the device hardware specifications.

2. The system asserts HRESET_B and TRST_B, causing all registers to be initialized to their default states and most I/O drivers to be three-stated (some clock, clock enabled, and system control signals are active).
3. The system applies a stable SYSCLK signal and stable PLL configuration inputs, and the device PLL begins locking to SYSCLK.
4. System negates HRESET after its required hold time and after POR configuration inputs have been valid for at least 4 SYSCLK cycles.
5. Chip enables I/O drivers.
6. The e500 PLL configuration inputs are applied, allowing the e500 PLLs to begin locking to the device clock (the CCB clock).
7. The CCB clock is cycled for approximately 50 μ s to lock the e500 PLLs.
8. The internal hard resets to the e500 cores are negated and soft resets are negated to the PLLs and other remaining I/O blocks. The PLLs begin to lock.
9. When PLL locking is completed, the enhanced local bus FCM is released provided NAND flash is configured as the boot device, as described in [Boot ROM location](#). Once the FCM finishes loading the pages from the NAND Flash device, the boot sequencer, if enabled, is allowed to progress, causing it to load configuration data from serial ROMs on the I2C1 interface, as described in [Boot sequencer configuration](#).
10. When the local bus FCM and boot sequencer complete, the PCI Express interfaces begin training and are released to accept external requests, and the boot vectors fetched by the e500 cores are allowed to proceed unless processor booting is further held off by POR configuration inputs as described in [Boot sequencer configuration](#).

The chip is now in its ready state.

11. The ASLEEP signal negates synchronized to a rising edge of SYSCLK, indicating the ready state of the system. The ready state for the e500 core is also indicated by the assertion of READY_P0/TRIG_OUT if TOSR[SEL] = 000. In this case, READY is asserted with the same rising edge of SYSCLK, to indicate that the e500 core has reached its ready state, allows external system monitors to know basic device status, for example, exactly when the e500 core emerges from reset, or if it is in a low-power mode.

This figure shows a timing diagram of the POR sequence.

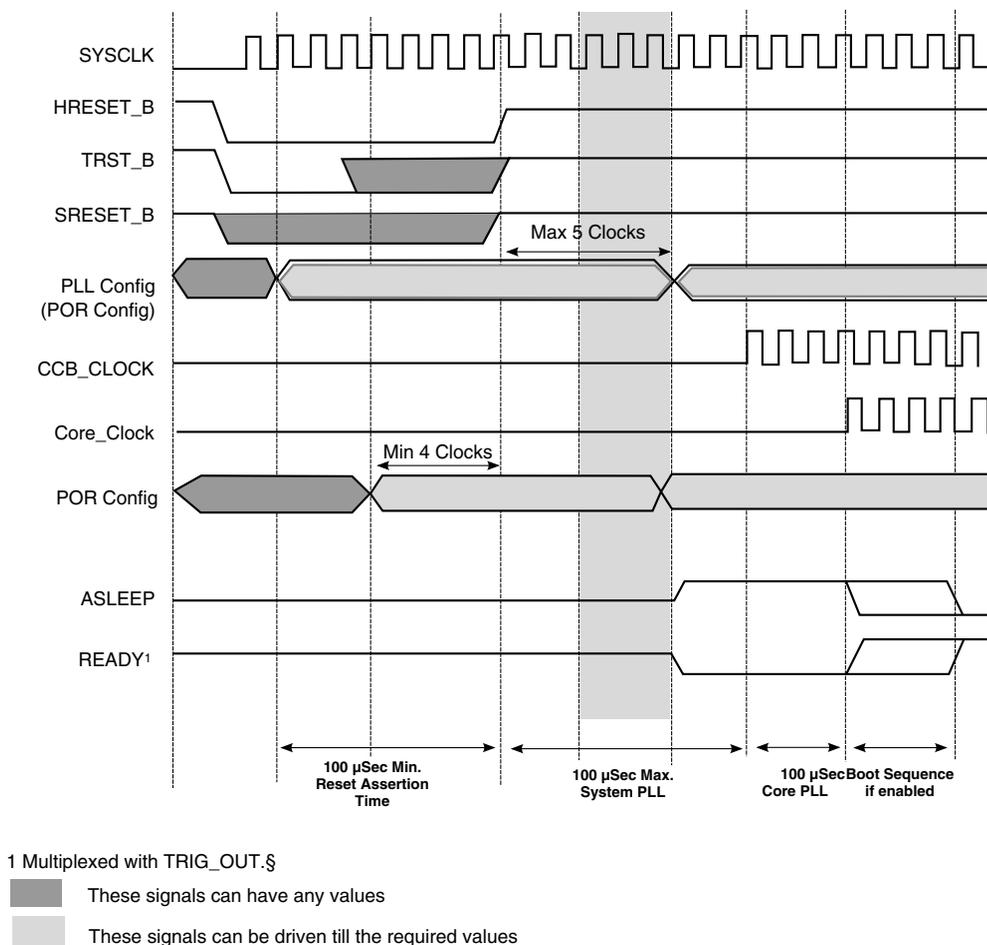


Figure 4-5. Power-on reset sequence

4.5.3 Power-on reset configuration

Various device functions are initialized by sampling certain signals during the assertion of HRESET_B.

The values of all these signals are sampled into registers while HRESET_B is asserted. These inputs are to be pulled high or low by external resistors. During HRESET_B, all other signal drivers connected to these signals must be in the high-impedance state.

Most POR configuration signals have internal pull-up resistors so that if the desired setting is high, there is no need for a pull-up resistor on the board. Other POR configuration signals do not use pull-ups and therefore must be pulled high or low. Refer to device hardware specifications for proper resistor values to be used for pulling POR configuration signals high or low.

This section describes the functions and modes configured by POR configuration signals. Note that many reset configuration settings are accessible to software through the following read-only memory-mapped registers described in [Global Utilities](#) :

- POR PLL status register (PORPLLSR)
- POR boot mode status register (PORBMSR)
- POR device status register (PORDEVSR)
- POR device status register 2 (PORDEVSR2)
- POR debug mode status register (PORDBGMSR)
- General-purpose POR configuration register (GPPORCR)-Reports the value on LAD[0:15] during POR (can be used to external system configuration)

NOTE

In the following tables, the binary value 0b0 represents a signal pulled down to GND and a value of 0b1 represents a signal pulled up to V_{DD} , regardless of the sense of the functional signal name on the signal.

4.5.3.1 System PLL ratio

The system PLL inputs establish the clock ratio between the SYSCLK input and the platform clock used by the chip.

The platform clock, also called the CCB clock, drives the L2 cache, the DDR SDRAM data rate, and the e500 core complex bus (CCB). There is no default value for this PLL ratio; these signals must be pulled to the desired values. See [Minimum frequency requirements](#), for optimal selection of this ratio with regard to available high-speed interface widths and frequencies. Note that the values latched on these signals during POR are accessible in the PORPLLSR (POR PLL status register), as described in [POR PLL ratio status register \(GUTS_PORPLLSR\)](#) .

The system PLL inputs are shown in this table.

Table 4-8. CCB clock PLL ratio

Functional signals	Reset configuration name	Value (Binary)	CCB Clock : SYSCLK Ratio
LA[29:31] No Default	cfg_sys_pll[0:2]	000	4:1
		001	5:1
		010	6:1
		011	8:1
		other	Reserved

4.5.3.2 DDR PLL Ratio

The DDR PLL inputs, shown in this table, establish the clock ratio between the DDRCLK input and the DDR controller clock.

The DDR controller clock drives the DDR data rate, which is twice the rate at which commands are issued on the DDR interface. This DDR controller clock domain is asynchronous to the platform clock or CCB clock domain and is sourced from a separate PLL than the rest of the platform, unless the DDR PLL encoding for synchronous mode operation is selected. When synchronous mode is selected, the DDR controller is driven by the CCB clock, which becomes the DDR data rate. There is no default value for this PLL ratio; these signals must be pulled to the desired values. Note that the encoded values latched on these signals during power-on reset-and not the actual values on the pins-are accessible in PORPLLSR (POR PLL status register), as described in [POR PLL ratio status register \(GUTS_PORPLLSR\)](#).

Table 4-9. DDR controller clock PLL ratio

Functional signals	Reset configuration name	Value (Binary)	DDR controller : DDRCLK Ratio
TSEC_1588_CLK_OUT, TSEC_1588_PULSE_OUT1, TSEC_1588_PULSE_OUT2 No Default	cfg_dds_pll[0:2]	000	Reserved
		001	4:1
		010	6:1
		011	8:1
		100	10:1
		101	12:1
		110	Reserved
		111	Synchronous mode

4.5.3.3 e500 core PLL ratios

The tables below describe the e500 core clock PLL inputs that program the core PLLs and establish the ratio between the e500 core clocks and the e500 core complex bus (CCB) clock.

There are no default values for these PLL ratios; these signals must be pulled to the desired values. Note that the values latched on these signals during POR are accessible through the memory-mapped PORPLLSR, as described in [POR PLL ratio status register \(GUTS_PORPLLSR\)](#).

Table 4-10. e500 core0 clock PLL ratios

Functional signals	Reset configuration name	Value (Binary)	e500 Core: CCB ClockRatio
LBCTL, LALE, LGPL2/LOE_B/LFRE_B No Default	cfg_core0_pll[0:2]	000	4:1
		001	9:2 (4.5:1)
		010	1:1
		011	3:2 (1.5:1)
		100	2:1
		101	5:2 (2.5:1)
		110	3:1
		111	7:2 (3.5:1)

Table 4-11. e500 Core1 Clock PLL Ratios

Functional Signals	Reset Configuration Name	Value (Binary)	e500 Core: CCB ClockRatio
LWE0_B, UART_SOUT1, READY_P1 No Default	cfg_core1_pll[0:2]	000	4:1
		001	9 : 2 (4.5:1)
		010	1:1
		011	3:2 (1.5:1)
		100	2:1
		101	5:2 (2.5:1)
		110	3:1
		111	7:2 (3.5:1)

4.5.3.4 Boot ROM location

The chip defines the default boot ROM address range to be 8 Mbytes at address 0x0_FF80_0000 to 0x0_FFFF_FFFF.

However, which peripheral interface handles these boot ROM accesses can be selected at power on.

The boot ROM location inputs, shown in this table, select the physical location of boot ROM. Accesses to the boot vector and the default boot ROM region of the local address map are directed to the interface specified by these inputs.

Table 4-12. Boot ROM location

Functional Signals	Reset Configuration Name	Value (Binary)	Meaning
TSEC1_TXD[6:4], TSEC1_TX_ER Default: 1111	cfg_rom_loc[0:3]	0000	PCI Express 1
		0001	PCI Express 2
		0010	Serial RapidIO 1
		0011	Serial RapidIO 2
		0100	DDR controller
		0101	PCI Express 3
		0110	On-chip boot ROM-SPI configuration
		0111	On-chip boot ROM-eSDHC configuration
		1000	Local bus FCM-8-bit NAND flash small page
		1001	Reserved
		1010	Local bus FCM-8-bit NAND flash large page
		1011	Reserved
		1100	Reserved
		1101	Local bus GPCM-8-bit ROM
		1110	Local bus GPCM-16-bit ROM
1111	Local bus GPCM-16-bit ROM		

Note that the values latched on these signals during POR are accessible through the memory-mapped PORBMSR (POR boot mode status register) described in [POR boot mode status register \(GUTS_PORBMSR\)](#).

See [Local address map example](#), for an example memory map that relies on the default boot ROM values. Also, see [Boot page translation register \(Reset_BPTR\)](#) for information on translation of the boot page.

4.5.3.5 Host/agent configuration

The host/agent reset configuration inputs, shown in this table, configure the chip to act as a host or as an agent of a master on another interface.

In host mode, the chip is immediately enabled to master transactions on the RapidIO interface. If the chip is an agent on the PCI Express or RapidIO interfaces, then the chip is disabled from mastering transactions on that interface until the external host enables it to do so. The external host does this by setting the control registers of the chip's interfaces appropriately. See details in the [PCI Express Interface Controller](#) and [Serial RapidIO Interface](#) respectively.

Note that the values latched on these signals during POR are accessible through the memory-mapped PORBMSR (POR boot mode status register) described in [POR boot mode status register \(GUTS_PORBMSR\)](#).

Table 4-13. Host/agent configuration

Functional signals	Reset configuration name	Value (Binary)	Meaning
LWET/LBST LA[18:19] Default (111)	cfg_host_agt[0:2]	000	The chip acts as an agent on all its PCI Express and SRIO interfaces.
		001	The chip acts as an agent on PCI Express 1 or host on Serial RapidIO 2. The chip acts as a host on PCI Express 2/Serial RapidIO 1. The chip acts as a host on PCI Express 3.
		010	The chip acts as a host on PCI Express 1 or agent on Serial RapidIO 2. The chip acts as an agent on PCI Express 2/Serial RapidIO 1. The chip acts as a host on PCI Express 3
		011	The chip acts as a host on PCI Express 1/Serial RapidIO 2. The chip acts as a host on PCI Express 2/Serial RapidIO 1. The chip acts as an agent on PCI Express 3
		100	The chip acts as an agent on PCI Express 1/Serial RapidIO 2. The chip acts as an agent on PCI Express 2/Serial RapidIO 1. The chip acts as a host on PCI Express 3
		101	The chip acts as an agent on PCI Express 1 or host on Serial RapidIO 2. The chip acts as a host on PCI Express 2/Serial RapidIO 1. The chip acts as an agent on PCI Express
		110	The chip acts as a host on PCI Express 1 or agent on Serial RapidIO 2. The chip acts as an agent on PCI Express 2/Serial RapidIO 1. The chip acts as an agent on PCI Express
		111	The chip acts as the host processor/root complex for all PCI Express/Serial RapidIO interfaces (default).

4.5.3.6 I/O port selection

The chip can be configured with different I/O ports active.

This table shows the configuration of I/O ports and bit rates (and required reference clocks) that are possible for the serial RapidIO, SGMII and PCI Express interfaces.

The POR configuration input `cfg_sgmi2` and `cfg_sgmi3`, in combination with `cfg_io_ports[0:3]`, determines whether SerDes lanes 2 and 3 are used by eTSEC2 and eTSEC3 respectively in SGMII mode.

Please note the following SerDes reference clock restrictions (as defined in [SerDes reference clock configuration](#)):

- For 100 MHz reference clock (PCI Express operation selected), `cfg_srds_refclk` must be set to 1.
- For 125 MHz reference clock (SRIO 3.125 Gbps operation selected), `cfg_srds_refclk` must be set to 0.

Table 4-14. I/O port selection

Functional signal	Reset configuration name	Value (Binary)	Meaning	SerDes reference clock speed
TSEC1_TXD[3:1], TSEC2_TX_ER Default (1111)	cfg_io_ports[0:3]	0000	PCI Express 1 (x1) (2.5Gbps) → SerDes lane 0 SerDes lanes 1-3 powered down.	100 MHz
		0001	SerDes lanes 0-3 powered down.	-
		0010	PCI Express 1 (x1) (2.5 Gbps) → SerDes lane 0 PCI Express 2 (x1) (2.5 Gbps) → SerDes lane 1 PCI Express 3 (x2) (2.5 Gbps) → SerDes lanes 2-3	100 MHz
		0011	Reserved	-
		0100	PCI Express 1 (x2) (2.5 Gbps) → SerDes lanes 0-1 PCI Express 3 (x2) (2.5 Gbps) → SerDes lanes 2-3	100 MHz
		0101	Reserved	-
		0110	PCI Express 1 (x4) (2.5 Gbps) → SerDes lanes 0-3	100 MHz
		0111	SRIO2 (1x) (3.125 Gbps) → SerDes lane 0 SRIO1 (1x) (3.125 Gbps) → SerDes lane 1 SerDes lanes 2-3 powered down.	125 MHz
		1000	SRIO2 (4x) (1.25 Gbps) → SerDes lanes 0-3	100 MHz 125 MHz
		1001	SRIO2 (4x) (2.5 Gbps) → SerDes lanes 0-3	100 MHz 125 MHz
		1010	SRIO2 (4x) (3.125 Gbps) → SerDes lanes 0-3	125 MHz
		1011	SRIO 2 (1x) (1.25 Gbps) → SerDes lane 0 SRIO 1 (1x) (1.25 Gbps) → SerDes lane 1 SGMII eTSEC2 (x1) (1.25 Gbps) → SerDes lane 2 (see Note1) SGMII eTSEC3 (x1) (1.25 Gbps) → SerDes lane 3 (see Note2)	100 MHz 125 MHz
		1100	SRIO 2 (1x) (2.5 Gbps) → SerDes lane 0 SRIO 1 (1x) (2.5 Gbps) → SerDes lane 1 SGMII eTSEC2 (x1) (1.25 Gbps) → SerDes lane 2 (see Note1) SGMII eTSEC3 (x1) (1.25 Gbps) → SerDes lane 3 (see Note2)	100 MHz 125 MHz
		1101	PCI Express 1 (x1) (2.5 Gbps) → SerDes lane 0 SRIO 1 (1x) (2.5 Gbps) → SerDes lane 1 SGMII eTSEC2 (x1) (1.25 Gbps) → SerDes lane 2 (see Note1) SGMII eTSEC3 (x1) (1.25 Gbps) → SerDes lane 3 (see Note2)	100 MHz
		1110	PCI Express 1 (x1) (2.5 Gbps) → SerDes lane 0 PCI Express 2 (x1) (2.5 Gbps) → SerDes lane 1 SGMII eTSEC2 (x1) (1.25 Gbps) → SerDes lane 2 (see Note1) SGMII eTSEC3 (x1) (1.25 Gbps) → SerDes lane 3 (see Note2)	100 MHz

Table continues on the next page...

Table 4-14. I/O port selection (continued)

Functional signal	Reset configuration name	Value (Binary)	Meaning	SerDes reference clock speed
TSEC1_TXD[3:1], TSEC2_TX_ER Default (1111)	cfg_IO_ports[0:3]	1111	PCI Express 1 (x2) (2.5 Gbps) → SerDes lanes 0-1 SGMII eTSEC2 (x1) (1.25 Gbps) → SerDes lane 2 (see Note1) SGMII eTSEC3 (x1) (1.25 Gbps) → SerDes lane 3 (see Note2)	100 MHz
Notes:				
¹ Port cfg_sgmii2 must also be logic 0 in addition for eTSEC2 to operate in SGMII mode.				
² Port cfg_sgmii3 must also be logic 0 in addition for eTSEC3 to operate in SGMII mode.				

4.5.3.7 CPU boot configuration

The CPU boot configuration input, shown in this table, specifies the boot configuration mode.

If `cfg_cpu0_boot` (LA27) is sampled low at reset, e500 core 0 is prevented from fetching boot code until configuration by an external master is complete. Similarly, `cfg_cpu1_boot` (LA16) can be used to gate off e500 core 1 from fetching boot code. The external master frees the core 0 and/or core 1 to boot by setting `EEBPCR[CPU0_EN]` and/or `EEBPCR[CPU1_EN]` in the ECM CCB port configuration register (`EEBPCR`). See [ECM CCB port configuration register \(ECM_EEBPCR\)](#) for more information.

Note that the value latched on these signals during POR are accessible through the memory-mapped `PORBMSR` (POR boot mode status register) described in [POR boot mode status register \(GUTS_PORBMSR\)](#).

Note also that the values latched on these signals during POR affect all of the PCI Express interfaces' configuration ready mode settings. See [Configuration Ready Register \(Configuration_Ready_Register\)](#). When the specific PCI Express module is configured for agent mode and both cores are configured to be in boot holdoff, the associated configuration register `PEX_CFG_READY[CFG_READY]` will have a reset value of 1, which indicates a configuration complete status to the transaction layer.

Table 4-15. CPU boot configuration

Functional signal	Reset configuration name	Value (Binary)	Meaning
LA27, LA16 No Default	cfg_cpu0_boot, cfg_cpu1_boot	00	CPU boot holdoff mode for both cores. The e500 cores are prevented from booting until configured by an external master.
		01	e500 core 1 is allowed to boot without waiting for configuration by an external master, while e500 core 0 is prevented from booting until configured by an external master or the other core.
		10	e500 core 0 is allowed to boot without waiting for configuration by an external master, while e500 core 1 is prevented from booting until configured by an external master or the other core.
		11	Both e500 cores are allowed to boot without waiting for configuration by an external master.

4.5.3.8 Boot sequencer configuration

The boot sequencer configuration options, shown in this table, allow the boot sequencer to load configuration data from the serial ROM located on the I²C1 port before the host tries to configure the chip.

These options also specify normal or extended I²C addressing modes. See [Boot sequencer mode](#) for more information on the boot sequencer.

Note that the values latched on these signals during POR are accessible through the memory-mapped PORBMSR (POR boot mode status register) described in [POR boot mode status register \(GUTS_PORBMSR\)](#).

Table 4-16. Boot sequencer configuration

Functional signal	Reset configuration name	Value (Binary)	Meaning
LGPL3/LFWP, LGPL5 Default (11)	cfg_boot_seq[0:1]	00	Reserved
		01	Normal I ² C addressing mode is used. Boot sequencer is enabled and loads configuration information from a ROM on the I ² C1 interface. A valid ROM must be present.
		10	Extended I ² C addressing mode is used. Boot sequencer is enabled and loads configuration information from a ROM on the I ² C1 interface. A valid ROM must be present.
		11	Boot sequencer is disabled. No I ² C ROM is accessed (default).

NOTE

When the boot sequencer is enabled, the processor core will be held in reset and thus prevented from fetching boot code until the boot sequencer has completed its task, regardless of the

state of the CPU boot configuration signal described in [CPU boot configuration](#).

4.5.3.9 DDR SDRAM type

DDR3 requires a different voltage level from DDR2.

This table describes the configuration of the DDR SDRAM type.

Table 4-17. DDR DRAM type

Functional signal	Reset configuration name	Value (Binary)	Meaning
TSEC2_TXD1 Default (1)	cfg_dram_type	0	DDR2 1.8 V, CKE low at reset
		1	DDR3 1.5 V, CKE low at reset (default)

4.5.3.10 SerDes reference clock configuration

As shown in this table, two options are available for the frequency of the input SerDes reference clock-either a 100-MHz or 125-MHz LVDS differential clock.

This one clock is applied to an internal PLL whose output creates the clocks used by all four SerDes lanes. The result is always a 1.25-Gbaud transmission/receive rate on each SGMII lane and a 2.5-Gbaud rate on each PCI Express lane.

Please note the following restrictions:

- For any SerDes configuration with PCI-Express active, a 100 MHz reference clock must be supplied.
- For SRIO 3.125-Gbaud operation, a 125-MHz reference clock must be supplied.

Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVS2 (POR device status 2 register) described in [POR device status register 2 \(GUTS_PORDEVS2\)](#).

Table 4-18. SerDes reference clock configuration

Functional signal	Reset configuration name	Value (Binary)	Meaning
TSEC_1588_ALARM_OUT1 Default (1)	cfg_srds_refclk	0	SerDes expects a 125 MHz reference clock frequency.
		1	SerDes expects a 100 MHz reference clock frequency (default).

4.5.3.11 eTSEC_n physical interface mode configuration

Table 4-19 through Table 4-21 show the eTSEC_n POR configuration summary.

This table shows the eTSEC1 POR configuration summary.

Table 4-19. eTSEC1 POR configuration summary

eTSEC1 configuration	cfg_tsec_reduce ¹	cfg_tsec1_ptcl
RMII	0	01
RGMII	0	10
RTBI	0	11
MII	1	01
GMII	1	10
TBI	1	11

1. cfg_tsec_reduce polarity matches the MPC8572E definition.

This table shows the eTSEC2 POR configuration summary.

Table 4-20. eTSEC2 POR configuration summary

eTSEC2 configuration	cfg_sgmi2 ¹	cfg_tsec_reduce ²	cfg_tsec1_ptcl	cfg_tsec2_ptcl
SGMII	0	-	-	-
RMII	1	0	-	01
RGMII	1	0	-	10
RTBI	1	0	-	11
MII	1	1	01 or 10 or 11	01
GMII	1	1	01 or 10 or 11	10
TBI	1	1	01 or 10 or 11	11

1. cfg_sgmi2 polarity matches the MPC8572E definition.

2. cfg_tsec_reduce polarity matches MPC8572E definition.

This table shows the eTSEC3 POR configuration summary.

Table 4-21. eTSEC3 POR configuration summary

eTSEC3 Configuration	cfg_sgmi3 ¹	cfg_tsec_reduce ²	cfg_tsec3_ptcl
SGMII	0	-	-
RMII	1	0	01
RGMII	1	0	10
RTBI	1	0	11

1. cfg_sgmi3 polarity matches the MPC8572E definition.
2. cfg_tsec_reduce polarity matches MPC8572E definition.

4.5.3.11.1 eTSEC2 SGMII mode

The eTSEC2 SGMII mode input, shown in this table, selects SGMII mode versus parallel mode for enhanced three-speed Ethernet controller (eTSEC) interface 2.

If eTSEC2 is configured to run in SGMII mode, none of the parallel mode configuration inputs are pertinent for configuring eTSEC2's interface. When operating in SGMII mode, the parallel interface pins normally used for eTSEC2 interface are not used, but rather the corresponding SGMII SerDes lane is used.

Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVSR (POR device status register) described in [POR device status register \(GUTS_PORDEVSR\)](#).

Table 4-22. eTSEC2 SGMII mode configuration

Function al signal	Reset configuration name	Value (Binary)	Meaning
LGPL1 Default (1)	cfg_sgmi2	0	eTSEC2 Ethernet interface operates in SGMII mode and uses SGMII SerDes lane 2 pins.
		1	eTSEC2 Ethernet interface operates in standard parallel interface mode and uses the TSEC2_* pins (default).

4.5.3.11.2 eTSEC3 SGMII mode

The eTSEC3 SGMII mode input, shown in this table, selects SGMII mode versus parallel mode for enhanced three-speed Ethernet controller (eTSEC) interface 3.

If eTSEC3 is configured to run in SGMII mode, none of the parallel mode configuration inputs are pertinent for configuring eTSEC3's interface. When operating in SGMII mode, the parallel interface pins normally used for eTSEC3 interface are not used, but rather the corresponding SGMII Serdes lane is used.

Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVSR (POR device status register) described in [POR device status register \(GUTS_PORDEVSR\)](#).

Table 4-23. eTSEC3 SGMII mode configuration

Functional signal	Reset configuration name	Value (Binary)	Meaning
TSEC_1588_ALARM_OUT2 Default (1)	cfg_sgmii3	0	eTSEC3 Ethernet interface operates in SGMII mode and uses SGMII SerDes lane 3 pins.
		1	eTSEC3 Ethernet interface operates in standard parallel interface mode and uses the TSEC3_* pins (default).

4.5.3.11.3 eTSEC1 and eTSEC2 width

The eTSEC width input, shown in this table, selects standard versus reduced width for three-speed Ethernet controller interfaces 1 and 2 operating in parallel mode.

This input only affects operation of eTSEC2 if it is not configured to operate in SGMII mode.

The value latched on this signal during POR is accessible through the memory-mapped PORDEVSR (POR device status register) described in [POR device status register \(GUTS_PORDEVSR\)](#).

Table 4-24. eTSEC1 width configuration

Functional signals	Reset configuration name	Value (Binary)	Meaning
EC_MDC Default (1)	cfg_tsec_reduce	0	eTSEC1 and eTSEC2 Ethernet interfaces operate in reduced pin mode (either RTBI, RGMII, or RMII mode).
		1	eTSEC1 and eTSEC2 Ethernet interfaces operate in their standard width TBI, GMII, or MII mode.

4.5.3.11.4 eTSEC1 protocol

The eTSEC1 protocol inputs, shown in this table, select the protocol (MII, GMII or TBI) used by the eTSEC1 controller when operating in parallel mode.

Note that the value latched on these signals during POR is accessible through the memory-mapped PORDEVSR (POR device status register) described in [POR device status register \(GUTS_PORDEVSR\)](#).

Table 4-25. eTSEC1 protocol configuration

Functional signals	Reset configuration name	Value (Binary)	Meaning
TSEC1_TXD0, TSEC1_TXD7 Default (11)	cfg_tsec1_prctl[0:1]	00	Reserved
		01	The eTSEC1 controller operates using the MII protocol (or RMII if configured in reduced mode as described in eTSEC1 and eTSEC2 width ”).
		10	The eTSEC1 controller operates using the GMII protocol (or RGMII if configured in reduced mode as described in eTSEC1 and eTSEC2 width ”).
		11	The eTSEC1 controller operates using the TBI protocol (or RTBI if configured in reduced mode as described in eTSEC1 and eTSEC2 width ”) (default).

4.5.3.11.5 eTSEC2 protocol

The eTSEC2 protocol inputs, shown in [Table 4-26](#), select the protocol (MII, GMII or TBI) used by the eTSEC2 controller when operating in parallel mode. This input only affects operation of eTSEC2 if it is not configured to operate in SGMII mode.

Note that the value latched on these signals during POR is accessible through the memory-mapped PORDEVSR (POR device status register) described in [POR device status register \(GUTS_PORDEVSR\)](#).

Table 4-26. eTSEC2 protocol configuration

Functional signals	Reset configuration name	Value (Binary)	Meaning
TSEC2_TXD0, TSEC2_TXD7 Default (11)	cfg_tsec2_prctl[0:1]	00	Reserved
		01	The eTSEC2 controller operates using the MII protocol (or RMII if configured in reduced mode as described in eTSEC1 and eTSEC2 width ”) if not configured to operate in SGMII mode.
		10	The eTSEC2 controller operates using the GMII protocol (or RGMII if configured in reduced mode as described in eTSEC1 and eTSEC2 width ”) if not configured to operate in SGMII mode.
		11	The eTSEC2 controller operates using the TBI protocol (or RTBI if configured in reduced mode as described in eTSEC1 and eTSEC2 width ”) if not configured to operate in SGMII mode. (default).

4.5.3.11.6 eTSEC3 protocol

The eTSEC3 protocol inputs, shown in [Table 4-27](#), select the protocol (RMII, RGMII or RTBI) used by the eTSEC3 controller when operating in parallel mode. This input only affects operation of eTSEC3 if it is not configured to operate in SGMII mode.

The value latched on these signals during POR is accessible through the memory-mapped PORDEVSR (POR device status register) described in [POR device status register \(GUTS_PORDEVSR\)](#).

Note that parallel mode operation for eTSEC3 is only available when the eTSEC1 and eTSEC2 are configured for reduced-mode width. See [Table 1-1](#) for more information.

Table 4-27. eTSEC3 protocol configuration

Functional signals	Reset configuration name	Value (Binary)	Meaning
UART_RTS0_B, UART_RTS1_B Default (11)	cfg_tsec3_prctl[0:1]	00	Reserved.
		01	The eTSEC3 controller operates using the RMII protocol if not configured to operate in SGMII mode.
		10	The eTSEC3 controller operates using the RGMII protocol if not configured to operate in SGMII mode.
		11	The eTSEC3 controller operates using the RTBI protocol if not configured to operate in SGMII mode (default).

4.5.3.12 RapidIO device ID

The RapidIO device ID inputs, shown in this table, specify the three lower-order bits of the device ID of the P2020 as used by RapidIO hosts.

Note that the 5 high-order RapidIO device ID bits cannot be set through the POR configuration inputs. They may be initialized by the boot sequencer or by the processor from boot ROM or by the RapidIO discovery process.

Table 4-28. RapidIO device ID

Functional signals	Reset configuration name	Meaning
TSEC2_TXD2	cfg_device_ID5	Device ID used for RapidIO hosts
TSEC2_TXD3	cfg_device_ID6	Device ID used for RapidIO hosts
TSEC2_TXD4	cfg_device_ID7	Device ID used for RapidIO hosts

functional description

If configured as a RapidIO host, the upper-order device ID bits default to zeros. If configured as a RapidIO agent, the upper-order device ID bits default to ones. Unconnected `cfg_device_IDn` inputs default to 1s regardless of the host/agent mode configuration.

Note that the value latched on this signal at POR is accessible through the memory-mapped PORDEVSR described in [POR device status register \(GUTS_PORDEVSR\)](#) as well as the memory-mapped BDIDCSR described in [Base device ID command and status register \(SRIO_BDIDCSR\)](#).

4.5.3.13 RapidIO system size

The RapidIO common transport specification defines two system sizes.

The large size uses 16-bit source and destination IDs allowing for 65,536 devices, while the small size uses 8-bit source and destination IDs, supporting 256 devices.

The RapidIO system size configuration shown in this table specifies which system size is to be used by the RapidIO controller on the chip.

Note that the value latched on this signal at POR is accessible through the memory-mapped PORDEVSR described in [POR device status register 2 \(GUTS_PORDEVSR2\)](#) as well as the memory-mapped PEFCAR described in [Processing element features capability register \(SRIO_PEFCAR\)](#).

Table 4-29. RapidIO system size

Functional signals	Reset configuration name	Value (Binary)	Meaning
LGPL0 Default (1)	cfg_rio_sys_size	0	Large system size (up to 65,536 devices)
		1	Small system size (up to 256 devices) (default)

4.5.3.14 Memory debug configuration

The memory debug configuration input, shown in this table, selects which debug outputs (DDR or eLBC memory controller) are driven onto the MSRCID and MDVAL debug signals.

Note that the value latched on this signal during POR is accessible through the memory-mapped PORDBGMSR (POR debug mode register) described in [POR debug mode status register \(GUTS_PORDBGMSR\)](#).

Table 4-30. Memory debug configuration

Functional signal	Reset configuration name	Value (Binary)	Meaning
DMA2_DACK0_B Default (1)	cfg_mem_debug	0	Debug information from the enhanced local bus controller (eLBC) is driven on the MSRCID and MDVAL signals
		1	Debug information from the DDR SDRAM controller is driven on the MSRCID and MDVAL signals (default).

4.5.3.15 DDR debug configuration

The DDR debug configuration input, shown in this table, enables a DDR memory controller debug mode in which the DDR SDRAM source ID field and data valid strobe are driven onto the ECC pins.

ECC checking and generation are disabled in this case. ECC signals driven from the SDRAMs must be electrically disconnected from the ECC I/O pins of the chip in this mode.

Table 4-31. DDR debug configuration

Functional signal	Reset configuration name	Value (Binary)	Meaning
DMA2_DDONE0_B Default (1)	cfg_ddr_debug	0	Debug information is driven on the ECC pins instead of normal ECC I/O. ECC signals from memory devices must be disconnected.
		1	Debug information is not driven on ECC pins. ECC pins function in their normal mode (default).

Note that the value latched on this signal during POR is accessible through the memory-mapped PORDBGMSR (POR debug mode register) described in [POR debug mode status register \(GUTS_PORDBGMSR\)](#).

4.5.3.16 General-purpose POR configuration

The eLBC address/data bus inputs, shown in this table, configure the value of the general-purpose POR configuration register defined in General-Purpose POR Configuration Register (GPPORCR).

functional description

This register is intended to facilitate POR configuration of user systems. A value placed on LAD[0:31] during POR is captured and stored (read only) in the GPPORCR. Software can then use this value to inform the operating system about initial system configuration. Typical interpretations include circuit board type, board ID number, or a list of available peripherals.

Table 4-32. General-purpose POR configuration

Functional signals	Reset configuration name	Value (Binary)	Meaning
LAD[0:15] Default (1)	cfg_gpinput[0:15]	-	General-purpose POR configuration vector to be placed in GPPORCR

4.5.3.17 Engineering use POR configuration

These POR configuration inputs may be used in the future to control functionality.

It is advised that boards are built with the ability to pull up or pull down these pins. Note that the value latched on these signals during POR are accessible through the PORDEVSR2, described in [POR device status register 2 \(GUTS_PORDEVSR2\)](#).

Table 4-33. Engineering use

Functional signals	Reset configuration name	Value (Binary)	Meaning
LA[20:22], UART_SOUT0, MSRCID1, MSRCID4, DMA1_DDONE_B Default (11111111)	cfg_eng_use[0:3, 5:7]	00000000-11111110	Reserved
		11111111	Default operation

4.5.3.18 eLBC ECC enable configuration

This POR configuration input is used to enable eLBC ECC checking on the external local bus interface on boot.

Table 4-34. eLBC ECC enable

Functional Signals	Reset Configuration Name	Value (Binary)	Meaning
MSRCID0 Default (1)	cfg_elbc_ecc	0	eLBC ECC checking is disabled.
		1	Default operation: eLBC ECC checking is enabled.

4.5.3.19 System speed

The SYSCLK speed configuration input configures internal logic for proper operation with the SYSCLK clock frequencies in use.

The default setting is appropriate for SYSCLK operating at or above 66 MHz. For low speed operation (SYSCLK below 66 MHz) this POR configuration input should be low during HRESET. If this configuration is not set properly, behavior of the system may be unreliable. Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVSR2, described in [POR device status register 2 \(GUTS_PORDEVSR2\)](#).

Table 4-35. System speed

Functional signals	Reset configuration name	Value (Binary)	Meaning
LA28 Default (1)	cfg_sys_speed	0	SYSCLK frequency is below 66 MHz.
		1	SYSCLK frequency is at or above 66 MHz.

4.5.3.20 Platform speed

The platform speed configuration input configures internal logic for proper operation with the platform clock frequencies in use. The default setting is appropriate for platform operating at or above 333 MHz. For low speed operation (platform below 333 MHz) this POR configuration input should be low during HRESET. If this configuration is not set properly, behavior of the system may be unreliable. Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVSR2, described in [POR device status register 2 \(GUTS_PORDEVSR2\)](#).

Table 4-36. Platform speed

Functional signals	Reset configuration name	Value (Binary)	Meaning
LA23 Default (1)	cfg_plat_speed	0	Platform clock frequency is below 333 MHz.
		1	Platform clock frequency is at or above 333 MHz.

4.5.3.21 Core 0 speed

The core 0 speed configuration input, shown in this table, configures internal logic for proper operation with the core 0 clock frequencies in use.

The default setting is appropriate for core 0 operating above 1000 MHz. For low speed operation (core 0 at or below 1000 MHz) this POR configuration input should be low during HRESET. If this configuration is not set properly, behavior of the system may be unreliable. Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVSR2, described in [POR device status register 2 \(GUTS_PORDEVSR2\)](#).

Table 4-37. Core 0 speed

Functional signals	Reset configuration name	Value (Binary)	Meaning
LA24	cfg_core0_speed	0	Core 0 clock frequency is less than or equal to 1000 MHz.
Default (1)		1	Core 0 clock frequency is greater than 1000 MHz.

4.5.3.22 Core 1 speed

The core 1 speed configuration input, shown in this table, configures internal logic for proper operation with the core 1 clock frequencies in use.

The default setting is appropriate for core 1 operating above 1000 MHz. For low speed operation (core 1 at or below 1000 MHz) this POR configuration input should be low during HRESET. If this configuration is not set properly, behavior of the system may be unreliable. Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVSR2, described in [POR Device Status Register 2 \(PORDEVSR2\)](#).

Table 4-38. Core 1 speed

Functional signals	Reset configuration name	Value (Binary)	Meaning
LA25	cfg_core1_speed	0	Core 1 clock frequency is less than or equal to 1000 MHz.
Default (1)		1	Core 1 clock frequency is greater than 1000 MHz.

4.5.3.23 DDR controller speed

The DDR controller speed configuration input, shown in this table, configures internal logic for proper operation with the DDR controller clock frequencies in use.

The DDR controller speed configuration input is only used by the internal logic when the DDR PLL ratio is not set to synchronous mode, see [DDR PLL Ratio](#).

The default setting is appropriate for the DDR controller operating at or above 500 MHz. For low speed operation (DDR controller below 500 MHz) this POR configuration input should be low during HRESET. If this configuration is not set properly, behavior of the system may be unreliable. Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVS2, described in [POR device status register 2 \(GUTS_PORDEVS2\)](#).

Table 4-39. DDR speed

Functional signals	Reset configuration name	Value (Binary)	Meaning
LA26 Default (1)	cfg_ddr_speed	0	DDR controller clock frequency is less than 500 MHz.
		1	DDR controller clock frequency is greater than or equal to 500 MHz.

4.5.3.24 SerDes PLL time-out enable

The SerDes PLL time-out enable input configures the reset sequence to enable the SerDes PLL time-out counter.

This configuration allows for the reset sequence to avoid hanging if a SerDes PLL reference clock is not provided.

When this counter is enabled, the reset sequence waits for the SerDes PLL to lock while the time-out counter has not expired. If the time-out counter expires before SerDes PLL locks, the reset sequence will proceed with the following modules disabled: SerDes, PCI-Express, SRIO, RMU, SGMII interface. The [Device disable control register \(GUTS_DEVDISR\)](#) reflects these modules as being disabled.

When this counter is disabled, the the reset sequence will wait indefinitely for the SerDes PLL to lock.

This table shows the SerDes PLL time-out enable.

Table 4-40. SerDes PLL time-out enable

Functional signals	Reset configuration name	Value (Binary)	Meaning
TRIG_OUT Default (1)	cfg_srds_pll_toe	0	Enable PLL lock time-out counter. The power-on-reset sequence waits for the SerDes PLL to lock while the time-out counter has not expired.
		1	Disable PLL lock time-out counter. The power-on-reset sequence waits indefinitely for the SerDes PLL to lock. (default)

4.5.3.25 eSDHC card-detect polarity select

The eSDHC card-detect polarity select pin sets the polarity of the eSDHC card-detect pin. Note that the value latched on this signal during POR is accessible through the memory-mapped PORDEVSR2, described in [POR device status register 2 \(GUTS_PORDEVSR2\)](#).

This table shows the eSDHC card-detect polarity select.

Table 4-41. SDHC card detect polarity select

Functional signals	Reset configuration name	Value (Binary)	Meaning
TSEC2_TXD[5] Default (1)	cfg_sdhc_cd_pol_sel	0	The eSDHC card-detect polarity is inverted.
		1	The eSDHC card-detect polarity is not inverted.(default)

4.5.4 Voltage selection configuration

The chip supports multiple supply voltages on its I/O supplies.

The signals used for these voltage selections are not POR configurations. This section describes the encoding used to select the voltage level for each I/O supply.

NOTE

Incorrect voltage select settings can lead to irreversible device damage. Please follow this section carefully.

For proper use, the voltage select device input signals LVDD_SEL, BVDD_VSEL[0:1], and CVDD_VSEL[0:1] must be statically tied to reflect the voltage applied on the LVDD, BVDD, and CVDD I/O supplies respectively.

Table 4-42. I/O Supply voltage select settings

Voltage select input	Voltage select setting ¹	Supply voltage	Interfaces effected
LVDD_VSEL	0	LVDD = 3.3 V	eTSEC1,2,3; Ethernet management; 1588
	1	LVDD = 2.5 V	
BVDD_VSEL[0:1]	00	BVDD = 3.3 V	Local Bus, GPIO[8:15]
	01	BVDD = 2.5 V	
	10	BVDD = 1.8 V	
	11	BVDD = 3.3 V	

Table continues on the next page...

Table 4-42. I/O Supply voltage select settings (continued)

Voltage select input	Voltage select setting ¹	Supply voltage	Interfaces effected
CVDD_VSEL[0:1]	00	CVDD = 3.3 V	USB, eSDHC, SPI
	01	CVDD = 2.5 V	
	10	CVDD = 1.8 V	
	11	CVDD = 3.3 V	

1. Logic 0 corresponds with a static tie to GND, while a logic 1 corresponds with a static tie to OVDD (3.3 V).

As an example, for local bus operation at 2.5 V, the BVDD_VSEL[0] and BVDD_VSEL[1] device inputs must be configured to 01 and thus be tied on the board to GND and OVDD respectively. For 2.5 V operation, tying BVDD_VSEL[0] and BVDD_VSEL[1] to anything except GND and OVDD respectively can lead to irreversible device damage.

4.5.5 Clocking

The following sections describe the clocking within the chip.

4.5.5.1 System clock and DDR controller clock

The chip takes a single input clock, SYSCLK, as its primary clock source for the e500 core and all of the devices and interfaces that operate synchronously with the core.

As shown in the figure below, the SYSCLK input (frequency) is multiplied up using a phase lock loop (PLL) to create the core complex bus (CCB) clock (also called the platform clock). The CCB clock is used by virtually all of the synchronous system logic, including the L2 cache, and other internal blocks such as the DMA and interrupt controller. The CCB clock also feeds the PLLs in the e500 cores and the PLL that create clocks for the local bus memory controller. Note that the divide-by-two CCB clock divider and the divide-by-*n* CCB clock divider, shown in the figure below, are located in the DDR and local bus blocks, respectively.

The DDR memory controller may use the platform clock and thus have operation of both DDR interfaces be synchronous with the platform. Alternately, an independent clock, DDRCLK, may be multiplied up using a separate PLL to create a unique DDR memory controller clock. In this case, the DDR controller operates asynchronous with respect to the platform clock.

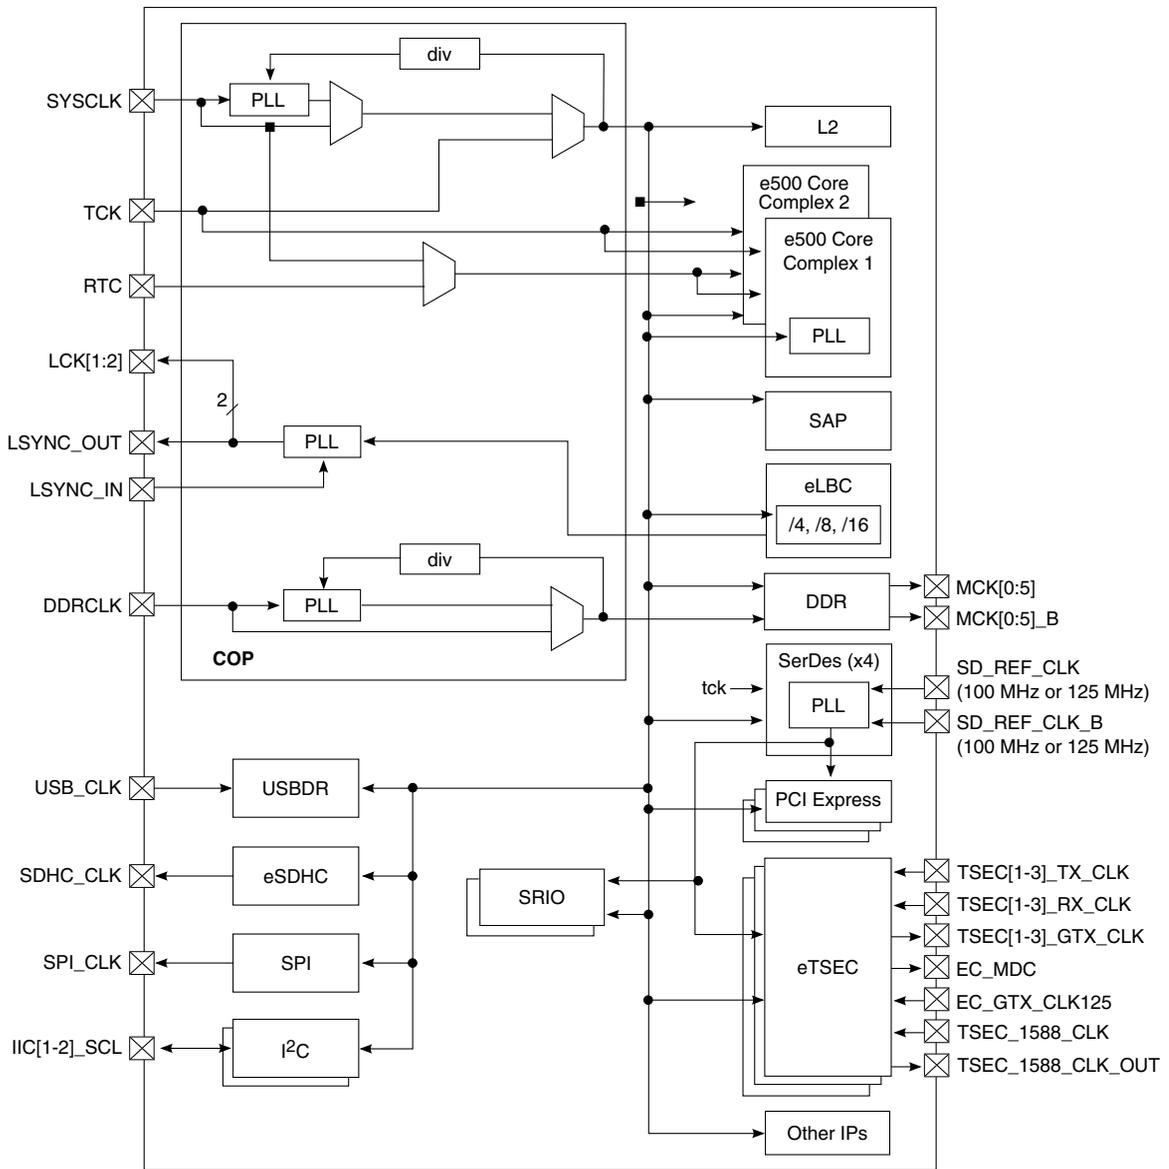


Figure 4-6. Clock subsystem block diagram

4.5.5.2 RapidIO, SGMII and PCI Express clocks

Clocks for these high speed interfaces are derived from a PLL in the SerDes block.

This PLL is driven by a reference clock ($\overline{\text{SD_REF_CLK}}$) whose input frequency is a function of the protocol and bit rate being used as shown in this table.

Table 4-43. High speed interface clocking

Interfaces	Bit rate (Gbps)	Reference clock frequency (MHz)
PCI Express	2.5	100
SGMII	1.25	100 or 125
Serial RapidIO	3.125	125
	2.5	100 or 125
	1.25	100 or 125

4.5.5.2.1 Minimum frequency requirements

[I/O port selection](#) describes various high-speed interface configuration options. CCB clock frequency must be considered for proper operation of such interfaces.

Please refer to the device hardware specifications, for specific supported frequencies.

4.5.5.3 Ethernet clocks

The Ethernet blocks operate asynchronously with respect to the rest of the device.

These blocks use receive and transmit clocks supplied by their respective PHY chips, plus a 125-MHz clock input for gigabit protocols. Data transfers are synchronized to the CCB clock internally.

4.5.5.4 Real time clock

As shown in this figure, the real time clock (RTC) input can optionally be used to clock the e500 core timer facilities.

RTC can also be used (optionally) by the chip programmable interrupt controller (PIC) global timer facilities. The RTC is separate from the e500 core clock and is intended to support relatively low frequency timing applications. The RTC frequency range is specified in the device hardware specifications, but the maximum value should not exceed one-quarter of the CCB frequency.

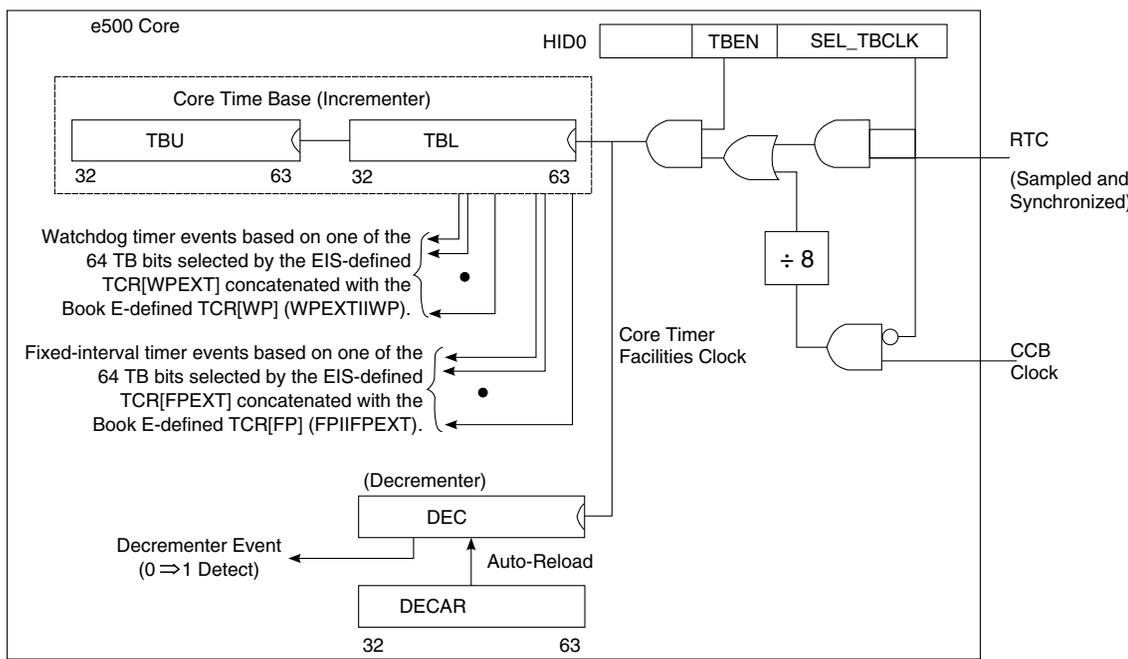
Before being distributed to the core time base, RTC is sampled and synchronized with the CCB clock.

The clock source for the core time base is specified by two fields in HID0: time base enable (TBEN), and select time base clock (SEL_TBCLK). If the time base is enabled, (HID0[TBEN] is set), the clock source is determined as follows:

- HID0[SEL_TBCLK] = 0, the time base is updated every 8 CCB clocks
- HID0[SEL_TBCLK] = 1, the time base is updated on the rising edge of RTC

The default source of the time base is the CCB clock divided by eight. For more details, see the *PowerPC e500 Core Complex Reference Manual*.

[Timer control register group n \(PIC_TCRn\)](#) provides additional information on the use of the RTC signal to clock the global timers in the PIC unit.



Note: The logic circuits shown depict functional relationships only; they do not represent physical implementation details.

Figure 4-7. RTC and core timer facilities clocking options

4.6 Initialization/applications information

Selecting on-chip ROM in boot ROM location (see [Table 4-12](#)), causes the e500 CPU to fetch data from the on-chip ROM. The on-chip ROM is selected by configuring the POR config pins *cfg_rom_loc[0:3]*. Two different configurations are provided for boot from the on-chip ROM—boot from eSPI and from eSDHC.

More detailed information is provided in the following sections:

- [eSDHC boot](#)

- eSPI boot ROM
- Default e500 addressing during system boot

4.6.1 eSDHC boot

This section provides an overview, list of features, and SD/MMC card data structures for the eSDHC boot.

4.6.1.1 eSDHC boot overview

The chip is capable of loading initialization code from a memory device that is connected to the eSDHC controller interface.

This device can be either a SD card or MMC card or other variants compatible with these devices. The term SD/MMC is used while referring to the memory device.

Boot from eSDHC is supported by the chip using an on-chip ROM which contains the basic eSDHC device driver and the code to perform block copy from SD/MMC to any target memory. Selecting on-chip ROM in boot ROM location (see [Table 4-12](#)) causes the e500 CPU to fetch data from the on-chip ROM. The on-chip ROM is selected by configuring the POR config pins *cfg_rom_loc[0:3]*. Prior to boot, the user must ensure that the SD/MMC card to boot from is inserted.

After the device has completed the reset sequence, the e500 core executes code from the internal on-chip ROM if the ROM location selects the on-chip ROM eSDHC boot configuration. The e500 core configures the eSDHC controller, enabling it to communicate with the external SD/MMC card. The SD/MMC card should contain a specific data structure with control words, device configuration information, and initialization code. The on-chip ROM boot code uses the information from the SD/MMC card content to configure the device and to copy the initialization code to a target memory device (for example, the DDR) through the eSDHC interface. After all the code has been copied, the e500 core starts to execute the code from the target memory device.

There are several different ways a user may utilize the eSDHC boot feature. The simplest way is for the on-chip ROM to copy an entire operating system boot image into system memory, and then jump to it to begin execution. However, this may be many megabytes and in some situations may be suboptimal because only 1-bit mode is used during booting.

A more advanced option is for the on-chip ROM to only copy a small user-customized subroutine, which configures the eSDHC in an optimal way. The user-customized subroutine then copies the rest of the boot code potentially much faster than the on-chip ROM software can achieve. For example, the user-customized subroutine may utilize 4-bit or 8-bit eSDHC interfaces, support new SD or MMC format revisions, or increase the external clock frequency based on knowledge of the exact frequency at which the chip is operating.

4.6.1.2 eSDHC boot features

The following list provides an overview of the key features:

- Provides mechanism to load initialization code from the following external devices:
 - SD memory cards, including the memory portion of SD Combo cards (up to and including version 2.0)
 - MMC, RS-MMC, and MMC*plus* (up to and including version 4.2)
 - SDHC cards (SD High Capacity, from 4 GByte to 32 GByte)
- Boot from the following devices is not supported
 - Locked (password-protected) SD/MMC cards
 - Secured Mode of SD cards (SD Card Specification Part 3: Security Specification)
- Simple data structure in SD/MMC card
- BOOT signature will be checked to validate that the SD/MMC card contains valid code
- Supports variable code length in SD/MMC card
- Flexible target memory device
- Supports target memory configuration controlled by the user
- Only 1-bit operation is supported for boot (even if the SD/MMC card supports 4 or 8-bit parallel access).
- Initial setting will use a serial clock below 400 kHz; the SD/MMC internal registers are read by initialization code and parsed to determine the optimal clock frequency supported by the SD/MMC card inserted.
- High speed cards are supported (up to 50 MHz SD and 52 MHz MMC).
- Control word will allow for user modification
- There must be precisely one device connected on the eSDHC bus. In particular, multiple MMC devices sharing the one bus are not supported.

- Compatible with FAT12/FAT16/FAT32 SD/MMC filesystems (provided ≤ 40 configuration words are used prior to copying user's code to system memory).
- Redundancy to support SD/MMC bad blocks, by searching for the BOOT signature in up to 24 blocks and trying the next block if the BOOT signature is not found or if a read CRC error is found.

4.6.1.3 SD/MMC card data structure

The SD/MMC card should contain the initialization code length in bytes, source address in the SD/MMC card, destination address in the target memory device, execution starting address, and multiple configuration words with pairs of target address and its respective data.

This figure shows the required SD/MMC card data structure.

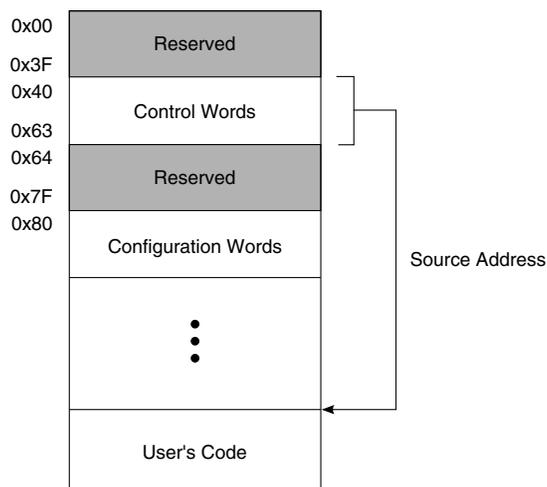


Figure 4-8. SD/MMC card data structure

This table describes the SD/MCC card data structure.

Table 4-44. SD/MMC card data structure

Address	Data bits [0:31]
0x00-0x3F	Reserved.
0x40-0x43	BOOT signature. This location should contain the value 0x424F_4F54, which is the ASCII encoding for BOOT. The boot loader code will search for this signature. If the value in this location doesn't match the BOOT signature, it means that the SD/MMC card does not contain a valid user code. In such case the boot loader code will disable the eSDHC and will issue a hardware reset request of the SoC by setting RSTCR[HRESET_REQ].
0x44-0x47	Reserved

Table continues on the next page...

Table 4-44. SD/MMC card data structure (continued)

Address	Data bits [0:31]
0x48-0x4B	User's code length. Number of bytes in the user's code to be copied. This must be a multiple of the SD/MMC card's block size (and the user's code zero-padded if necessary to achieve that length). User's code length \leq 2 GBytes.
0x4C-0x4F	Reserved
0x50-0x53	Source Address. Contains the starting address of the user's code as an offset from the SD/MMC card starting address. In Standard Capacity SD/MMC Cards, the 32-bit source address specifies the memory address in byte address format. This must be a multiple of the SD/MMC card's block size. In High Capacity SD Cards (> 2 GByte), the 32-bit source address specifies the memory address in block address format. Block length is fixed to 512 bytes as per the SD High Capacity specification.
0x54-0x57	Reserved
0x58-0x5B	Target Address. Contains the target address in the system's local memory address space in which the user's code will be copied. This is a 32-bit effective address. The core is configured in such a way that the 36-bit real address is equal to this (with 4 most significant bits zero).
0x5C-0x5F	Reserved
0x60-0x63	Execution Starting Address. Contains the jump address in the system's local memory address space into the user's code first instruction to be executed. This is a 32-bit effective address. The core is configured in such a way that the 36-bit real address is equal to this (with 4 most significant bits zero).
0x64-0x67	Reserved
0x68-0x6B	N. Number of Config Address/Data pairs. Must be $1 \leq N \leq 1024$ (but is recommended to be as small as possible).
0x6C-0x7F	Reserved.
0x80-0x83	Config Address 1
0x84-0x87	Config Data 1
0x88-0x8B	Config Address 2
0x8C-0x8F	Config Data 2
...	
$0x80 + 8 \times (N - 1)$	Config Address N ¹
$0x80 + 8 \times (N - 1) + 4$	Config Data N (final Config Data N optional)
~	
~	
~	
	User's code. Note that user's code must start on a 512-byte boundary.

1. $N \leq 40$ if compatibility with FAT12/FAT16/FAT32 file systems is required. See [Notes on compatibility with FAT12/FAT16/FAT32 filesystems](#)" for details.

4.6.1.3.1 Configuration words section

The configuration words section is comprised of Config Address and Config Data pairs of adjacent 32-bit fields.

Table 4-47. Config address field description, CNT = 1

Bits	Name	Description
0	EC	<p>End Configuration. Indicates the end of the configuration stage. Valid only if bit CNT is set.</p> <p>0 Not the last Config Address field.</p> <p>1 The Last Config Address field. The e500 core will stop the configuration stage and start to copy the user's code.</p> <p>NOTE: This must be set for Config Address Word N and must not be set for Config Address words prior to Config Address Word N.</p>
1	DLY	<p>Delay. Instruct the e500 core to perform delay according to the number that is specified in the adjacent Config Data field. The adjacent Config Data field provides the delay measured in terms of the number of 8 CCB clocks. Valid only if bit CNT is set.</p> <p>0 No delay</p> <p>1 Delay</p>
2-30	-	Reserved. Must be zero.
31	CNT	<p>Control. Select between Address mode and Control mode.</p> <p>0 Address mode</p> <p>1 Control mode</p> <p>NOTE: When CNT= 1, bits 0-29 select the control instruction. Only one bit in the range of bits 0-29 can be set at any specific control instruction. A control instruction with bits 0-29 all cleared is also illegal.</p>

4.6.1.3.2 Notes on compatibility with FAT12/FAT16/FAT32 filesystems

Depending upon application, compatibility may be desired between the SD/MMC Card data structure defined here and the FAT12, FAT16, or FAT32 filesystems (documented in SD Card Specifications Part 2-File System Specification v2.0, among other places).

This compatibility is possible, but imposes a limit on the number of configuration words that can be parsed by the processor prior to fetching the user's code. Compatibility is achieved by ensuring that the entire data structure of control words and configuration words is contained within the first 446-byte (0x1BE) master boot record code area of the filesystem.

Given that Configuration words start at address 0x80, and all configuration words (except the last one with EC = 1 to end the configuration) occupy 8 bytes per configuration address/data pair, this imposes the limit of a maximum of 40 configuration address words. More configuration words can be used in applications for which compatibility with the FAT master boot record is not required. If exactly 40 configuration address words are used and FAT12/FAT16/FAT32 compatibility is required, then the final configuration data word must be omitted to ensure that the data structure fits in less than 446 bytes.

Note that FAT12, FAT16, and FAT32 standards impose additional requirements on the data structures that must be present on the SD/MMC card, such as partition tables and a fixed signature word at the end of the master boot record. These features are not interpreted or required by the eSDHC boot process and are outside the scope of this document.

Furthermore, FAT12 and FAT16 define a boot sector with defined fields in the first 0x36 addressable bytes (which does not conflict with the SD/MMC card data structure for boot from SD/MMC defined in this document). Therefore FAT12 and FAT16 filesystems are completely compatible with the defined data structure, even if they also contain a FAT boot sector. However, FAT32 defines a boot sector with defined fields in the first 0x52 addressable bytes. Therefore, FAT32 file system compatibility is only possible if used in a system in which this boot sector information is not required.

Also note that the user code is copied from one sequential area of SD/MMC card memory space specified by the Source Address. The boot ROM software does not look for or parse any file allocation table, and furthermore, the boot ROM software assumes that the user code is in one contiguous range of memory addresses.

4.6.1.4 eSDHC controller initial configuration

The eSDHC controller configuration is used by the boot ROM software.

After the boot from eSDHC has finished, the user can change this configuration for other uses of the eSDHC interface. The boot ROM software also changes some of this configuration automatically, depending upon the features supported by the SD/MMC card that is connected.

The eSDHC controller is initially configured to operate in the following configuration:

- Address invariant mode (eSDHC.PROTCTL[EMODE] = 10)
- SDHC_DAT[3] does not monitor card insertion. The GPIO4/SDHC_CD pin is used for card detect (eSDHC.PROTCTL[D3CD] = 0 and Global_Uilities. PMUXCR[SDHC_CD] = 1).
- 1-bit Mode (eSDHC.PROTCTL[DTW] = 00)
- For SD clock frequency selection, refer to example provided in SYSCTL[SCLKFS] in [System control \(eSDHC_SYSCTL\)](#).
- There must be precisely one device connected on the eSDHC bus (and this device must be inserted prior to boot). Multiple MMC devices sharing the one bus are not supported.
- The bus operates in push-pull mode (The chip pads drive both logic 0 and logic 1 as appropriate). If a MMC card is to be connected, then weak external pull-ups are

required on the SDHC_CMD and SDHC_DAT[0:7] pins in order to interface with the MMC open-drain mode during initialization.

- The eSDHC DMA engine is not used for control or configuration word accesses; instead, all eSDHC data transfers are initiated by the processor core polling eSDHC.PRSTAT[BRR] and accessing data through the DATPORT register (XFERTYP[DMAEN] = 0). The eSDHC DMA engine is used for user code accesses.

4.6.1.5 eSDHC controller boot sequence

The code in the eSDHC Boot ROM configuration performs the following sequence of events:

1. The eSDHC controller is configured as per [eSDHC controller initial configuration](#).
2. Card-detect.
3. The SD/MMC card is reset.
4. SD/MMC card voltage validation is performed.
5. SD/MMC card identification.
6. With CMD9, the CSD (Card-Specific Data) register of the SD/MMC card is read.
7. Based on the values returned from the SD/MMC card's CSD register, the eSDHC's registers are updated to reflect the maximum clock frequency jointly supported by the eSDHC controller and the SD/MMC card connected to it.
8. The eSDHC begins reading the SD/MMC data structure from the card.
9. The eSDHC begins fetching the user code from the card.
10. If either the BOOT signature is not found at memory offset 0x40 or when reading the control and configuration words or the user's code a read CRC error is detected, then it may be due to a bad block on the SD/MMC memory card. To counteract this and provide error resilience, if this occurs the eSDHC returns to step 8, fetching data from an address 0x200 greater than the previously fetched address. For example, if there have been i failed attempts, then on the following try the BOOT signature is checked at offset $0x40 + i \times 0x200$. If this sequence fails 24 times, then the system boot is deemed to have failed.
11. The processor core waits until the user code DMA transfer is complete.
12. The processor core jumps to the execution starting address to begin execution of the user's code.

4.6.1.6 eSDHC boot error handling

If at any stage the boot loader code detects an error and cannot continue, it disables the eSDHC and issues a hardware reset request of the chip by setting RSTCR[HRESET_REQ].

This may occur in any of the following scenarios:

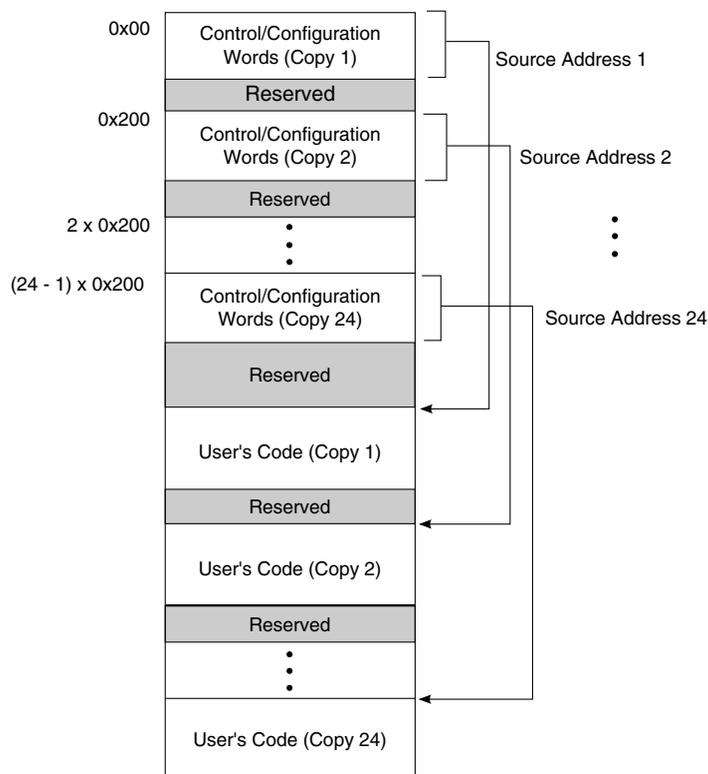
- BOOT signature not found at offset 0x40 or CRC error on any of the data read by the eSDHC 24 times.
- Timeout while waiting for the SD/MMC card to respond at any stage.
- No card inserted.
- Incorrect type of card inserted that is not supported for boot (such as CE-ATA).
- There is no common protocol, voltage, or frequency mutually supported by the SD/MMC card and the eSDHC.
- The eSDHC reads as far as the source address (specified by the control word of the SD/MMC data structure) without seeing a EC = 1 Configuration Word.

The boot loader code supports redundancy, which allows boot to succeed even in the presence of SD/MMC bad blocks. It does this by searching for the BOOT signature in up to 24 locations. When either the BOOT signature is not found or a read CRC error is found, the boot loader tries the next block. Each location tried is at a fixed offset of 512 bytes (0x200) from the previous (unsuccessful) offset, irrespective of the actual block size of the SD/MMC card.

For reference, [Figure 4-9](#) provides an example SD/MMC memory card data structure that can be used for maximum SD/MMC card data redundancy.

Note that if $0x40 + 8 \times (N - 1) + 4 \geq 0x200$ (where N is the number of configuration words), care needs to be taken to ensure that the configuration words at $0x40 + i \times 0x200$ (for all $2 \leq i \leq 24$) must not contain the BOOT signature. This ensures that the boot loader code does not mistakenly detect a BOOT signature. This also reduces the number of copies of boot code that can be used on one device.

Each copy of the control/configuration words are generally identical except for the pointer to the source address (offset 0x50) of the SD/MMC card, which may be different for each copy. If the user's code section is sufficiently large that 24 copies of it do not fit in the capacity of the SD/MMC card (or if the SD/MMC card capacity must also be utilized for functional features other than system boot), it may still be desirable to support up to 24 copies of the control/configuration words, but only have them reference a limited number of user's code sections. This is also possible.



Note: User's code must begin on 512-byte boundary and its length must be a multiple of 512 bytes.

Figure 4-9. SD/MMC card data structure for maximum redundancy

4.6.2 eSPI boot ROM

This section provides an overview, list of features, and SD/MMC card data structures for the eSPI boot.

4.6.2.1 eSPI boot overview

The chip is capable of loading initialization code from a memory device that is connected to the eSPI controller interface.

This device can be either an EEPROM or a serial flash with an eSPI-compatible interface. The term EEPROM will be used when referring to the memory device.

The eSPI controller supports RapidS full clock cycle operation and Winbond dual output read serial interface, but these modes are not enabled for boot by the on-chip ROM.

Boot from eSPI is supported by the chip using an on-chip ROM which contains the basic eSPI device driver and the code to perform block copy from eSPI EPROM to any target memory. Selecting on-chip ROM in boot ROM location, see [Table 4-12](#), causes the e500 CPU to fetch data from the on-chip ROM. The on-chip ROM is selected by configuring the POR config pins `cfg_rom_loc[0:3]`.

After the device has completed the reset sequence, the e500 core starts to execute code from the internal on-chip ROM if the ROM location selects the on-chip ROM eSPI boot configuration. The e500 core configures the eSPI controller, enabling it to communicate with the external EEPROM. The EEPROM should contain a specific data structure with control words, device configuration information, and initialization code. The on-chip ROM boot code uses the information from the EEPROM content to configure the device and to copy the initialization code to a target memory device (for example, the DDR) through the eSPI interface. After all the code has been copied, the e500 core starts to execute the code from the target memory device.

There are several different ways a user may utilize the eSPI boot feature. The simplest is for the on-chip ROM to copy an entire operating system boot image into system memory and then jump to it to begin execution. However, this may be many megabytes and in some situations may suboptimal.

A more advanced option is for the on-chip ROM to only copy a small user-customized subroutine, which configures the eSPI in an optimal way. The user-customized subroutine then copies the rest of the boot code potentially much faster than the on-chip ROM software can achieve. For example, the user-customised subroutine may utilize Atmel RapidS or Winbond dual output eSPI modes.

4.6.2.2 eSPI boot features

The following list describes the key features:

- Provides mechanism to load initialization code from external eSPI EEPROM
- Simple data structure in eSPI EEPROM
- BOOT signature will be checked to validate that the EEPROM contains valid code
- Supports variable code length in EEPROM
- Flexible target memory device
- Supports target memory configuration controlled by the user
- Supports standard eSPI interface EEPROMs with read instruction code 0x03 followed by a 2-byte address (16-bit addressable EEPROMs) or 3-byte address (24-bit addressable EEPROMs)
- Initial setting will generate a serial clock below 5 MHz; the control word will allow for user modification of clock frequency.

4.6.2.3 EEPROM data structure

The EEPROM should contain the initialization code length in bytes, source address in the eSPI EEPROM, destination address in the target memory device, execution starting address, and multiple configuration words with pairs of target address and its respective data.

This figure shows the required eSPI EEPROM data structure.

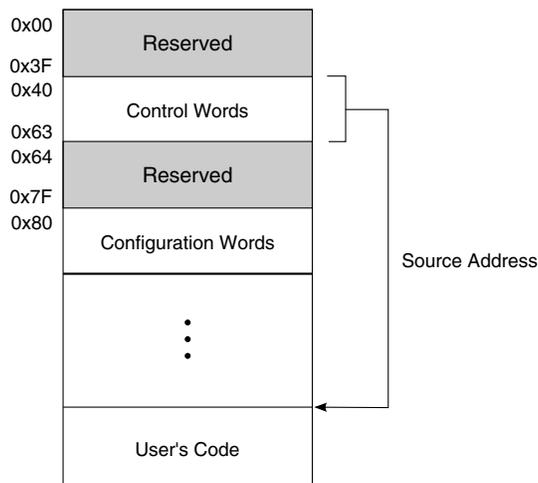


Figure 4-10. eSPI EEPROM data structure

This table describes the eSPI EEPROM data structure.

Table 4-48. eSPI EEPROM data structure

Address	Data bits [0:31]
0x00-0x3F	Reserved.
0x40-0x43	BOOT signature. This location should contain the value 0x424F_4F54, which is the ASCII encoding for BOOT. The eSPI loader code will search for this signature, initially in 24-bit addressable mode. If the value in this location does not match the BOOT signature, the EEPROM is re-accessed in 16-bit mode. If the value in this location still does not match the BOOT signature, it means that the eSPI device does not contain a valid user code. In such cases, the eSPI loader code disables the eSPI and issues a hardware reset request of the chip by setting RSTCR[HRESET_REQ].
0x44-0x47	Reserved
0x48-0x4B	User's code length. Number of bytes in the user's code to be copied. Must be a multiple of 4. $4 \leq \text{User's code length} \leq 2 \text{ GBytes}$.
0x4C-0x4F	Reserved
0x50-0x53	Source Address. Contains the starting address of the user's code as an offset from the EEPROM starting address. In 24-bit addressing mode, the 8 most significant bits of this should be written to as zero, because the EEPROM is accessed with a 3-byte (24-bit) address. In 16-bit addressing mode, the 16 most significant bits of this should be written to as zero.

Table continues on the next page...

Table 4-48. eSPI EEPROM data structure (continued)

Address	Data bits [0:31]
0x54-0x57	Reserved
0x58-0x5B	Target Address. Contains the target address in the system's local memory address space in which the user's code will be copied to. This is a 32-bit effective address. The core is configured in such a way that the 36-bit real address is equal to this (with 4 most significant bits zero).
0x5C-0x5F	Reserved
0x60-0x63	Execution Starting Address. Contains the jump address in the system's local memory address space into the user's code first instruction to be executed. This is a 32-bit effective address. The core is configured in such a way that the 36-bit real address is equal to this (with 4 most significant bits zero).
0x64-0x67	Reserved
0x68-0x6B	N. Number of Config Address/Data pairs. Must be ≤ 1024 (but is recommended to be as small as possible).
0x6C-0x7F	Reserved.
0x80-0x83	Config Address 1
0x84-0x87	Config Data 1
0x88-0x8B	Config Address 2
0x8C-0x8F	Config Data 2
...	
$0x80 + 8 \times (N - 1)$	Config Address N
$0x80 + 8 \times (N - 1)$ + 4	Config data N (Final Config Data N optional)
~	
~	
~	
	User's code

4.6.2.3.1 EEPROM configuration words section

The configuration words section is comprised of Config Address and Config Data pairs of adjacent 32-bit fields.

These are typically used to configure the local access windows and the target memory controller's registers. They are therefore system-dependent, as they need to be aware of the type and configuration of memory in a particular system.

The Config Address field has two modes that are selected by the least significant bit in the field (CNT). If the CNT bit is clear, then the 30 most significant bits are used to form the address pointer and the config data contains the data to be written to this address. If the CNT bit is set then the 30 most significant bits are used for control instruction. This flexible structure allows the user to configure any 4-byte aligned memory mapped register, perform control instructions, and specify the end of the configuration stage.

Table 4-51. Config Address Field Description, CNT = 1 (continued)

Bits	Name	Description
1	DLY	Delay. Instruct the e500 core to perform delay according to the number that is specified in the adjacent Config Data field. The adjacent Config Data field provides the delay measured in terms of the number of 8 CCB clocks. Valid only if bit CNT is set. 0 No delay. 1 Delay.
2	CF	Change frequency. Instruct the e500 core to perform sequence of operations to setup the eSPI CS0 mode register with the frequency related (PM and DIV16) bits as defined by the user. The adjacent Config Data field will be written to the eSPI mode register. Software will use DIV16 and PM bits and mask all other bits such that they will not change. Software will perform the necessary steps which are required by the eSPI controller before and after changing the eSPI mode register. This only takes effect after all of the configuration and control words have been read.
3-29	-	Reserved. Must be zero.
30	-	Reserved. Must be zero.
31	CNT	Control. Select between address mode and control mode. 0 Address mode 1 Control mode NOTE: When CNT = 1, bits 0-29 select the control instruction. Only one bit in the range of bits 0-29 can be set at any specific control instruction. A control instruction with bits 0-29 all cleared is also illegal.

4.6.2.4 eSPI controller configuration

The eSPI controller configuration is used by the eSPI boot ROM software.

After the boot from eSPI has finished, the user can change this configuration for other uses of the eSPI interface.

The eSPI controller is configured to operate in master mode. The eSPI chip select 0 (SPI_CS[0]) must be connected to the EEPROM \overline{CS} and selectively enables the EEPROM.

This figure shows the external signal connection.

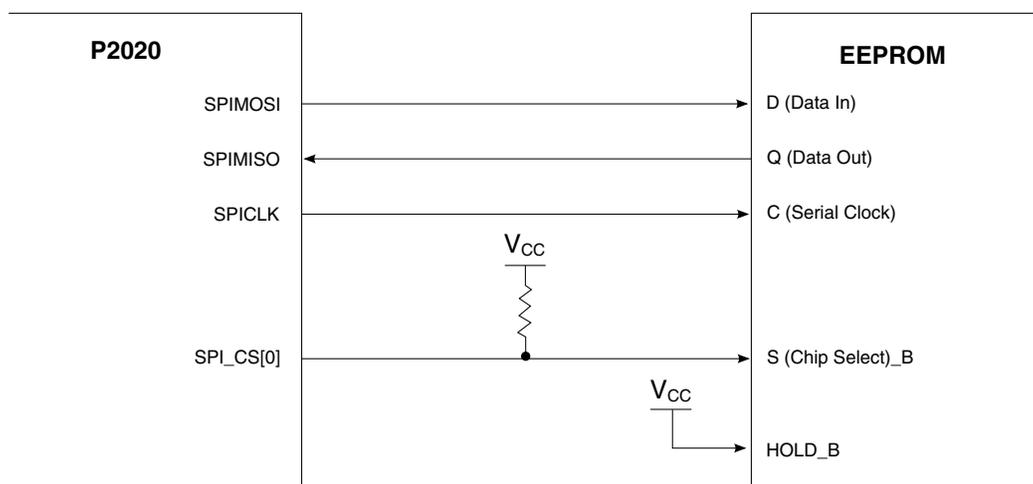


Figure 4-11. External signal connection

The eSPI controller is configured by the on-chip ROM code. The controller is configured as follows:

- Data is shifted out data on SPIMOSI during the falling edge of SPICLK. It samples data in from SPIMISO during the rising edge of SPICLK
- The clock is low when the line is idle
- It uses 8-bit length characters
- The platform clock is divided by 64. For example, when the platform clock is configured to 533 MHz, the SPICLK will run at 8.33 MHz. (Note that frequency setting can be changed by using the CF control word, as explained in the [Configuration words section](#).)
- The MSB is sent and received first.

For default eSPI CS0 mode register (SPMODE0) configuration, see eSPI CS0 Mode Register (SPMODE0). During ROM boot, the ROM code needs to change the resetvalue of eSPI CS0 mode register to 0x3117_2210.

The ROM code will initially use the eSPI controller to generate standard read instruction code 0x03 followed by a 3-byte address for every non-sequential read operation (reading from a location which is not sequential to the last byte read). For sequential read operation, toggling the eSPI clock will cause the eSPI EEPROM to present the content of the next address location. The serial EEPROM must have an eSPI compatible interface with read instruction code 0x03 followed by a 2 or 3-byte address.

16-bit addressable EEPROM memories are supported and detected automatically by the boot code. This is accomplished by the boot code trying 16-bit mode if it fails to find the "BOOT" signature when in 24-bit mode.

This figure shows the read instruction timing diagram for normal (not Atmel RapidS or Winbond Dual Output) modes of operation with a 24-bit addressable eSPI memory. With 16-bit addressable eSPI memories, only a 16-bit address is transmitted, and valid data is received from the EEPROM on the 24th SPICLK cycle rather than the 32nd SPICLK cycle for 24-bit addressable memories.

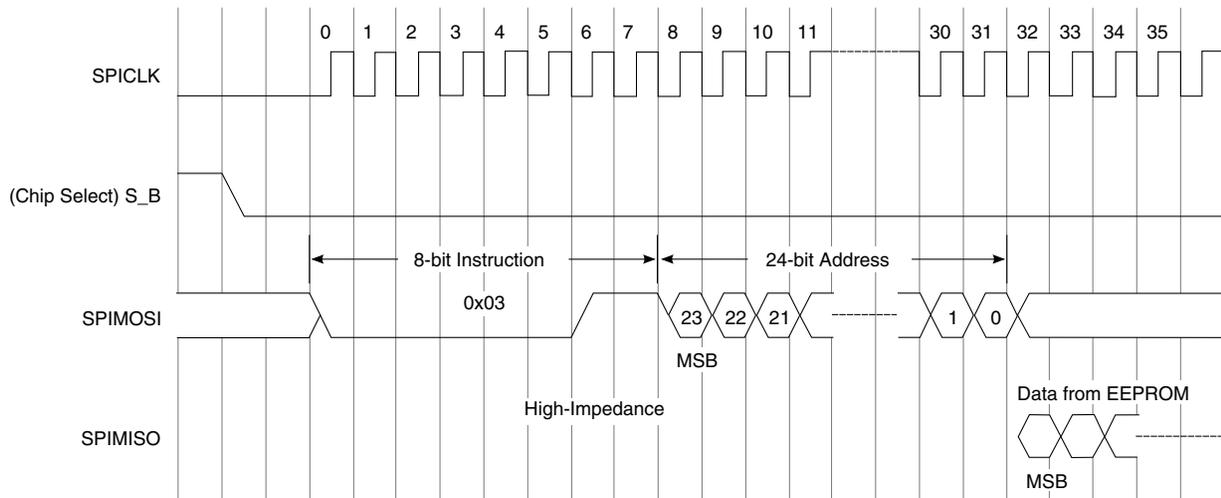


Figure 4-12. Read instruction timing diagram (24-bit addressable eSPI memory)

4.6.3 Default e500 addressing during system boot

During boot from the on-chip ROM (for boot targets of either eSPI or SD/MMC), the user specifies 32-bit addresses for several fields (target address for copying the user's code, and the execution starting address).

This section describes how these 32-bit effective addresses are translated into 36-bit real addresses and the associated address translation and mapping.

The L2 cache remains disabled as per its power-on-reset state. The e500 level 1 and level 2 MMU configuration is left as per defaults, with the exception that the TLBI Entry 0 is replaced with the following:

- V = 1 (valid)
- TS = 0 (address space 0)
- TID = 0x00 (global)
- EPN[32-51] = 0x0_0000
- RPN[32-51] = 0x0_0000
- SIZE[0-3] = 1011 (4 Gbyte)

- SX/SR/SW = 111 (Full supervisor mode access allowed)
- UX/UR/UW = 000 (No user mode access allowed)
- WIMGE = 01110 (Cache-inhibited, Memory coherency required, Guarded)
- X0-X1 = 00
- U0-U3 = 0
- IPROT = 1 (Page is protected from invalidation)

This configuration results in a 32-bit byte address with a 0-bit effective page number. Therefore, the 36-bit real address is equal to the 32-bit effective address, with the 4 MSbits of the 36-bit real address equal to 0.

The on-chip ROM code does not setup any local access windows. Access to CCSR address space (and therefore by extension, also access to the on-chip ROM) does not require a local access window. It is the user's responsibility to setup a local access window through a control word address/data pair for the desired target address and execution starting address (which will typically be in either DDR or local bus memory space).

Note that any such local access window configured at this time must have the 4 MSbits of the address equal to 0. This is due to the 32-bit addressing enabled by the e500 MMUs as described above.

The user can reconfigure the system in the user code portion based on system requirements.

Chapter 5

e500 Core Integration Details

e500 core integration and the core complex bus (CCB) describes hardware aspects of that integration and provides links to chapters that discuss functionality in which core and SoC operations interact.

5.1 Overview

This chapter describes how the core is integrated into the SoC.

[e500 core integration and the core complex bus \(CCB\)](#) describes hardware aspects of that integration and provides links to chapters that discuss functionality in which core and SoC operations interact. Such topics include reset, power management, interrupt management, and debug.

This chapter also lists SoC-specific details of the core's programming model. For example, the e500 programming model defines the processor version register (PVR), system version register (SVR), and special-purpose registers (SPRs) that respectively identify the version and revision of the core and of the integrated device. These values are provided in [Summary of core integration details](#), and additional links are provided to other chapters that provide a context for a discussion of these registers. The section [Register model integration details](#) in [Table 5-1](#) describes a few aspects of the core programming model that has SoC-specific behavior that cannot be fully understood by reading the *e500 Reference Manual* alone.

General information about e500 core functionality can be found in this chapter and in the following documentation:

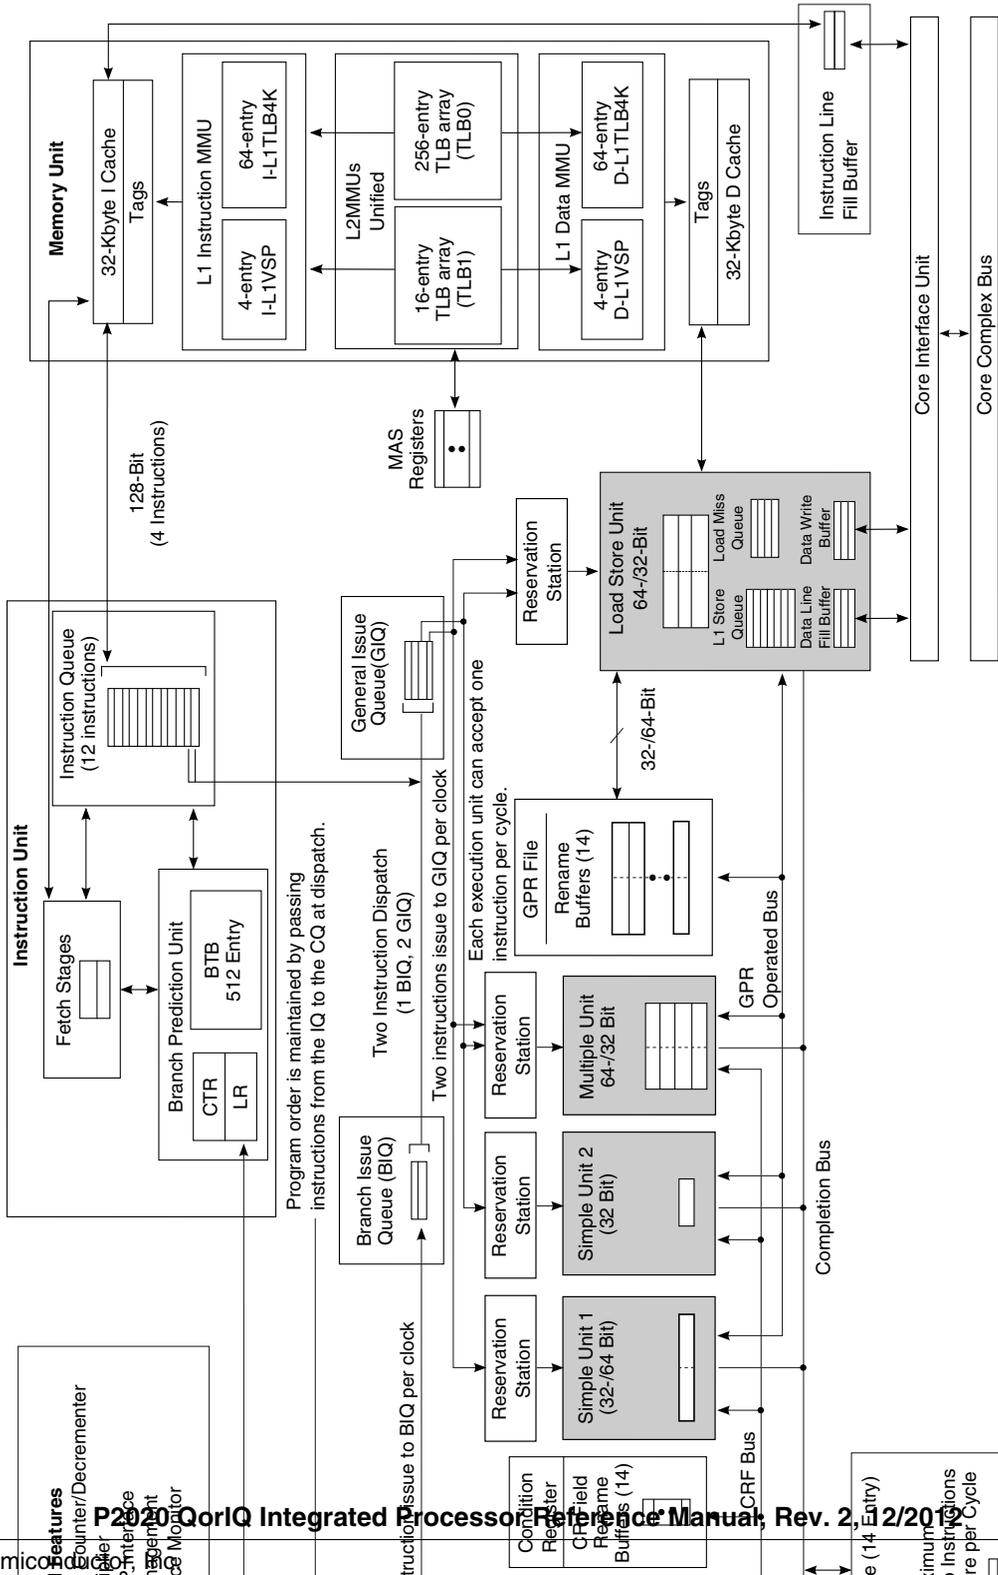
- The *PowerPC e500 Core Family Reference Manual* (referred to here as the *e500 Reference Manual*) provides detailed information about the functions and features of the core. In particular, it describes details about how architecture-defined features are implemented.

- The "e500 Core Complex Bus (CCB) and System Integration" chapter in the *e500 Reference Manual* describes core-to-SoC integration issues from the perspective of the e500 core.
- The *EREF: A Programmer's Reference Manual for Freescale Embedded Processors (Including the e200 and e500 Families)* describes in detail features defined by the Power ISA and Freescale EIS (referred to generically as the architecture). Unless otherwise stated in the *e500 Reference Manual*, e500 features are implemented as defined by the architecture and described in the *EREF*.
- How the integrated device implements e500 core features, including specific registers and register fields, is summarized in [Summary of core integration details](#).

5.2 e500 core block diagram

The chip core complex block diagram shows how the functional units operate independently and in parallel.

Note that this is a conceptual diagram that does not attempt to show how these features are physically implemented.



5.3 e500 core integration and the core complex bus (CCB)

The CCB is the hardware interface between the core and the SoC. With a few exceptions, the user cannot access these internal signals directly.

This figure shows a selection of CCB signals that are discussed in various places in this manual and in the core reference manual, because understanding how they work helps to understand the functionality of this chip.

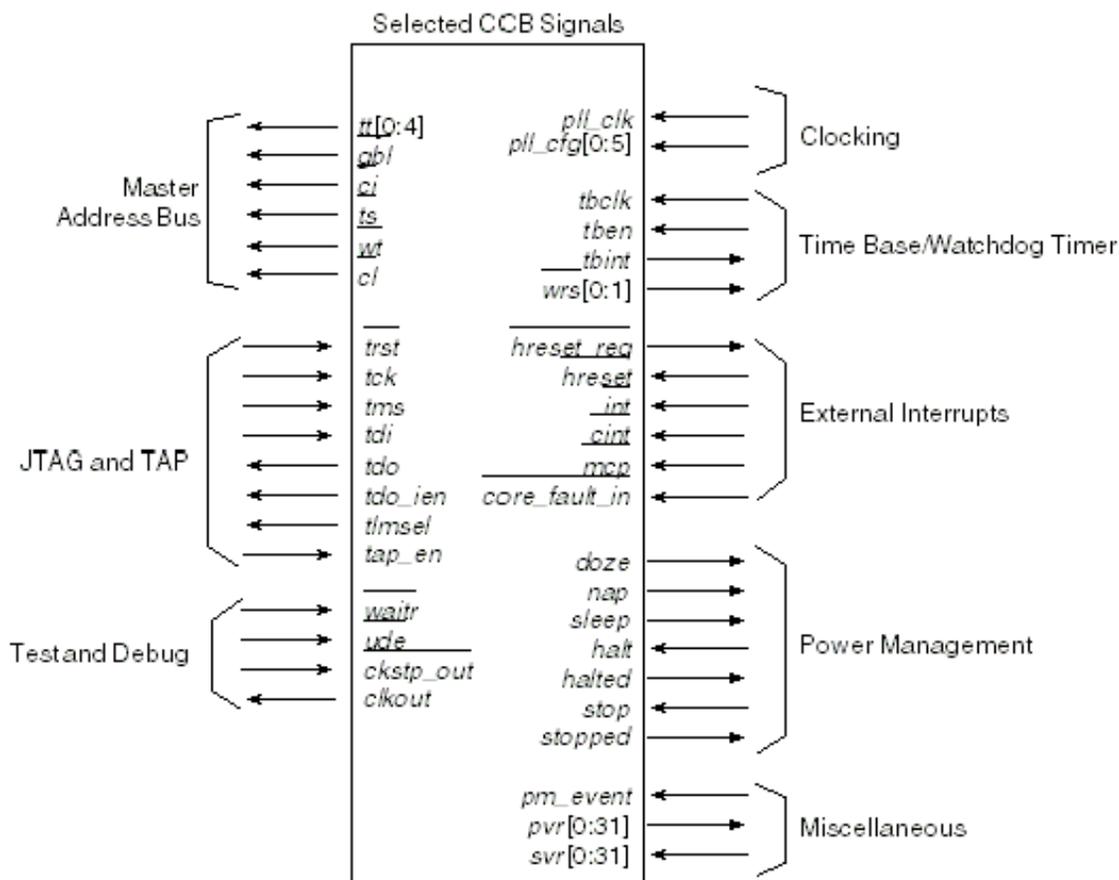


Figure 5-2. e500 core integration

Aspects of the e500 chip integration are summarized as follows:

- Reset

The core directs and coordinates device's hard and soft resets and the power-on reset (POR) sequence, power-on reset configuration, and initialization. Core integration of the reset signals shown in [Figure 5-2](#) is described in the chapters [Reset](#), [Clocking](#), [and Initialization](#) and [Global Utilities](#).

- Clocking and timers

Integration details of the CCB clocking signals are described in [Reset, Clocking, and Initialization](#). Additional details regarding the timer configuration are described in [Summary of core integration details](#).

- Cache and memory-mapped SRAM

The e500 cache implementation interacts with the SoC's L2 cache. In particular, the core implements a number of instructions that interact with the L2 cache implementation, which are described in the *e500 Reference Manual* and in the *EREF*. [L2 Look-Aside Cache/SRAM](#) describes the SoC's L2 cache. [Figure 5-2](#) shows the e500 signals that interface with the L2 cache.

- e500 coherency module (ECM)

The ECM, described in [e500 Coherency Module](#), facilitates communication between the core, the L2 cache, and the other blocks that comprise the coherent memory domain of the SoC.

The ECM provides a mechanism for I/O-initiated transactions to snoop the core complex bus (CCB) of the core to maintain coherency across cacheable local memory. It also provides a flexible, easily expandable switch-type structure for core- and I/O-initiated transactions to be routed (dispatched) to target modules on the device. The CCB is described in the "Core Complex Bus (CCB) and System Integration" chapter of the *e500 Reference Manual*.

- Interrupts

The e500 core handles the hardware interrupts that are generated from SoC peripheral logic, typically by error conditions encountered from within blocks on the integrated device. [Programmable Interrupt Controller \(PIC\)](#) describes the programmable interrupt controller, which prioritizes interrupt requests to the core. [Figure 5-2](#) shows the e500 signals that interface with the PIC.

- Power management

The core's HID0[NAP, DOZE, SLEEP] can be used to assert *nap*, *doze*, and *sleep* output signals to initiate power-saving modes at the integrated device level. [Figure 5-2](#) shows the e500 signals, which interact with the SoC level power management logic described in [Global Utilities](#).

- System debug

The architecture defines many features for software and hardware debug that interact with the SoC. [Debug Features and Watchpoint Facility](#) describes the debug features and watchpoint monitor. [Figure 5-2](#) shows the e500 signals that interface with the debug block. The core reference manual describes how the architecture-defined debug resources are implemented on the e500 core.

The "Core Complex Bus (CCB) and System Integration" chapter of the core reference manual describes the signals that are shown in the figure above and other aspects of core integration.

5.4 Summary of core integration details

The summary of core integration details is organized into two sections: General feature integration details and register model integration details.

This table summarizes details of the QorIQ -specific implementation of the core, which is organized into the following sections:

- [General feature integration details](#) summarizes integration-specific details by functionality.
- [Register model integration details](#) summarizes how integration-specific details are reflected in the SoC's implementation of the core register model.

Table 5-1. Differences between the e500 core and the QorIQ core implementation

Feature	QorIQ implementation
General feature integration details	
Cache protocol	The write-through L2 cache implemented on the SoC does not support MESI cache protocol.
Clocking	Internal clock multipliers ranging from 1 to 8 times the bus clock, including integer and half-mode multipliers. The integrated device supports multipliers of 2, 2.5, 3, and 3.5. See the table entry, HID1 Implementation, for further details.
R1 and R2 data bus parity	R1 and R2 data bus parity are disabled on QorIQ devices. HID1[R1DPE,R2DPE] are reserved.
Dynamic bus snooping	The QorIQ devices do not perform dynamic bus snooping as described here. That is, when the e500 core is in core-stopped state (which is the state of the core when the QorIQ device is in either the nap or sleep state), the core is not awakened to perform snoops on global transactions.
Device specific definition for TCR[WRC]	QorIQ devices define values for 00, 01, 10, and 11, as described in Register Model Integration Details in this table.

Table continues on the next page...

Table 5-1. Differences between the e500 core and the QorIQ core implementation (continued)

Feature	QorIQ implementation
SPE and floating-point categories	<p>The SPE (which includes the embedded vector and scalar floating-point instructions) will not be implemented in the next generation of QorIQ devices. Freescale Semiconductor strongly recommends that use of these instructions be confined to libraries and device drivers. Customer software that uses these instructions at the assembly level or that uses SPE or floating-point intrinsics will require rewriting for upward compatibility with next generation QorIQ devices.</p> <p>The e500v2 core implements SPE double-precision floating-point instructions.</p> <p>Freescale Semiconductor offers a libcfsl_e500 library that uses SPE instructions. Freescale Semiconductor will also provide future libraries to support next generation QorIQ devices.</p> <p>Note that in the Power ISA, MSR[SPE] and ESR[SPE] are renamed to MSR[SPV] and ESR[SPV].</p>
Register Model Integration Details	
HID0 implementation	SEL_TBCLK. Selects time base clock. If this bit is set and the time base is enabled, the time base is based on the TBCLK input, which on the QorIQ devices is RTC.
HID1 Implementation	<p>HID1[PLL_CFG] is implemented as two subfields:</p> <p>PLL_MODE (HID1[32-33]):</p> <p>Read-only for integrated devices.</p> <p>1 1 Fixed value for this device</p> <p>PLL_CFG, (HID1[34-39]): The following clock ratios are supported:</p> <p>0001_00 Ratio of 2:1</p> <p>0001_01 Ratio of 5:2 (2.5:1)</p> <p>0001_10 Ratio of 3:1</p> <p>0001_11 Ratio of 7:2 (3.5:1)</p> <p>NEXEN, R1DPE, R2DPE, MPXTT, MSHARS, SSHAR, ATS, and MID are not implemented.</p> <p>On QorIQ devices, ABE must be set to ensure that cache and TLB management instructions operate properly on the L2 cache.</p>
	<p>HID1[RFXE].</p> <p>If RFXE is 0, conditions that cause the assertion of core_fault_in_B cannot directly cause the e500 to generate a machine check; however, QorIQ devices must be configured to detect and enable such conditions. The following bullets describe how error bits should be configured:</p> <ul style="list-style-type: none"> • ECM mapping errors: EEER[LAEE] must be set. See ECM error enable register (ECM_EEER). • L2 multiple-bit ECC errors: L2ERRDIS[MBECCDIS] must be cleared to ensure that error can be detected. L2ERRINTEN[MBECCINTEN] must be set. See L2 error registers. • DDR multiple-bit ECC errors: ERR_DISABLE[MBED] must be zero and ERR_INT_EN[MBEE] must be set and DDR_SDRAM_CFG[ECC_EN] must be one to ensure that an interrupt is generated. See DDR memory map/register definition. • PCI. The appropriate parity detect and master-abort bits in ERR_DR must be cleared and the corresponding enable bits in ERR_EN must be set to ensure that an interrupt is generated. • Local bus controller parity errors. LTEDR[PAR] must be cleared and LTEIR[PARI] must be set to ensure that an parity errors can generate an interrupt. See Transfer error interrupt register (eLBC_LTEIR), and Transfer error attributes register (eLBC_LTEATR).
PIR value	The PIR value corresponds to the core number on multicore QorIQ devices.
PVR value	<p>The PVR reset value is 0x80nn_nnnn. See Table 5-2 for specific values.</p> <p>PVR[VERSION] = 0x80nn</p> <p>PVR[REVISION] = 0xnxxx</p>

Table continues on the next page...

Table 5-1. Differences between the e500 core and the QorIQ core implementation (continued)

Feature	QorIQ implementation
SVR value	The SVR reset value is 0x80nn_nnnn. See Table 5-2 for specific values.
TCR (timer control register)	<p>TCR[WRC] is defined more specifically for the implementation of the core in the integrated device.</p> <p>Watchdog timer reset control. This value is written into TSR[WRS] when a watchdog event occurs. WRC may be set by software but cannot be cleared by software, except by a software-induced reset. Once written to a non-zero value, WRC may no longer be altered by software.</p> <p>00 No watchdog timer reset will occur</p> <p>01 Second timeout generates machine check</p> <p>10 Second timeout generates HRESET_REQ output externally</p> <p>11 Second timeout automatically resets the given processor core</p>

5.4.1 Processor version register (PVR) and system version register (SVR)

The table below matches the revision codes in the processor version register (PVR) with the core revision and the SoC revision; it includes a cross-reference to the section in global utilities that contains the corresponding SVR values. Note that the SVR and PVR can be accessed both as SPRs through the e500 core (see the e500 Reference Manual) and as memory-mapped registers defined by the integrated device.

Table 5-2. Device revision level cross-reference

SoC revision	Core revision	Processor version register (PVR)	System version register (SVR)
1.0	4.0	0x8021_1040	System version register (GUTS_SVR)
2.0	5.0	0x8021_1050	System version register (GUTS_SVR)
2.1	5.1	0x8021_1051	System version register (GUTS_SVR)

Chapter 6

L2 Look-Aside Cache/SRAM

This chapter describes the organization of the on-chip L2/SRAM, cache coherency rules, cache line replacement algorithm, cache control instructions, and various cache operations.

6.1 Introduction

This chapter describes the organization of the on-chip L2/SRAM, cache coherency rules, cache line replacement algorithm, cache control instructions, and various cache operations.

It also describes the interaction between the L2/SRAM and the e500 core complex.

6.2 L2 cache overview

The integrated 512 KB L2 cache is organized as 2048 eight-way sets of 32-byte cache lines based on 36-bit physical addresses.

This figure shows the L2 cache/SRAM configuration.

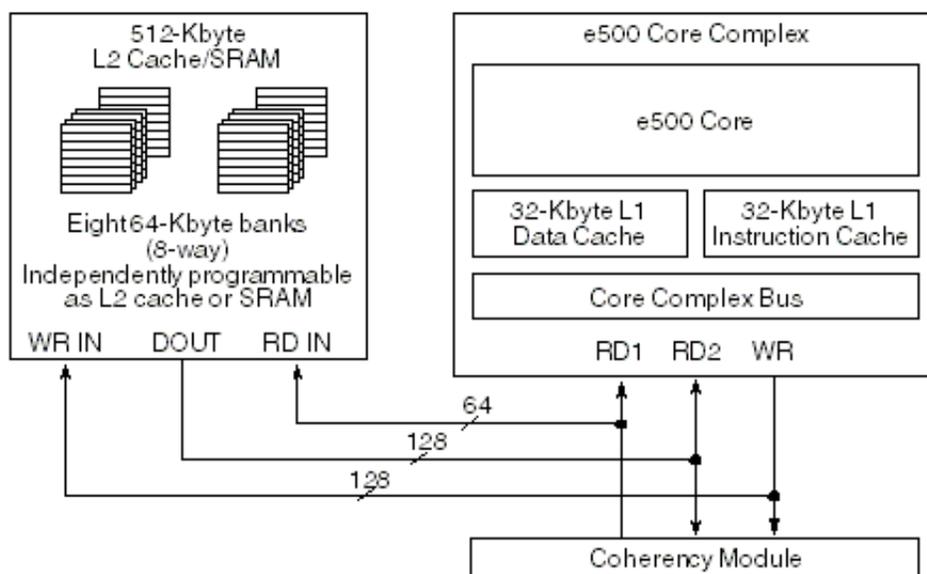


Figure 6-1. L2 cache/SRAM configuration

The SRAM can be configured with memory-mapped registers as externally accessible memory-mapped SRAM in addition to or instead of cache. The L2 cache can operate in the following modes, described in [L2 cache and SRAM organization](#):

- Full cache mode (512 KB cache).
- Full memory-mapped SRAM mode (512 KB SRAM mapped as a single 512 KB block or two 256 KB blocks)
- Partial SRAM and partial cache mode, in which 1/8, 1/4, or 1/2 the total on-chip memory can be allocated to 1 or 2 SRAM regions.

6.2.1 L2 cache and SRAM features

The L2 cache includes the following characteristics:

- Supports 36-bit address space
- Write-through, front-side cache
 - Front-side design provides easier cache access for the I/O masters, such as Ethernet or RapidIO
 - Write-through design is more efficient on the processor bus for front-side caches
- Valid, locked, and stale states (no modified state)
- Two input data buses (64 and 128 bits wide) and one output data bus (128 bits wide)
- All accesses are fully pipelined and non-blocking (allows hits under misses)
- 512-KB array organized as 2048 eight-way sets of 32-byte cache lines
- Eight-way set associativity with a pseudo-LRU (7-bit) replacement algorithm.
 - High level of associativity yields good performance even with many lines locked or used as SRAM regions

- I/O devices can store data into the cache in a process called 'stashing.'
 - Stashing is indicated for global I/O writes either by a transaction attribute or by a programmable memory range
 - Regions of the cache can be reserved exclusively for stashing to prevent pollution of processor cache regions.
- Processor L2 cache regions are configurable to allocate instructions, data, or both.
 - External masters can force data to be allocated into the cache through programmed memory ranges or special transaction types (stashing).
 - 1, 2, or 4 ways can be configured for stashing only
- Data ECC on 64-bit boundaries (single-error correction, double-error detection)
- Tag arrays use 20 tag bits and 1 tag parity bit per line.
- Multiple cache locking methods supported
 - Individual line locks are set and cleared using e500 cache locking APU instructions-Data Cache Block Touch and Lock Set (**dcbtls**), Data Cache Block Touch for Store and Lock Set (**dcbtstls**), and Instruction Cache Block Touch and Lock Set (**icbtls**).
 - A lock attribute can be attached to write operations.
 - Individual line locks are set and cleared through core-initiated instructions, by external reads or writes, or by accesses to programmed memory ranges defined in L2 cache external write address registers (L2CEWAR_n).
 - The entire cache can be locked by setting configuration register appropriately.
- Lock clearing methods
 - Individual locks can be cleared by cache-locking APU instructions-Data Cache Block Lock Clear (**dcble**) and Instruction Cache Block Lock Clear (**icble**)-or by a snooped flush unless entire cache is locked.
 - Flash clearing of all instruction and/or data locks is done by writes to configuration registers.
 - An unlock attribute can be attached to a read instruction.
- Error injection modes supported for testing error handling

SRAM features include the following:

- SRAM regions are created by configuring 1, 2, 4 or 8 ways of each set to be reserved for memory-mapped SRAM.
- Regions can reside at any location in the memory map aligned to the SRAM size.
- SRAM memory is byte addressable; for accesses of less than a cache line, ECC is updated using read-modify-write transactions.
- I/O devices access SRAM regions by marking transactions as snoopable (global).

Table 6-1 lists the possible L2 cache/SRAM configurations.

Table 6-1. Available L2 cache/SRAM configurations(512-KB SRAM)

Cache	Stash-only Region	SRAM Region 1	SRAM Region 2
512 KB (8 ways)	-	-	-
448 KB (7 ways)	-	64 KB	-
	64 KB	-	-
384 KB (6 ways)	-	128 KB	-
	-	64 KB	64 KB
	64 KB	64 KB	-
	128 KB	-	-
320 KB (5 ways)	64 KB	128 KB	-
	-	64 KB	64 KB
	128 KB	64 KB	-
256 KB (4 ways)	-	256 KB	-
	-	128 KB	128 KB
	128 KB	128 KB	-
	-	64 KB	64 KB
	256 KB	-	-
192 KB (3 ways)	64 KB	256 KB	-
	-	128 KB	128 KB
	256 KB	64 KB	-
128 KB (2 ways)	128 KB	256 KB	-
	-	128 KB	128 KB
	256 KB	128 KB	-
	-	64 KB	64 KB
-	-	512 KB	-
	-	256 KB	256 KB
	256 KB	256 KB	-
	-	128 KB	128 KB

6.3 L2 cache and SRAM organization

The on-chip memory array has eight banks, each containing 256 sets of eight cache blocks (or "ways"). Each block consists of 32 bytes of data and a tag.

This figure shows the organization of the cache.

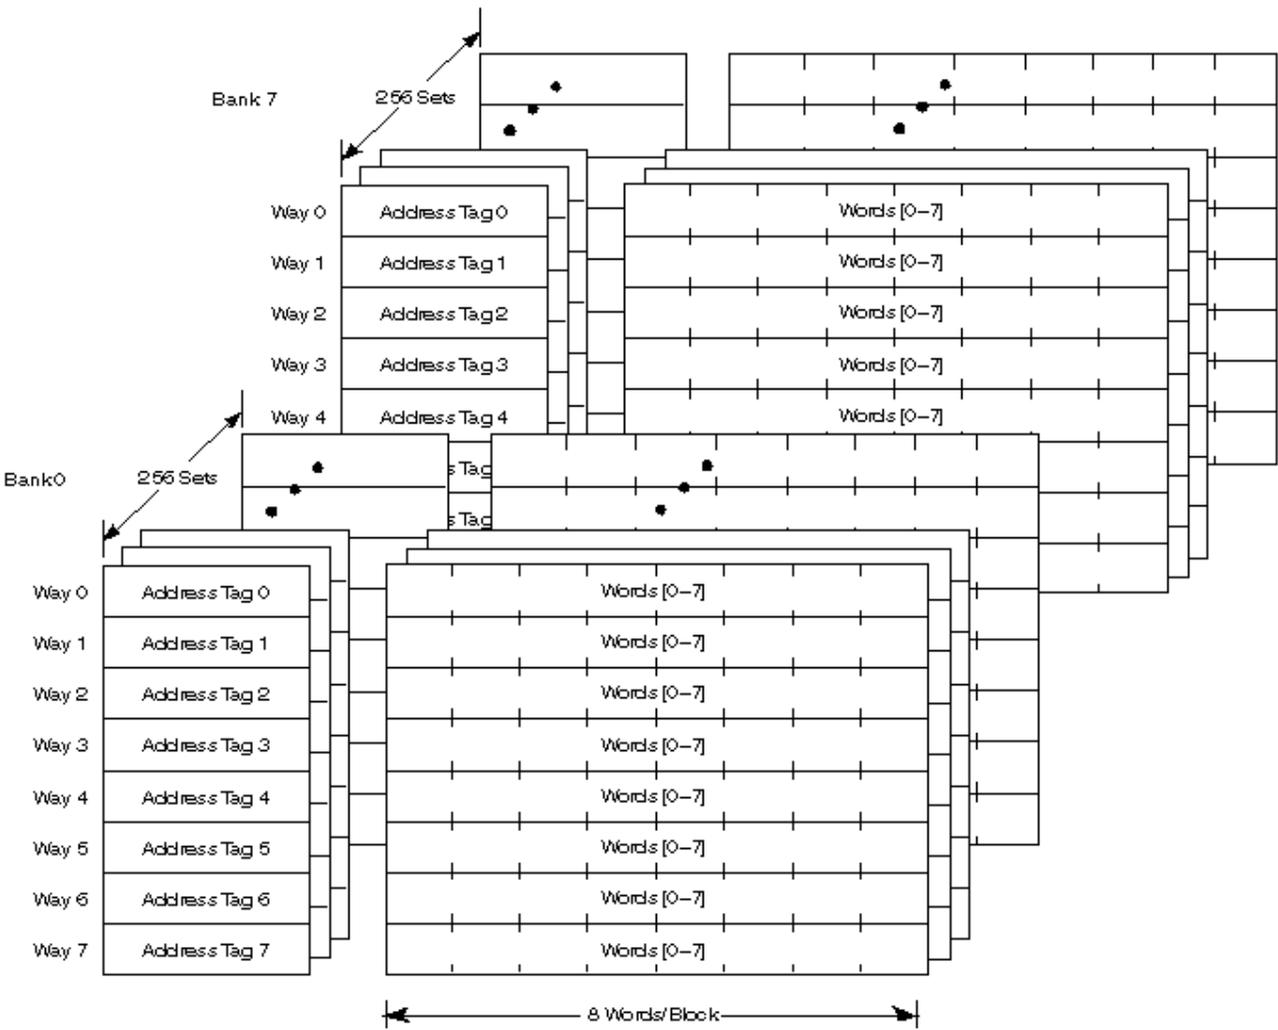


Figure 6-2. Cache organization

6.3.1 Accessing the on-chip array as an L2 cache

This figure shows how physical address bits are used to access the L2 cache.

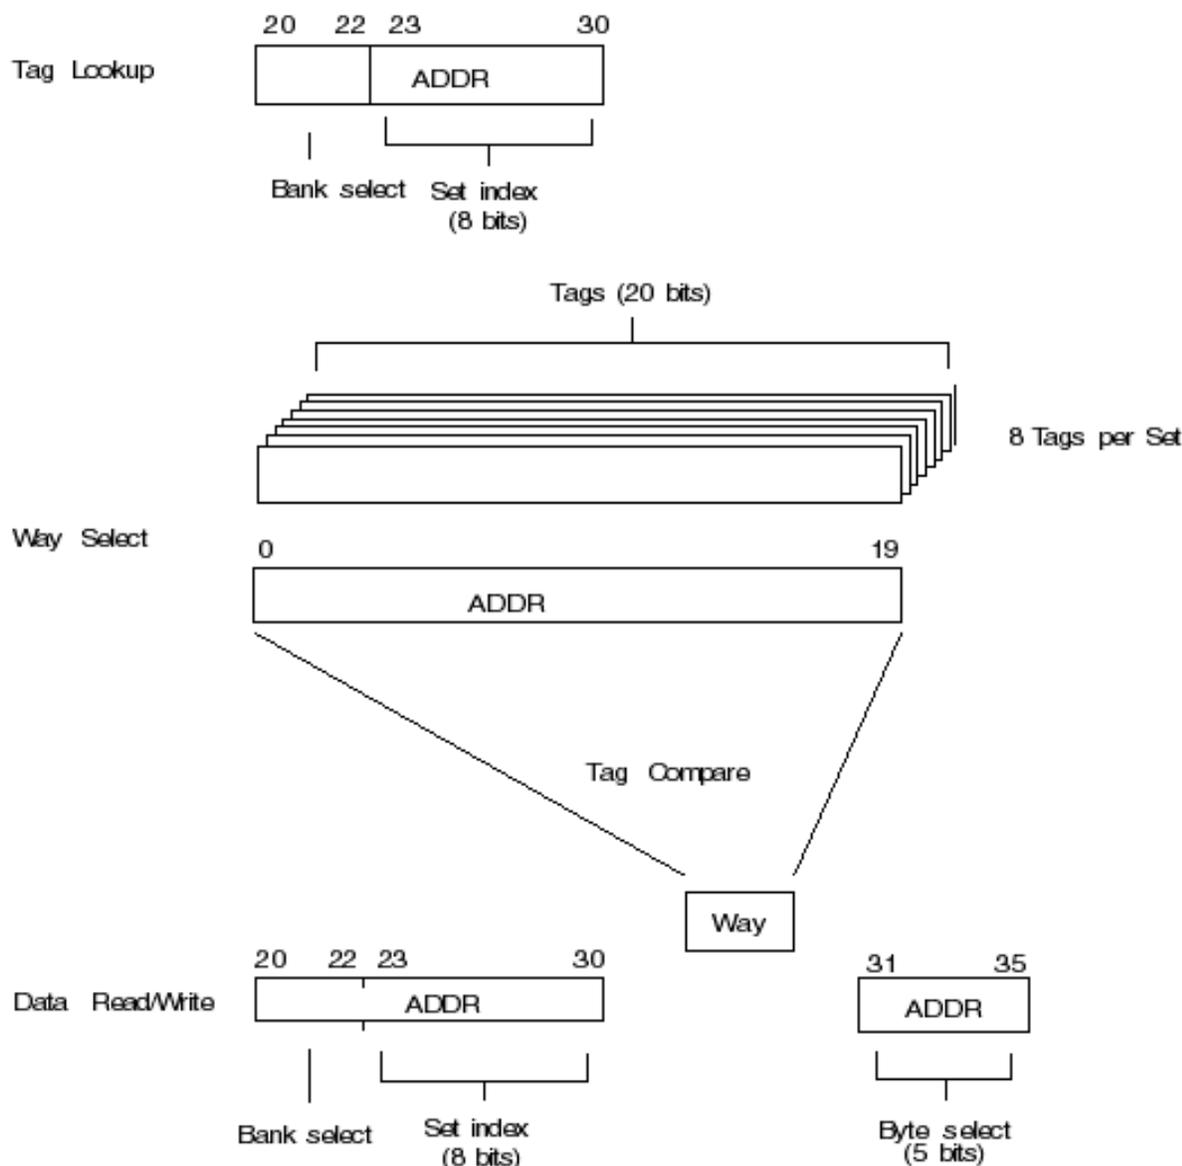


Figure 6-3. Physical address usage for L2 cache accesses

Physical address bits 20-30 identify the bank and set of the tag and data. Physical address bits 0- 19 are compared against the tags of all eight ways. A match of a valid tag selects a 32-byte block of data (or way) within the set. Physical address bits 31-35 identify the byte or bytes of data within the block.

6.3.2 Accessing the on-chip array as an SRAM

When all or part of the array is dedicated to memory mapped SRAM, individual ways of each set are reserved for that purpose.

SRAM accesses use physical address bits 17-19 in conjunction with the SRAM mode to select a way of the indexed set.

This figure shows the physical address usage for SRAM accesses.

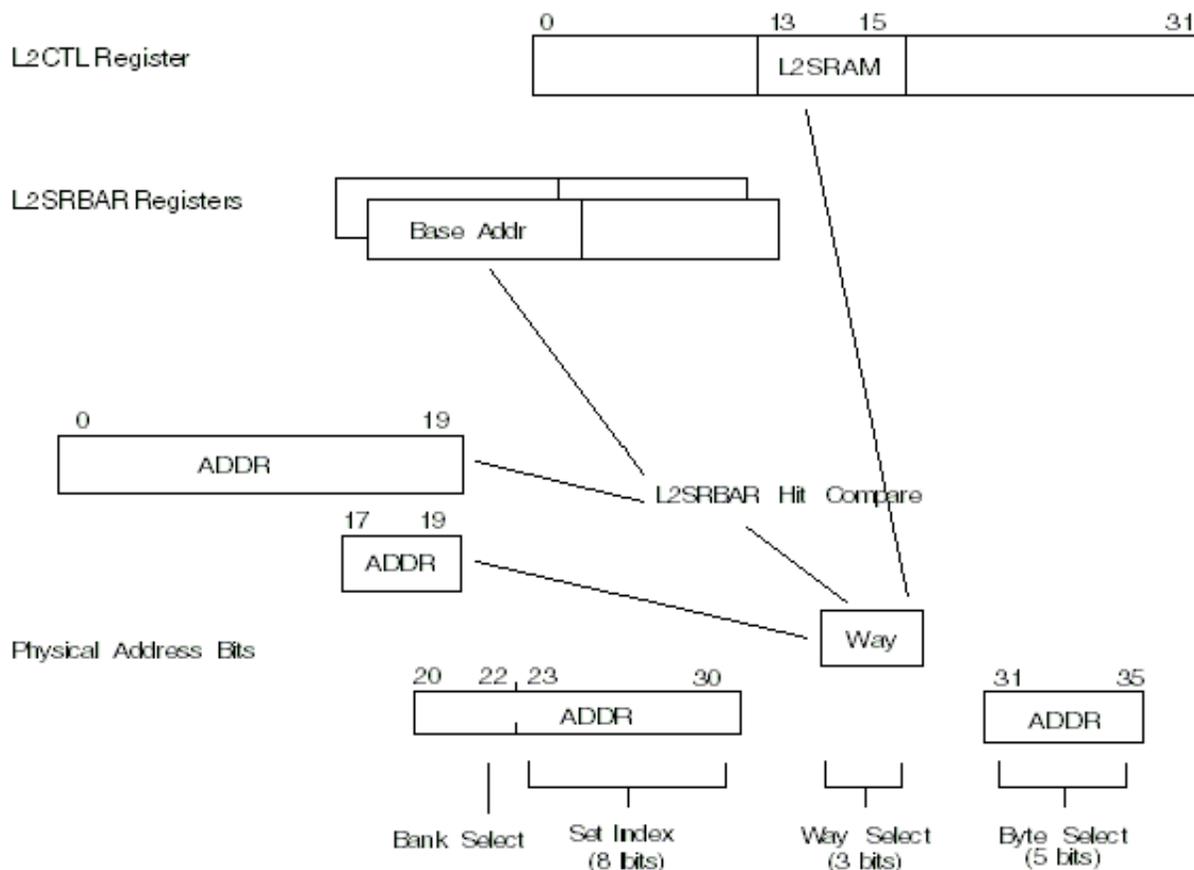


Figure 6-4. Physical address usage for SRAM accesses

The mapping of address bits and SRAM mode to a way select is shown in this table. SRAM size is reflected in L2CTL[L2SIZ].

Table 6-2. Way selection for SRAM accesses

Description	L2SRAM	L2SRBAR 0 Hit	L2SRBAR 1 Hit	Addr[17-19]	Way Select
No SRAM	000	-	-	-	-
Entire Array is SRAM (single 512 KB SRAM if L2SIZ=512 KB)	001	1	0	000	0
				001	1
				010	2
				011	3
				100	4
				101	5
				110	6
				111	7

Table continues on the next page...

Table 6-2. Way selection for SRAM accesses (continued)

Description	L2SRAM	L2SRBAR 0 Hit	L2SRBAR 1 Hit	Addr[17-19]	Way Select
One-half of array is an SRAM (single 256 KB SRAM if L2SIZ=512 KB)	010	1	0	x00	0
				x01	1
				x10	2
				x11	3
Both halves of array are SRAM (two 256 KB SRAM if L2SIZ=512 KB)	011	1	0	x00	0
				x01	1
				x10	2
				x11	3
		0	1	x00	4
				x01	5
				x10	6
				x11	7
One quarter of the array is SRAM (single 128 KB SRAM if L2SIZ=512 KB)	100	1	0	xx0	0
				xx1	1
Two quarters of the array are SRAMs (single 128 KB SRAM if L2SIZ=512 KB)	101	1	0	xx0	0
				xx1	1
		0	1	xx0	2
				xx1	3
One-eighth of the array is an SRAM (single 64 KB SRAM if L2SIZ=512 KB)	110	1	0	-	0
Two-eighths of the array are SRAM (single 64 KB SRAM if L2SIZ=512 KB)	111	1	0	-	0
		0	1	-	1

6.3.3 Connection of the on-chip memory to the system

The e500 core connects to the L2 cache and the system interface through the high-speed core complex bus (CCB).

The e500 core and the L2 cache connect to the rest of the integrated device through the e500 coherency module (ECM). This figure shows the data connections of the e500 core and L2/SRAM. The e500 core can simultaneously read 128 bits of data from the L2/SRAM, read 64 bits of data from the system interface, and write 128 bits of data to the L2/SRAM and/or system interface.

The L2/SRAM can be accessed by the e500 core or the system interface through the ECM. The L2 cache does not initiate transactions. [Figure 6-5](#) shows the data bus connections of the e500 core and L2/SRAM.

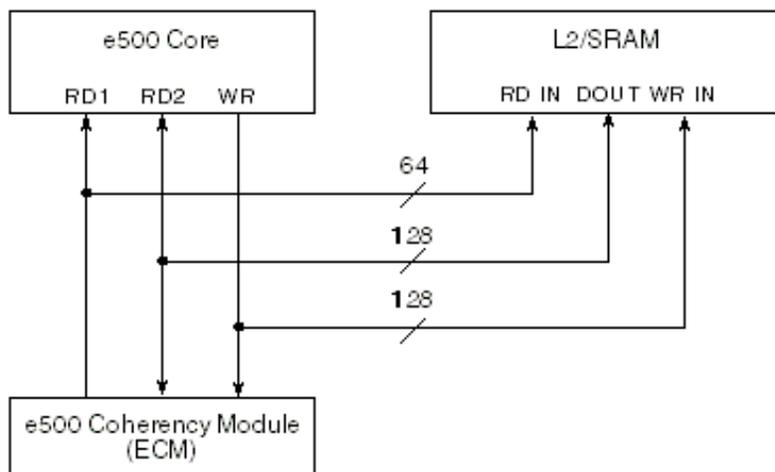


Figure 6-5. Data bus connection of CCB

This figure shows address connections of the e500 core and L2/SRAM.

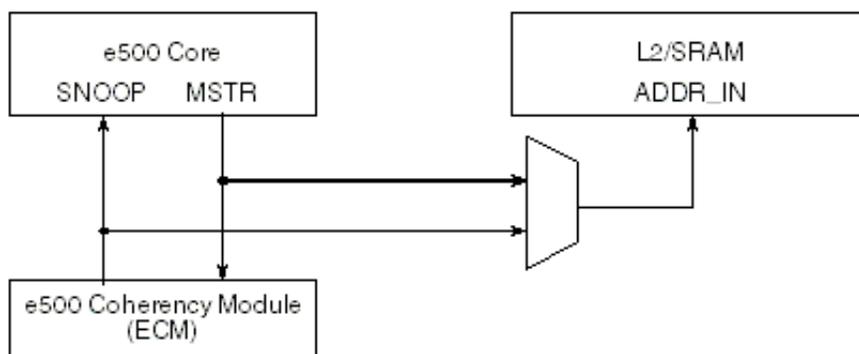


Figure 6-6. Address bus connection of CCB

In SRAM mode, if a non-cache-line read or write transaction is not preceded by a cache-line write, an ECC error occurs; such a non-cache-line write transaction cannot be allocated in the L2.

6.4 L2 cache external write registers

The device supports allocating and locking L2 cache lines from external agents such as PCI. This functionality is called stashing.

Four sets of registers are provided to support this feature; each set has three registers that specify a programmed memory range that can be locked with a snoop write transaction. All three registers in a set must be configured in order to use an external write address.

These registers are the L2 cache external write address registers 0-3 (L2CEWAR n), the L2 cache external write address registers extended address 0-3 (L2CEWAREA n), and the L2 cache external write control registers 0-3 (L2CEWCR n). L2CEWAR n contain the lower 24 bits of the external write base address and L2CEWAREA n contain the upper 4 bits. The base address specified in the address registers must be naturally aligned to the window size in the corresponding control register.

Further details on the locations and fields of these registers are given in the following sections.

6.5 L2 memory-mapped SRAM registers

The L2 memory-mapped SRAM base address registers 0-1 (L2SRBAR n) and the L2 memory-mapped SRAM base address registers extended address 0-1 (L2SRBAREA n), control the memory-mapped SRAM mode functionality. Together, these two pairs of registers define memory blocks that can be mapped into the L2 cache.

Specified SRAM base addresses must be aligned to the size of the SRAM region. If L2CTL[L2SRAM] specifies one memory-mapped SRAM block, its base address must be written to the pair L2SRBAR0 and L2SRBAREA0; if it specifies two memory-mapped SRAM blocks, L2SRBAR0 and L2SRBAREA0 are used for the first SRAM block and L2SRBAR1 and L2SRBAREA1 are used for the second block.

6.6 L2 error registers

L2 error detection, reporting, and injection allow flexible handling of ECC and parity errors in the L2 data and tag arrays.

When the device detects an L2 error, the appropriate bit in the error detect register (L2ERRDET) is set. Error detection is disabled by setting the corresponding bit in the error disable register (L2ERRDIS).

The address and attributes of the first detected error are also saved in the error capture registers (L2ERRADDR, L2ERRATTR, L2CAPTDATAHI, L2CAPTDATALO, and L2CAPTACC). Subsequent errors set error bits in the error detection registers, but information is saved only for the first one. Error reporting (by generating an interrupt) is enabled by setting the corresponding bit in the error interrupt enable register (L2ERRINTEN). Note that the error detect bit is set regardless of the state of the interrupt enable bit. When an error is detected, if error detection is enabled the L2 cache/SRAM

always asserts an internal error signal with read data to prevent the L1 caches and architectural registers from being loaded with corrupt data. If error detection is disabled, the detected error bit is not set and no internal signal is asserted.

The L2 error detect register (L2ERRDET) is implemented as a bit-reset type register. Reading from this register occurs normally; however, write operations can clear but not set bits. A bit is cleared whenever the register is written and the data in the corresponding bit location is a 1. For example, to clear bit 6 and not affect any other bits in the register, the value 0x0200_0000 is written to the register.

Note that in SRAM mode, if a non-cache-line read or write transaction is not preceded by a cache-line write, an ECC error occurs; such a non-cache-line write transaction cannot be allocated in the L2.

6.7 L2_Cache memory map/register definition

The following table shows the memory map for the L2/SRAM registers.

L2_Cache memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_0000	L2 control register (L2_Cache_L2CTL)	32	R/W	2000_0000h	6.7.1/194
2_0004	L2 cache way allocation for processors (L2_Cache_L2CWAP)	32	R/W	0000_0000h	6.7.2/199
2_0010	L2 cache external write address register n (L2_Cache_L2CEWAR0)	32	R/W	0000_0000h	6.7.3/201
2_0014	L2 cache external write address register extended address n (L2_Cache_L2CEWAREA0)	32	R/W	0000_0000h	6.7.4/201
2_0018	L2 cache external write control register n (L2_Cache_L2CEWCR0)	32	R/W	0000_0000h	6.7.5/202
2_0020	L2 cache external write address register n (L2_Cache_L2CEWAR1)	32	R/W	0000_0000h	6.7.3/201
2_0024	L2 cache external write address register extended address n (L2_Cache_L2CEWAREA1)	32	R/W	0000_0000h	6.7.4/201
2_0028	L2 cache external write control register n (L2_Cache_L2CEWCR1)	32	R/W	0000_0000h	6.7.5/202
2_0030	L2 cache external write address register n (L2_Cache_L2CEWAR2)	32	R/W	0000_0000h	6.7.3/201
2_0034	L2 cache external write address register extended address n (L2_Cache_L2CEWAREA2)	32	R/W	0000_0000h	6.7.4/201

Table continues on the next page...

L2_Cache memory map (continued)

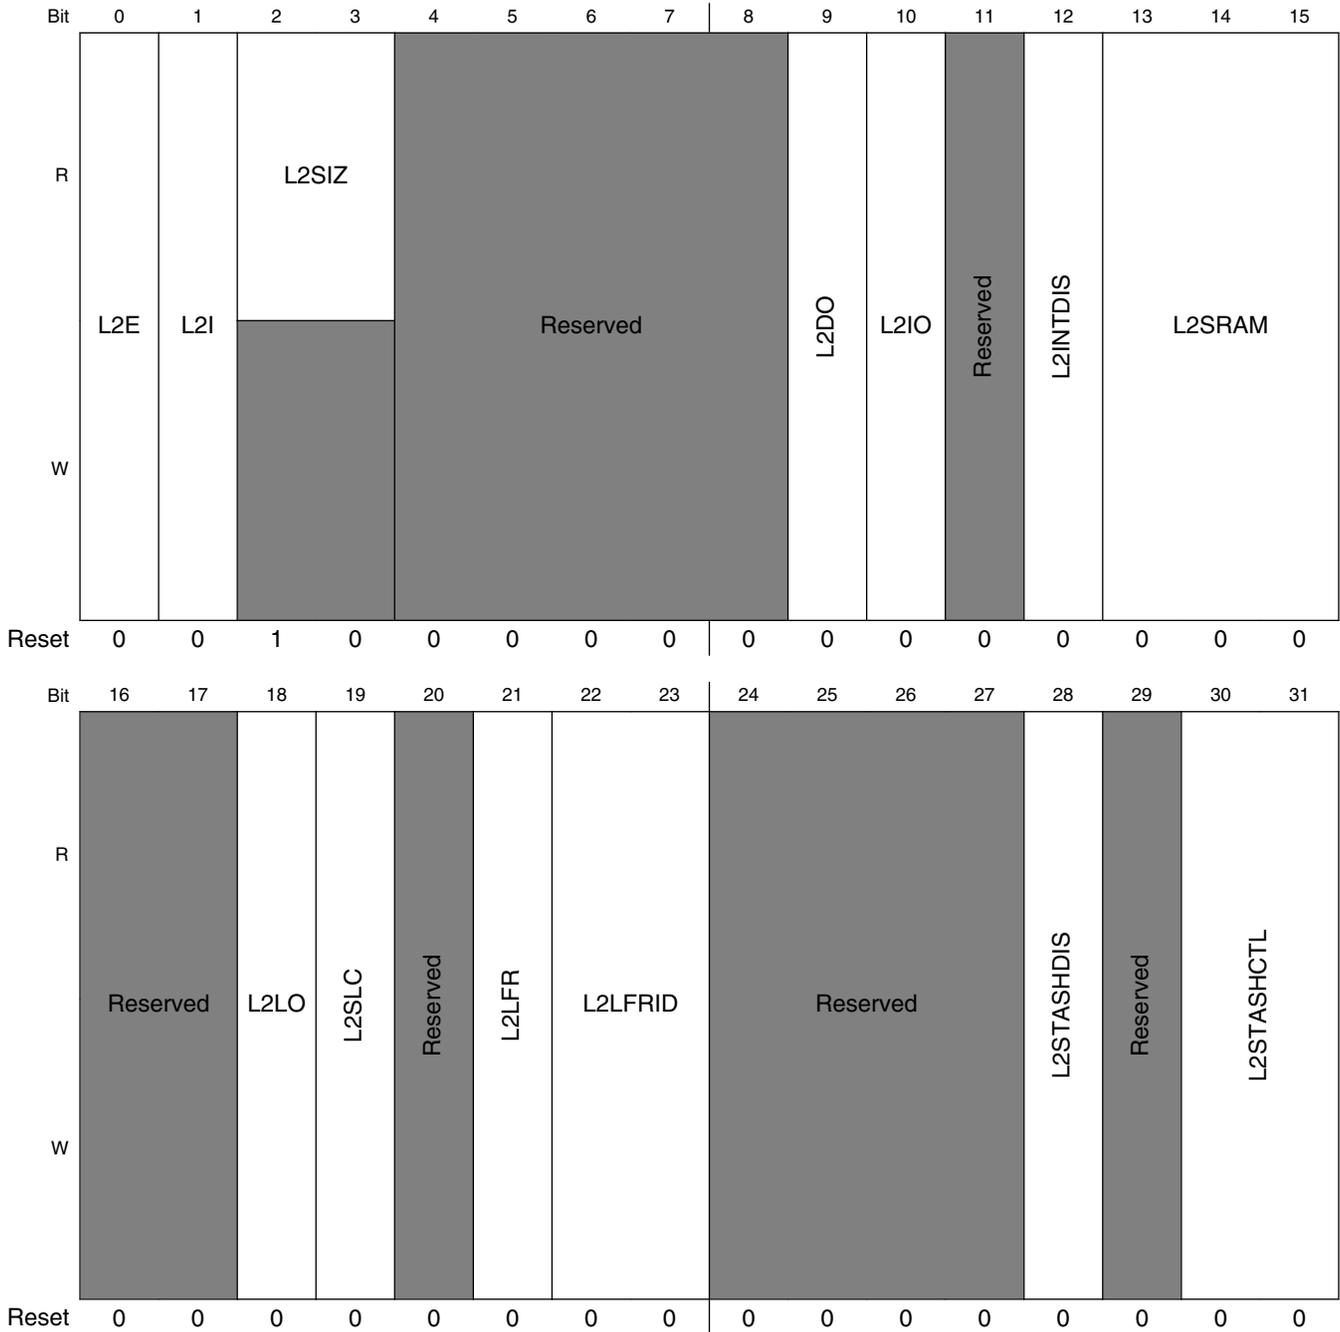
Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_0038	L2 cache external write control register n (L2_Cache_L2CEWCR2)	32	R/W	0000_0000h	6.7.5/202
2_0040	L2 cache external write address register n (L2_Cache_L2CEWAR3)	32	R/W	0000_0000h	6.7.3/201
2_0044	L2 cache external write address register extended address n (L2_Cache_L2CEWAREA3)	32	R/W	0000_0000h	6.7.4/201
2_0048	L2 cache external write control register n (L2_Cache_L2CEWCR3)	32	R/W	0000_0000h	6.7.5/202
2_0100	L2 memory-mapped SRAM base address register n (L2_Cache_L2SRBAR0)	32	R/W	0000_0000h	6.7.6/203
2_0104	L2 memory-mapped SRAM base address register extended address n (L2_Cache_L2SRBAREA0)	32	R/W	0000_0000h	6.7.7/204
2_0108	L2 memory-mapped SRAM base address register n (L2_Cache_L2SRBAR1)	32	R/W	0000_0000h	6.7.6/203
2_010C	L2 memory-mapped SRAM base address register extended address n (L2_Cache_L2SRBAREA1)	32	R/W	0000_0000h	6.7.7/204
2_0E00	L2 error injection mask high register (L2_Cache_L2ERRINJHI)	32	R/W	0000_0000h	6.7.8/205
2_0E04	L2 error injection mask low register (L2_Cache_L2ERRINJLO)	32	R/W	0000_0000h	6.7.9/206
2_0E08	L2 error injection tag/ECC control register (L2_Cache_L2ERRINJCTL)	32	R/W	0000_0000h	6.7.10/206
2_0E20	L2 error data high capture register (L2_Cache_L2CAPTDATAHI)	32	R	0000_0000h	6.7.11/207
2_0E24	L2 error data low capture register (L2_Cache_L2CAPTDATALO)	32	R	0000_0000h	6.7.12/208
2_0E28	L2 error syndrome register (L2_Cache_L2CAPTECC)	32	R	0000_0000h	6.7.13/208
2_0E40	L2 error detect register (L2_Cache_L2ERRDET)	32	w1c	0000_0000h	6.7.14/210
2_0E44	L2 error disable register (L2_Cache_L2ERRDIS)	32	R/W	0000_0000h	6.7.15/212
2_0E48	L2 error interrupt enable register (L2_Cache_L2ERRINTEN)	32	R/W	0000_0000h	6.7.16/214
2_0E4C	L2 error attributes capture register (L2_Cache_L2ERRATTR)	32	R/W	0000_0000h	6.7.17/216
2_0E50	L2 error address capture register low (L2_Cache_L2ERRADDRLO)	32	R	0000_0000h	6.7.18/218
2_0E54	L2 error address capture register high (L2_Cache_L2ERRADDRH)	32	R	0000_0000h	6.7.19/218
2_0E58	L2 error control register (L2_Cache_L2ERRCTL)	32	R/W	0000_0000h	6.7.20/219

6.7.1 L2 control register (L2_Cache_L2CTL)

The L2 control register (L2CTL), shown in the figure below, controls configuration and operation of the L2/SRAM array. The sequence for modifying L2CTL is as follows:

1. **mbar**
2. **isync**
3. **stw** (WIMG = 01xx) CCSRBAR + 0x2_0000
4. **lwz** (WIMG = 01xx) CCSRBAR + 0x2_0000
5. **mbar**

Address: 2_0000h base + 0h offset = 2_0000h



L2_Cache_L2CTL field descriptions

Field	Description
0 L2E	<p>L2 enable. Used to enable the L2 array (cache or memory-mapped SRAM). Note that L2I can be set regardless of the value of L2E.</p> <p>0 The L2 SRAM (cache and memory-mapped SRAM) is disabled and is not accessed for reads, snoops, or writes. Setting the L2 flash invalidate bit (L2I) is allowed. 1 The L2 SRAM (cache or memory-mapped SRAM) is enabled.</p>
1 L2I	<p>L2 flash invalidate. Data to memory-mapped SRAM are unaffected by the flash invalidate. The hardware automatically clears L2I when the invalidate is complete.</p> <p>0 The L2 status and LRU bits are not being cleared. 1 Setting L2I invalidates the L2 cache globally by clearing all the L2 status bits, as well as the LRU algorithm. Memory-mapped SRAM is unaffected.</p>
2–3 L2SIZ	<p>L2 SRAM size (read only). Indicates the total available size of on-chip memory array (to be configured as cache or memory-mapped SRAM).</p> <p>00 Reserved 01 Reserved 10 512 KB 11 Reserved</p>
4–8 -	<p>This field is reserved. Reserved</p>
9 L2DO	<p>L2 data-only. Reserved in full memory-mapped SRAM mode. L2DO may be changed while the L2 is enabled or disabled. Note that if L2DO and L2IO are both set, no new lines are allocated into the L2 cache for any processor transactions, and processor writes and castouts that hit existing data in the cache invalidate those lines rather than updating them.</p> <p>0 The L2 cache allocates entries for instruction fetches that miss in the L2. 1 The L2 cache allocates entries for processor data loads that miss in the L2 and for processor L1 castouts but does not allocate entries for instruction fetches that miss in the L2. Instruction accesses that hit in the L2, data accesses, and accesses from the system (including I/O stash writes) are unaffected.</p>
10 L2IO	<p>L2 instruction-only. Reserved in full memory-mapped SRAM mode. Causes the L2 cache to allocate lines for instruction cache transactions only. L2IO may be changed while the L2 is enabled or disabled. Note that if L2DO and L2IO are both set, no new lines are allocated into the L2 cache for any processor transactions, and processor writes and castouts that hit existing data in the cache invalidate those lines rather than updating them.</p> <p>0 The L2 cache entries are allocated for data loads that miss in the L2 and for processor L1 castouts. 1 The L2 cache allocates entries for instruction fetch misses, but does not allocate entries for processor data transactions. Data accesses that hit in the L2, instruction accesses, and accesses from the system (including I/O stash writes) are unaffected.</p>
11 -	<p>This field is reserved. Reserved</p>
12 L2INTDIS	<p>Cache read intervention disable. Reserved for full memory-mapped SRAM mode. Used to disable cache read intervention. May be changed while the L2 is enabled or disabled.</p>

Table continues on the next page...

L2_Cache_L2CTL field descriptions (continued)

Field	Description
	<p>0 Cache intervention is enabled. The ECM ensures that if a data read from another device hits in the L2 cache, it is serviced from the L2 cache.</p> <p>1 Cache intervention is disabled</p>
13–15 L2SRAM	<p>L2 SRAM configuration. Determines the L2 cache/memory-mapped SRAM allocation of the on-chip memory array. SRAM size depends on the value of L2SIZ. Since L2SIZ is 512 KB, SRAM can have sizes from 64 KB to 512 KB.</p> <p>For one SRAM region L2SRBAR0 is used and for two SRAM regions L2SRBAR0 and L2SRBAR1 are used. Regions of the array that are not allocated to SRAMs are used as cache memory.</p> <p>To change these bits, the L2 must be disabled (L2CTL[L2E] = 0).</p> <p>Note that when setting L2SRAM after cache has been enabled, L2I should be set as well. The fields can be set simultaneously, and this step is not needed if SRAM size is getting smaller.</p> <p>000 No SRAM. Entire array is cache. 001 Entire array is a single SRAM (512 -KB SRAM for L2SIZ = 512 KB) 010 One half of the array is an SRAM (256 -KB SRAM for L2SIZ = 512 KB) 011 Both halves of the array are SRAMs (two 256 -KB SRAMs for L2SIZ = 512 KB) 100 One quarter of the array is an SRAM (one 128 -KB SRAM for L2SIZ = 512 KB) 101 Two quarters of the array are SRAMs (two 128 -KB SRAMs for L2SIZ = 512 KB) 110 One eighth of the array is an SRAM (one 64 -KB SRAM for L2SIZ = 512 KB) 111 Two eighths of the array are SRAMs (two 64 -KB SRAMs for L2SIZ = 512 KB)</p>
16–17 -	<p>This field is reserved. Reserved</p>
18 L2LO	<p>L2 cache lock overflow. Reserved in full memory-mapped SRAM mode. This sticky bit is set if an overlock condition is detected in the L2 cache. A lock overflow is triggered either by executing instruction or data cache block touch and lock set instructions or by performing L2 cache external writes with lock set. If all ways are locked and an attempt to stash is made, the stash is not allocated.</p> <p>0 The L2 cache did not encounter a lock overflow. L2LO is cleared only by software. 1 The L2 cache encountered a lock overflow condition.</p>
19 L2SLC	<p>L2 snoop lock clear. This sticky bit is set if a snoop invalidated a locked data cache line. Note that the lock bit for that line is cleared whenever the line is invalidated. L2SLC is reserved in full memory-mapped SRAM mode.</p> <p>0 A snoop did not invalidate a locked L2 cache line. L2SLC is cleared only by software. 1 The L2 cache encountered a snoop that invalidated a locked line.</p>
20 -	<p>This field is reserved. Reserved</p>
21 L2LFR	<p>L2 cache lock bits flash reset. The L2 cache must be enabled (L2CTL[L2E] = 1) for reset to occur. This field is reserved in full memory-mapped SRAM mode.</p> <p>0 The L2 cache lock bits are not cleared or the clear operation completed. 1 A reset operation is issued that clears each L2 cache line's lock bits. Depending on the L2LFRID value, data or instruction locks, or both, can be reset. Cache access is blocked during this time. After L2LFR is set, the L2 cache unit automatically clears L2LFR when the reset operation is complete (if L2CTL[L2E] is set).</p>
22–23 L2LFRID	<p>L2 cache lock bits flash reset select instruction or data. Indicates whether data or instruction lock bits or both are reset.</p> <p>00 Not used</p>

Table continues on the next page...

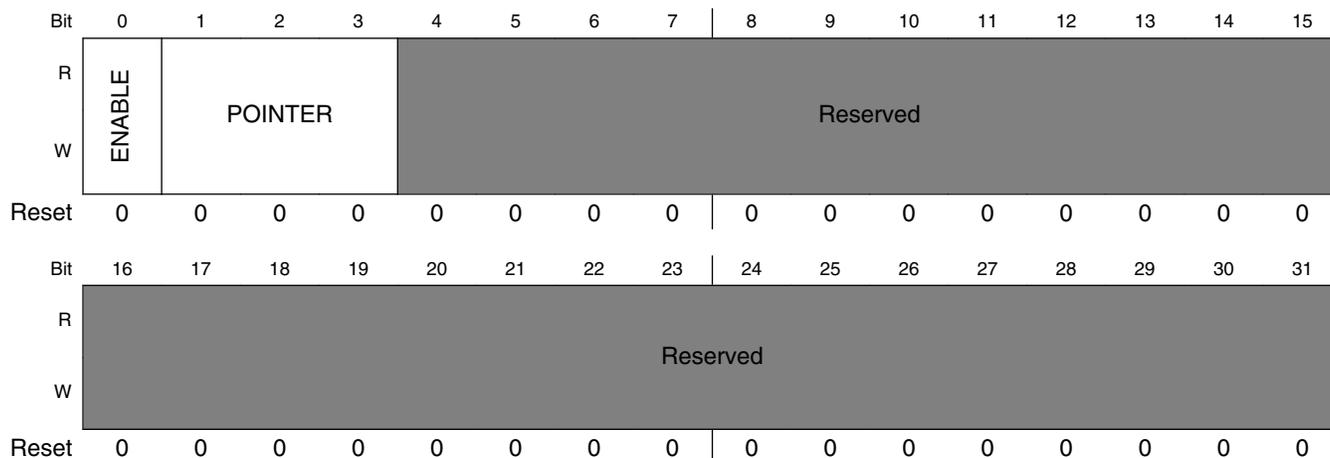
L2_Cache_L2CTL field descriptions (continued)

Field	Description
	01 Reset data locks if L2LFR = 1. 10 Reset instruction locks if L2LFR = 1. 11 Reset both data and instruction locks if L2LFR = 1.
24–27 -	This field is reserved. Reserved
28 L2STASHDIS	L2 stash allocate disable. Disables allocation of lines for stashing. This bit does NOT affect the updating of lines that are already resident in the cache and have the stash attribute set or hit a stash range. Such lines are updated even if this bit is set. NOTE: To change this bit, the L2 must be disabled (L2CTL[L2E] = 0). 0 The L2 cache allocates lines for global writes that hit in a stash range or that have the stashing attribute set. 1 The L2 does not allocate lines for stashed writes.
29 -	This field is reserved. Reserved
30–31 L2STASHCTL	L2 stash configuration. This field reserves regions of the cache for stash-only operation. That is, blocks of each cache set are reserved so that they can only be allocated for stash data. If such a region is created, processor reads and writes are not allocated into this region; it can only be populated by stash writes. Similarly, stash writes are only allocated into this region. This prevents processor and stashed I/O data from polluting one another. Like L2SRAM configuration, stash-only regions subtract from the amount of the on-chip memory that is available to the processor as cache. If the L2SRAM configuration uses the entire on-chip memory array as SRAM, then no stash-only region can be created. To change these bits, the L2 must be disabled (L2CTL[L2E] = 0). This field has no effect if the L2STASHDIS bit is set. 00 No stash-only region. Stashed writes are allocated across the entire cache and can evict processor data and can be evicted by processor data. 01 One half of the array is a stash-only cache (way4, way5, way6 & way7 of each set) 10 One quarter of the array is a stash-only cache (way6 & way7 of each set) 11 One eighth of the array is a stash-only cache (way7 of each set)

6.7.2 L2 cache way allocation for processors (L2_Cache_L2CWAP)

The L2 cache way allocation for processors (L2CWAP) register, shown in the figure below, allows the L2 cache to be programmed to distribute eight ways between two processor cores. A pointer is used to set the ways allocated between processor 0 and processor 1.

Address: 2_0000h base + 4h offset = 2_0004h



L2_Cache_L2CWAP field descriptions

Field	Description
0 ENABLE	L2 cache way pointer enable. The L2 cache uses the pointer value only when ENABLE is set. 0 The L2 way pointer is disabled. No ways are allocated only for processor 0 or for processor 1. All ways (except ways allocated for SRAM and stashing) are shared between processors. 1 The L2 way pointer is enabled. Ways are allocated for processors based on way pointer.
1-3 POINTER	L2 cache way pointer. This pointer is used to allocate ways between processor 0 and processor 1. To change this field the L2 must be disabled (L2CTL[L2E] = 0). If ways have been allocated for SRAM (L2CTL[13-15]) or stashing (L2CTL[30-31]), way allocation is affected as noted in the following examples. In general, when allocating ways among SRAM, stashing, and processors, the controller gives SRAM first preference, then stashing, then processors. Suppose 2 ways are allocated for SRAM and 2 ways are allocated for stashing and POINTER = 100... Way0 -> SRAM Way1 -> SRAM Way2 -> Processor 0 Way3 -> Processor 0 Way4 -> Processor 1 Way5 -> Processor 1

Table continues on the next page...

L2_Cache_L2CWAP field descriptions (continued)

Field	Description
	<p>Way6 -> Stash Way7 -> Stash Suppose 4 ways are allocated for SRAM and 2 ways are allocated for stashing and POINTER = 100... Way0 -> SRAM Way1 -> SRAM Way2 -> SRAM Way3 -> SRAM Way4 -> Processor 1 Way5 -> Processor 1 Way6 -> Stash Way7 -> Stash Suppose 2 ways are allocated for SRAM and no ways are allocated for stashing and POINTER = 100... Way0 -> SRAM Way1 -> SRAM Way2 -> Processor 0/Stash Way3 -> Processor 0/Stash Way4 -> Processor 1/Stash Way5 -> Processor 1/Stash Way6 -> Processor 1/Stash Way7 -> Processor 1/Stash</p> <p>The following description for way pointer values assumes that no ways are allocated for L2SRAM (L2CTL[13-15]) and L2STASH (L2CTL[30-31]).</p> <p>000 Way0-way7 are allocated for processor 1 and no ways are allocated for processor 0. 001 Way0 is allocated for processor 0 and way1-way7 are allocated for processor 1. 010 Way0-way1 are allocated for processor 0 and way2-way7 are allocated for processor 1. 011 Way0-way2 are allocated for processor 0 and way3- way7 are allocated for processor 1. 100 Way0-way3 are allocated for processor 0 and way4-way7 are allocated for processor 1. 101 Way0-way4 are allocated for processor 0 and way5-way7 are allocated for processor 1. 110 Way0-way5 are allocated for processor 0 and way6-way7 are allocated for processor 1. 111 Way0-way6 are allocated for processor 0 and way7 is allocated for processor 1.</p>
4-31 -	This field is reserved. Reserved

6.7.3 L2 cache external write address register n (L2_Cache_L2CEWAR n)

The L2CEWAR n registers contain the lower 24 bits of the 28-bit L2 cache external write base address. Each of these registers has identical fields, as shown in the figure below.

Address: 2_0000h base + 10h offset + (16d × i), where $i=0d$ to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	ADDR															Reserved																
W	ADDR															Reserved																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

L2_Cache_L2CEWAR n field descriptions

Field	Description
0–23 ADDR	Contains the lower 24 bits of the 28-bit L2 cache external write base address. Note that the upper 4 bits of the base address are in L2CEWAREA n [ADDR].
24–31 -	This field is reserved. Reserved

6.7.4 L2 cache external write address register extended address n (L2_Cache_L2CEWAREA n)

The L2 cache external write address registers extended address (L2CEWAREA n), shown in the figure below, contain the upper 4 bits of the 28-bit L2 cache external write base address.

Address: 2_0000h base + 14h offset + (16d × i), where $i=0d$ to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															ADDR																
W	Reserved															ADDR																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

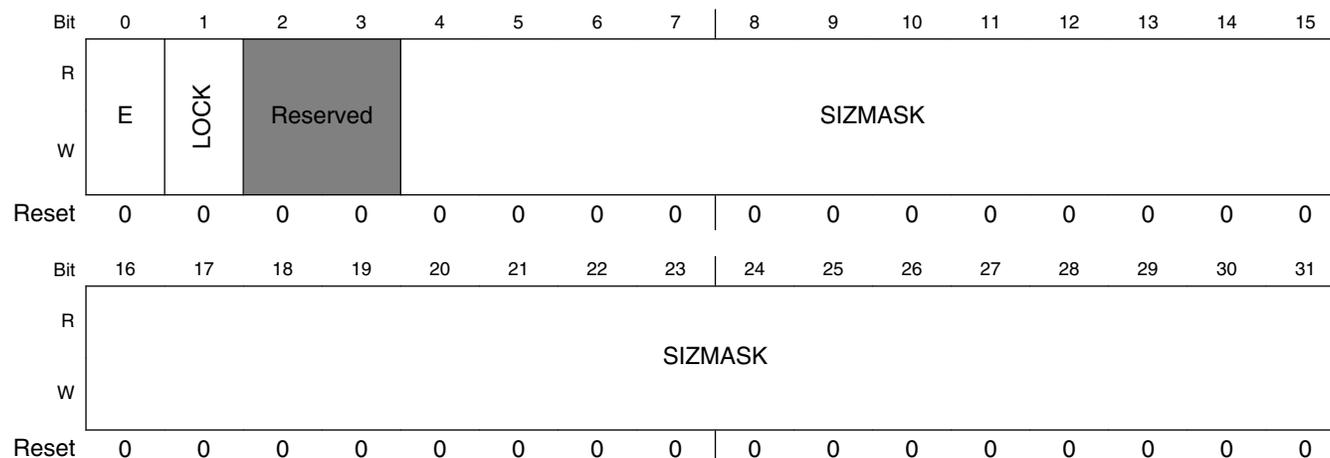
L2_Cache_L2CEWAREA n field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 ADDR	Contains the upper 4 bits of the L2 cache external write base address. Note that the rest of the base address is in L2CEWAR n [ADDR].

6.7.5 L2 cache external write control register n (L2_Cache_L2CEWCRn)

The L2CEWAR *n* /L2CEWAREA *n* address registers work with the L2 cache external write control registers 0-3 (L2CEWCR *n*), shown in the figure below, to control cache external write functionality.

Address: 2_0000h base + 18h offset + (16d × i), where i=0d to 3d



L2_Cache_L2CEWCRn field descriptions

Field	Description
0 E	External write enable. An external write matching the address window defined by L2CEWAR <i>n</i> / L2CEWAREA <i>n</i> /L2CEWCR <i>n</i> is allocated or updated in the L2 cache. 0 External writes for the L2CEWAR <i>n</i> /L2CEWAREA <i>n</i> /L2CEWCR <i>n</i> set are disabled. 1 External writes are enabled for the L2CEWAR <i>n</i> /L2CEWAREA <i>n</i> /L2CEWCR <i>n</i> set.
1 LOCK	Lock lines in the targeted cache. An external write matching the address window defined by L2CEWAR <i>n</i> / L2CEWAREA <i>n</i> /L2CEWCR <i>n</i> is locked in the L2 cache when it is allocated or updated. 0 The locked bit is not set when a line is allocated unless explicitly specified by transaction attributes. 1 Cache lines are allocated as locked. A hit to a valid, unlocked line sets the lock.
2-3 -	This field is reserved. Reserved
4-31 SIZMASK	Mask size. Defines the size of the naturally aligned address region for cache external writes. The address region must be aligned to a boundary that is a multiple of the mask size. Any value not listed below is illegal and produces boundedly undefined results. 1111 1111 1111 1111 1111 1111 1111 256 bytes 1111 1111 1111 1111 1111 1111 1110 512 bytes 1111 1111 1111 1111 1111 1111 1100 1 KB 1111 1111 1111 1111 1111 1111 1000 2 KB 1111 1111 1111 1111 1111 1111 0000 4 KB

Table continues on the next page...

L2_Cache_L2CEWCR n field descriptions (continued)

Field	Description
1111 1111 1111 1111 1111 1110 0000	8 KB
1111 1111 1111 1111 1111 1100 0000	16 KB
1111 1111 1111 1111 1111 1000 0000	32 KB
1111 1111 1111 1111 1111 0000 0000	64 KB
1111 1111 1111 1111 1110 0000 0000	128 KB
1111 1111 1111 1111 1100 0000 0000	256 KB
1111 1111 1111 1111 1000 0000 0000	512 KB
1111 1111 1111 1111 0000 0000 0000	1 MB
1111 1111 1111 1110 0000 0000 0000	2 MB
1111 1111 1111 1100 0000 0000 0000	4 MB
1111 1111 1111 1000 0000 0000 0000	8 MB
1111 1111 1111 0000 0000 0000 0000	16 MB
1111 1111 1110 0000 0000 0000 0000	32 MB
1111 1111 1100 0000 0000 0000 0000	64 MB
1111 1111 1000 0000 0000 0000 0000	128 MB
1111 1111 0000 0000 0000 0000 0000	256 MB
1111 1110 0000 0000 0000 0000 0000	512 MB
1111 1100 0000 0000 0000 0000 0000	1 GB
1111 1000 0000 0000 0000 0000 0000	2 GB
1111 0000 0000 0000 0000 0000 0000	4 GB
1110 0000 0000 0000 0000 0000 0000	8 GB
1100 0000 0000 0000 0000 0000 0000	16 GB
1000 0000 0000 0000 0000 0000 0000	32 GB
0000 0000 0000 0000 0000 0000 0000	64 GB

6.7.6 L2 memory-mapped SRAM base address register n (L2_Cache_L2SRBAR n)

The L2 memory-mapped SRAM base address registers (L2SRBAR n) contain the lower 18 bits of the 22-bit SRAM base address.

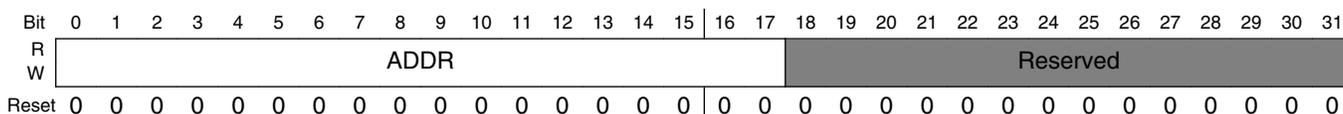
SRAM Partition	Bits Required for SRAM Offset	Bits Used for Actual Base Address
64 KB	16	20 (0-19)
128 KB	17	19 (0-18)
256 KB	18	18 (0-17)
512 KB	19	17 (0-16)

When enabled, the windows defined in L2SRBAR n and L2SRBAREA n supersede all other mappings of these addresses for processor and global (snoopable) I/O transactions. Therefore, SRAM windows must never overlap configuration space as defined by CCSRBAR (see [Configuration, control, and status registers base address register.](#)) Overlapping SRAM and local access windows is discouraged because processor and

L2_Cache memory map/register definition

snooperable I/O transactions would map to the SRAM while non-snoopered I/O transactions would be mapped by the local access windows. Only if all accesses to the SRAM address range are snooperable can results be consistent if SRAM and local access windows overlap.

Address: 2_0000h base + 100h offset + (8d × i), where i=0d to 1d



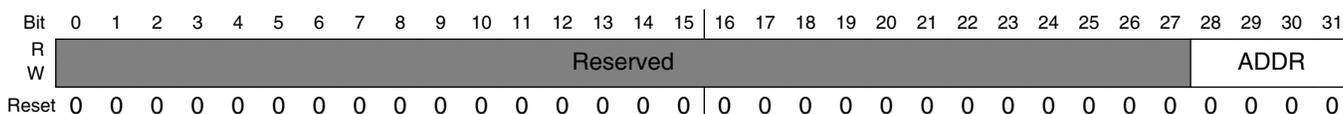
L2_Cache_L2SRBARn field descriptions

Field	Description
0–17 ADDR	Contains the lower 18 bits of the 22-bit L2 memory-mapped SRAM base address; the upper 4 bits are contained in L2SRBAREA n [ADDR]. (Note that some of these bits may not be needed, depending on how the L2 cache is partitioned.) The combined base address from L2SRBAREA n [ADDR] L2SRBAR n [ADDR] is used as shown in L2 memory-mapped SRAM base address register n (L2_Cache_L2SRBARn) Unused bits of the base address are masked off by the hardware
18–31 -	This field is reserved. Reserved

6.7.7 L2 memory-mapped SRAM base address register extended address n (L2_Cache_L2SRBAREAn)

The L2 memory-mapped SRAM base address registers extended address (L2SRBAREAn), shown in the figure below, contain the upper 4 bits of the L2 cache SRAM base address.

Address: 2_0000h base + 104h offset + (8d × i), where i=0d to 1d



L2_Cache_L2SRBAREAn field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 ADDR	Contains the upper 4 bits of the L2 cache SRAM base address. Note that the 18 low-order bits of the base address are contained in L2SRBAR n [ADDR].

6.7.8 L2 error injection mask high register (L2_Cache_L2ERRINJHI)

The L2 cache includes support for injecting errors into the L2 data, data ECC, or tag. This may be used to test error recovery software by deterministically creating error scenarios.

The preferred method for error injection is to set all data pages to cache-inhibited (MMU TLB entry I = 1) except a scratch page, set L2CTL[L2DO] to prevent allocation of instruction accesses, and invalidate the L2 by setting L2CTL[L2I] = 1. The following code sequence triggers an error, then detects it (A is an address in the scratch page):

```

dcbz A           | allocates the line in the L1 in the modified state

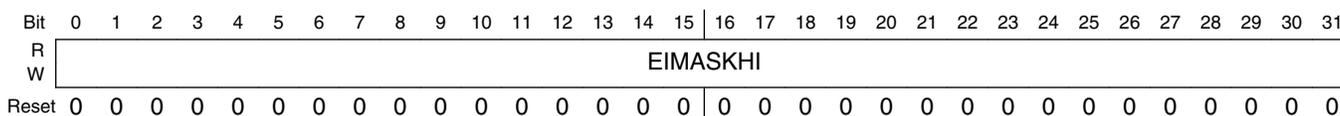
dcbt1s_L2 A     | forces the line from the L1 and allocates the line in the L2

lwz A
    
```

Data or tag errors are injected into the line, according to the error injection settings in L2ERRINJHI, L2ERRINJLO, and L2ERRINJCTL, at allocation. The final load detects and reports the error (if enabled) and allows software to examine the offending data, address, and attributes.

Note that error injection enable bits in L2ERRINJCTL must be cleared by software and the L2 must be invalidated (by setting L2CTL[L2I]) before resuming L2 normal operation. The figure below shows the L2 error injection mask high register (L2ERRINJHI).

Address: 2_0000h base + E00h offset = 2_0E00h



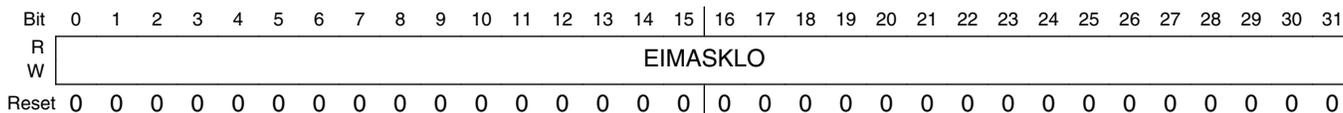
L2_Cache_L2ERRINJHI field descriptions

Field	Description
0–31 EIMASKHI	Error injection mask/high word. A set bit corresponding to a data path bit causes that bit on the data path to be inverted on cache/SRAM writes if L2ERRINJCTL[DERRIEN] = 1.

6.7.9 L2 error injection mask low register (L2_Cache_L2ERRINJLO)

The figure below shows the L2 error injection mask low register (L2ERRINJLO).

Address: 2_0000h base + E04h offset = 2_0E04h



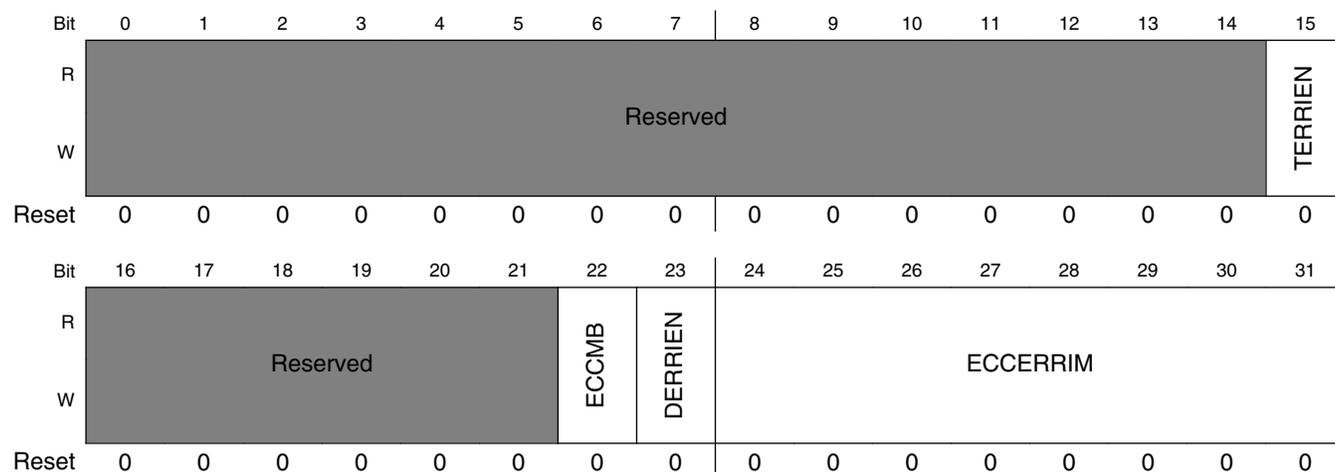
L2_Cache_L2ERRINJLO field descriptions

Field	Description
0–31 EIMASKLO	Error injection mask/low word. A set bit corresponding to a data path bit causes that bit on the data path to be inverted on SRAM writes if L2ERRINJCTL[DERRIEN] = 1.

6.7.10 L2 error injection tag/ECC control register (L2_Cache_L2ERRINJCTL)

The figure below shows the L2 error injection mask control register (L2ERRINJCTL).

Address: 2_0000h base + E08h offset = 2_0E08h



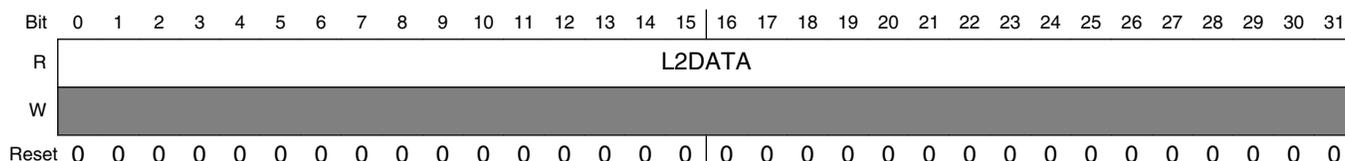
L2_Cache_L2ERRINJCTL field descriptions

Field	Description
0–14 -	This field is reserved. Reserved
15 TERRIEN	L2 tag array error injection enable 0 No tag errors are injected. 1 All subsequent entries written to the L2 tag array have the parity bit inverted.
16–21 -	This field is reserved. Reserved
22 ECCMB	ECC mirror byte enable. 0 ECC byte mirroring is disabled 1 The most significant data path byte is mirrored onto the ECC byte if DERRIEN = 1.
23 DERRIEN	L2 data array error injection enable: NOTE: If both ECC mirror byte and data error injection are enabled, ECC mask error injection is performed on the mirrored ECC. 0 No data errors are injected. 1 Subsequent entries written to the L2 data array have data or ECC bits inverted as specified in the data and ECC error injection masks and/or data path byte mirrored onto ECC as specified by ECC mirror byte enable.
24–31 ECCERRIM	Error injection mask for the ECC bits. A set bit corresponding to an ECC bit causes that bit to be inverted on SRAM writes if DERRIEN = 1.

6.7.11 L2 error data high capture register (L2_Cache_L2CAPTDATAHI)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error capture data high register (L2CAPTDATAHI).

Address: 2_0000h base + E20h offset = 2_0E20h



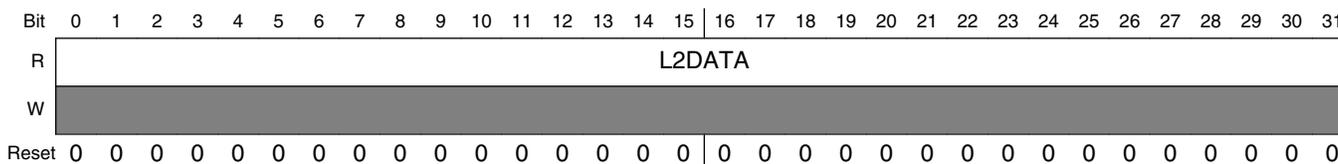
L2_Cache_L2CAPTDATAHI field descriptions

Field	Description
0–31 L2DATA	L2 data high word

6.7.12 L2 error data low capture register (L2_Cache_L2CAPTDATALO)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error capture data low register (L2CAPTDATALO).

Address: 2_0000h base + E24h offset = 2_0E24h



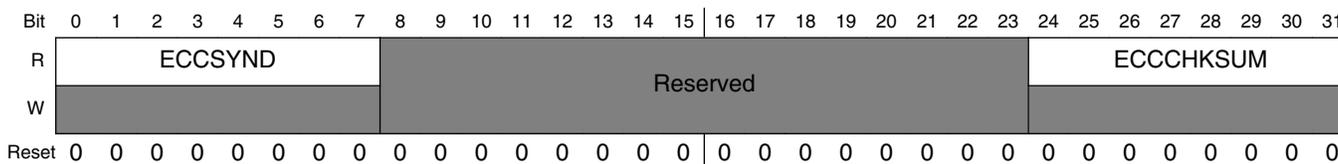
L2_Cache_L2CAPTDATALO field descriptions

Field	Description
0–31 L2DATA	L2 data low word

6.7.13 L2 error syndrome register (L2_Cache_L2CAPTECC)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error syndrome register (L2CAPTECC).

Address: 2_0000h base + E28h offset = 2_0E28h



L2_Cache_L2CAPTECC field descriptions

Field	Description
0–7 ECCSYND	The calculated ECC syndrome of the failing double word

Table continues on the next page...

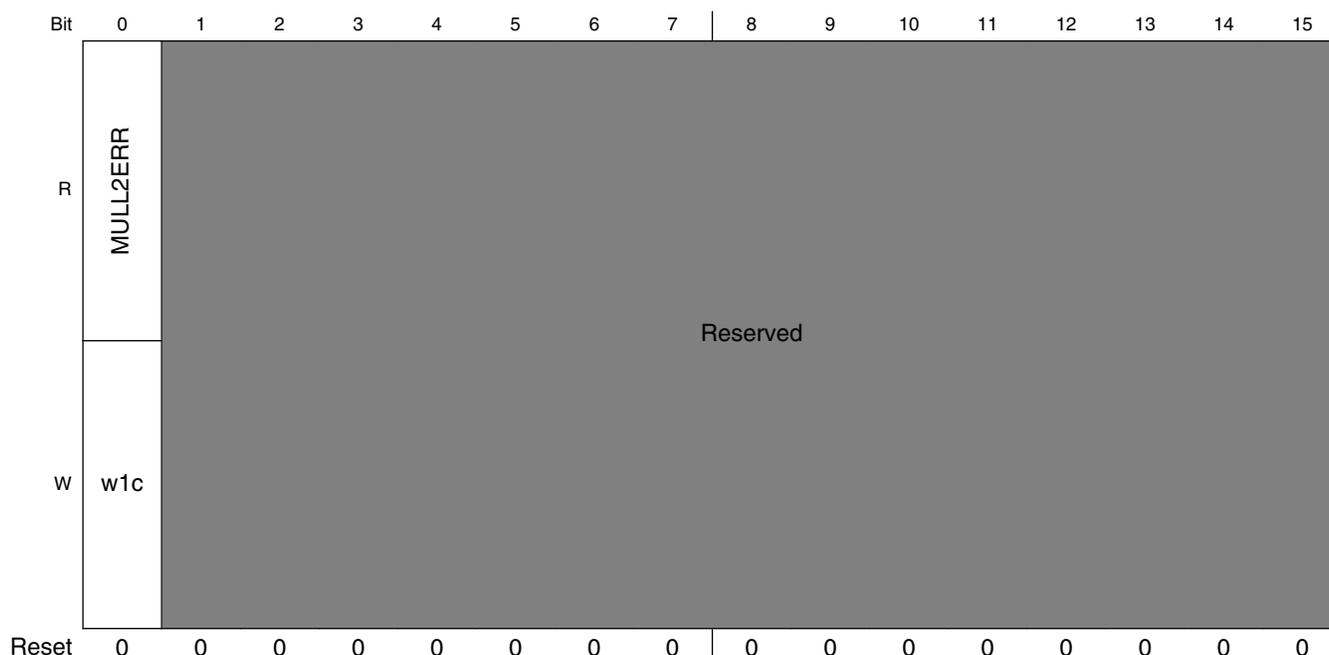
L2_Cache_L2CAPTECC field descriptions (continued)

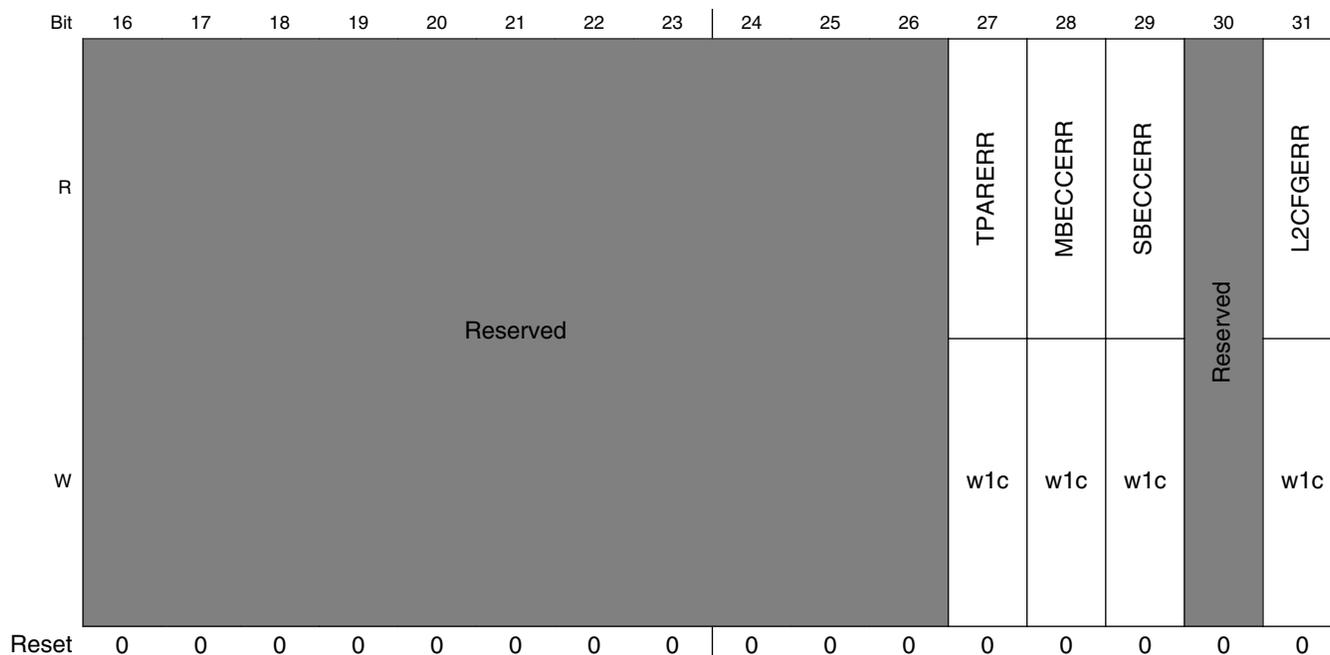
Field	Description
8–23 -	This field is reserved. Reserved
24–31 ECCCHKSUM	The data path ECC of the failing double word

6.7.14 L2 error detect register (L2_Cache_L2ERRDET)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error detect register (L2ERRDET).

Address: 2_0000h base + E40h offset = 2_0E40h





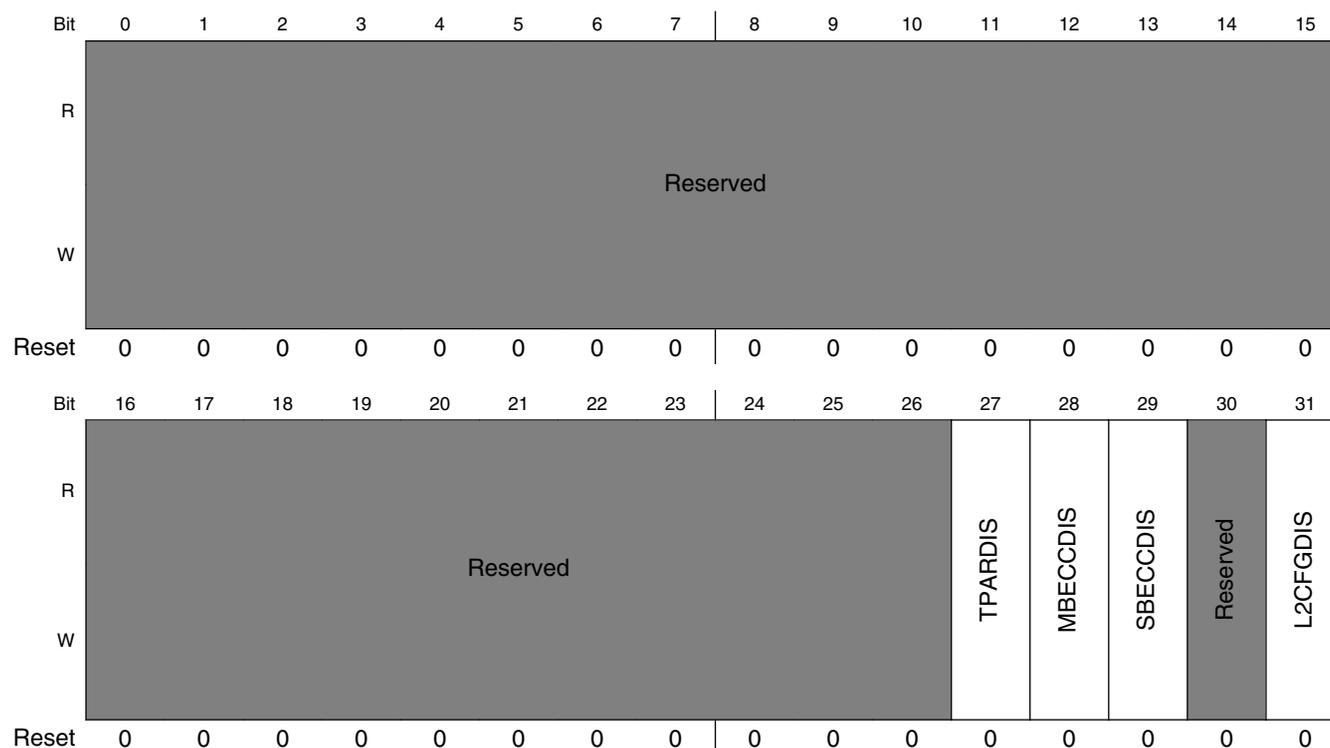
L2_Cache_L2ERRDET field descriptions

Field	Description
0 MULL2ERR	Multiple L2 errors (write 1 to clear) 0 Multiple L2 errors of the same type were not detected 1 Multiple L2 errors of the same type were detected
1–26 -	This field is reserved. Reserved
27 TPARERR	Tag parity error (write 1 to clear) Note that if an L2 cache tag parity error occurs on an attempt to write a new line, the L2 cache must be Flash invalidated. L2 functionality is not guaranteed if Flash invalidation is not performed after a tag parity error. 0 Tag parity error was not detected 1 Tag parity error was detected
28 MBECCERR	Multiple-bit ECC error (write 1 to clear) 0 Multiple-bit ECC errors were not detected 1 Multiple-bit ECC errors were detected
29 SBECCERR	Single-bit ECC error (write 1 to clear) 0 Single-bit ECC error was not detected 1 Single-bit ECC error was detected.
30 -	This field is reserved. Reserved
31 L2CFGERR	L2 configuration error (write 1 to clear) 0 L2 configuration errors were not detected 1 L2 illegal configuration error detected. Reports inconsistencies between the L2SRAM, L2STASHDIS and L2STASHCTL fields of the L2 control register (L2CTL)

6.7.15 L2 error disable register (L2_Cache_L2ERRDIS)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error disable register (L2ERRDIS).

Address: 2_0000h base + E44h offset = 2_0E44h



L2_Cache_L2ERRDIS field descriptions

Field	Description
0–26 -	This field is reserved. Reserved
27 TPARDIS	Tag parity error disable 0 Tag parity error detection enabled 1 Tag parity error detection disabled
28 MBECCDIS	Multiple-bit ECC error disable. Note that non-correctable read errors may cause the assertion of <i>core_fault_in</i> , which causes the core to generate a machine check interrupt, unless it is disabled (by clearing HID1[RFXE]). If RFXE is zero and this error occurs, MBECCDIS must be cleared and L2ERRINTEN[MBECCINTEN] must be set to ensure that an interrupt is generated. NOTE: In normal operation, ECC must be enabled.

Table continues on the next page...

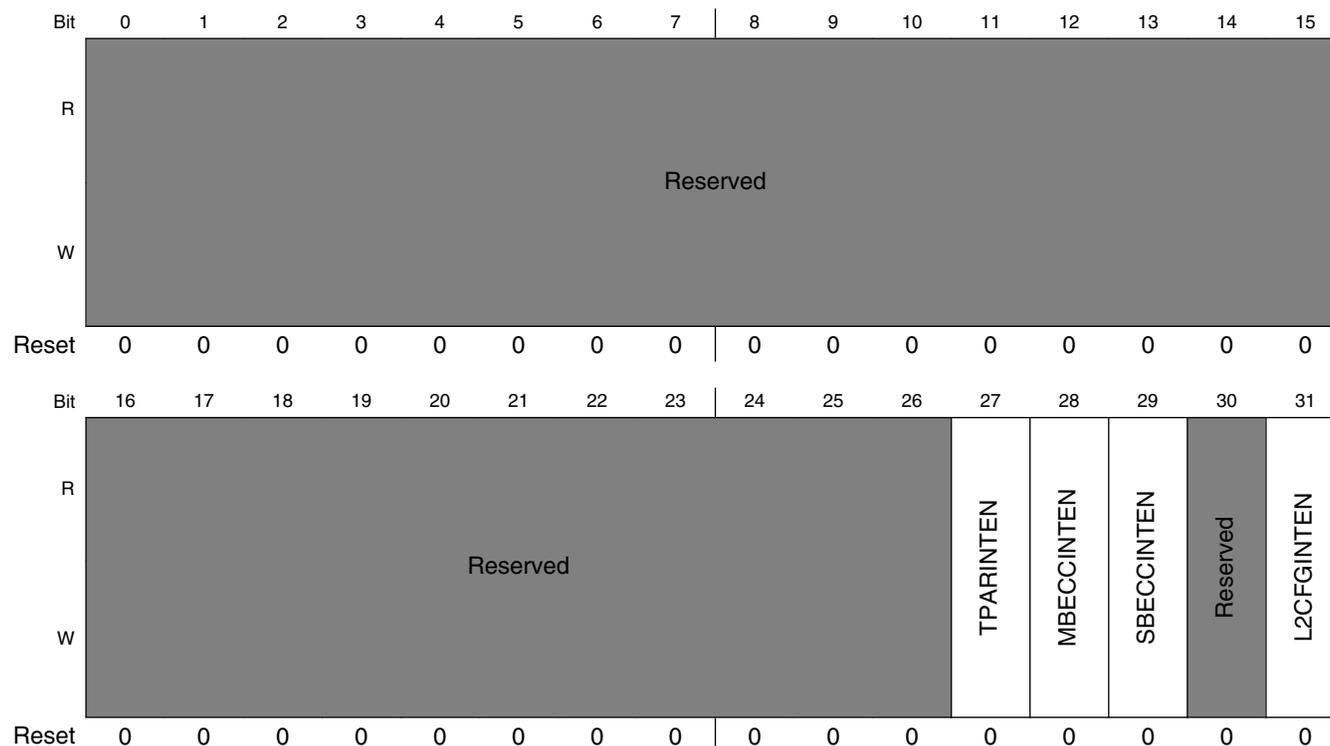
L2_Cache_L2ERRDIS field descriptions (continued)

Field	Description
	0 Multiple-bit ECC error detection enabled 1 Multiple-bit ECC error detection disabled
29 SBECCDIS	Single-bit ECC error disable NOTE: In normal operation , ECC must be enabled. 0 Single-bit ECC error detection enabled 1 Single-bit ECC error detection disabled
30 -	This field is reserved. Reserved
31 L2CFGDIS	L2 configuration error disable 0 L2 configuration error detection enabled 1 L2 configuration error detection disabled

6.7.16 L2 error interrupt enable register (L2_Cache_L2ERRINTEN)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error interrupt enable register (L2ERRINTEN). When an enabled error condition exists, the L2 signals an interrupt to the core through the internal $\overline{\text{int}}$ signal.

Address: 2_0000h base + E48h offset = 2_0E48h



L2_Cache_L2ERRINTEN field descriptions

Field	Description
0–26 -	This field is reserved. Reserved
27 TPARINTEN	Tag parity error reporting enable 0 Tag parity error reporting disabled 1 Tag parity error reporting enabled
28 MBECCINTEN	Multiple-bit ECC error reporting enable. Note that non-correctable read errors may cause the assertion of <i>core_fault_in</i> , which causes the core to generate a machine check interrupt, unless it is disabled (by

Table continues on the next page...

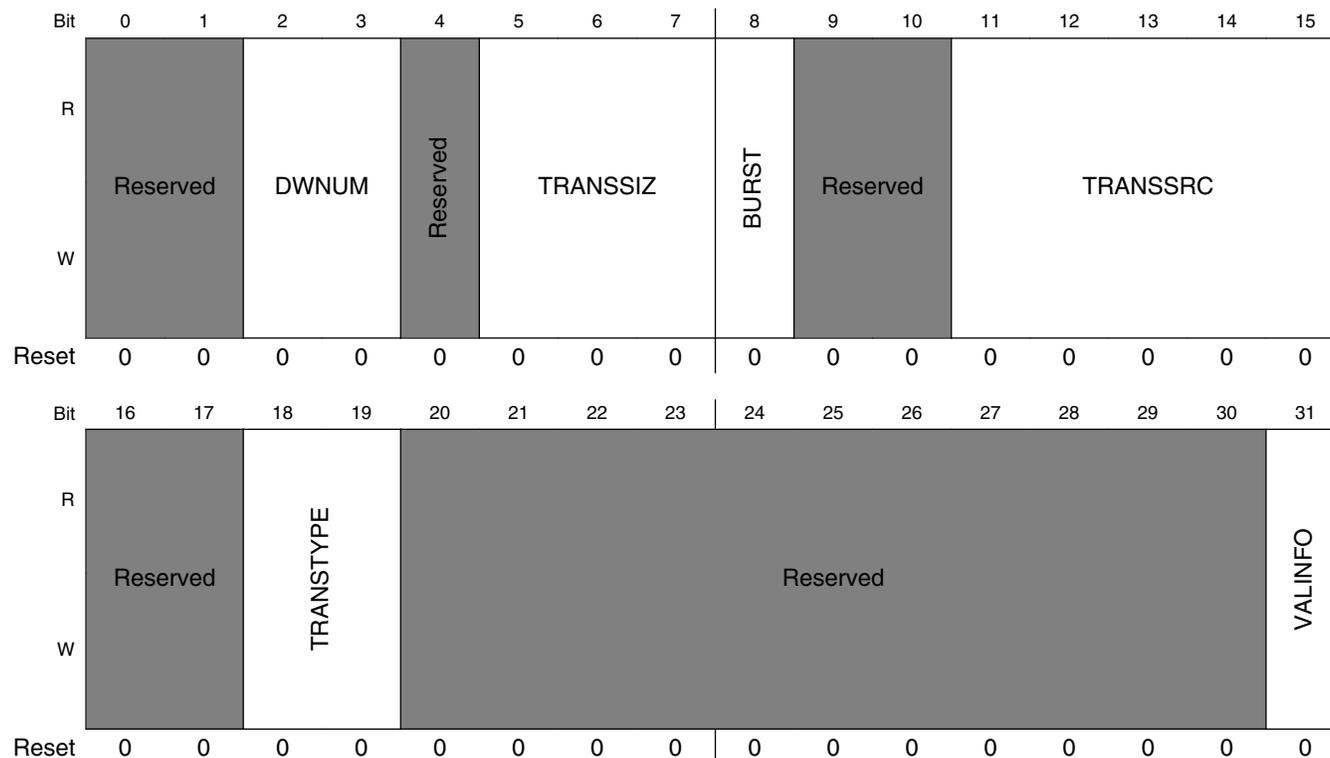
L2_Cache_L2ERRINTEN field descriptions (continued)

Field	Description
	clearing HID1[RFXE]). If RFXE is zero and this error occurs, L2ERRDIS[MBECCDIS] must be cleared and MBECCINTEN must be set to ensure that an interrupt is generated. 0 Multiple-bit ECC error reporting disabled 1 Multiple-bit ECC error reporting enabled
29 SBECCINTEN	Single-bit ECC error reporting enable 0 Single-bit ECC error reporting disabled 1 Single-bit ECC error reporting enabled
30 -	This field is reserved. Reserved
31 L2CFGINTEN	L2 configuration error reporting enable 0 L2 configuration error reporting disabled 1 L2 configuration error reporting enabled

6.7.17 L2 error attributes capture register (L2_Cache_L2ERRATTR)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error attributes capture register (L2ERRATTR).

Address: 2_0000h base + E4Ch offset = 2_0E4Ch



L2_Cache_L2ERRATTR field descriptions

Field	Description
0-1 -	This field is reserved. Reserved
2-3 DWNUM	Double-word number of the detected error (data ECC errors only)
4 -	This field is reserved. Reserved
5-7 TRANSSIZ	Transaction size for detected error Single-beat Burst

Table continues on the next page...

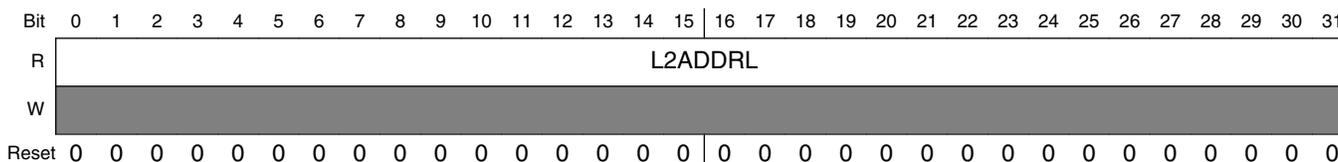
L2_Cache_L2ERRATTR field descriptions (continued)

Field	Description
	000 8 bytes Reserved 001 1 byte 16 bytes 010 2 bytes 32 bytes 011 3 bytes Reserved 100 4 bytes Reserved 101 5 bytes Reserved 110 6 bytes Reserved 111 7 bytes Reserved
8 BURST	Burst transaction for detected error 0 Single-beat (≤ 64 bits) transaction 1 Burst transaction
9–10 -	This field is reserved. Reserved
11–15 TRANSSRC	Transaction source for detected error 00000 External (system logic) 10000 Processor (instruction) 10001 Processor (data)
16–17 -	This field is reserved. Reserved
18–19 TRANSTYPE	Transaction type for detected error 00 Snoop (tag/status read) 01 Write 10 Read 11 Read-modify-write
20–30 -	This field is reserved. Reserved
31 VALINFO	L2 capture registers valid 0 L2 capture registers contain no valid information or no enabled errors were detected. 1 L2 capture registers contain information of the first detected error which has reporting enabled. Software must clear this bit to unfreeze error capture so error detection hardware can overwrite the capture address/data/attributes for a newly detected error.

6.7.18 L2 error address capture register low (L2_Cache_L2ERRADDRL)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error address capture register low (L2ERRADDRL).

Address: 2_0000h base + E50h offset = 2_0E50h



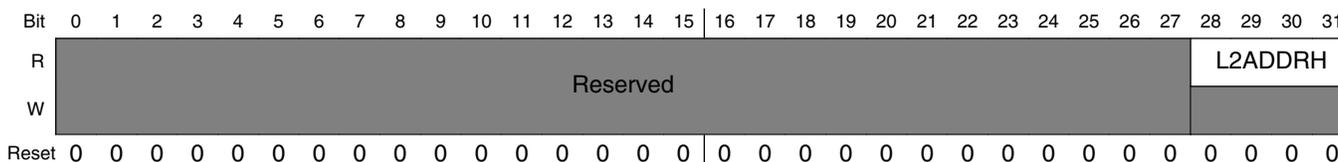
L2_Cache_L2ERRADDRL field descriptions

Field	Description
0–31 L2ADDRL	L2 address bits 4-35 corresponding to detected error

6.7.19 L2 error address capture register high (L2_Cache_L2ERRADDRH)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled (L2CTL[L2E] = 1). The figure below shows the L2 error address capture register high (L2ERRADDRH).

Address: 2_0000h base + E54h offset = 2_0E54h



L2_Cache_L2ERRADDRH field descriptions

Field	Description
0–27 -	This field is reserved. Reserved

Table continues on the next page...

L2_Cache_L2ERRADDRH field descriptions (continued)

Field	Description
28–31 L2ADDRH	L2 address bits 0-3 corresponding to detected error

6.7.20 L2 error control register (L2_Cache_L2ERRCTL)

The error control and capture registers control detection and reporting of tag parity, ECC and L2 configuration errors. L2 configuration errors are illegal combinations of L2 size and block size and are detected when the L2 is enabled ($L2CTL[L2E] = 1$). The figure below shows the L2 error control register (L2ERRCTL).

Address: 2_0000h base + E58h offset = 2_0E58h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved							L2CTHRESH								Reserved							L2CCOUNT									
W	0																															
Reset	0																															

L2_Cache_L2ERRCTL field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–15 L2CTHRESH	L2 cache threshold Threshold value for the number of ECC single-bit errors that are detected before reporting an error condition.
16–23 -	This field is reserved. Reserved
24–31 L2CCOUNT	L2 count Counts ECC single-bit errors detected. If L2CCOUNT equals the ECC single-bit error trigger threshold, an error is reported if single-bit error reporting is enabled.

6.8 External writes to the L2 cache (cache stashing)

Data from an I/O master can be allocated into the L2 cache while simultaneously being written to memory.

External (stashed) writes can be performed from any I/O master, such as the following:

- Ethernet
- RapidIO

- PCI/PCI-Express
- DMA

Stashing is controlled either by an attribute from the initiator of a write or by address range registers in the L2 cache. New cache lines are allocated for full-cache-line writes (unless the line is already resident in the cache). Sub-cache-line write data is stashed only if the line is already valid in the cache. For these sub-cache-line writes, a read-modify-write process is used to merge the write data with the valid data already in the cache.

For information on how to initiate cache stashing from an I/O master, see the respective chapters for the I/O masters that support stashing.

For address range based control of stashing, the L2 cache external write address registers 0-3 (L2CEWAR n) and the L2 cache external write address registers extended address 0-3 (L2CEWAREA n) are used with the L2 cache external write control registers 0-3 (L2CEWCR n) to control the cache stashing functionality. Each register set (for example L2CEWAR0, L2CEWAREA0, and L2CEWCR0) specifies a programmed memory range that can be allocated and optionally locked with a global write transaction. The address register must be naturally aligned to the window size in the corresponding control register. For more information, see [L2 cache external write registers](#).

Note that stashing can occur regardless of whether the L1 cache is enabled or whether the cache-inhibited bit in the MMU is set for the page.

6.8.1 Stash-only cache regions

In order to prevent stashed I/O data from polluting processor data in the L2 cache (and vice versa), it is possible to create stash-only regions. This is controlled by the L2STASHCTL field of L2CTL.

See [L2 control register \(L2_Cache_L2CTL\)](#).

If a stash-only region is created, then that region of the cache is only used for stashed I/O data, and stashed I/O data does not cause the eviction of processor data; they are kept in separate ways of each set. The processor may allocate data into the ways of the cache that are not allocated to SRAM or stash-only memory. Replacement within the stash-only region and the processor region is governed by a pseudo-LRU algorithm modified by masks that allow only applicable ways of a cache set to be considered for replacement.

6.9 L2 cache timing

Table 6-42 shows the timing of back-to-back loads that miss in the L1 data cache and hit in the L2 cache, assuming the core is running at 2 1/2 times the L2 cache frequency.

The L2 returns the 128 bits containing the requested data (critical quad word) first. This data is forwarded to the result register before the full cache line reloads the L1.

Table 6-42. Fastest read timing-hit in L2

Core clocks	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	
e500 core load 1	to D-cache	D-cache miss	to CIU	CIU Q												to CIU	LSU DLB	LSU reads command	LSU reads out data	Result bus							

Table continues on the next page...

Table 6-42. Fastest read timing-hit in L2 (continued)

Core clocks	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	
e500 core load 2		to D-cache	D-cache miss	to CIU	CIU Q																to CIU	LSU	LSU reads command	LSU reads out data	Result bus		
CCB clocks	<1	2	3		4		5		6		7		8		9		10		11		12		13		14		
CCB addr bus load 1		BG ¹	TS ²				AACK ³		HIT		DATA-COMING																
CCB addr bus load 2							BG	TS				AACK	HIT														
CCB data bus load 1															DATA	DATA											
CCB data bus load 2																				DATA	DATA						

1. BG-Bus grant
2. TS-Transfer start
3. AACK-Address acknowledge

6.10 L2 cache and SRAM coherency

This section explains the rules of cache and memory-mapped SRAM coherency.

The term 'snoop transaction' refers to transactions initiated by the system logic or by I/O traffic, as opposed to e500 core-initiated transactions.

6.10.1 L2 cache coherency rules

L2 cache coherency rules are as follows:

- The L2 is non-inclusive of the L1-valid L1 lines may be valid or invalid in the L2.
- The L2 cache holds no modified data. Data is in one of four states-invalid, exclusive, exclusive locked, and stale.
- The L2 allocates entries for data cast out or pushed (non-global, non-write-through write with kill) from the L1 caches.

- Lines for e500 core-initiated burst read transactions are allocated as exclusive in the L2.
- The L2 supports I/O devices reading data from valid lines in the L2 cache (data intervention) if L2CTL[L2INTDIS] = 0. An optional unlock attribute causes I/O reads to clear a lock when the read is performed.
- The L2 cache does not respond to cache-inhibited read transactions.
- e500 core-initiated, cache-inhibited store transactions invalidate the line when they hit on a valid L2 line. If the line is locked, it goes to the stale state. For other write transactions the cache-inhibited bit is ignored.
- Non-burst cacheable write transactions from the e500 core (generated by write-through cacheable stores) update a valid L2 cache line through a read-modify-write operation.
- e500 core cast out transactions that hit on a stale line in the L2 cache cause a data update of the line and a change to the valid locked state for that line.
- An e500 core-initiated, cacheable, non-write-through store that misses in the L1 and hits on a line in the L2 invalidates that line in the L2. If the line is marked exclusive locked, the L2 marks the line as stale.
- Transactions that hit a stale L2 cache line that would cause an allocate if they miss cause a data update of the line (when data arrives from memory) and a change to the line's valid locked state. Data is not supplied by the L2 cache for the read in this case.
- The following transactions kill the data and the respective locks when they hit a valid L2 line:
 - **dcbf**
 - **dcbi**
- The L2 cache supports mixed cache external writes and core-initiated writes to the same addresses if the core-initiated writes are marked coherency-required, caching allowed, not write-through (WIMG = 001x) and the external writes are marked coherency-required, caching-allowed.
- The L2 cache supports writes to the L2 cache from peripheral devices or from I/O controllers through snoop write transactions with addresses that hit in a programmed memory range. Full cache line (32-byte) write transactions update the data for a valid line in the L2 and if the line is not valid in the L2, a line is allocated. Sub-cache line write transactions update the data only for valid L2 cache lines through read-modify-write operations.
- The L2 cache supports burst writes that lock an L2 cache line from peripheral devices or from I/O controllers through write transactions with addresses that hit in a programmed memory range that has the lock attribute set.
- The L2 cache supports burst writes that allocate and/or lock an L2 cache line from peripheral devices or I/O controllers through a write allocate transaction. See the system logic programming model (for example, that of the DMA controller) for details on how to set the transaction type for cache external writes to the L2.

6.10.2 Memory-mapped SRAM coherency rules

Memory-mapped SRAM coherency rules are as follows:

- External (non-core-initiated) accesses to memory-mapped SRAM must be marked coherency-required. External accesses to memory-mapped SRAM marked coherency-not-required may cause an address unavailable error.
- Accesses to memory-mapped SRAM are cacheable only in the corresponding e500 L1 caches. External accesses must be marked cache-inhibited or be performed with non-caching transactions.

6.11 L2 cache locking

The caches can be locked and cleared using the following methods:

- Cache locking methods
 - Individual line locks are set and cleared using instructions defined by the e500 cache locking APU, which is part of the Freescale Embedded Implementation Standards (EIS). These instructions include Data Cache Block Touch and Lock Set (**dcbtls**), Data Cache Block Touch for Store and Lock Set (**dcbtstls**), and Instruction Cache Block Touch and Lock Set (**icbtls**). For detailed information about these instructions, see the *PowerPC e500 Core Reference Manual*.
 - A lock attribute can be attached to write operations.
 - Individual line locks are set and cleared through core-initiated instructions, by external reads or writes, or by accesses to programmed memory ranges defined in L2 cache external write address registers (L2CEWAR_n).
 - The entire cache can be locked by setting configuration registers appropriately
- Methods for clearing locks
 - Individual locks can be cleared by cache locking APU instructions (Instruction Cache Block Lock Clear (**icble**) and Data Cache Block Lock Clear (**dcble**)) or by snoop flush unless the entire cache is locked.
 - Flash clearing of all instruction and/or data locks can be done by writes to configuration registers.
 - An unlock attribute can be attached to I/O read operations.

6.11.1 Locking the entire L2 cache

The entire L2 cache can be locked by setting $L2CTL[L2DO] = 1$ and $L2CTL[L2IO] = 1$. This has the effect of preventing any further allocation of new lines in the cache by core requests.

If there are lines in the cache that are not valid, they cannot be used by core requests until the cache is unlocked. While the cache is locked, read requests are serviced as normal, and snooping continues as normal to maintain coherency. Lines invalidated to satisfy coherency requirements cannot be reallocated by core requests while the cache remains locked. The L2 cache can be unlocked by clearing $L2CTL[L2IO]$ and/or $L2CTL[L2DO]$. Note that $L2CTL[L2DO]$ and $L2CTL[L2IO]$ have no effect on cache external write allocations or memory-mapped SRAM.

Note that this form of cache locking does not use the lock bits of the cache and cannot be cleared by resetting the cache or lock bits.

6.11.2 Locking programmed memory ranges

A programmed memory range can be locked with a snoop write transaction that matches a cache external write address range (specified by $L2CEWAR_n/L2CEWAREA_n$ and $L2CEWCR_n$).

There are no clearing of locks through the programmed address ranges. Locks can be cleared using clear lock instructions, flushes, read-and-clear-lock snoop (RWNITC with clear lock attribute), or flash clear locks.

6.11.3 Locking selected lines

Individual lines are locked when the L2 receives one of the following burst transactions:

- **icbtls** (CT = 1)-Instruction Cache Block Touch and Lock Set instruction
- **dcbtls** (CT = 1)-Data Cache Block Touch and Lock Set instruction
- **dcbtstls** (CT = 1)-Data Cache Block Touch for Store and Lock Set instruction
- Snoop burst write-If the address hits on a programmed cache external write space with the lock attribute set, or if the write allocate transaction type is used
- Snoop non-burst write-If the address hits on a programmed cache external write space with the lock attribute set

Note that the core complex broadcasts these instructions to the L2 if the CT field in the instruction specifies the L2 cache (CT = 1). When the L2 cache is specified, data is not placed in the L1, only the L2. If the L1 cache is specified (CT = 0), the L2 does not lock the line, and the data is placed in the L1 (and locked).

When the touch lock set L2 instruction (**dcbtls** or **dcbtstls**) hits are modified in the L1 cache, the modified data is allocated into the L2 cache (and written back to main memory) and a data lock is set. The L1 line state transitions to invalid.

Note that if the L2 receives a request to allocate and lock a line, but all lines in the selected way are locked, the requested L2 line is not allocated and the L2 cache lock overflow bit (L2CTL[L2LO]) is set.

Lines invalidated to satisfy coherency requirements cannot be reallocated while the cache remains locked.

6.11.4 Clearing locks on selected lines

Individual locks in the L2 are cleared by a lock clear (**icblc** or **dcblc**, CT = 1) instruction.

This directs the L2 cache to clear a lock on that line if it hits in the L2 cache. Both data and instruction locks are cleared by the **icblc** and **dcblc** instructions.

Note that the lock on a line is cleared if the line is invalidated by a snooped flush transaction, and the line in the cache is available for allocation of a new line of instruction or data unless the entire cache is locked.

6.11.5 Flash clearing of instruction and data locks

Locks for instructions and data are recorded separately in the L2 cache, and they can be flash cleared separately by writing the appropriate value to the L2 cache control register (L2CTL[L2LFR] and L2CTL[L2LFRID]).

Flash invalidating of the L2 (setting L2CTL[L2I]) clears all locks on both instructions and data.

Note that flash clearing is the only way to clear data locks without clearing instruction locks, or to clear instruction locks without clearing data locks. All instructions and snoop transactions that clear locks clear both data and instruction locks.

6.11.6 Locks with stale data

If data is locked in the L2 and either the e500 core performs a cacheable copyback store or a **dcbtst** misses in the L1, the L2 invalidates the line; however, the L2 clears the valid bit for the data, the lock remains, and the line cannot be victimized.

If the e500 core casts out modified data or pushes it in response to a non-flush snoop, the L2 updates the data and sets the valid bit again, maintaining the lock and keeping the data in the cache hierarchy.

6.12 PLRU L2 replacement policy

Line replacement is determined using a pseudo least-recently-used (PLRU) algorithm.

There is a valid bit (V0-V7) for each line. To determine the replacement victim (the line to be cast out), there are seven PLRU bits (P0-P6) for each set. PLRU bits are updated every time a new line is allocated and every time an existing line is read by the processor, updated by a write, or invalidated.

Figure 6-43 shows the binary decision tree used to generate the victim line. The eight ways of the L2 cache are labeled W0-W7; the seven PLRU bits are labeled P0-P6.

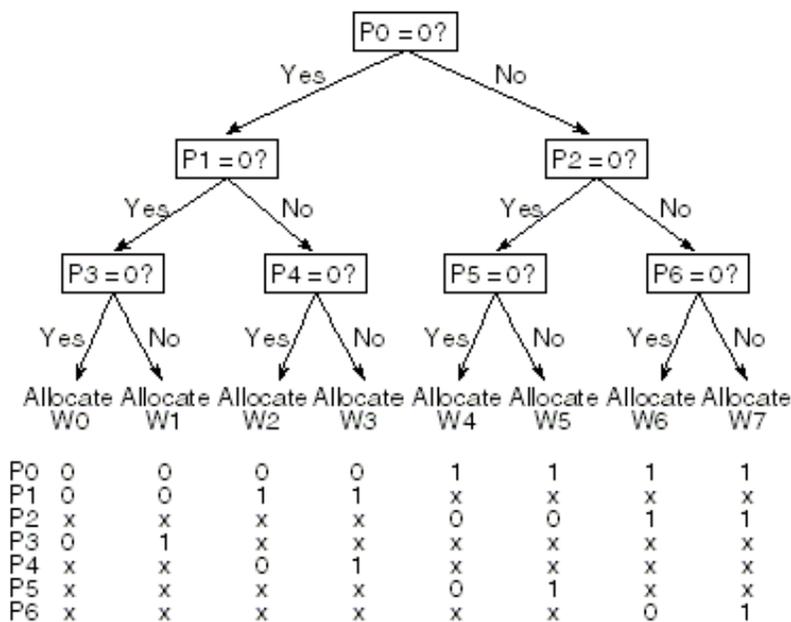


Figure 6-43. L2 cache line replacement algorithm

6.12.1 PLRU bit update considerations

PLRU bit updates depend on which cache way was last accessed, as summarized in [Table 6-43](#).

Table 6-43. PLRU bit update algorithm

Last Way Accessed	PLRU Bits						
	P0	P1	P2	P3	P4	P5	P6
0	1	1	-	1	-	-	-
1	1	1	-	0	-	-	-
2	1	0	-	-	1	-	-
3	1	0	-	-	0	-	-
4	0	-	1	-	-	1	-
5	0	-	1	-	-	0	-
6	0	-	0	-	-	-	1
7	0	-	0	-	-	-	0

When an L2 line is invalidated, the PLRU bits are updated, marking the corresponding way as least-recently used. This causes the invalidated way to be selected as the next victim.

6.12.2 Allocation of lines

The general PLRU algorithm described above must be modified to take into account special features of the L2 cache; namely SRAM regions, line locking, and stash-only regions.

Each of these features reserves ways within each cache set such that some ways are not eligible for allocation/victimization by the general LRU algorithm.

To preserve the state of the ways that are set aside for other special functions, the PLRU pointers are modified by a mask that is a function of the L2 configuration registers, the lock bits in the cache status array, and initiator of the transaction. The mask effectively points the PLRU algorithm away from ways that are not to be considered for replacement.

L2 cache lines are locked through the status array lock bits. There are two lock bits for each way of each set (2048 sets by eight ways). These bits are set or cleared through special L2 controller commands. There are two sets of lock bits, one for instructions (I0-I7) and one for data (D0-D7) for every line. The lock bits act as a mask over the PLRU bits to determine victim selection. The PLRU bits are updated regardless of line locking.

Lock bits are used at allocate time to steer the PLRU algorithm away from selecting locked victims. In the following discussion, the eight lock bits for a particular set are called L0-L7.

Where Lock Way i : $L_i = D_i \mid I_i$, $i=0\dots7$ (D_i = data lock, I_i = instruction lock)

An effective value of each PLRU bit is calculated as follows:

$$\begin{aligned}
 P0_eff &= f(P0, L0, L1, L2, L3, L4, L5, L6, L7) = (L0 \ \& \ L1 \ \& \ L2 \ \& \ L3) \mid (P0 \ \& \ \sim(L4 \ \& \ L5 \ \& \ L6 \ \& \ L7)) \\
 P1_eff &= f(P1, L0, L1, L2, L3) = (L0 \ \& \ L1) \mid (P1 \ \& \ \sim(L2 \ \& \ L3)) \\
 P2_eff &= f(P2, L4, L5, L6, L7) = (L4 \ \& \ L5) \mid (P2 \ \& \ \sim(L6 \ \& \ L7)) \\
 P3_eff &= f(P3, L0, L1) = L0 \mid (P3 \ \& \ \sim L1) \\
 P4_eff &= f(P4, L2, L3) = L2 \mid (P4 \ \& \ \sim L3) \\
 P5_eff &= f(P5, L4, L5) = L4 \mid (P5 \ \& \ \sim L5) \\
 P6_eff &= f(P6, L6, L7) = L6 \mid (P6 \ \& \ \sim L7)
 \end{aligned}$$

These effective PLRU bits are used to select a victim, as indicated in [Table 6-44](#).

Table 6-44. PLRU-based victim selection mechanism

Way selected	Effective PLRU state (binary)	Reduced logic equation (using effective PLRU bits)
W0	00x0xxx	$\sim P0 \ \& \ \sim P1 \ \& \ \sim P3$
W1	00x1xxx	$\sim P0 \ \& \ \sim P1 \ \& \ P3$
W2	01xx0xx	$\sim P0 \ \& \ P1 \ \& \ \sim P4$
W3	01xx1xx	$\sim P0 \ \& \ P1 \ \& \ P4$
W4	1x0xx0x	$P0 \ \& \ \sim P2 \ \& \ \sim P5$
W5	1x0xx1x	$P0 \ \& \ \sim P2 \ \& \ P5$
W6	1x1xxx0	$P0 \ \& \ P2 \ \& \ \sim P6$
W7	1x1xxx1	$P0 \ \& \ P2 \ \& \ P6$

6.13 L2 cache operation

This section describes the behavior of the L1 and L2 cache in response to various operations and configurations.

6.13.1 Initialization

This section describes L2 cache initialization and memory-mapped SRAM initialization.

6.13.1.1 L2 cache initialization

After power-on reset the valid bits in the L2 cache status array are in random states. Therefore, it is necessary to perform a flash invalidate before using the array as an L2 cache.

This is done by writing a one to the L2I field of the L2 control register (L2CTL). This can be done before or simultaneously with the write that enables the L2 cache. That is, the L2E and L2I bits of L2CTL can be set simultaneously. The L2I bit clears automatically, so no further writes are necessary.

6.13.1.2 Memory-mapped SRAM initialization

After power-on reset the contents of the data and ECC arrays are random, so all SRAM data must be initialized before it is read.

If the cache is initialized by the processor or any other device that uses sub-cache-line transactions, ECC error checking should be disabled during the initialization process to avoid false ECC errors generated during the read-modify-write process used for sub-cache-line writes to the SRAM array. This is done by setting the multi- and single-bit ECC error disable bits of the L2 error disable register (L2ERRDIS[MBECCDIS, SBECCDIS]). See [L2 error disable register \(L2_Cache_L2ERRDIS\)](#). If the array is initialized by a DMA engine using cache-line writes, ECC checking can remain enabled during the initialization process.

6.13.2 Flash invalidation of the L2 cache

The L2 cache may be completely invalidated by setting the L2I bit of the L2 control register (L2CTL).

Note that no data is lost in this process because the L2 cache is a write-through cache and contains no modified data. Flash invalidation of the cache is necessary when the cache is initially enabled and may be necessary to recover from some error conditions such as a tag parity error.

The invalidation process requires several cycles to complete. The L2I bit remains set during this procedure and is then cleared automatically when the procedure is complete. The L2 cache controller issues retries for all transactions on the e500 core complex bus while the flash invalidation process is in progress.

Note that the contents of memory-mapped SRAM regions of the data array are unaffected by a flash invalidation of the L2 cache regions of the array.

6.13.3 Managing errors

This section describes ECC and tag parity errors.

6.13.3.1 ECC errors

An individual soft error that causes a single- or multi-bit ECC error can be cleared from the L2 array simply by performing a **dcbf** instruction on the address captured in the L2ERRADDR register. This invalidates the line in the L2 cache.

When the load that caused the ECC error is performed again, the data is reallocated into the L2 with ECC bits set properly again.

If the threshold for single bit errors set in the L2ERRCTL register is exceeded, then the L2 cache should be flash invalidated to clear out all single-bit errors.

Note that no data is lost by **dcbf**s or flash invalidates, since the L2 cache is write-through and contains no modified data.

6.13.3.2 Tag parity errors

A tag parity error must be fixed by flash invalidating the L2 cache.

Note that a **dcbf** operation to the address that caused the error to be reported is not sufficient since a tag parity error is seen as an L2 miss and does not cause invalidation of the bad tag. Proper L2 operation cannot be guaranteed if an L2 tag parity error is not repaired by a flash invalidation of the entire cache.

6.13.4 L2 cache states

The L2 status array uses four bits for each line to determine the status of the line.

Different combinations of these bits result in different L2 states. The status bits are as follows:

- Valid (V)
- Instruction locked (IL)
- Data locked (DL)
- Stale (T)

Table 6-45 shows L2 cache states. Note that these conventions are also used in Table 6-46.

Table 6-45. L2 cache states

V	T	IL	DL	L2 states
0	x	x	x	Invalid (I)
1	0	0	0	Exclusive (E)
1	0	0	1	Exclusive data locked (EDL)
1	0	1	0	Exclusive instruction locked (EIL)
1	0	1	1	Exclusive instruction and data locked (EL)
1	1	0	0	Stale (data invalid, locks invalid) (T)
1	1	0	1	Stale (data invalid, dlock valid) (TDL)
1	1	1	0	Stale (data invalid, ilock valid) (TIL)
1	1	1	1	Stale (data invalid, locks valid) (TL)

6.13.5 L2 state transitions

Table 6-46 lists state transitions for all e500 core-initiated transactions that change the L2 cache state.

Core-initiated transactions caused when the core executes **msync**, **mbar**, **tlbivax**, or **tlbsync** do not change the L2 cache state.

The table does not list initial L1 states for transactions that hit in the L1 (iL1 or dL1) and are not sent to the L2.

In the table, the heading 'L2 hit' indicates that the L2 provides (on a read) or captures (on a write) data for an existing line. Some entries list two final L1 states. L2 touch instructions never allocate into iL1 or dL1.

Note that if the L2 SRAM is disabled, the L2 initial and final states are always invalid and the L2 never hits. Similarly, if the L2 SRAM is in full memory-mapped SRAM mode, the L2 initial and final states are always invalid and the L2 never hits for addresses not in the memory-mapped SRAM address range. The L2 always hits for addresses in the enabled memory-mapped SRAM address ranges.

Table 6-46. State transitions due to core-initiated transactions

Source of transaction	Initial states		L2 Hit	Final states		Comments
	L1	L2		L1	L2	
Cacheable instruction fetch icbtl_L1	iL1 I	I/T	No	I/V	Same	L2CTL[L2DO] = 1. L2 touch instructions not allocated in L1
		I	No	I/V	E	L2CTL[L2DO] = 0
icbt_L2	dL1 I,E	E/EL	Yes	I/V	Same	-
		T	No	I/V	EL	L2CTL[L2DO] = 0. Restore locked line in L2 with valid data from bus
icbtl_L2	dL1 I,E	I/T	No	I	Same	L2CTL[L2DO] = 1
		E	Yes	I	I	L2CTL[L2DO] = 1
		EL	Yes	I	T	L2CTL[L2DO] = 1
		I	No	I	EL	L2CTL[L2DO] = 0
		E	Yes	I	EL	L2CTL[L2DO] = 0
		EL	Yes	I	Same	L2CTL[L2DO] = 0
		T	No	I	EL	L2CTL[L2DO] = 0. Restore locked line in L2 with valid data from bus
Cache-inhibited instruction fetch	N/A	N/A	No	N/A	N/A	No L1/L2 effect
Cacheable load (4-state) Cacheable lwarx (4-state) dcbt_L1 (4-state) dcbtl_L1 (4-state)	dL1 I	I/T	No	E	Same	L2CTL[L2IO] = 1
		E	Yes	E	I	L2CTL[L2IO] = 1
		EL	Yes	E	T	L2CTL[L2IO] = 1
		I	No	E	E	L2CTL[L2IO] = 0
		E/EL	Yes	E	Same	L2CTL[L2IO] = 0
		T	No	EL	EL	L2CTL[L2IO] = 0. Restore locked line in L2 with valid data from bus
Cache-inhibited load	N/A	N/A	No	N/A	N/A	No L1/L2 effect
Cache-inhibited lwarx	N/A	N/A	No	N/A	N/A	No L2 effect
Writeback Store	dL1 I	I/T	No	M	Same	L2 allocates when a line is cast out of L1.
		E	Yes	M	I	-
		EL	Yes	M	T	-
Writeback stwcx	dL1 I	I/T	No	M	Same	-
		E	Yes	M	I	-
		EL	Yes	M	T	-

Table continues on the next page...

Table 6-46. State transitions due to core-initiated transactions (continued)

Source of transaction	Initial states		L2 Hit	Final states		Comments
	L1	L2		L1	L2	
Cacheable load (3-state) Cacheable lwarx (3-state) dcbt_L1 (3-state) dcbtIs_L1 (3-state)	dL1 I	I	No	E/I	Same	L2CTL[L2IO] = 1
		T	No	E/I	Same	L2CTL[L2IO] = 1
		E	Yes	E/I	I	L2CTL[L2IO] = 1
		EL	Yes	E/I	T	L2CTL[L2IO] = 1
		I	No	E/I	E	L2CTL[L2IO] = 0
dcbt_L2 dcbtst_L2	dL1 I,E	E/EL	Yes	E/I	Same	L2CTL[L2IO] = 0
		T	No	E/I	EL	L2CTL[L2IO] = 0. Restore locked line with valid data from bus
dcbtst_L1 dcbtstIs_L1	dL1 I	I/T	No	E	Same	-
		E	Yes	E	I	-
		EL	Yes	E	T	-
dcbtIs_L2 dcbtstIs_L2	dL1 I,E	I	No	I	Same	L2CTL[L2IO] = 1
		T	No	I	Same	L2CTL[L2IO] = 1
		E	Yes	I	I	L2CTL[L2IO] = 1
		EL	Yes	I	T	L2CTL[L2IO] = 1
		I	No	I	EL	L2CTL[L2IO] = 0
		E/EL	Yes	I	EL	L2CTL[L2IO] = 0
		T	No	I	EL	L2CTL[L2IO] = 0. Restore locked line with valid data from bus
Write-through store	dL1 I,E,M	I/T	No	same	I	-
		E/EL	Yes	same	Same	Read-modify-write
Cache-inhibited store	N/A	I/E	No	N/A	I	Invalidate line
		EL/T	No	N/A	T	Invalidate data, keep lock
Cache-inhibited stwcx	N/A	I/E	No	N/A	I	Invalidate line
		EL/T	No	N/A	T	Invalidate data, keep lock
dcbLc_L2 icbLc_L2	dL1 I,E,M	I/E	No	same	Same	-
		EL	No	same	E	-
		T	No	same	I	-

Table continues on the next page...

Table 6-46. State transitions due to core-initiated transactions (continued)

Source of transaction	Initial states		L2 Hit	Final states		Comments
	L1	L2		L1	L2	
Victim castout dcbt_L2 icbt_L2 dcbtst_L2	dL1 M	I/T	No	I	Same	L2CTL[L2IO] = 1. If software sharing cache lines between instructions and data wishes to capture instruction lines in L2 with L2CTL[L2IO] = 1, it must perform dcbst to flush the line out of the dL1 before fetching it into L2.
		I	No	I	E	L2CTL[L2IO] = 0
		E/EL	No	I	I/T	L2CTL[L2IO] = 1.
			Yes	I	Same	L2CTL[L2IO] = 0.
		T	Yes	I	EL	L2CTL[L2IO] = 0.
dcbtIs_L2 icbtIs_L2 dcbtstIs_L2	dL1 M	I	No	I	EL	An icbtIs_L2 that hits modified in L1 cannot be distinguished from dcbtIs_L2 and sets the L2 dlock bit. If software shares cache lines between instructions and data and wishes to set locks in L2, it must perform dcbst to flush the line out of the dL1 before locking it in L2.
		E/EL/T	Yes	I	EL	-
Snoop push	dL1 M	I/E	No	I/E	I	-
		EL/T	No	I/E	T	Invalidate data, keep lock
dcbf dcbst	dL1 M	I/E/EL	No	I	I	-
dcbz dcba	dL1 I	I/E	No	M	I	-
		EL	No	M	T	-
dcbi	dL1 I,E,M	I/ E/EL/T	No	I	I	-
dcbf dcbst	dL1 I,E	I/ E/EL/T	No	I	I	-
icbi	iL1 I,V	I/ E/EL/T	No	I	I	-

Table 6-47 lists L2 cache state transitions for all system-initiated (non-core) transactions that change the L2. The transaction types and attributes listed follow MPX bus nomenclature, with the addition of write allocate (burst write with L2 cache allocation). Table 6-47 accounts for changes caused by L1 snoop pushes triggered by snoops, listed in Table 6-46.

Table 6-47. State transitions due to system-initiated transactions

Transaction type	\overline{wt}	\overline{ci}	\overline{gbl}	Initial L2 state	Final L2 state	Comments				
Clean IKill	x	x	0	I/E/EL/T	Same	-				
Flush	x	x	0	I/E/EL/T	I	-				
Write allocate	x	1	0	I/E/EL/T	EL	Allocate and lock regardless of cache external write (CEW) window				
				I/E	E	Allocate regardless of CEW window				
				EL/T	EL	-				
	x	0	0	I/E	I	No allocate if cache-inhibited				
				EL/T	T	Invalidate data, keep lock				
WWK 32-byte WWF 32-byte WWF atomic	x	1	0	I/E/EL/T	I	Miss in cache external write windows				
				I	E/EL	Hit in cache external write window				
				EL	Same	Hit in cache external write window				
				E	E/EL	Hit in cache external write window				
				T	EL	Hit in cache external write window				
				I/E	I	Invalidate line				
				EL/T	T	Invalidate data, keep lock				
				< 32-byte WWF < 32-byte WWF atomic	x	1	0	I/E	I	Miss in cache external write windows
								EL/T	T	Miss in cache external write windows.
								I/T	Same	Hit in CEW window but need burst data
EL	Same	Hit in cache external write window								
E	E/EL	Hit in cache external write window. Set lock if CEW lock attribute set.								
I/E	I	Invalidate line								
Read Read atomic	1	1	0	I/T	Same	-				
				E	Same	-				
				EL	Same	-				
				x	0	0	N/A	N/A	No L1/L2 effect	
				RWNITC	1	1	0	I/T	Same	-
								E	Same	-
EL	Same	-								
0	1	0	I					Same	Read-and-clear-lock	
				EL	E	-				
				T	I	-				
				x	0	0	N/A	N/A	No L1/L2 effect	

Table continues on the next page...

Table 6-47. State transitions due to system-initiated transactions (continued)

Transaction type	\overline{wt}	\overline{ci}	\overline{gbl}	Initial L2 state	Final L2 state	Comments
Kill RWITM RWITM atomic RClaim	x	1	0	I/E	I	-
				EL/T	T	Invalidate data, keep lock
	x	0	0	I/E/EL/T	Same	-

6.14 Error checking and correcting (ECC)

The L2 cache supports error checking and correcting (ECC) for the data path between the core master and system memory.

It detects all double-bit errors, detects all multi-bit errors within a nibble, and corrects all single-bit errors. Other errors may be detected, but are not guaranteed to be corrected or detected. Multiple-bit errors are always reported when error reporting is enabled. When a single-bit error occurs, the single-bit error counter register is incremented, and its value compared to the single-bit error trigger register. An error is reported when these values are equal. The single-bit error registers can be programmed such that minor memory faults are corrected and ignored, but double- or multi-bit errors generate an interrupt.

The syndrome encodings for the ECC code are shown in [Table 6-48](#) and [Table 6-49](#).

Table 6-48. L2 cache ECC syndrome encoding

Data bit	Syndrome bit								Data bit	Syndrome bit							
	0	1	2	3	4	5	6	7		0	1	2	3	4	5	6	7
0	•	•						•	32			•	•				•
1	•		•					•	33			•		•			•
2	•			•				•	34	•		•		•			
3	•				•			•	35		•	•		•			
4	•	•				•			36			•	•		•		
5	•		•			•			37			•		•	•		
6	•			•		•			38	•		•		•	•		•
7	•				•	•			39		•	•		•	•		•
8	•	•						•	40			•	•			•	
9	•		•					•	41			•		•		•	
10	•			•				•	42	•		•		•		•	•

Table continues on the next page...

Table 6-48. L2 cache ECC syndrome encoding (continued)

Data bit	Syndrome bit								Data bit	Syndrome bit							
	0	1	2	3	4	5	6	7		0	1	2	3	4	5	6	7
11	•				•		•		43		•	•		•		•	•
12	•	•				•	•	•	44			•	•		•	•	•
13	•		•			•	•	•	45			•		•	•	•	•
14	•			•		•	•	•	46	•		•		•	•	•	
15	•				•	•	•	•	47		•	•		•	•	•	
16		•	•					•	48		•				•	•	
17		•		•				•	49			•			•	•	
18		•			•			•	50				•		•	•	
19	•	•			•				51	•					•	•	
20		•	•			•			52		•				•		•
21		•		•		•			53			•			•		•
22		•			•	•			54				•		•		•
23	•	•			•	•		•	55	•					•		•
24		•	•				•		56		•					•	•
25		•		•			•		57			•				•	•
26		•			•		•		58				•			•	•
27	•	•			•		•	•	59	•						•	•
28		•	•			•	•	•	60				•	•		•	
29		•		•		•	•	•	61	•			•	•		•	•
30		•			•	•	•	•	62		•		•	•		•	•
31	•	•			•	•	•		63			•	•	•		•	•

Table 6-49. L2 cache ECC syndrome encoding (check bits)

Check bit	Syndrome bit							
	0	1	2	3	4	5	6	7
0	•							
1		•						
2			•					
3				•				
4					•			
5						•		
6							•	
7								•

Chapter 7

e500 Coherency Module

The e500 coherency module (ECM) provides a flexible switching structure for routing e500- and I/O-initiated transactions to target modules on the device.

7.1 Introduction

The e500 coherency module (ECM) provides a flexible switching structure for routing e500- and I/O-initiated transactions to target modules on the device.

Figure 7-1 shows a high-level block diagram of the ECM.

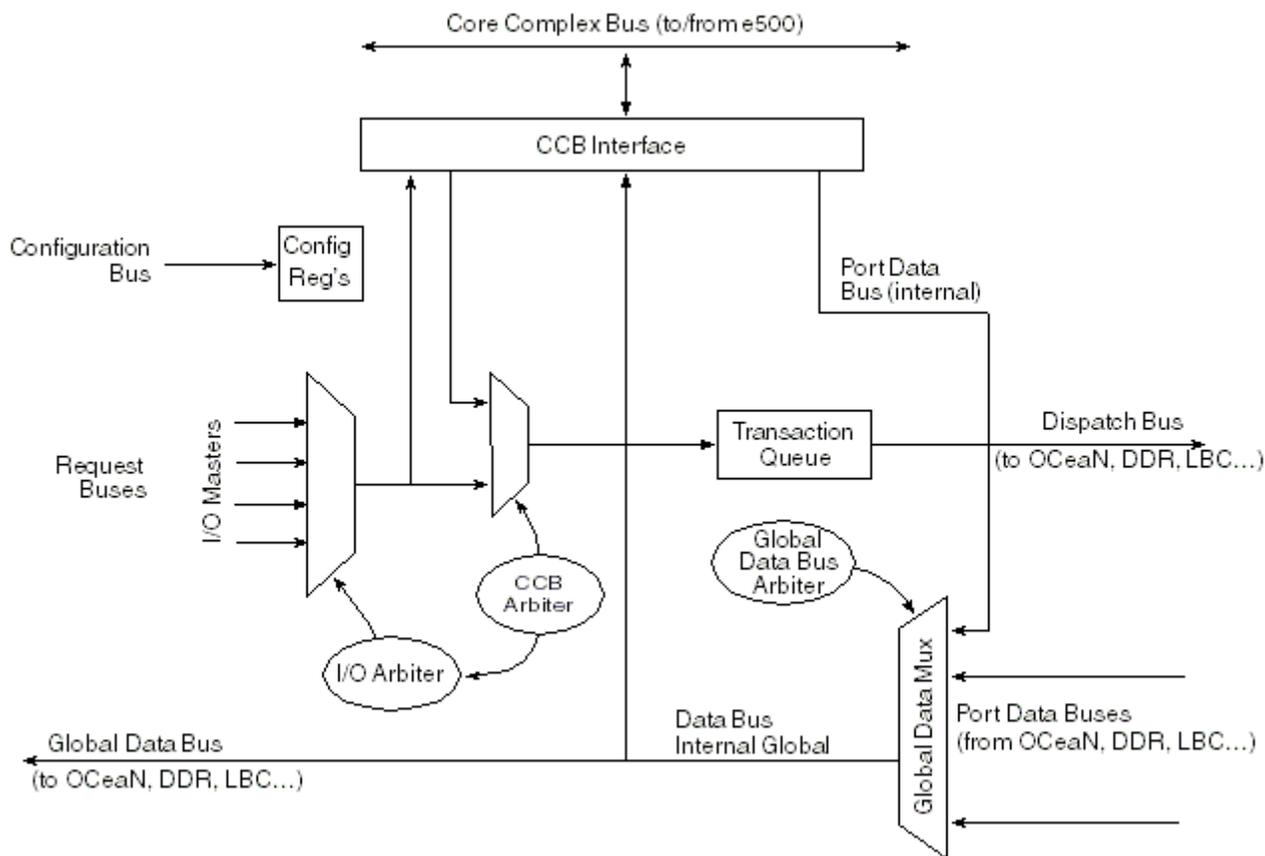


Figure 7-1. e500 coherency module block diagram

7.1.1 Overview

The ECM routes transactions initiated by the e500 cores to the appropriate target interface on the device.

In a manner analogous to a bridging router in a local area network, the ECM forwards I/O-initiated transactions that are tagged with the global attribute onto the core complex bus (CCB). This allows on-chip caches to snoop these transactions as if they were locally initiated and to take actions to maintain coherency across cacheable memory.

7.1.2 Features

The ECM includes these distinctive features:

- Support for two e500 cores and an L2/SRAM on the CCB, including a CCB arbiter.
- Sources a 64-bit data bus for returning read data from the ECM to the e500 cores and routing write data from the ECM to the L2/SRAM. It sinks a 128-bit data bus for receiving data from the L2/SRAM and two 128-bit write data buses—one from each of the e500 cores.
- Four connection points for I/O initiating (mastering into the device) interfaces. The ECM supports five connection points for I/O targets. The DDR memory controllers, enhanced local bus, OCeaN targets, and configuration register access block all have a target port connection to the ECM.
- Split transaction support—separate address and data tenures allow for pipelining of transactions and out-of-order data tenures between initiators and targets.
- Proper ordering of I/O-initiated transactions.
- Speculative read bus for low-latency dispatch of reads to the DDR controllers.
- Low-latency paths for returning read data from DDR to the e500 cores.
- Error registers trap transactions with invalid addresses. Errors can be programmed to generate interrupts to the e500 core, as described in the following sections:
 - [ECM error detect register \(ECM_EEDR\)](#)
 - [ECM error enable register \(ECM_EEER\)](#)
 - [ECM error attributes capture register \(ECM_EEATR\)](#)
 - [ECM error low address capture register \(ECM_EELADR\)](#)
 - [ECM error high address capture register \(ECM_EEHADR\)](#)
- Errors from reading I/O devices terminate with data sent to the master with a corrupt attribute. If the master is the e500 core, the ECM asserts *core_fault_in* to the core, which causes the core to generate a machine check interrupt, unless it is disabled (by clearing HID1[RFXE]). If RFXE is zero and one of these errors occurs, appropriate interrupts must be enabled to ensure that an interrupt is generated.

7.2 ECM memory map/register definition

The table below shows the register memory map for the ECM.

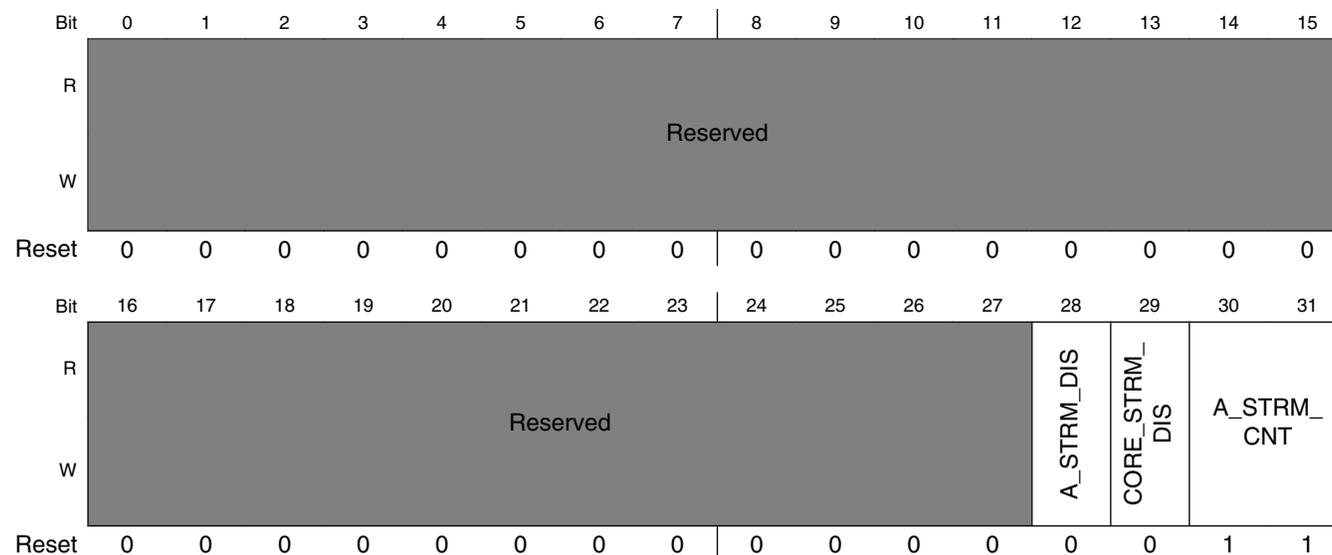
ECM memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
1000	ECM CCB address configuration register (ECM_EEBACR)	32	R/W	0000_0003h	7.2.1/242
1010	ECM CCB port configuration register (ECM_EEBPCR)	32	R/W	See section	7.2.2/243
1BF8	ECM IP Block Revision Register 1 (ECM_EIPBRR1)	32	R	0001_0000h	7.2.3/245
1BFC	ECM IP Block Revision Register 2 (ECM_EIPBRR2)	32	R	0000_0000h	7.2.4/245
1E00	ECM error detect register (ECM_EEDR)	32	w1c	0000_0000h	7.2.5/246
1E08	ECM error enable register (ECM_EEER)	32	R/W	0000_0000h	7.2.6/247
1E0C	ECM error attributes capture register (ECM_EEATR)	32	R	0000_0000h	7.2.7/248
1E10	ECM error low address capture register (ECM_EELADR)	32	R	0000_0000h	7.2.8/250
1E14	ECM error high address capture register (ECM_EEHADR)	32	R	0000_0000h	7.2.9/250

7.2.1 ECM CCB address configuration register (ECM_EEBACR)

The ECM CCB address configuration register controls arbitration and streaming policies for the CCB.

Address: 1000h base + 0h offset = 1000h



ECM_EEBACR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved

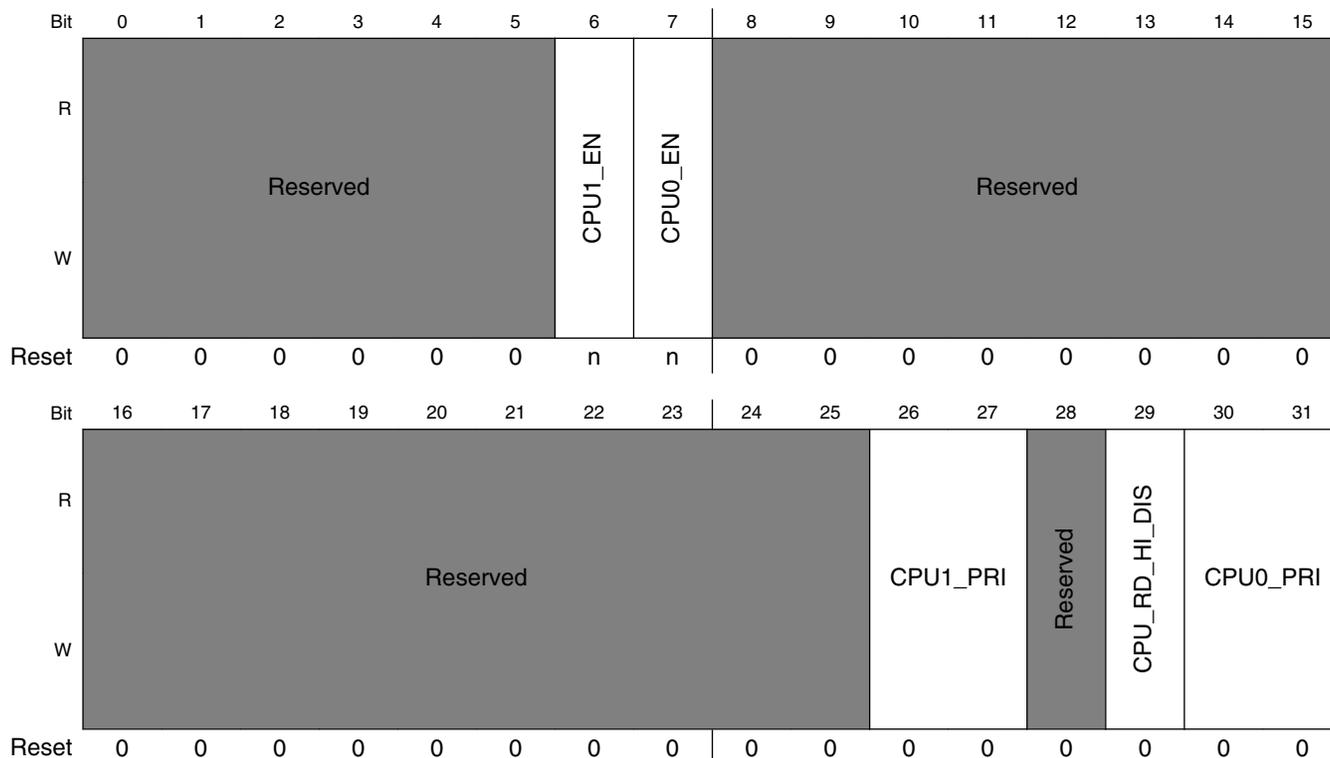
Table continues on the next page...

ECM_EEBACR field descriptions (continued)

Field	Description
28 A_STRM_DIS	Controls whether the ECM allows any streaming to occur. 0 Streaming is enabled. 1 Streaming is disabled.
29 CORE_STRM_DIS	With A_STRM_DIS, controls whether the e500 cores can stream commands onto the CCB . A_STRM_DIS and CORE_STRM_DIS must both be cleared for the e500 cores to be enabled to stream their address tenures that they master. 0 Stream address tenures initiated by the e500 cores , provided A_STRM_DIS is cleared. 1 Streaming of address tenures initiated by the e500 cores not allowed.
30-31 A_STRM_CNT	Stream count. Specifies the maximum number of transactions that any master can stream (issue sequentially without preemption) on the CCB following an initial transaction. 00 Reserved 01 One transaction can be streamed with the initial transaction. 10 Two transactions can be streamed with the initial transaction. 11 Three transactions can be streamed with the initial transaction. Default.

7.2.2 ECM CCB port configuration register (ECM_EEBPCR)

Address: 1000h base + 10h offset = 1010h



ECM_EEBPCR field descriptions

Field	Description
0–5 -	This field is reserved. Reserved
6 CPU1_EN	<p>CPU1 port enable</p> <p>Controls boot holdoff mode when the device is an agent of an external host or when the other core is used to perform initialization prior to allowing this core to boot. Specifies whether the e500 core (CPU) port1 is enabled to run transactions on the CCB. The CPU boot configuration power-on reset signal (cfg_cpu1_boot) determines the initial value of this bit. If the signal is sampled as a logic 1 at the negation of reset, CPU1 is enabled to boot at the end of the POR sequence. Otherwise, CPU1 cannot fetch its boot vector until an external host or another core sets the CPU1_EN bit.</p> <p>After this bit is set, it should not be cleared by software. It is not intended to dynamically enable and disable CPU operation. It is only intended to end boot holdoff mode. See CPU boot configuration , for more information.</p> <p>The settings are as follows:</p> <p>0 Boot holdoff mode. CPU1 arbitration is disabled on the CCB and no bus grants are issued. 1 CPU1 is enabled and receives bus grants in response to bus requests for the boot vector.</p>
7 CPU0_EN	<p>CPU 0 port enable</p> <p>Controls boot holdoff mode when the device is an agent of an external host or when the other core is used to perform initialization prior to allowing this core to boot . Specifies whether the e500 core (CPU) port0 is enabled to run transactions on the CCB. The CPU boot configuration power-on reset pin (cfg_cpu0_boot) determines the initial value of this bit. If the pin is sampled as a logic 1 at the negation of reset,CPU0 is enabled to boot at the end of the POR sequence. Otherwise,CPU0 cannot fetch its boot vector until an external host or another coresets the CPU0_EN bit.</p> <p>After this bit is set, it should not be cleared by software. It is not intended to dynamically enable and disable CPU operation. It is only intended to end boot holdoff mode. See CPU boot configuration , for more information.</p> <p>The settings are as follows:</p> <p>0 Boot holdoff mode. CPU0 arbitration is disabled on the CCB and no bus grants are issued. 1 CPU0 is enabled and receives bus grants in response to bus requests for the boot vector.</p>
8–25 -	This field is reserved. Reserved
26–27 CPU1_PRI	<p>Identifies the priority level of the e500 core 1 (CPU) port . This priority level is used to determine whether a particular port's bus request can cause the CCB Arbiter to terminate another port's streaming of address tenures.</p> <p>00 Lowest priority level 01 Second lowest priority level 10 Highest priority level 11 Reserved</p>
28 -	This field is reserved. Reserved
29 CPU_RD_HI_DIS	<p>Identifies which read queue of DDR targets is assigned to the e500 core (CPU) ports' read transactions (in understressed system).</p> <p>0 Read high queue (higher bandwidth DDR queue) is assigned for the e500 cores' read transactions 1 Read low queue (lower bandwidth DDR queue) is assigned for the e500 cores' read transactions</p>

Table continues on the next page...

ECM_EEBPCR field descriptions (continued)

Field	Description
30–31 CPU0_PRI	Specifies the priority level of the e500 core 0 (CPU) port . This priority level is used to determine whether a particular port's bus request can cause the CCB arbiter to terminate another port's streaming of address tenures. 00 Lowest priority level 01 Second lowest priority level 10 Highest priority level 11 Reserved

7.2.3 ECM IP Block Revision Register 1 (ECM_EIPBRR1)

Address: 1000h base + BF8h offset = 1BF8h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	IP_ID															IP_MJ								IP_MN								
W	[Shaded]																															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ECM_EIPBRR1 field descriptions

Field	Description
0–15 IP_ID	IP block ID
16–23 IP_MJ	Major revision
24–31 IP_MN	Minor revision

7.2.4 ECM IP Block Revision Register 2 (ECM_EIPBRR2)

Address: 1000h base + BFCh offset = 1BFCh

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								IP_INT								Reserved								IP_CFG							
W	[Shaded]																															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ECM_EIPBRR2 field descriptions

Field	Description
0–7 -	This field is reserved. Reserved

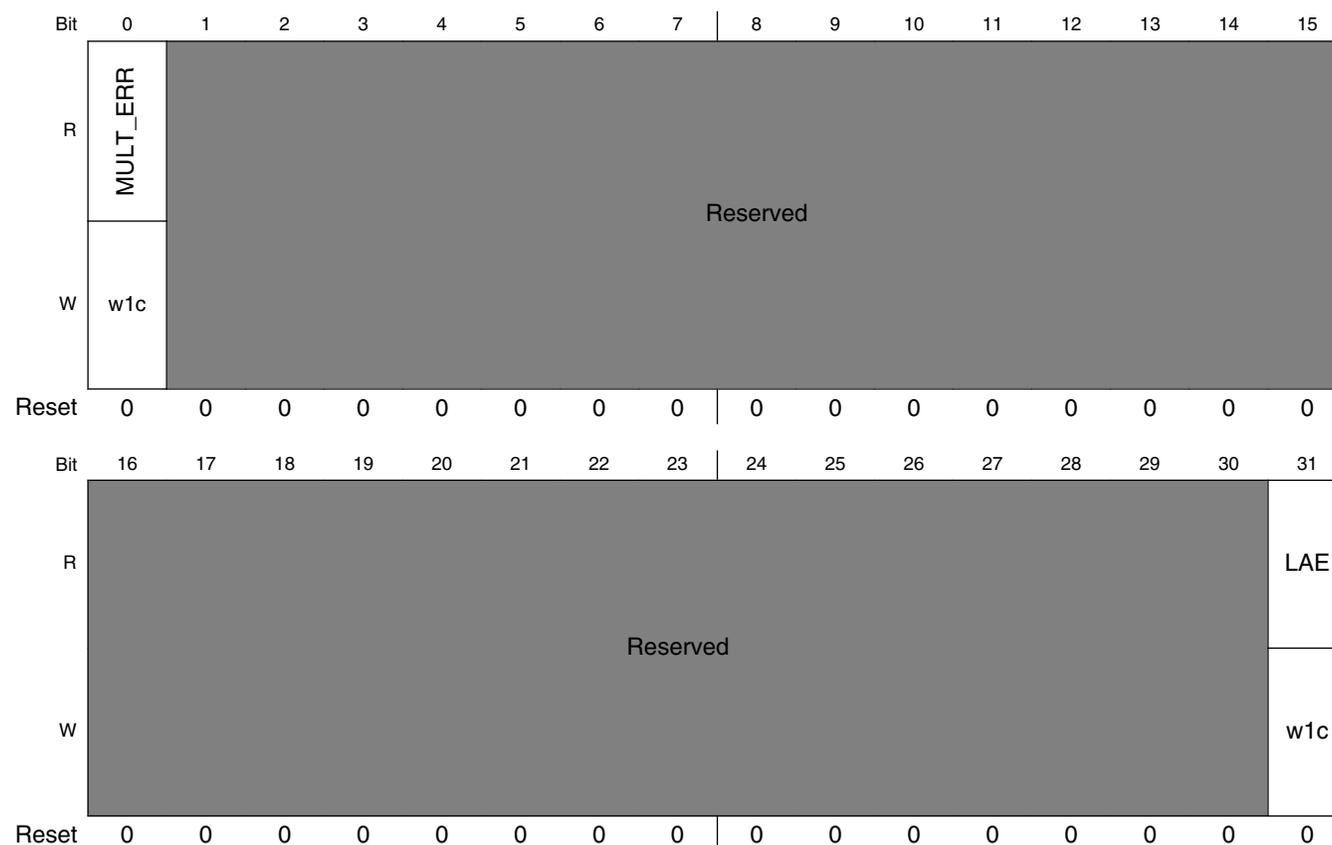
Table continues on the next page...

ECM_EIPBRR2 field descriptions (continued)

Field	Description
8–15 IP_INT	IP block integration options
16–23 -	This field is reserved. Reserved
24–31 IP_CFG	IP block configuration options

7.2.5 ECM error detect register (ECM_EEDR)

Address: 1000h base + E00h offset = 1E00h



ECM_EEDR field descriptions

Field	Description
0 MULT_ERR	Multiple error Indicates the occurrence of multiple errors of the same type. Write 1 to clear. 0 Multiple errors of the same type were not detected. 1 Multiple errors of the same type were detected.

Table continues on the next page...

ECM_EEDR field descriptions (continued)

Field	Description
1–30 -	This field is reserved. Reserved
31 LAE	Local access error Write 1 to clear. The following two cases can generate LAEs: <ul style="list-style-type: none"> Transaction does not map to any target. In this case the ECM injects read responses (with the corrupt attribute set) and write data is dropped. Note that a read that attempts to access an unmapped target causes the assertion of <i>core_fault_in</i>, which causes the core to generate a machine check interrupt, unless it is disabled (by clearing HID1[RFXE]). If RFXE is zero and this error occurs, EEER[LAE] must be set to ensure that an interrupt is generated. Source and target IDs indicate that an OCN port initiated a transaction that targets an OCN port. This loopback behavior can result from programming errors where inbound ATMU window targets are inconsistent with targets configured in the local access windows for a given address range. For this type of LAE, the dispatch (to OCN target in this case) is not screened off; the LAE error is reported, but the transaction is still sent to its OCN target. 0 Local access error has not occurred. 1 Local access error occurred.

7.2.6 ECM error enable register (ECM_EEER)

The ECM error enable register (EEER) enables the reporting of error conditions to the e500 core through the internal $\overline{\text{int}}$ interrupt signal.

Address: 1000h base + E08h offset = 1E08h

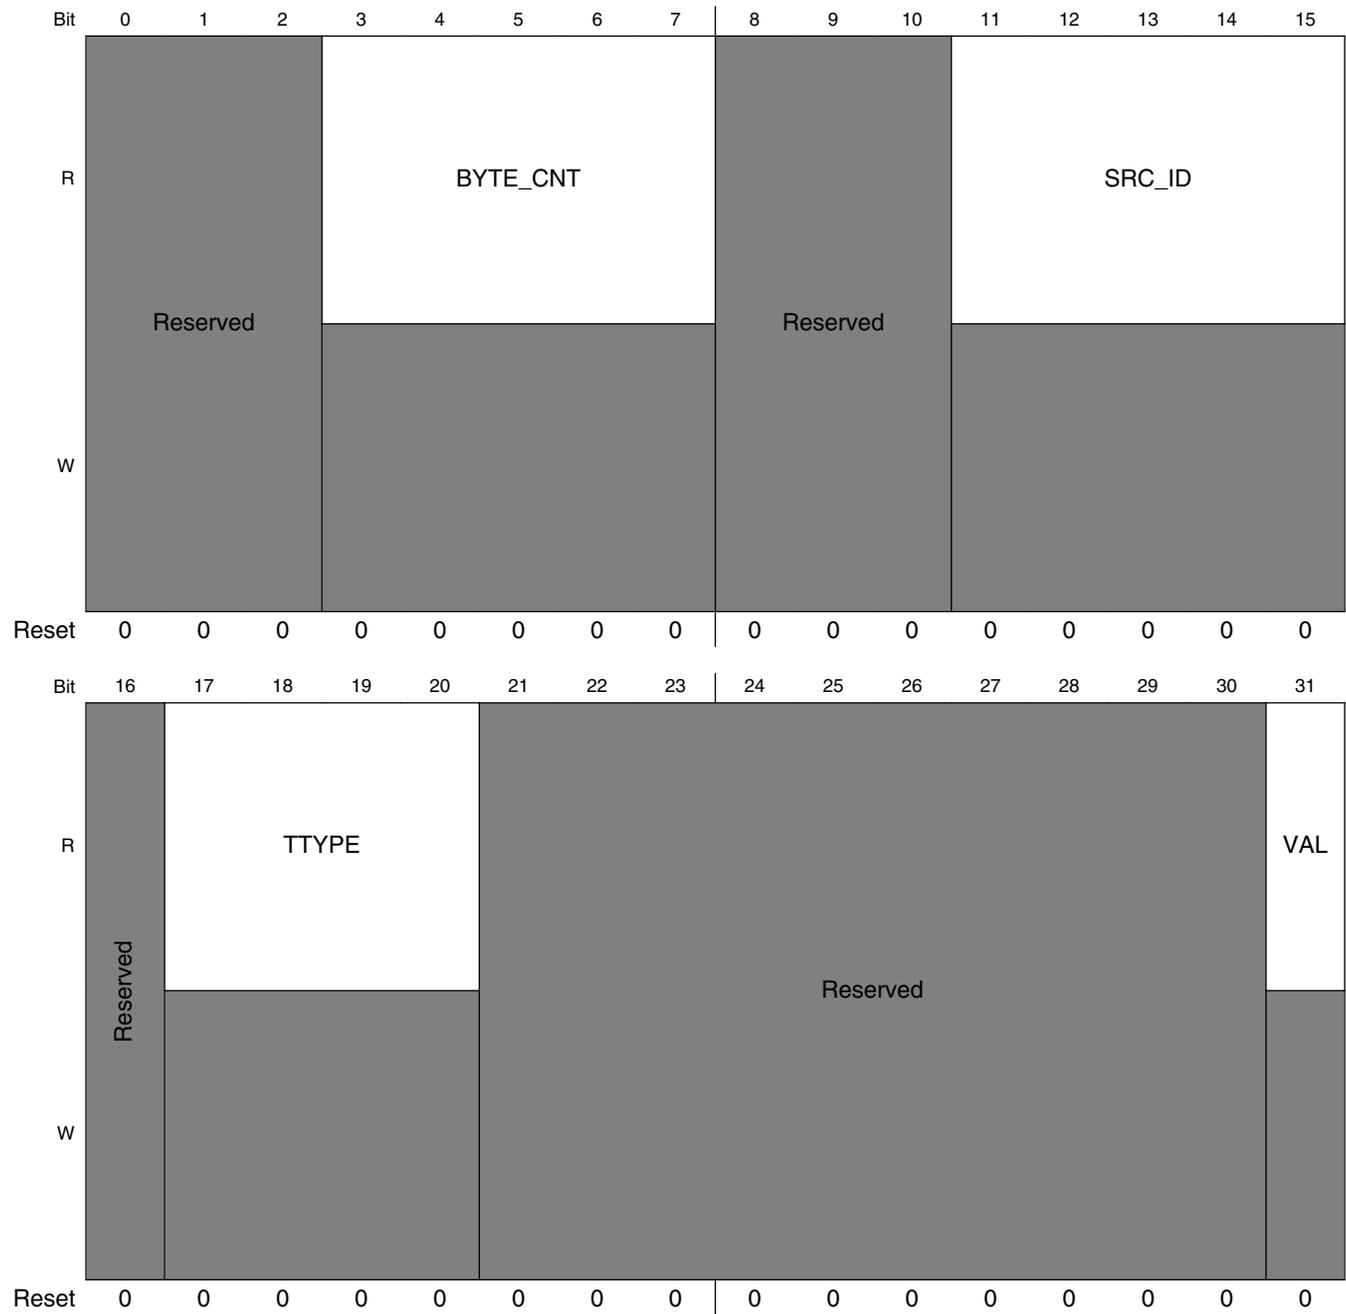
Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	Reserved																
W	LAE																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

ECM_EEER field descriptions

Field	Description
0–30 -	This field is reserved. Reserved
31 LAEE	Local access error enable Note that a read that attempts to access an unmapped target causes the assertion of <i>core_fault_in</i> , which causes the core to generate a machine check interrupt, unless it is disabled (by clearing HID1[RFXE]). If HID1[RFXE] is zero and this error occurs, LAEE must be set to ensure that an interrupt is generated. 0 Disable reporting local access errors as interrupts. 1 Enable reporting local access errors as interrupts .

7.2.7 ECM error attributes capture register (ECM_EEATR)

Address: 1000h base + E0Ch offset = 1E0Ch



ECM_EEATR field descriptions

Field	Description
0-2 -	This field is reserved. Reserved

Table continues on the next page...

ECM_EEATR field descriptions (continued)

Field	Description
3–7 BYTE_CNT	Byte count. Specifies the transaction byte count. 00000 32 bytes 00001 1 byte 00010 2 bytes 00100 4 bytes 01000 8 bytes 10000 16 bytes
8–10 -	This field is reserved. Reserved
11–15 SRC_ID	Source ID. Specifies the source device mastering the transaction. Patterns not shown are reserved. 00000 PCI Express 3 00001 PCI Express 2 00010 PCI Express 1 00101 USB 00111 Security 01010 Boot sequencer 01011 eSDHC 01100 Serial RapidIO 1 01101 Serial RapidIO 2 10000 Processor 0 (instruction) 10001 Processor 0 (data) 10010 Processor 1 (instruction) 10011 Processor 1 (data) 10101 DMA 1 10110 DMA 2 11000 eTSEC 1 11001 eTSEC 2 11010 eTSEC 3 11100 RapidIO message unit 11101 RapidIO doorbell unit 11110 RapidIO port-write unit
16 -	This field is reserved. Reserved
17–20 TTYPE	Transaction type. Defined as follows: All other settings are reserved. 0000 Write 0010 Write with allocate 0011 Write with allocate with lock 0100 Address only transaction 1000 Read 1001 Read with unlock 1100 Read with clear atomic 1101 Read with set atomic

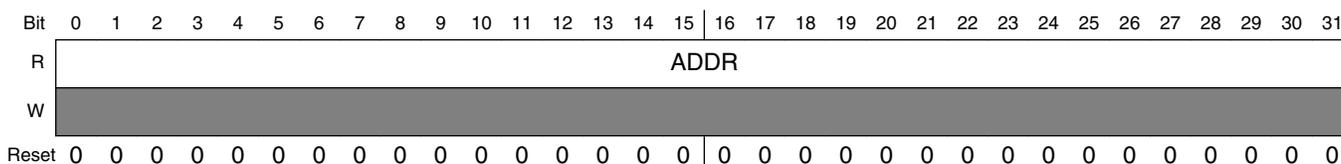
Table continues on the next page...

ECM_EEATR field descriptions (continued)

Field	Description
	1110 Read with decrement atomic 1111 Read with increment atomic
21–30 -	This field is reserved. Reserved
31 VAL	Register data valid. 0 ECM error attribute capture register does not contain valid information. 1 ECM error attribute capture register contains valid information.

7.2.8 ECM error low address capture register (ECM_EELADR)

Address: 1000h base + E10h offset = 1E10h

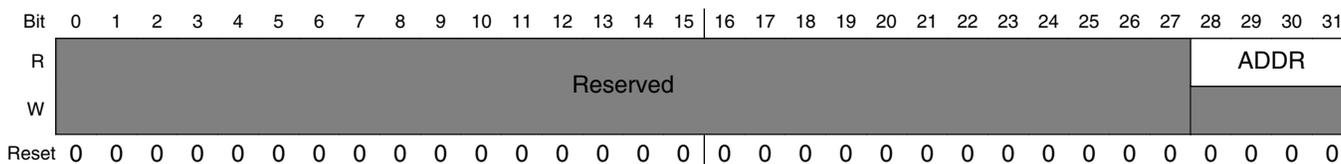


ECM_EELADR field descriptions

Field	Description
0–31 ADDR	Address. Specifies the lower-order 32 bits of the 36-bit address of the transaction. Qualified by EEATR[VAL].

7.2.9 ECM error high address capture register (ECM_EEHADR)

Address: 1000h base + E14h offset = 1E14h



ECM_EEHADR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 ADDR	Address. Specifies the high-order 4 bits of the 36-bit address of the transaction. Qualified by EEATR[VAL].

7.3 Functional description

The following is a general discussion of ECM operation.

7.3.1 I/O arbiter

Figure 7-1 shows the I/O arbiter block that manages I/O-initiated address tenure requests arriving on the request buses.

Four request buses compete for access to the ECM, which can only process one request at a time.

The ECM uses two factors to select the winning request bus: the primary factor is requested bandwidth and the secondary factor is longest waiting/least recently granted status. By default all requesters start requesting low levels of bandwidth. A starvation avoidance algorithm ensures that low bandwidth requesters make forward progress in the presence of high bandwidth requesters. The transaction from the winning request bus competes with e500 core requests for the CCB and entry into the transaction queue.

7.3.2 CCB arbiter

Figure 7-1 shows the CCB arbiter block coordinating the entry of new transactions into the ECM's transaction queue .

It handles arbitration for requests to use the CCB from the e500 cores and the winning request bus and consequently controls when these new transactions can enter the transaction queue.

Because the CCB bus operates most efficiently when it streams commands from one initiator, the CCB arbiter alternates grants between streams of transactions from the e500 cores and from the winner of the I/O arbiter. The length of a stream (number of back-to-back transactions) is limited by the A_STRM_CNT field in the EEBACR register. However, the arbiter also uses the priority of the requests to limit streaming. If the priority of a new request is higher than that of a stream in progress, then the higher priority transaction interrupts the other stream. The priority of e500 transactions is set by the CPU[0/1]_PRI field in EEBPCR register. Depending how the CPU_RD_HI_DIS field in EEBPCR register is set, read transactions from the e500 cores are initially assigned to either the higher- or lower-bandwidth queue of the DDR target.

7.3.3 Transaction queue

The ECM's transaction queue performs four basic functions: arbitration across the e500 cores and I/O masters, target mapping and dispatching, enforcement of ordering, and enforcement of coherency.

The address of each transaction is compared against each local access window, and the transaction is then routed to the appropriate target interface associated with the local access window that the address hits within. Even though the CCB and ECM allow the pipelining of transactions, the address tenures of all transactions issued from I/O masters (masters other than the e500 cores) may still be ordered. For those transactions accessing address space marked as snoopable, or space that may be cached by the e500 cores, the ECM enforces coherency, snooping those transactions on the CCB, and taking castouts from the e500 cores as is necessary.

7.3.4 Global data multiplexor

The global data mux allows initiators of write transactions to route data to their targets and read targets to return data to the initiators.

See [Figure 7-1](#) for an illustration of how the global data multiplexor takes data bus connections and multiplexes them onto one 128-bit global data bus.

7.3.5 CCB interface

This interface formats CCB address tenures for the ECM transaction queue. It also contains the queueing and buffering needed to manage outstanding CCB data tenures.

The buffers receive e500 core-initiated write and I/O-initiated read data (that hit in the L2/SRAM module) from the e500 write (128-bit wide) and read (128-bit wide) data buses and route them through the global data mux to the global data bus. The buffers also receive e500 core-initiated read and I/O-initiated write data (that hit in the L2/SRAM module) from the global data bus and forward them onto the CCB data bus (64 bits).

[Figure 7-1](#) shows the CCB interface for both CCB address and data tenures.

7.4 Initialization/application information

If either e500 core is used to initialize the device, the applicable CPU boot configuration power-on reset pin should be pulled high to initially set `EEBPCR[CPU0_EN]` or `EEBPCR[CPU1_EN]`. See [Reset, Clocking, and Initialization](#) for more information on power-up reset initialization.

If any device other than the e500 cores is used to initialize the device, the CPU boot configuration power-on reset pins should be pulled low to initially clear `EEBPCR[CPU0_EN]` and `EEBPCR[CPU1_EN]`. This prevents the e500 cores from accessing any configuration registers or local memory space during initialization. However, in any such system, one step near the end of the initialization routine must set `EEBPCR[CPU0_EN]` and `EEBPCR[CPU1_EN]` to re-enable the e500 cores. Note that for basic functionality, `EEBPCR[CPU0/1_EN]` is the only field that must be written (provided a device other than the e500 cores is used to initialize the device) in the ECM.

`EEBPCR[CPU0_PRI]` and `EEBPCR[CPU1_PRI]` specify the priority level associated with all e500 core initiated transactions. These values allow users running time-critical applications to adjust the average response latency of transactions initiated by the cores compared to those initiated by I/O masters. These priority levels affect whether e500 core requests can interrupt the streaming of address tenures initiated by (the ECM on behalf of) I/O masters or the other e500 core. Only transactions with a priority greater than the current CCB transaction can interrupt streaming. The higher the core's priority, the lower the average latency needed for it to obtain bus grants from the ECM because it can interrupt lower priority streaming. The default value of zero gives all core-initiated transactions the lowest priority, which prevents the cores from interrupting I/O master transaction streams.

`EEBACR[A_STRM_CNT]` allows users to balance response latency with throughput and should prove useful in tuning systems with multiple time-critical tasks. The default value of 0b11 causes the ECM to attempt to stream as many as four transactions initiated from the same CCB master. Decreasing this value decreases the maximum number of transactions that may be streamed together from any one CCB master. Decreasing this value can decrease throughput for high priority transactions, but may decrease latency for lower priority transactions from another CCB master. Note that the e500 cores must also have streaming enabled (through `HID1[ASTME]` in each core) for the CCB to stream.



Chapter 8

DDR Memory Controller

The fully programmable DDR SDRAM controller supports most JEDEC standard x8, x16, or x32 DDR3/3L memories available. In addition, unbuffered and registered DIMMs are supported.

8.1 Introduction

The fully programmable DDR SDRAM controller supports most JEDEC standard x8, x16, or x32 DDR2 and DDR3 memories available.

In addition, unbuffered and registered DIMMs are supported. However, mixing different memory types or unbuffered and registered DIMMs in the same system is not supported. Built-in error checking and correction (ECC) ensures very low bit-error rates for reliable high-frequency operation. Dynamic power management and auto-precharge modes simplify memory system design. A large set of special features, including ECC error injection, support rapid system debug.

NOTE

In this chapter, the word 'bank' refers to a physical bank specified by a chip select; 'logical bank' refers to one of the four or eight sub-banks in each SDRAM chip. A sub-bank is specified by the 2 or 3 bits on the bank address (MBA) pins during a memory access.

The figure below is a high-level block diagram of the DDR memory controller with its associated interfaces. [Functional description](#) contains detailed figures of the controller.

features

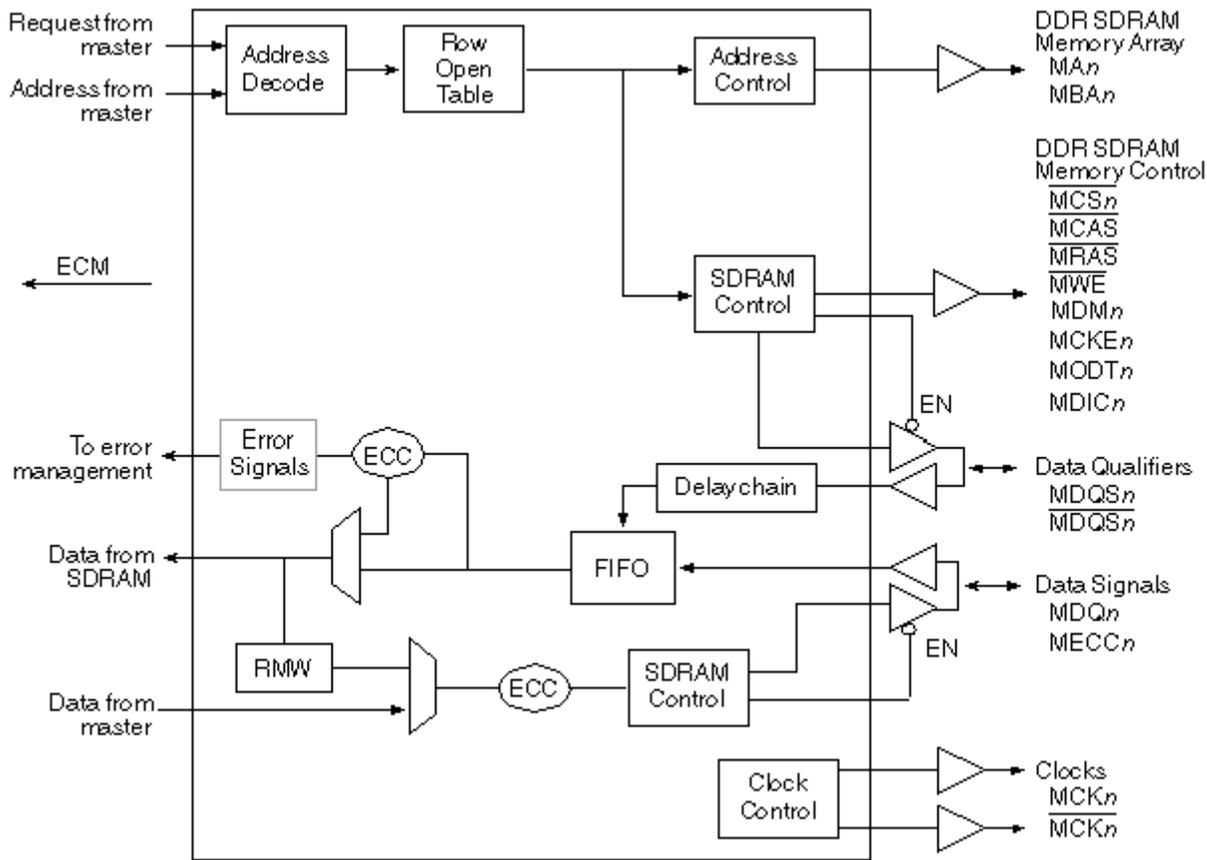


Figure 8-1. DDR memory controller simplified block diagram

8.2 Features

The DDR memory controller includes these distinctive features:

- Support for DDR2 and DDR3 SDRAM
- 64-/72-bit or 32-/40-bit SDRAM data bus for DDR2 and DDR3
- Support for up to 32GB of memory
- Support for DRAM clocking asynchronously from the platform clock
- Programmable settings for meeting all SDRAM timing parameters
- Support for the following SDRAM configurations:
 - As many as four physical banks (chip selects), each bank independently addressable
 - 64-Mbit to 8-Gbit devices depending on internal device configuration with x8/x16/x32 data ports (no direct x4 support)
 - Unbuffered and registered DIMMs
- Chip select interleaving support
- Partial array self refresh support

- Support for data mask signals and read-modify-write for sub-double-word writes. Note that a read-modify-write sequence is only necessary when ECC is enabled.
- Support for double-bit error detection and single-bit error correction ECC (8-bit check word across 32 or 64-bit data)
- Support for address parity for registered DIMMs
- Open page management (dedicated entry for each logical bank)
- Automatic DRAM initialization sequence or software-controlled initialization sequence
- Automatic DRAM data initialization
- Write leveling supported for DDR3 memories
- Support for up to eight posted refreshes
- Support for error injection

8.2.1 Modes of operation

The DDR memory controller supports the following modes:

- Dynamic power management mode. The DDR memory controller can reduce power consumption by negating the SDRAM CKE signal when no transactions are pending to the SDRAM.
- Auto-precharge mode. Clearing `DDR_SDRAM_INTERVAL[BSTOPRE]` causes the memory controller to issue an auto-precharge command with every read or write transaction. Auto-precharge mode can be enabled for separate chip selects by setting `CSn_CONFIG[APn_EN]`.

8.3 External signal descriptions

This section provides descriptions of the DDR memory controller's external signals.

It describes each signal's behavior when the signal is asserted or negated and when the signal is an input or an output.

8.3.1 Signals overview

Memory controller signals are grouped as follows:

- Memory interface signals
- Clock signals

The table below shows all signals of DDR controller. The device hardware specification shows detailed multiplexing of these signal. Note that availability of all signals depends upon device configurations.

Table 8-1. DDR memory interface signal summary

Name	Function/description	Reset	Pins	I/O
MAPAR_ERR_B	Address parity error	One	1	I
MAPAR_OUT	Address parity out	Zero	1	O
MDQ[0:63]	Data bus	All zeros	64	I/O
MDQS[0:8]	Data strobes	All zeros	9	I/O
MDQS_B[0:8]	Data strobes (complement)	All ones	9	I/O
MECC[0:7]	Error checking and correcting	All zeros	8	I/O
MCAS_B	Column address strobe	One	1	O
MA[15:0]	Address bus	All zeros	16	O
MBA[2:0]	Logical bank address	All zeros	3	O
MCS_B[0:3]	Chip selects	All ones	4	O
MWE_B	Write enable	One	1	O
MRAS_B	Row address strobe	One	1	O
MDM[0:8]	Data mask	All zeros	9	O
MCK[0:5]	DRAM clock outputs	All zeros	6	O
MCK_B[0:5]	DRAM clock outputs (complement)	All zeros	6	O
MCKE[0:3]	DRAM clock enable	All zeros	4	O
MODT[0:3]	DRAM on-die termination external control	All zeros	4	O
MDIC[0:1]	Driver impedance calibration	10	2	I/O

This table shows the memory address signal mappings.

Table 8-2. Memory address signal mappings

Signal name (Outputs)		JEDEC DDR DIMM signals (Inputs)
msb	MA15	A15
	MA14	A14
	MA13	A13
	MA12	A12
	MA11	A11
	MA10	A10 (AP for DDR) ¹
	MA9	A9
	MA8	A8 (alternate AP for DDR) ²
	MA7	A7
	MA6	A6
	MA5	A5
	MA4	A4
	MA3	A3
	MA2	A2
MA1	A1	
lsb	MA0	A0
msb	MBA2	MBA2
	MBA1	MBA1
lsb	MBA0	MBA0

1. Auto-precharge for DDR signaled on A10 when DDR_SDRAM_CFG[PCHB8] = 0
2. Auto-precharge for DDR signaled on A8 when DDR_SDRAM_CFG[PCHB8] = 1

8.3.2 Detailed signal descriptions

The following sections describe the DDR SDRAM controller input and output signals, the meaning of their different states, and relative timing information for assertion and negation.

8.3.2.1 Memory interface signals

This table describes the DDR controller memory interface signals.

Table 8-3. Memory interface signals-detailed signal descriptions

Signal	I/O	Description	
MDQ[0:63]	I/O	Data bus. Both input and output signals on the DDR memory controller.	
	O	As outputs for the bidirectional data bus, these signals operate as described below.	
		State meaning	Asserted/Negated-Represent the value of data being driven by the DDR memory controller.
		Timing	Assertion/Negation-Driven coincident with corresponding data strobes (MDQS) signal. High impedance-No READ or WRITE command is in progress; data is not being driven by the memory controller or the DRAM.
	I	As inputs for the bidirectional data bus, these signals operate as described below.	
		State Meaning	Asserted/Negated-Represents the state of data being driven by the external DDR SDRAMs.
Timing		Assertion/Negation-The DDR SDRAM drives data during a READ transaction. High impedance-No READ or WRITE command in progress; data is not being driven by the memory controller or the DRAM.	
MDQS[0:8]/ MDQS_B[0:8]	I/O	Data strobes. Inputs with read data, outputs with write data.	
	O	As outputs, the data strobes are driven by the DDR memory controller during a write transaction.	
		State Meaning	Asserted/Negated-Driven high when positive capture data is transmitted and driven low when negative capture data is transmitted. Centered in the data "eye" for writes; coincident with the data eye for reads. Treated as a clock. Data is valid when signals toggle. See Table 8-63 for byte lane assignments.
		Timing	Assertion/Negation (for DDR2/DDR3)-If a WRITE command is registered at clock edge n , and the latency is programmed in TIMING_CFG_2[WR_LAT] to be m clocks, data strobes at the DRAM assert coincident with the data on clock edge $n + m$. See the JEDEC DDR SDRAM specification for more information.
	I	As inputs, the data strobes are driven by the external DDR SDRAMs during a read transaction. The data strobes are used by the memory controller to synchronize data latching.	
		State Meaning	Asserted/Negated-Driven high when positive capture data is received and driven low when negative capture data is received. Centered in the data eye for writes; coincident with the data eye for reads. Treated as a clock. Data is valid when signals toggle. See Table 8-63 for byte lane assignments.
Timing		Assertion/Negation-If a READ command is registered at clock edge n , and the latency is programmed in TIMING_CFG_1[CASLAT] to be m clocks, data strobes at the DRAM assert coincident with the data on clock edge $n + m$. See the JEDEC DDR SDRAM specification for more information.	

Table continues on the next page...

Table 8-3. Memory interface signals-detailed signal descriptions (continued)

Signal	I/O	Description	
MECC[0:7]	I/O	Error checking and correcting codes. Input and output signals for the DDR controller's bidirectional ECC bus. MECC[0:5] function in both normal and debug modes.	
	O	As normal mode outputs, the ECC signals represent the state of ECC driven by the DDR controller on writes. As debug mode outputs MECC[0:5] provide source ID and data-valid information. See Error checking and correcting (ECC) and Debug information on debug pins for more details.	
		State Meaning	Asserted/Negated-Represents the state of ECC being driven by the DDR controller on writes.
		Timing	Assertion/Negation-Same timing as MDQ High impedance-Same timing as MDQ
	I	As inputs, the ECC signals represent the state of ECC driven by the SDRAM devices on reads.	
		State Meaning	Asserted/Negated-Represents the state of ECC being driven by the DDR SDRAMs on reads.
Timing		Assertion/Negation-Same timing as MDQ High impedance-Same timing as MDQ	
MA[15:0]	O	Address bus. Memory controller outputs for the address to the DRAM. MA[15:0] carry 16 of the address bits for the DDR memory interface corresponding to the row and column address bits. MA0 is the lsb of the address output from the memory controller.	
		State Meaning	Asserted/Negated-Represents the address driven by the DDR memory controller. Contains different portions of the address depending on the memory size and the DRAM command being issued by the memory controller. See Table 8-4 for a complete description of the mapping of these signals.
		Timing	Assertion/Negation-The address lines are only driven when the controller has a command scheduled to issue on the address/CMD bus; otherwise they will be at high-Z. It is valid when a transaction is driven to DRAM (when MCS_B n is active). High impedance-When the memory controller is disabled
MBA[2:0]	O	Logical bank address. Outputs that drive the logical (or internal) bank address pins of the SDRAM. Each SDRAM supports four or eight addressable logical sub-banks. Bit zero of the memory controller's output bank address must be connected to bit zero of the SDRAM's input bank address. MBA0, the least-significant bit of the three bank address signals, is asserted during the mode register set command to specify the extended mode register.	
		State Meaning	Asserted/Negated-Selects the DDR SDRAM logical (or internal) bank to be activated during the row address phase and selects the SDRAM internal bank for the read or write operation during the column address phase of the memory access. Table 8-4 describes the mapping of these signals in all cases.
		Timing	Assertion/Negation-Same timing as MA _n High impedance-Same timing as MA _n

Table continues on the next page...

Table 8-3. Memory interface signals-detailed signal descriptions (continued)

Signal	I/O	Description
MCAS_B	O	Column address strobe. Active-low SDRAM address multiplexing signal. MCAS_B is asserted for read or write transactions and for mode register set, refresh, and precharge commands.
		State Meaning Asserted-Indicates that a valid SDRAM column address is on the address bus for read and write transactions. See Table 8-74 for more information on the states required on MCAS_B for various other SDRAM commands. Negated-The column address is not guaranteed to be valid.
		Timing Assertion/Negation-Assertion and negation timing is directed by the values described in DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0) , DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1) , DDR SDRAM timing configuration 2 (DDR_TIMING_CFG_2) , and DDR SDRAM timing configuration 3 (DDR_TIMING_CFG_3) . High impedance-MCAS_B is tri-stated when memory controller is idle.
MRAS_B	O	Row address strobe. Active-low SDRAM address multiplexing signal. Asserted for activate commands. In addition; used for mode register set commands and refresh commands.
		State Meaning Asserted-Indicates that a valid SDRAM row address is on the address bus for read and write transactions. See Table 8-74 for more information on the states required on MRAS_B for various other SDRAM commands. Negated-The row address is not guaranteed to be valid.
		Timing Assertion/Negation-Assertion and negation timing is directed by the values described in DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0) , DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1) , DDR SDRAM timing configuration 2 (DDR_TIMING_CFG_2) , and DDR SDRAM timing configuration 3 (DDR_TIMING_CFG_3) . High impedance-MRAS_B is tri-stated when memory controller is idle.
MCS_B[0:3]	O	Chip selects. Four chip selects supported by the memory controller.
		State Meaning Asserted-Selects a physical SDRAM bank to perform a memory operation as described in Chip select n memory bounds (DDR_CS_n_BNDS) , and Chip select n configuration (DDR_CS_n_CONFIG) . The DDR controller asserts one of the MCS_B[0:3] signals to begin a memory cycle. Negated-Indicates no SDRAM action during the current cycle.
		Timing Assertion/Negation-Asserted to signal any new transaction to the SDRAM. The transaction must adhere to the timing constraints set in TIMING_CFG_0-TIMING_CFG_3 . High impedance-Always driven unless the memory controller is disabled.
MWE_B	O	Write enable. Asserted when a write transaction is issued to the SDRAM. This is also used for mode registers set commands and precharge commands.
		State Meaning Asserted-Indicates a memory write operation. See Table 8-74 for more information on the states required on MWE_B for various other SDRAM commands. Negated-Indicates a memory read operation.
		Timing Assertion/Negation-Similar timing as MRAS_B and MCAS_B. Used for write commands. High impedance-MWE_B is tri-stated when memory controller is idle.

Table continues on the next page...

Table 8-3. Memory interface signals-detailed signal descriptions (continued)

Signal	I/O	Description
MDM[0:8]	O	DDR SDRAM data output mask. Masks unwanted bytes of data transferred during a write. They are needed to support sub-burst-size transactions (such as single-byte writes) on SDRAM where all I/O occurs in multi-byte bursts. MDM0 corresponds to the most significant byte (MSB) and MDM7 corresponds to the LSB, while MDM8 corresponds to the ECC byte. Table 8-63 shows byte lane encodings.
		State Meaning Asserted-Prevents writing to DDR SDRAM. Asserted when data is written to DRAM if the corresponding byte(s) should be masked for the write. Note that the MDM n signals are active-high for the DDR controller. MDM n is part of the DDR command encoding. Negated-Allows the corresponding byte to be read from or written to the SDRAM.
		Timing Assertion/Negation-Same timing as MDQx as outputs. High impedance-Always driven unless the memory controller is disabled.
MODT[0:3]	O	On-Die termination. Memory controller outputs for the ODT to the DRAM. MODT[0:3] represents the on-die termination for the associated data, data masks, ECC, and data strobes.
		State Meaning Asserted/Negated-Represents the ODT driven by the DDR memory controller.
		Timing Assertion/Negation-Driven in accordance with JEDEC DRAM specifications for on-die termination timings. It is configured through the CS n _CONFIG[ODT_RD_CFG] and CS n _CONFIG[ODT_WR_CFG] fields. High impedance-Always driven.
MDIC[0:1]	I/O	Driver impedance calibration. Note that the MDIC signals require the use of resistors; MDIC0 must be pulled to GND, while MDIC1 must be pulled to GV _{DD} . See DDR Control Driver Register 1 (DDR_DDRCDR_1) .
		State Meaning These pins are used for automatic calibration of the DDR IOs.
		Timing These are driven for four DRAM cycles at a time while the DDR controller is executing the automatic driver compensation.
MAPAR_ERR_B	I	Address parity error. Reflects whether an address parity error has been detected by the DRAM. This signal is active low.
		State Meaning Asserted-An error has been detected. Negated-An error has not been detected.
		Timing Assertion/Negation-are driven by the registered DIMMs one DRAM cycle after the parity bit has been driven by the memory controller. This error signal should be held valid for two DRAM cycles.
MAPAR_OUT_B	O	Address parity out. Driven by the memory controller as the parity bit calculated across the address and command bits. Even parity is used, and parity is not calculated for the MCKE[0:3], MODT[0:3], or MCS_B[0:3] signals.
		State Meaning Asserted-The parity bit is high. Negated-The parity bit is low.
		Timing Assertion/Negation-are issued one DRAM cycle after the chip select for each command.

8.3.2.2 Clock interface signals

The table below contains the detailed descriptions of the clock signals of the DDR controller.

Table 8-4. Clock signals-detailed signal descriptions

Signal	I/O	Description	
MCK[0:5] , MCK_B [0:5]	O	DRAM clock outputs and their complements.	
		State Meaning	Asserted/Negated-The JEDEC DDR SDRAM specifications require true and complement clocks. A clock edge is seen by the SDRAM when the true and complement cross.
		Timing	Assertion/Negation-Timing is controlled by the DDR_CLK_CNTL register at offset 0x130.
MCKE[0:3]	O	Clock enable. Output signals used as the clock enables to the SDRAM. MCKE[0:3] can be negated to stop clocking the DDR SDRAM. The MCKE signals should be connected to the same rank of memory as the corresponding MCS_B and MODT signals. For example, MCKE[0] should be connected to the same rank of memory as MCS_B[0] and MODT[0].	
		State Meaning	Asserted-Clocking to the SDRAM is enabled. Negated-Clocking to the SDRAM is disabled and the SDRAM should ignore signal transitions on MCK or MCK_B. MCK/MCK_B are don't cares while MCKE[0:3] are negated.
		Timing	Assertion/Negation-Asserted when DDR_SDRAM_CFG[MEM_EN] is set. Can be negated when entering dynamic power management or self refresh. Are asserted again when exiting dynamic power management or self refresh. High impedance-Always driven.

8.4 DDR memory map/register definition

The following table shows the register memory map for the DDR memory controller .

DDR memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2000	Chip select n memory bounds (DDR_CS0_BNDS)	32	R/W	0000_0000h	8.4.1/266
2008	Chip select n memory bounds (DDR_CS1_BNDS)	32	R/W	0000_0000h	8.4.1/266
2010	Chip select n memory bounds (DDR_CS2_BNDS)	32	R/W	0000_0000h	8.4.1/266
2018	Chip select n memory bounds (DDR_CS3_BNDS)	32	R/W	0000_0000h	8.4.1/266
2080	Chip select n configuration (DDR_CS0_CONFIG)	32	R/W	0000_0000h	8.4.2/267
2084	Chip select n configuration (DDR_CS1_CONFIG)	32	R/W	0000_0000h	8.4.2/267
2088	Chip select n configuration (DDR_CS2_CONFIG)	32	R/W	0000_0000h	8.4.2/267
208C	Chip select n configuration (DDR_CS3_CONFIG)	32	R/W	0000_0000h	8.4.2/267
20C0	Chip select n configuration 2 (DDR_CS0_CONFIG_2)	32	R/W	0000_0000h	8.4.3/269

Table continues on the next page...

DDR memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
20C4	Chip select n configuration 2 (DDR_CS1_CONFIG_2)	32	R/W	0000_0000h	8.4.3/269
20C8	Chip select n configuration 2 (DDR_CS2_CONFIG_2)	32	R/W	0000_0000h	8.4.3/269
20CC	Chip select n configuration 2 (DDR_CS3_CONFIG_2)	32	R/W	0000_0000h	8.4.3/269
2100	DDR SDRAM timing configuration 3 (DDR_TIMING_CFG_3)	32	R/W	0000_0000h	8.4.4/270
2104	DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)	32	R/W	0011_0105h	8.4.5/272
2108	DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)	32	R/W	0000_0000h	8.4.6/276
210C	DDR SDRAM timing configuration 2 (DDR_TIMING_CFG_2)	32	R/W	0000_0000h	8.4.7/280
2110	DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG)	32	R/W	0300_0000h	8.4.8/284
2114	DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2)	32	R/W	0000_0000h	8.4.9/287
2118	DDR SDRAM mode configuration (DDR_DDR_SDRAM_MODE)	32	R/W	0000_0000h	8.4.10/290
211C	DDR SDRAM mode configuration 2 (DDR_DDR_SDRAM_MODE_2)	32	R/W	0000_0000h	8.4.11/291
2120	DDR SDRAM mode control (DDR_DDR_SDRAM_MD_CNTL)	32	R/W	0000_0000h	8.4.12/291
2124	DDR SDRAM interval configuration (DDR_DDR_SDRAM_INTERVAL)	32	R/W	0000_0000h	8.4.13/295
2128	DDR SDRAM data initialization (DDR_DDR_DATA_INIT)	32	R/W	0000_0000h	8.4.14/295
2130	DDR SDRAM clock control (DDR_DDR_SDRAM_CLK_CNTL)	32	R/W	0200_0000h	8.4.15/296
2148	DDR training initialization address (DDR_DDR_INIT_ADDR)	32	R/W	0000_0000h	8.4.16/297
214C	DDR training initialization extended address (DDR_DDR_INIT_EXT_ADDR)	32	R/W	0000_0000h	8.4.17/298
2160	DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4)	32	R/W	0000_0000h	8.4.18/299
2164	DDR SDRAM timing configuration 5 (DDR_TIMING_CFG_5)	32	R/W	0000_0000h	8.4.19/302
2170	DDR ZQ calibration control (DDR_DDR_ZQ_CNTL)	32	R/W	0000_0000h	8.4.20/304
2174	DDR write leveling control (DDR_DDR_WRLVL_CNTL)	32	R/W	0000_0000h	8.4.21/306
217C	DDR Self Refresh Counter (DDR_DDR_SR_CNTR)	32	R/W	0000_0000h	8.4.22/309
2180	DDR Register Control Words 1 (DDR_DDR_SDRAM_RCW_1)	32	R/W	0000_0000h	8.4.23/310
2184	DDR Register Control Words 2 (DDR_DDR_SDRAM_RCW_2)	32	R/W	0000_0000h	8.4.24/311
2190	DDR write leveling control 2 (DDR_DDR_WRLVL_CNTL_2)	32	R/W	0000_0000h	8.4.25/312
2194	DDR write leveling control 3 (DDR_DDR_WRLVL_CNTL_3)	32	R/W	0000_0000h	8.4.26/314
2B20	DDR Debug Status Register 1 (DDR_DDRDSR_1)	32	R	0000_0000h	8.4.27/317
2B24	DDR Debug Status Register 2 (DDR_DDRDSR_2)	32	R	0000_0000h	8.4.28/318
2B28	DDR Control Driver Register 1 (DDR_DDRCDR_1)	32	R/W	0000_0000h	8.4.29/318
2B2C	DDR Control Driver Register 2 (DDR_DDRCDR_2)	32	R/W	0000_0000h	8.4.30/322
2BF8	DDR IP block revision 1 (DDR_DDR_IP_REV1)	32	R	See section	8.4.31/323

Table continues on the next page...

DDR memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2BFC	DDR IP block revision 2 (DDR_DDR_IP_REV2)	32	R	0000_0300h	8.4.32/323
2E00	Memory data path error injection mask high (DDR_DATA_ERR_INJECT_HI)	32	R/W	0000_0000h	8.4.33/324
2E04	Memory data path error injection mask low (DDR_DATA_ERR_INJECT_LO)	32	R/W	0000_0000h	8.4.34/324
2E08	Memory data path error injection mask ECC (DDR_ERR_INJECT)	32	R/W	0000_0000h	8.4.35/325
2E20	Memory data path read capture high (DDR_CAPTURE_DATA_HI)	32	R/W	0000_0000h	8.4.36/326
2E24	Memory data path read capture low (DDR_CAPTURE_DATA_LO)	32	R/W	0000_0000h	8.4.37/326
2E28	Memory data path read capture ECC (DDR_CAPTURE_ECC)	32	R/W	0000_0000h	8.4.38/327
2E40	Memory error detect (DDR_ERR_DETECT)	32	w1c	0000_0000h	8.4.39/328
2E44	Memory error disable (DDR_ERR_DISABLE)	32	R/W	0000_0000h	8.4.40/330
2E48	Memory error interrupt enable (DDR_ERR_INT_EN)	32	R/W	0000_0000h	8.4.41/332
2E4C	Memory error attributes capture (DDR_CAPTURE_ATTRIBUTES)	32	R/W	0000_0000h	8.4.42/334
2E50	Memory error address capture (DDR_CAPTURE_ADDRESS)	32	R/W	0000_0000h	8.4.43/336
2E54	Memory error extended address capture (DDR_CAPTURE_EXT_ADDRESS)	32	R/W	0000_0000h	8.4.44/336
2E58	Single-Bit ECC memory error management (DDR_ERR_SBE)	32	R/W	0000_0000h	8.4.45/337

8.4.1 Chip select n memory bounds (DDR_CS_n_BNDS)

The chip select bounds registers (CS_n_BNDS) define the starting and ending address of the memory space that corresponds to the individual chip selects. Note that the size specified in CS_n_BNDS should equal the size of physical DRAM. Also, note that EA_n must be greater than or equal to SA_n.

If chip select interleaving is enabled, all fields in the lower interleaved chip select are used, and the other chip selects' bounds registers are unused. For example, if chip selects 0 and 1 are interleaved, all fields in CS₀_BNDS are used, and all fields in CS₁_BNDS are unused.

Address: 2000h base + 0h offset + (8d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_CS_n_BNDS field descriptions

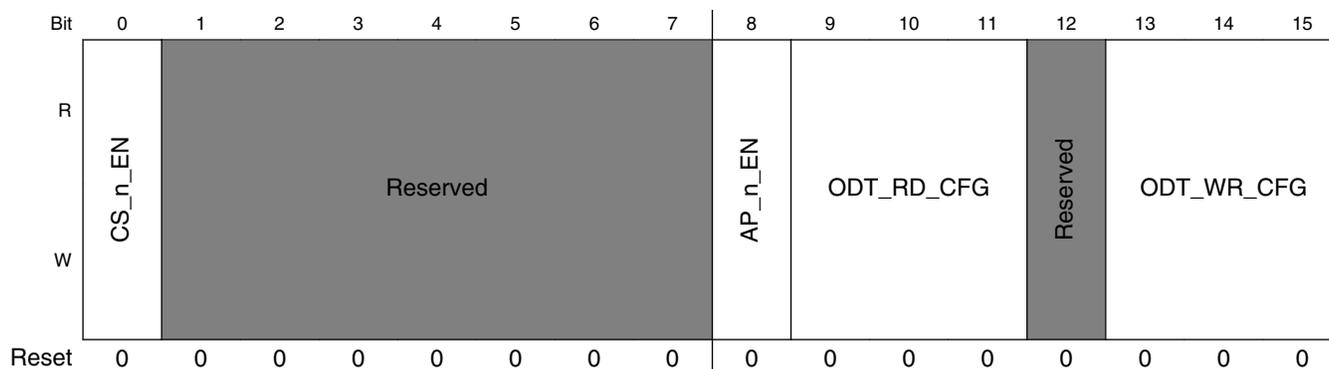
Field	Description
0–3 -	This field is reserved. Reserved
4–15 SAn	Starting address for chip select (bank) <i>n</i> . This value is compared against the 12 msbs of the 36-bit address.
16–19 -	This field is reserved. Reserved
20–31 EAn	Ending address for chip select (bank) <i>n</i> . This value is compared against the 12 msbs of the 36-bit address.

8.4.2 Chip select *n* configuration (DDR_CS_n_CONFIG)

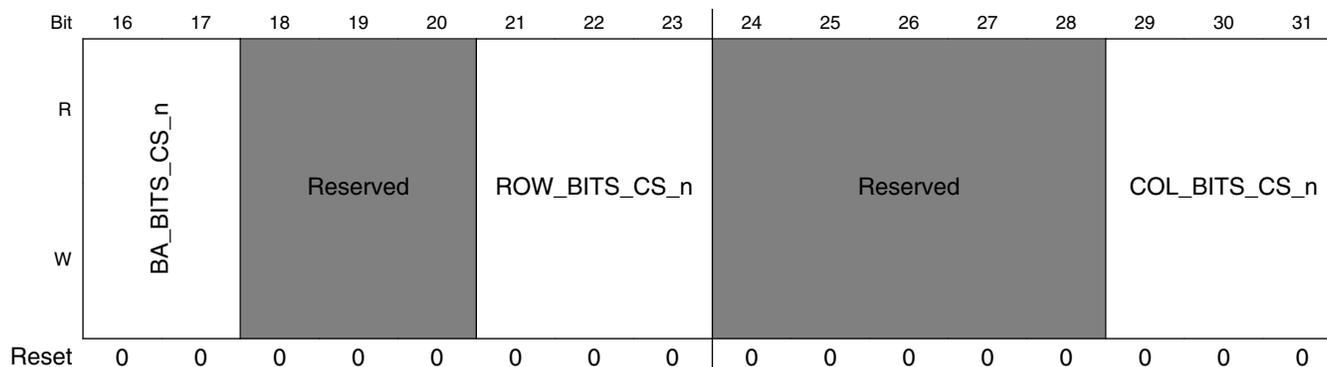
The chip select configuration (CS *n* _CONFIG) registers enable the DDR chip selects and set the number of row and column bits used for each chip select. These registers should be loaded with the correct number of row and column bits for each SDRAM. Because CS *n* _CONFIG[ROW_BITS_CS_ *n* , COL_BITS_CS_ *n*] establish address multiplexing, the user should take great care to set these values correctly.

If chip select interleaving is enabled, then all fields in the lower interleaved chip select are used, and the other registers' fields are unused, with the exception of the ODT_RD_CFG and ODT_WR_CFG fields. For example, if chip selects 0 and 1 are interleaved, all fields in CS0_CONFIG are used, but only the ODT_RD_CFG and ODT_WR_CFG fields in CS1_CONFIG are used.

Address: 2000h base + 80h offset + (4d × i), where i=0d to 3d



DDR memory map/register definition



DDR_CS_n_CONFIG field descriptions

Field	Description
0 CS_n_EN	Chip select <i>n</i> enable 0 Chip select <i>n</i> is not active 1 Chip select <i>n</i> is active and assumes the state set in CS _n _BNDS.
1–7 -	This field is reserved. Reserved
8 AP_n_EN	Chip select <i>n</i> auto-precharge enable 0 Chip select <i>n</i> is only auto-precharged if global auto-precharge mode is enabled (DDR_SDRAM_INTERVAL[BSTOPRE] = 0). 1 Chip select <i>n</i> always issues an auto-precharge for read and write transactions.
9–11 ODT_RD_CFG	ODT for reads configuration. Note that CAS latency plus additive latency must be at least 3 cycles for ODT_RD_CFG to be enabled. ODT should only be used with DDR2or DDR3 memories. 000 Never assert ODT for reads 001 Assert ODT only during reads to CS _n 010 Assert ODT only during reads to other chip selects 011 Assert ODT only during reads to other DIMM modules. It is assumed that CS0 and CS1 are on the same DIMM module, whereas CS2 and CS3 are on a separate DIMM module. 100 Assert ODT for all reads 101-111 Reserved
12 -	This field is reserved. Reserved
13–15 ODT_WR_CFG	ODT for writes configuration. Note that write latency plus additive latency must be at least 3 cycles for ODT_WR_CFG to be enabled. ODT should only be used with DDR2or DDR3 memories. 000 Never assert ODT for writes 001 Assert ODT only during writes to CS <i>n</i> 010 Assert ODT only during writes to other chip selects 011 Assert ODT only during writes to other DIMM modules. It is assumed that CS0 and CS1 are on the same DIMM module, whereas CS2 and CS3 are on a separate DIMM module. 100 Assert ODT for all writes 101-111 Reserved
16–17 BA_BITS_CS_n	Number of bank bits for SDRAM on chip select <i>n</i> . These bits correspond to the sub-bank bits driven on MBA <i>n</i> in DDR SDRAM address multiplexing . 00 2 logical bank bits

Table continues on the next page...

DDR_CS_n_CONFIG field descriptions (continued)

Field	Description
	01 3 logical bank bits 10-11 Reserved
18–20 -	This field is reserved. Reserved
21–23 ROW_BITS_CS_ n	Number of row bits for SDRAM on chip select <i>n</i> . See DDR SDRAM address multiplexing for details. 000 12 row bits 001 13 row bits 010 14 row bits 011 15 row bits 100 16 row bits 101-111 Reserved
24–28 -	This field is reserved. Reserved
29–31 COL_BITS_CS_n	Number of column bits for SDRAM on chip select <i>n</i> . For DDR, the decoding is as follows: 000 8 column bits 001 9 column bits 010 10 column bits 011 11 column bits 100-111 Reserved

8.4.3 Chip select n configuration 2 (DDR_CS_n_CONFIG_2)

The chip select configuration (CS_n_CONFIG_2) registers enable the partial array self refresh address decode in each chip select.

If chip select interleaving is enabled, then all fields in the lower interleaved chip select are used, and the other registers' fields are unused.

Address: 2000h base + C0h offset + (4d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W	Reserved				PASR_CFG		Reserved																									
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_CS_n_CONFIG_2 field descriptions

Field	Description
0–4 -	This field is reserved. Reserved
5–7 PASR_CFG	Partial array self refresh config. Controls the bits that are placed on MA[2:0] during the write to the EMRS(2) register when the automatic hardware DRAM initialization is used (DDR_SDRAM_CFG[BI] is cleared when DDR_SDRAM_CFG[MEM_EN] is set). If this field is a non-zero value, then it overrides the

Table continues on the next page...

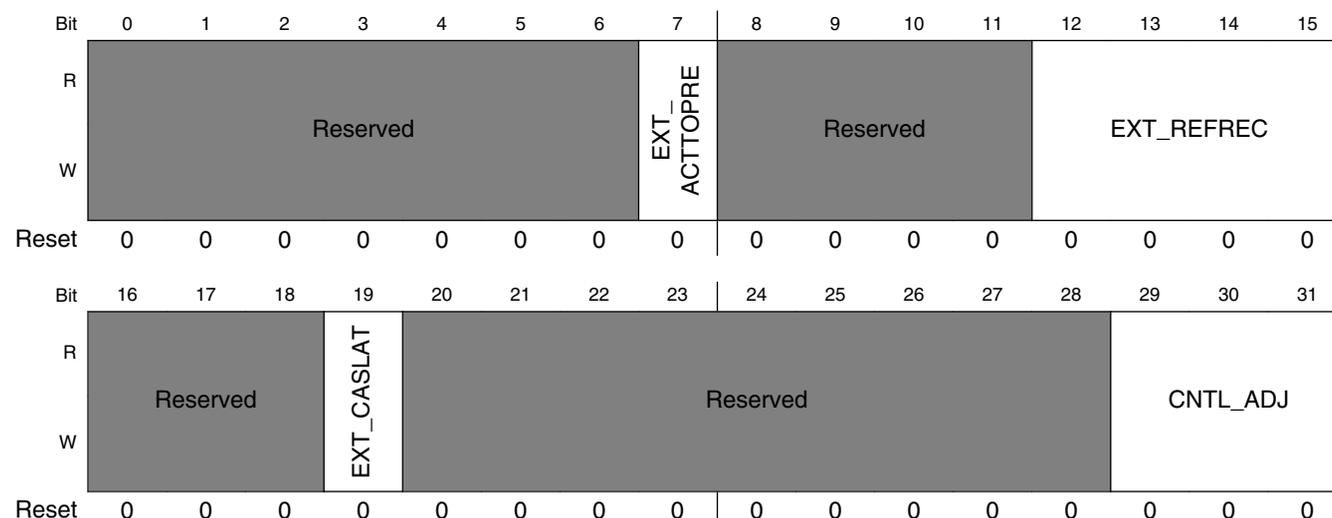
DDR_CS_n_CONFIG_2 field descriptions (continued)

Field	Description
	<p>least significant 3 bits in DDR_SDRAM_MODE_2[ESDMODE2] during the automatic initialization for chip select <i>n</i> .</p> <p>In addition, if a non-zero value is programmed in this field, then the address decode for chip select <i>n</i> is optimized for partial array self refresh, as shown in DDR SDRAM address multiplexing</p> <p>000 Partial array self refresh is disabled 001-111 Partial array self refresh is enabled per JEDEC specifications. Overriding the least significant 3 bits of EMRS or EMRS(2) is only supported for DDR2 and DDR3 memory types.</p>
8–31 -	This field is reserved. Reserved

8.4.4 DDR SDRAM timing configuration 3 (DDR_TIMING_CFG_3)

DDR SDRAM timing configuration register 3 sets the extended refresh recovery time, which is combined with TIMING_CFG_1[REFREC] to determine the full refresh recovery time.

Address: 2000h base + 100h offset = 2100h



DDR_TIMING_CFG_3 field descriptions

Field	Description
0–6 -	This field is reserved. Reserved
7 EXT_ACTTOPRE	Extended Activate to precharge interval (t_{RAS}). Determines the number of clock cycles from an activate command until a precharge command is allowed. This field is concatenated with TIMING_CFG_1[ACTTOPRE] to obtain a 5-bit value for the total activate to precharge. Note that a 5-bit value of 0_0000 is the same as a 5-bit value of 1_0000. Both values represent 16 cycles.

Table continues on the next page...

DDR_TIMING_CFG_3 field descriptions (continued)

Field	Description
	0 0 clocks 1 16 clocks
8–11 -	This field is reserved. Reserved
12–15 EXT_REFREC	<p>Extended refresh recovery time (t_{RFC}). Controls the number of clock cycles from a refresh command until an activate command is allowed. This field is concatenated with TIMING_CFG_1[REFREC] to obtain an 8-bit value for the total refresh recovery. Note that hardware adds an additional 8 clock cycles to the final, 8-bit value of the refresh recovery. $t_{RFC} = \{EXT_REFREC \parallel REFREC\} + 8$, such that t_{RFC} is calculated as follows:</p> <p>0000 0 clocks 0001 16 clocks 0010 32 clocks 0011 48 clocks 0100 64 clocks 0101 80 clocks 0110 96 clocks 0111 112 clocks 1000 128 clocks 1001 144 clocks 1010 160 clocks 1011 176 clocks 1100 192 clocks 1101 208 clocks 1110 224 clocks 1111 240 clocks</p>
16–18 -	This field is reserved. Reserved
19 EXT_CASLAT	<p>Extended MCAS latency from READ command. Number of clock cycles between registration of a READ command by the SDRAM and the availability of the first output data. If a READ command is registered at clock edge n and the latency is m clocks, data is available nominally coincident with clock edge $n + m$. This field is concatenated with TIMING_CFG_1[CASLAT] to obtain a 5-bit value for the total CAS latency. Note that if this bit is set, then 8 clocks are added to the programmed value in TIMING_CFG_1[CASLAT].</p> <p>0 0 clocks 1 8 clocks</p>
20–28 -	This field is reserved. Reserved
29–31 CNTL_ADJ	<p>Control Adjust. Controls the amount of delay to add to the lightly loaded control signals with respect to all other DRAM address and command signals. The signals affected by this field are MODT[0:3], \overline{MCS}[0:3], and MCKE[0:3]</p> <p>000 MODT[0:3], \overline{MCS} [0:3], and MCKE[0:3] are launched aligned with the other DRAM address and control signals. 001 MODT[0:3], \overline{MCS} [0:3], and MCKE[0:3] are launched 1/4 DRAM cycle later than the other DRAM address and control signals. 010 MODT[0:3], \overline{MCS} [0:3], and MCKE[0:3] are launched 1/2 DRAM cycle later than the other DRAM address and control signals.</p>

Table continues on the next page...

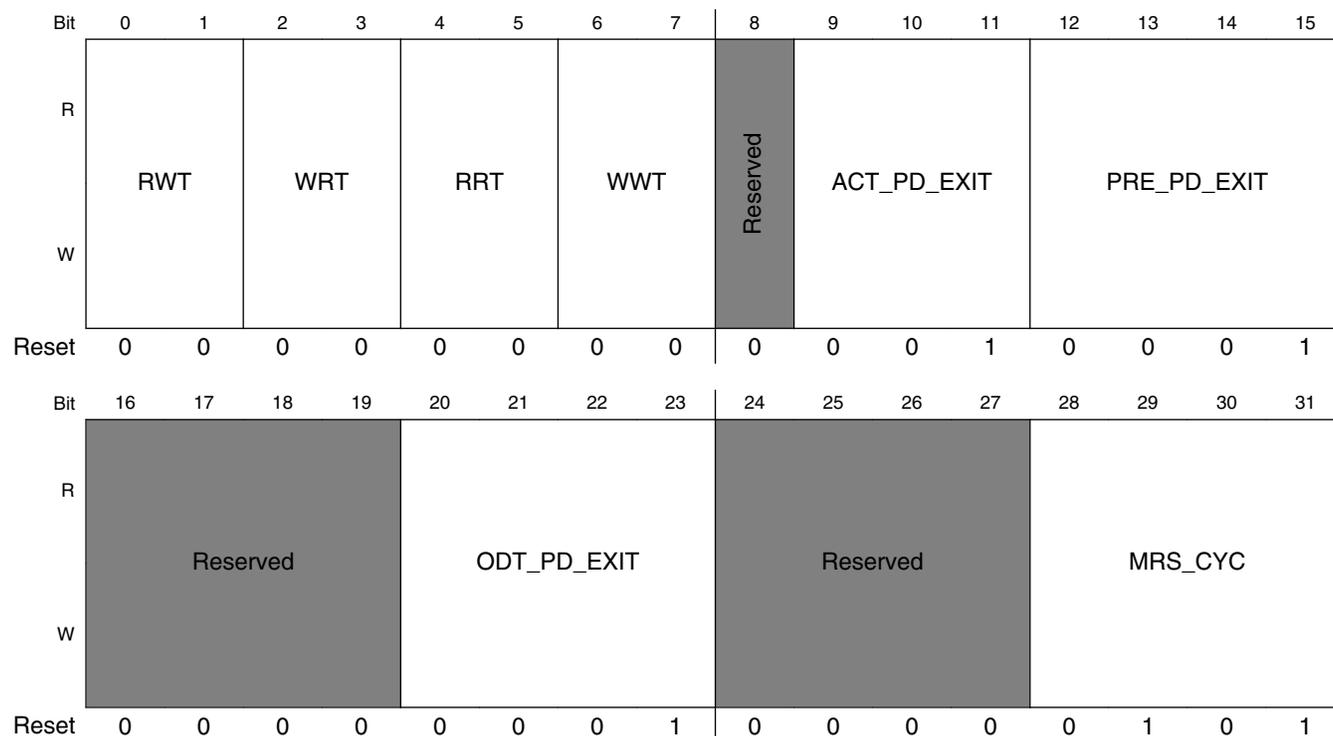
DDR_TIMING_CFG_3 field descriptions (continued)

Field	Description
011	MODT[0:3], \overline{MCS} [0:3], and MCKE[0:3] are launched 3/4 DRAM cycles later than the other DRAM address and control signals.
100	MODT[0:3], \overline{MCS} [0:3], and MCKE[0:3] are launched 1 DRAM cycles later than the other DRAM address and control signals.
101	MODT[0:3], \overline{MCS} [0:3], and MCKE[0:3] are launched 5/4 DRAM cycles later than the other DRAM address and control signals.
110-111	Reserved

8.4.5 DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)

DDR SDRAM timing configuration register 0 sets the number of clock cycles between various SDRAM control commands.

Address: 2000h base + 104h offset = 2104h



DDR_TIMING_CFG_0 field descriptions

Field	Description
0-1 RWT	Read-to-write turnaround ($\overline{t_{RTW}}$). Specifies how many extra cycles are added between a read to write turnaround. If 0 clocks is chosen, then the DDR controller uses a fixed number based on the CAS latency and write latency. Choosing a value other than 0 adds extra cycles past this default calculation. As a default the DDR controller determines the read-to-write turnaround as $CL - WL + BL \div 2 + 2$. In this

Table continues on the next page...

DDR_TIMING_CFG_0 field descriptions (continued)

Field	Description
	<p>equation, CL is the CAS latency rounded up to the next integer, WL is the programmed write latency, and BL is the burst length.</p> <p>00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks</p>
2–3 WRT	<p>Write-to-read turnaround. Specifies how many extra cycles are added between a write to read turnaround. If 0 clocks is chosen, then the DDR controller uses a fixed number based on the, read latency, and write latency. Choosing a value other than 0 adds extra cycles past this default calculation. As a default, the DDR controller determines the write-to-read turnaround as $WL - CL + BL \div 2 + 1$. In this equation, CL is the CAS latency rounded down to the next integer, WL is the programmed write latency, and BL is the burst length.</p> <p>00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks</p>
4–5 RRT	<p>Read-to-read turnaround. Specifies how many extra cycles are added between reads to different chip selects. As a default, 3 cycles are required between read commands to different chip selects. Extra cycles may be added with this field.</p> <p>NOTE: If 8-beat bursts are enabled, then 5 cycles are the default. Note that DDR2 does not support 8-beat bursts.</p> <p>00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks</p>
6–7 WWT	<p>Write-to-write turnaround. Specifies how many extra cycles are added between writes to different chip selects. As a default, 2 cycles are required between write commands to different chip selects. Extra cycles may be added with this field.</p> <p>NOTE: If 8-beat bursts are enabled, then 4 cycles are the default. Note that DDR2 does not support 8-beat bursts.</p> <p>00 0 clocks 01 1 clock 10 2 clocks 11 3 clocks</p>
8 -	<p>This field is reserved. Reserved</p>
9–11 ACT_PD_EXIT	<p>Active powerdown exit timing (DDR2: t_{XARD} and t_{XARDS} and DDR3 : t_{XP}). Specifies how many clock cycles to wait after exiting active powerdown before issuing any command.</p> <p>000 Reserved 001 2 clocks 010 3 clocks 011 4 clocks 100 5 clocks 101 6 clocks</p>

Table continues on the next page...

DDR_TIMING_CFG_0 field descriptions (continued)

Field	Description
	110 7 clocks 111 8 clocks
12–15 PRE_PD_EXIT	Precharge powerdown exit timing (t_{XP}). Specifies how many clock cycles to wait after exiting precharge powerdown before issuing any command. 0000 Reserved 0001 2 clocks 0010 3 clocks 0011 4 clocks 0100 5 clocks 0101 6 clocks 0110 7 clocks 0111 8 clocks 1000 9 clocks 1001 10 clocks 1010 11 clocks 1011 12 clocks 1100 13 clocks 1101 14 clocks 1110 15 clocks 1111 16 clocks
16–19 -	This field is reserved. Reserved
20–23 ODT_PD_EXIT	ODT powerdown exit timing (DDR2: t_{AXPD} , DDR3 : set to 1). Specifies how many clocks must pass after exiting powerdown before ODT may be asserted. ODT_PD_EXIT must be greater than TIMING_CFG_5[RODT_ON] when using RODT_ON overrides and must be greater than TIMING_CFG_5[WODT_ON] when using WODT_ON overrides. 0000 0 clock 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks
24–27 -	This field is reserved. Reserved

Table continues on the next page...

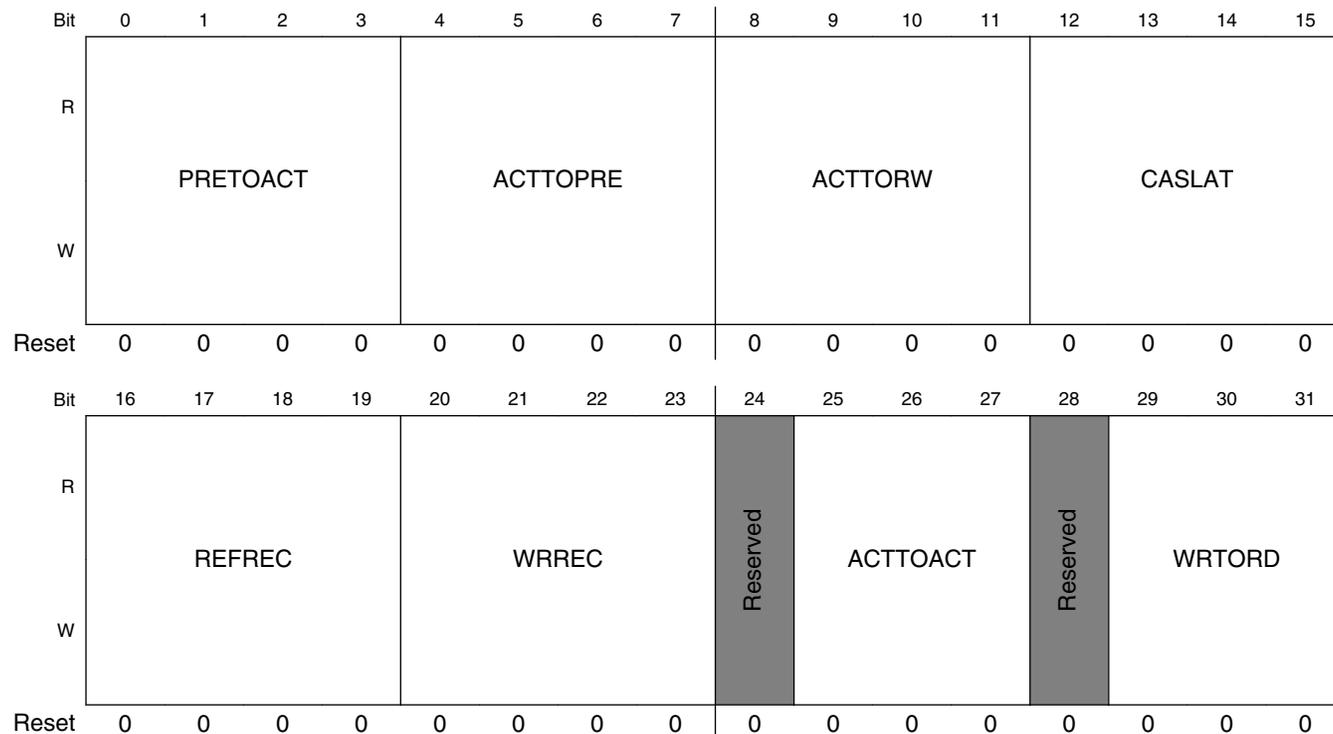
DDR_TIMING_CFG_0 field descriptions (continued)

Field	Description
28–31 MRS_CYC	Mode register set cycle time (t_{MRD}). Specifies the number of cycles that must pass after a Mode Register Set command until any other command. 0000 Reserved 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks

8.4.6 DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)

DDR SDRAM timing configuration register 1 sets the number of clock cycles between various SDRAM control commands.

Address: 2000h base + 108h offset = 2108h



DDR_TIMING_CFG_1 field descriptions

Field	Description
0-3 PRETOACT	Precharge-to-activate interval (t_{RP}). Determines the number of clock cycles from a precharge command until an activate or refresh command is allowed.
0000	Reserved
0001	1 clock
0010	2 clocks
0011	3 clocks
0100	4 clocks
0101	5 clocks
0110	6 clocks
0111	7 clocks
1000	8 clocks
1001	9 clocks
1010	10 clocks
1011	11 clocks

Table continues on the next page...

DDR_TIMING_CFG_1 field descriptions (continued)

Field	Description
	1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks
4–7 ACTTOPRE	Activate to precharge interval (t_{RAS}). Determines the number of clock cycles from an activate command until a precharge command is allowed. This field is concatenated with TIMING_CFG_3[EXT_ACTTOPRE] to obtain a 5-bit value for the total activate to precharge time. Note that the decode of 0000-0011 is equal to 16-19 clocks when TIMING_CFG_3[EXT_ACTTOPRE] = 0, but it is equal to 0-3 clocks when TIMING_CFG_3[EXT_ACTTOPRE] = 1. 0000 16 clocks 0001 17 clocks 0010 18 clocks 0011 19 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks ... 1111 15 clocks
8–11 ACTTORW	Activate to read/write interval for SDRAM (t_{RCD}). Controls the number of clock cycles from an activate command until a read or write command is allowed. 0000 Reserved 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks
12–15 CASLAT	MCAS latency from READ command. Number of clock cycles between registration of a READ command by the SDRAM and the availability of the first output data. If a READ command is registered at clock edge n and the latency is m clocks, data is available nominally coincident with clock edge $n + m$. This field is concatenated with TIMING_CFG_3[EXT_CASLAT] to obtain a 5-bit value for the total CAS latency. This value must be programmed at initialization as described in DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2) . 0000 Reserved 0001 1 clock 0010 1.5 clocks

Table continues on the next page...

DDR_TIMING_CFG_1 field descriptions (continued)

Field	Description
	0011 2 clocks 0100 2.5 clocks 0101 3 clocks 0110 3.5 clocks 0111 4 clocks 1000 4.5 clocks 1001 5 clocks 1010 5.5 clocks 1011 6 clocks 1100 6.5 clocks 1101 7 clocks 1110 7.5 clocks 1111 8 clocks
16–19 REFREC	Refresh recovery time (t_{RFC}). Controls the number of clock cycles from a refresh command until an activate command is allowed. This field is concatenated with TIMING_CFG_3[EXTREFREC] to obtain a 7-bit value for the total refresh recovery. Note that hardware adds an additional 8 clock cycles to the final, 7-bit value of the refresh recovery, such that t_{RFC} is calculated as follows: $t_{RFC} = \{EXT_REFREC \parallel REFREC\} + 8$. 0000 8 clocks 0001 9 clocks 0010 10 clocks 0011 11 clocks ... 1111 23 clocks
20–23 WRREC	Last data to precharge minimum interval (t_{WR}). Determines the number of clock cycles from the last data associated with a write command until a precharge command is allowed. If DDR_SDRAM_CFG_2[OBC_CFG] is set, then this field needs to be programmed to ($t_{WR} + 2$ cycles). 0000 Reserved 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks
24 -	This field is reserved. Reserved

Table continues on the next page...

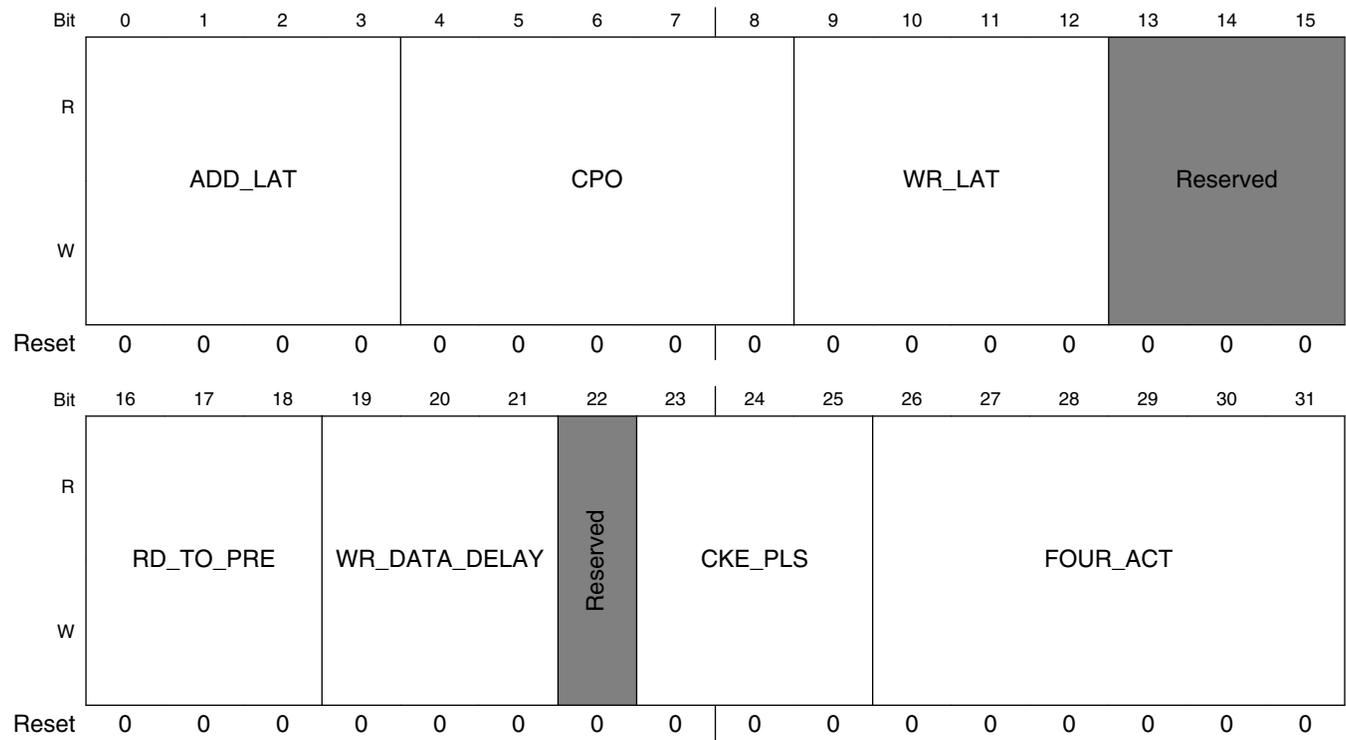
DDR_TIMING_CFG_1 field descriptions (continued)

Field	Description
25–27 ACTTOACT	Activate-to-activate interval (t_{RRD}). Number of clock cycles from an activate command until another activate command is allowed for a different logical bank in the same physical bank (chip select). 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
28 -	This field is reserved. Reserved
29–31 WRTORD	Last write data pair to read command issue interval (t_{WTR}). Number of clock cycles between the last write data pair and the subsequent read command to the same physical bank. If <code>DDR_SDRAM_CFG_2[OBC_CFG]</code> is set, then this field needs to be programmed to ($t_{WTR} + 2$ cycles). 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks

8.4.7 DDR SDRAM timing configuration 2 (DDR_TIMING_CFG_2)

DDR SDRAM timing configuration 2 sets the clock delay to data for writes.

Address: 2000h base + 10Ch offset = 210Ch



DDR_TIMING_CFG_2 field descriptions

Field	Description
0-3 ADD_LAT	Additive latency. The additive latency must be set to a value less than TIMING_CFG_1[ACTTORW]. (DDR2-specific) 0000 0 clocks 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks
4-8 CPO	MCAS -to-preamble override. Defines the number of DRAM cycles between when a read is issued and when the corresponding DQS preamble is valid for the memory controller. For these decodings, "READ_LAT" is equal to the CAS latency plus the additive latency. 00000 READ_LAT + 1 00001 Reserved 00010 READ_LAT 00011 READ_LAT + 1/4 00100 READ_LAT + 1/2 00101 READ_LAT + 3/4 00110 READ_LAT + 1 00111 READ_LAT + 5/4 01000 READ_LAT + 3/2 01001 READ_LAT + 7/4 01010 READ_LAT + 2 01011 READ_LAT + 9/4 01100 READ_LAT + 5/2 01101 READ_LAT + 11/4 01110 READ_LAT + 3 01111 READ_LAT + 13/4 10000 READ_LAT + 7/2 10001 READ_LAT + 15/4 10010 READ_LAT + 4 10011 READ_LAT + 17/4 10100 READ_LAT + 9/2 10101 READ_LAT + 19/4 10110 READ_LAT + 5 10111 READ_LAT + 21/4 11000 READ_LAT + 11/2

Table continues on the next page...

DDR_TIMING_CFG_2 field descriptions (continued)

Field	Description
	11001 READ_LAT + 23/4 11010 READ_LAT + 6 11011 READ_LAT + 25/4 11100 READ_LAT + 13/2 11101 READ_LAT + 27/4 11110 READ_LAT + 7 11111 Automatic Calibration (recommended)
9–12 WR_LAT	Write latency. Note that the total write latency for DDR2/ DDR3 is equal to WR_LAT + ADD_LAT. If a write latency of 1 is desired, then the additive latency must also be set to at least 1 cycle. 0000 Reserved 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks
13–15 -	This field is reserved. Reserved
16–18 RD_TO_PRE	Read to precharge (t_{RTP}). For DDR2, with a non-zero ADD_LAT value, takes a minimum of ADD_LAT + t_{RTP} cycles between read and precharge. If DDR_SDRAM_CFG_2[OBC_CFG] is set, then this field needs to be programmed to ($t_{RTP} + 2$ cycles). 000 Reserved 001 1 cycle 010 2 cycles 011 3 cycles 100 4 cycles 101 5 cycles 110 6 cycles 111 7 cycles
19–21 WR_DATA_DELAY	Write command to write data strobe timing adjustment. Controls the amount of delay applied to the data and data strobes for writes. See DDR SDRAM write timing adjustments , for details. The write preamble typically is driven high for 1/2 DRAM cycle, and then it is driven low for 1/2 DRAM cycle. However, for WR_DATA_DELAY settings of 0 clocks and 1/4 clocks, the write preamble is driven low for the entire DRAM cycle. If the preamble needs to switch high first (to meet DDR3 specifications), then these values should not be used. 000 0 clock delay

Table continues on the next page...

DDR_TIMING_CFG_2 field descriptions (continued)

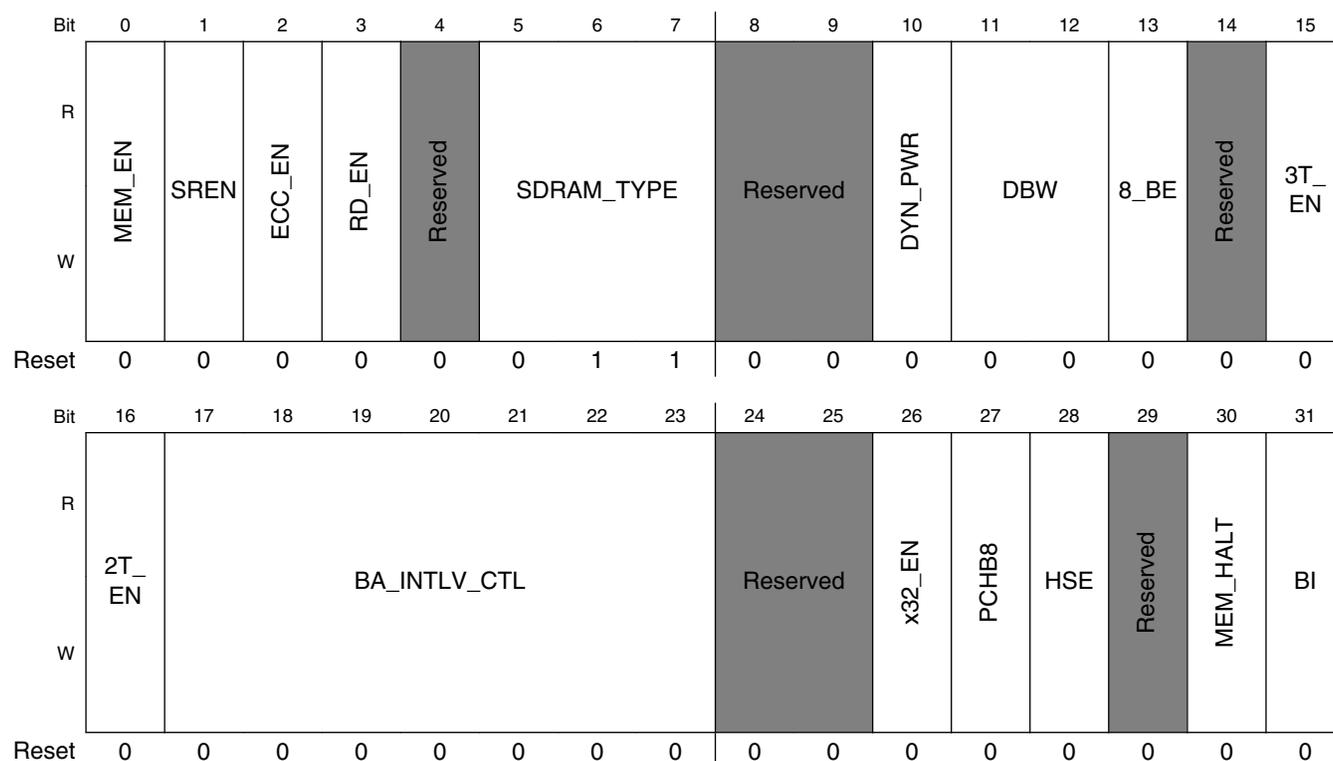
Field	Description
	001 1/4 clock delay 010 1/2 clock delay 011 3/4 clock delay 100 1 clock delay 101 5/4 clock delay 110 3/2 clock delay 111 Reserved
22 -	This field is reserved. Reserved
23–25 CKE_PLS	Minimum CKE pulse width (t_{CKE}). 000 Reserved 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
26–31 FOUR_ACT	Window for four activates (t_{FAW}). This is applied to DDR2/ DDR3 with eight logical banks only. 000000 Reserved 000001 1 cycle 000010 2 cycles 000011 3 cycles 000100 4 cycles ... 011110 30 cycles 011111 31 cycles 100000 32 cycles 100001-111111 Reserved

- For CPO decodings other than 00000 and 11111, READ_LAT is rounded up to the next integer value.

8.4.8 DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG)

The DDR SDRAM control configuration register enables the interface logic and specifies certain operating features such as self refreshing, error checking and correcting, registered DIMMs, and dynamic power management.

Address: 2000h base + 110h offset = 2110h



DDR_DDR_SDRAM_CFG field descriptions

Field	Description
0 MEM_EN	DDR SDRAM interface logic enable. 0 SDRAM interface logic is disabled. 1 SDRAM interface logic is enabled. Must not be set until all other memory configuration parameters have been appropriately configured by initialization code.
1 SREN	Self refresh enable (during sleep). 0 SDRAM self refresh is disabled during sleep. Whenever self-refresh is disabled, the system is responsible for preserving the integrity of SDRAM during sleep. 1 SDRAM self refresh is enabled during sleep.
2 ECC_EN	ECC enable. Note that non-correctable read errors may cause the assertion of <i>core_fault_in</i> , which causes the core to generate a machine check interrupt unless it is disabled (by clearing HID1[RFXE]). If

Table continues on the next page...

DDR_DDR_SDRAM_CFG field descriptions (continued)

Field	Description
	RFXE is zero and this error occurs, ERR_DISABLE[MBED] must be zero, and ERR_INT_EN[MBEE] and ECC_EN must be one to ensure an interrupt is generated. 0 No ECC errors are reported. No ECC interrupts are generated. 1 ECC is enabled.
3 RD_EN	Registered DIMM enable. Specifies the type of DIMM used in the system. NOTE: RD_EN and 2T_EN must not both be set at the same time. 0 Indicates unbuffered DIMMs. 1 Indicates registered DIMMs.
4 -	This field is reserved. Reserved
5–7 SDRAM_TYPE	Type of SDRAM device to be used. This field is used when issuing the automatic hardware initialization sequence to DRAM through Mode Register Set and Extended Mode Register Set commands. Patterns not shown are reserved. 011 DDR2 SDRAM 111 DDR3 SDRAM
8–9 -	This field is reserved. Reserved
10 DYN_PWR	Dynamic power management mode 0 Dynamic power management mode is disabled. 1 Dynamic power management mode is enabled. If there is no ongoing memory activity, the SDRAM CKE signal is negated.
11–12 DBW	DRAM data bus width. 00 64-bit bus is used 01 32-bit bus is used 10 Reserved 11 Reserved
13 8_BE	8-beat burst enable. NOTE: DDR2 must use 4-beat bursts when using 32/64-bit bus mode NOTE: DDR3 must use 8-beat bursts when using 32-bit bus mode and 4-beat bursts when using 64-bit bus mode 0 4-beat bursts are used on the DRAM interface. 1 8-beat bursts are used on the DRAM interface.
14 -	This field is reserved. Reserved
15 3T_EN	Enable 3T timing. This field cannot be set if DDR_SDRAM_CFG[2T_EN] is also set. This field cannot be used with a 32-bit bus if 4-beat bursts are used. 0 1T timing is enabled if 2T_EN is cleared. The DRAM command/address are held for only 1 cycle on the DRAM bus. 1 3T timing is enabled. The DRAM command/address are held for 3 full cycles on the DRAM bus for every DRAM transaction. However, the chip select is only held for the third cycle.

Table continues on the next page...

DDR_DDR_SDRAM_CFG field descriptions (continued)

Field	Description
16 2T_EN	<p>Enable 2T timing. This field should not be set if DDR_SDRAM_CFG[3T_EN] is set.</p> <p>NOTE: RD_EN and 2T_EN must not both be set at the same time.</p> <p>0 1T timing is enabled if 3T_EN is cleared. The DRAM command/address is held for only 1 cycle on the DRAM bus.</p> <p>1 2T timing is enabled. The DRAM command/address are held for 2 full cycles on the DRAM bus for every DRAM transaction. However, the chip select is only held for the second cycle.</p>
17–23 BA_INTLV_CTL	<p>Bank (chip select) interleaving control. Set this field only if you wish to use bank interleaving.</p> <p>(' x ' denotes a don't care bit value. All unlisted field values are reserved.)</p> <p>xx 00100 External memory banks 0 through 3 are all interleaved together</p> <p>0000000 No external memory banks are interleaved</p> <p>1000000 External memory banks 0 and 1 are interleaved</p> <p>0100000 External memory banks 2 and 3 are interleaved</p> <p>1100000 External memory banks 0 and 1 are interleaved together and banks 2 and 3 are interleaved together</p>
24–25 -	<p>This field is reserved.</p> <p>Reserved</p>
26 x32_EN	<p>x32 enable.</p> <p>0 Either x8 or x16 discrete DRAM chips are used. In this mode, each data byte has a dedicated corresponding data strobe.</p> <p>1 x32 discrete DRAM chips are used. In this mode, DQS0 is used to capture DQ[0:31], DQS4 is used to capture DQ[32:63] and DQS8 is used to capture ECC[0:7].</p>
27 PCHB8	<p>Precharge bit 8 enable.</p> <p>0 MA[10] is used to indicate the auto-precharge and precharge all commands.</p> <p>1 MA[8] is used to indicate the auto-precharge and precharge all commands.</p> <p>If x32_EN is cleared, then PCHB8 should be cleared as well .</p>
28 HSE	<p>Global half-strength override.</p> <p>Sets I/O driver impedance to half strength. This impedance is used by the MDIC, address/command, data, and clock impedance values, but only if automatic hardware calibration is disabled and the corresponding group's software override is disabled in the DDR control driver register(s) described in DDR Control Driver Register 1 (DDR_DDRCDR_1) . This bit should be cleared if using automatic hardware calibration.</p> <p>0 I/O driver impedance is configured to full strength.</p> <p>1 I/O driver impedance is configured to half strength.</p>
29 -	<p>This field is reserved.</p> <p>Reserved</p>
30 MEM_HALT	<p>DDR memory controller halt. When this bit is set, the memory controller does not accept any new data read/write transactions to DDR SDRAM until the bit is cleared again. This can be used when bypassing initialization and forcing MODE REGISTER SET commands through software.</p> <p>0 DDR controller accepts new transactions.</p> <p>1 DDR controller finishes any remaining transactions, and then it remains halted until this bit is cleared by software.</p>

Table continues on the next page...

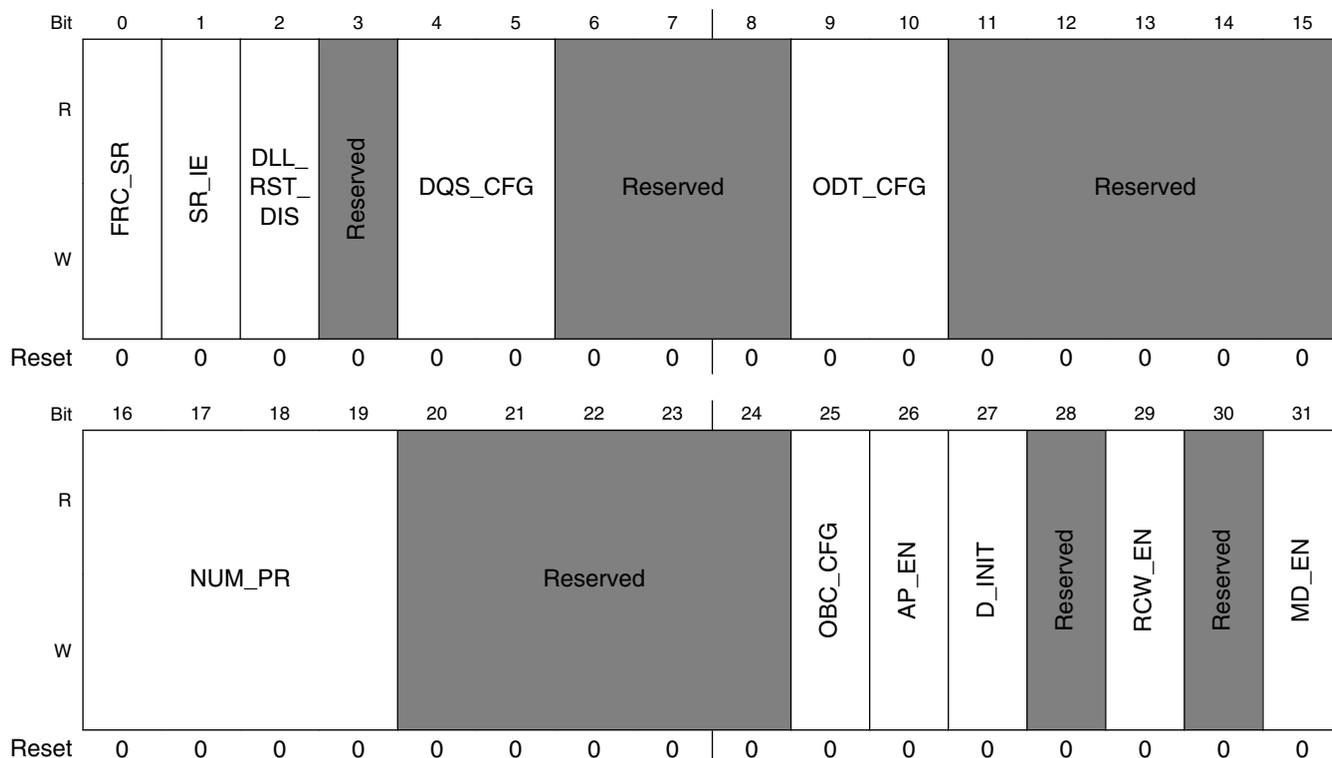
DDR_DDR_SDRAM_CFG field descriptions (continued)

Field	Description
31 BI	Bypass initialization
0	DDR controller cycles through initialization routine based on SDRAM_TYPE
1	Initialization routine is bypassed. Software is responsible for initializing memory through DDR_SDRAM_MODE2 register. If software is initializing memory, then the MEM_HALT bit can be set to prevent the DDR controller from issuing transactions during the initialization sequence. Note that the DDR controller does not issue a DLL reset to the DRAMs when bypassing the initialization routine, regardless of the value of DDR_SDRAM_CFG[DLL_RST_DIS]. If a DLL reset is required, then the controller should be forced to enter and exit self-refresh after the controller is enabled. See DDR training initialization address (DDR_DDR_INIT_ADDR) for details on avoiding ECC errors in this mode.

8.4.9 DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2)

The DDR SDRAM control configuration register 2 provides more control configuration for the DDR controller.

Address: 2000h base + 114h offset = 2114h



DDR_DDR_SDRAM_CFG_2 field descriptions

Field	Description
0 FRC_SR	Force self-refresh 0 DDR controller operates in normal mode. 1 DDR controller enters self-refresh mode.
1 SR_IE	Self-refresh interrupt enable The DDR controller can be placed into self refresh mode by forcing the PIC to assert $\overline{IRQ_OUT}$. This is considered a 'panic interrupt' by the DDR controller, and it enters self refresh as soon as possible. DDR_SDRAM_CFG[SREN] must also be set if the panic interrupt is used. 0 DDR controller does not enter self-refresh mode if panic interrupt is asserted. 1 DDR controller enters self-refresh mode if panic interrupt is asserted.
2 DLL_RST_DIS	DLL reset disable. If self-refresh is to be used with DDR3 , this bit should be set The DDR controller typically issues a DLL reset to the DRAMs when exiting self refresh. However, this function may be disabled by setting this bit during initialization. 0 DDR controller issues a DLL reset to the DRAMs when exiting self refresh. 1 DDR controller does not issue a DLL reset to the DRAMs when exiting self refresh.
3 -	This field is reserved. Reserved
4-5 DQS_CFG	DQS configuration 00 Reserved 01 Differential DQS signals are used. 10 Reserved 11 Reserved
6-8 -	This field is reserved. Reserved
9-10 ODT_CFG	ODT configuration This field defines how ODT is driven to the on-chip IOs. See DDR Control Driver Register 1 (DDR_DDRCDR_1) , which defines the termination value that is used. 00 Never assert ODT to internal IOs 01 Assert ODT to internal IOs only during writes to DRAM 10 Assert ODT to internal IOs only during reads to DRAM 11 Always keep ODT asserted to internal IOs
11-15 -	This field is reserved. Reserved.
16-19 NUM_PR	Number of posted refreshes This determines how many posted refreshes, if any, can be issued at one time. Note that if posted refreshes are used, then this field, along with DDR_SDRAM_INTERVAL[REFINT], must be programmed such that the maximum t_{RAS} specification cannot be violated. 0000 Reserved 0001 1 refresh is issued at a time 0010 2 refreshes is issued at a time 0011 3 refreshes is issued at a time

Table continues on the next page...

DDR_DDR_SDRAM_CFG_2 field descriptions (continued)

Field	Description
	... 1000 8 refreshes is issued at a time 1001-1111 Reserved
20-24 -	This field is reserved. Reserved
25 OBC_CFG	On-The-Fly Burst Chop Configuration Determines whether on-the-fly Burst Chop is used. This bit should only be set if DDR3 memories are used. If on-the-fly Burst Chop mode is not used with DDR3 memories, then fixed Burst Chop mode may be used if the proper turnaround times are programmed into TIMING_CFG_0 and TIMING_CFG_4. DDR_SDRAM_CFG[8_BE] should be cleared for both on-the-fly Burst Chop mode or fixed Burst Chop mode when using a 64-bit data bus with DDR3 memories. 0 On-the-fly Burst Chop mode is disabled. Fixed burst lengths as defined in DDR_SDRAM_CFG[8_BE] are used. If fixed Burst Chop is used (with DDR3 memories), then DDR_SDRAM_CFG[8_BE] should be cleared. 1 On-the-fly Burst Chop mode is used. DDR_SDRAM_CFG[8_BE] should be cleared for on-the-fly Burst Chop mode. DDR_SDRAM_CFG[DBW] should also be cleared for on-the-fly Burst Chop mode
26 AP_EN	Address Parity Enable Determines whether address parity is generated and checked for the address and control signals when using registered DIMMs. If address parity is used, the MAPAR_OUT and MAPAR_ERR pins are used to drive the parity bit and to receive errors from the open-drain parity error signal. Even parity is used, and parity is generated for the MA[15:0], MBA[2:0], $\overline{M}RAS$, $\overline{M}CAS$, $\overline{M}WE$ signals. Parity does not generate for the MCKE[0:3], MODT[0:3], or $\overline{M}CS$ [0:3] signals. Note that address parity should not be used for non-zero values of TIMING_CFG_3[CNTL_ADJ]. 0 Address parity is not used 1 Address parity is used
27 D_INIT	DRAM data initialization This bit is set by software, and it is cleared by hardware. If software sets this bit before the memory controller is enabled, the controller automatically initializes DRAM after it is enabled. This bit is automatically cleared by hardware once the initialization is completed. This data initialization bit should only be set when the controller is idle. 0 There is not data initialization in progress, and no data initialization is scheduled 1 The memory controller initializes memory once it is enabled. This bit remains asserted until the initialization is complete. The value in DDR_DATA_INIT register is used to initialize memory.
28 -	This field is reserved. Reserved
29 RCW_EN	Register Control Word Enable If DDR3 registered DIMMs are used, it may be necessary to write the register control words before issuing commands to DRAM. If this bit is set, the controller will write the register control words after DDR_SDRAM_CFG[MEM_EN] is set, unless DDR_SDRAM_CFG[BI] is set. The register control words are written with the values in DDR_SDRAM_RCW_1 and DDR_SDRAM_RCW_2. 0 Register control words will not be automatically written during DRAM initialization 1 Register control words are automatically written during DRAM initialization. This bit should only be set if DDR3 registered DIMMs are used, and the default settings need to be modified.
30 -	This field is reserved. Reserved

Table continues on the next page...

DDR_DDR_SDRAM_CFG_2 field descriptions (continued)

Field	Description
31 MD_EN	<p>Mirrored DIMM Enable</p> <p>Some DDR3 DIMMs are mirrored, where certain MA and MBA pins are mirrored on one side of the DIMM. When this bit is set, the controller will know to swap these signals before transmitting to the DRAM. The controller will assume that CS1 and CS3 are the 'mirrored' ranks of memory. The following signals are mirrored (MBA[0] vs MBA[1]; MA[3] vs MA[4]; MA[5] vs MA[6]; MA[7] vs MA[8]).</p> <p>0 Mirrored DIMMs are not used 1 Mirrored DIMMs are used</p>

8.4.10 DDR SDRAM mode configuration (DDR_DDR_SDRAM_MODE)

The DDR SDRAM mode configuration register sets the values loaded into the DDR's mode registers.

Address: 2000h base + 118h offset = 2118h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDR_SDRAM_MODE field descriptions

Field	Description
0–15 ESDMODE	<p>Extended SDRAM mode</p> <p>Specifies the initial value loaded into the DDR SDRAM extended mode register. The range and meaning of legal values is specified by the DDR SDRAM manufacturer.</p> <p>When this value is driven onto the address bus (during the DDR SDRAM initialization sequence), MA[0] presents the lsb of ESDMODE, which, in the big-endian convention shown in DDR SDRAM mode configuration (DDR_DDR_SDRAM_MODE), corresponds to ESDMODE[15]. The msb of the SDRAM extended mode register value must be stored at ESDMODE[0]. The value programmed into this field is also used for writing MR1 during write leveling for DDR3, although the bits specifically related to the write leveling scheme are handled automatically by the DDR controller. Even if DDR_SDRAM_CFG[Bi] is set, this field is still used during write leveling. The write leveling enable bit should be cleared by software in this field.</p>
16–31 SDMODE	<p>SDRAM mode</p> <p>Specifies the initial value loaded into the DDR SDRAM mode register. The range of legal values is specified by the DDR SDRAM manufacturer.</p> <p>When this value is driven onto the address bus (during DDR SDRAM initialization), MA[0] presents the lsb of SDMODE, which, in the big-endian convention shown in DDR SDRAM mode configuration (DDR_DDR_SDRAM_MODE), corresponds to SDMODE[15]. The msb of the SDRAM mode register value must be stored at SDMODE[0]. Because the memory controller forces SDMODE[7] to certain values depending on the state of the initialization sequence, (for resetting the SDRAM's DLL) the corresponding bits of this field are ignored by the memory controller. Note that SDMODE[7] is mapped to MA[8].</p>

8.4.11 DDR SDRAM mode configuration 2 (DDR_DDR_SDRAM_MODE_2)

The DDR SDRAM mode 2 configuration register sets the values loaded into the DDR's extended mode 2 and 3 registers (for DDR2 and DDR3).

Address: 2000h base + 11Ch offset = 211Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	ESDMODE2															ESDMODE3																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDR_SDRAM_MODE_2 field descriptions

Field	Description
0–15 ESDMODE2	<p>Extended SDRAM mode 2</p> <p>Specifies the initial value loaded into the DDR SDRAM extended 2 mode register. The range and meaning of legal values is specified by the DDR SDRAM manufacturer.</p> <p>When this value is driven onto the address bus (during the DDR SDRAM initialization sequence), MA[0] presents the lsb bit of ESDMODE2, which, in the big-endian convention shown in DDR SDRAM mode configuration 2 (DDR_DDR_SDRAM_MODE_2), corresponds to ESDMODE2[15]. The msb of the SDRAM extended mode 2 register value must be stored at ESDMODE2[0].</p>
16–31 ESDMODE3	<p>Extended SDRAM mode 3</p> <p>Specifies the initial value loaded into the DDR SDRAM extended 3 mode register. The range of legal values of legal values is specified by the DDR SDRAM manufacturer.</p> <p>When this value is driven onto the address bus (during DDR SDRAM initialization), MA[0] presents the lsb of ESDMODE3, which, in the big-endian convention shown in DDR SDRAM mode configuration 2 (DDR_DDR_SDRAM_MODE_2), corresponds to ESDMODE3[15]. The msb of the SDRAM extended mode 3 register value must be stored at ESDMODE3[0].</p>

8.4.12 DDR SDRAM mode control (DDR_DDR_SDRAM_MD_CNTL)

The DDR SDRAM mode control register allows the user to carry out the following tasks:

- Issue a mode register set command to a particular chip select
- Issue an immediate refresh to a particular chip select
- Issue an immediate precharge or precharge all command to a particular chip select
- Force the CKE signals to a specific value

Before issuing a command via the DDR_SDRAM_MD_CNTL register, the DDR interface should be idle. This can be done by setting DDR_SDRAM_CFG[MEM_HALT] and disabling refreshes by clearing DDR_INTERVAL[REFINT]. If there are memory contents that need to be preserved during this time, then software should also force any required refresh commands while DDR_INTERVAL[REFINT] is cleared.

The following table shows how DDR_SDRAM_MD_CNTL fields should be set for each of the tasks described above:

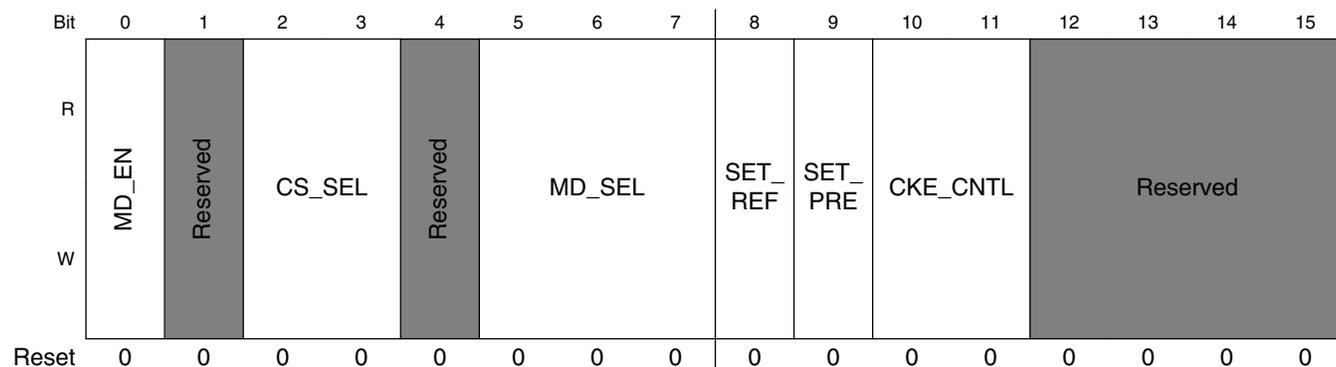
NOTE

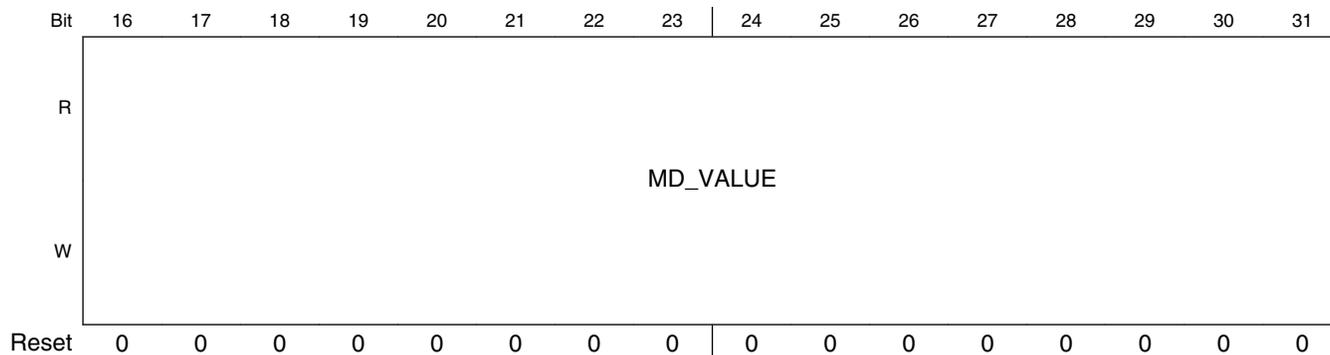
Note that MD_EN, SET_REF, and SET_PRE are mutually exclusive; only one of these fields can be set at a time.

Table 8-29. Settings of DDR_SDRAM_MD_CNTL Fields

Field	Mode Register Set	Refresh	Precharge	Clock Enable Signals Control
MD_EN	1	0	0	-
SET_REF	0	1	0	-
SET_PRE	0	0	1	-
CS_SEL	Chooses chip select (CS)			-
MD_SEL	Select mode register. See DDR SDRAM mode control (DDR_DDR_SDRAM_MD_CNTL) .	-	Selects logical bank	-
MD_VALUE	Value written to mode register	-	Only bit five is significant. See DDR SDRAM mode control (DDR_DDR_SDRAM_MD_CNTL) .	-
CKE_CNTL	0	0	0	See DDR SDRAM mode control (DDR_DDR_SDRAM_MD_CNTL) .

Address: 2000h base + 120h offset = 2120h





DDR_DDR_SDRAM_MD_CNTL field descriptions

Field	Description
0 MD_EN	<p>Mode enable</p> <p>Setting this bit specifies that valid data in MD_VALUE is ready to be written to DRAM as one of the following commands:</p> <ul style="list-style-type: none"> • MODE REGISTER SET • EXTENDED MODE REGISTER SET • EXTENDED MODE REGISTER SET 2 • EXTENDED MODE REGISTER SET 3 <p>The specific command to be executed is selected by setting MD_SEL. In addition, the chip select must be chosen by setting CS_SEL. MD_EN is set by software and cleared by hardware once the command has been issued.</p> <p>0 Indicates that no mode register set command needs to be issued. 1 Indicates that valid data contained in the register is ready to be issued as a mode register set command.</p>
1 -	<p>This field is reserved. Reserved</p>
2-3 CS_SEL	<p>Select chip select</p> <p>Specifies the chip select that is driven active due to any command forced by software in DDR_SDRAM_MD_CNTL.</p> <p>00 Chip select 0 is active 01 Chip select 1 is active 10 Chip select 2 is active 11 Chip select 3 is active</p>
4 -	<p>This field is reserved. Reserved</p>
5-7 MD_SEL	<p>Mode register select</p> <p>MD_SEL specifies one of the following:</p> <ul style="list-style-type: none"> • During a mode select command, selects the SDRAM mode register to be changed • During a precharge command, selects the SDRAM logical bank to be precharged. A precharge all command ignores this field. • During a refresh command, this field is ignored. <p>NOTE: MD_SEL contains the value that is presented onto the memory bank address pins (MBA_n) of the DDR controller.</p>

Table continues on the next page...

DDR_DDR_SDRAM_MD_CNTL field descriptions (continued)

Field	Description
	000 MR 001 EMR 010 EMR2 011 EMR3
8 SET_REF	Set refresh Forces an immediate refresh to be issued to the chip select specified by DDR_SDRAM_MD_CNTL[CS_SEL]. This bit is set by software and cleared by hardware once the command has been issued. 0 Indicates that no refresh command needs to be issued. 1 Indicates that a refresh command is ready to be issued.
9 SET_PRE	Set precharge Forces a precharge or precharge all to be issued to the chip select specified by DDR_SDRAM_MD_CNTL[CS_SEL]. This bit is set by software and cleared by hardware once the command has been issued. 0 Indicates that no precharge all command needs to be issued. 1 Indicates that a precharge all command is ready to be issued.
10–11 CKE_CNTL	Clock enable control Allows software to globally clear or set all CKE signals issued to DRAM. Once software has forced the value driven on CKE, that value continues to be forced until software clears the CKE_CNTL bits. At that time, the DDR controller continues to drive the CKE signals to the same value forced by software until another event causes the CKE signals to change (such as, self refresh entry/exit, power down entry/exit). 00 CKE signals are not forced by software. 01 CKE signals are forced to a low value by software. 10 CKE signals are forced to a high value by software. 11 Reserved
12–15 -	This field is reserved. Reserved
16–31 MD_VALUE	Mode register value This field, which specifies the value that is presented on the memory address pins of the DDR controller during a mode register set command, is significant only when this register is used to issue a mode register set command or a precharge or precharge all command. For a mode register set command, this field contains the data to be written to the selected mode register. For a precharge command, only bit five is significant: 0 Issue a precharge command; MD_SEL selects the logical bank to be precharged 1 Issue a precharge all command; all logical banks are precharged

8.4.13 DDR SDRAM interval configuration (DDR_DDR_SDRAM_INTERVAL)

The DDR SDRAM interval configuration register sets the number of DRAM clock cycles between bank refreshes issued to the DDR SDRAMs. In addition, the number of DRAM cycles that a page is maintained after it is accessed is provided here.

Address: 2000h base + 124h offset = 2124h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	REFINT															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved		BSTOPRE													
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDR_SDRAM_INTERVAL field descriptions

Field	Description
0–15 REFINT	Refresh interval Represents the number of memory bus clock cycles between refresh cycles. Depending on DDR_SDRAM_CFG_2[NUM_PR], some number of rows are refreshed in each DDR SDRAM physical bank during each refresh cycle. The value for REFINT depends on the specific SDRAMs used and the interface clock frequency. Refreshes are not issued when the REFINT is set to all 0s.
16–17 -	This field is reserved. Reserved
18–31 BSTOPRE	Precharge interval Sets the duration (in memory bus clocks) that a page is retained after a DDR SDRAM access. If BSTOPRE is zero, the DDR memory controller uses auto-precharge read and write commands rather than operating in page mode. This is called global auto-precharge mode.

8.4.14 DDR SDRAM data initialization (DDR_DDR_DATA_INIT)

The DDR SDRAM data initialization register provides the value that is used to initialize memory if DDR_SDRAM_CFG2[D_INIT] is set.

Address: 2000h base + 128h offset = 2128h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
R																																		
W	INIT_VALUE																																	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		

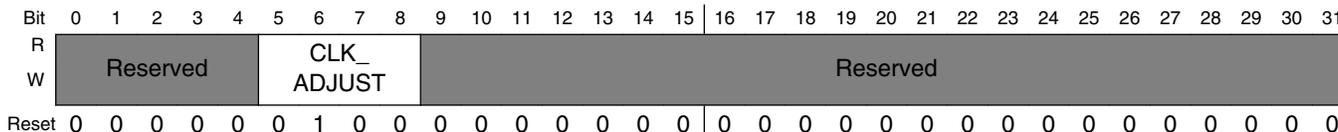
DDR_DDR_DATA_INIT field descriptions

Field	Description
0–31 INIT_VALUE	Initialization value. Represents the value that DRAM is initialized with if DDR_SDRAM_CFG2[D_INIT] is set.

8.4.15 DDR SDRAM clock control (DDR_DDR_SDRAM_CLK_CNTL)

The DDR SDRAM clock control configuration register provides a 1/8-cycle clock adjustment.

Address: 2000h base + 130h offset = 2130h



DDR_DDR_SDRAM_CLK_CNTL field descriptions

Field	Description
0–4 -	This field is reserved. Reserved
5–8 CLK_ADJUST	Clock adjust 0000 Clock is launched aligned with address/command 0001 Clock is launched 1/8 applied cycle after address/command 0010 Clock is launched 1/4 applied cycle after address/command 0011 Clock is launched 3/8 applied cycle after address/command 0100 Clock is launched 1/2 applied cycle after address/command 0101 Clock is launched 5/8 applied cycle after address/command 0110 Clock is launched 3/4 applied cycle after address/command 0111 Clock is launched 7/8 applied cycle after address/command 1000 Clock is launched 1 applied cycle after address/command 1001-1111 Reserved
9–31 -	This field is reserved. Reserved

8.4.16 DDR training initialization address (DDR_DDR_INIT_ADDR)

The DDR SDRAM initialization address register provides the address that is used for the data strobe to data skew adjustment and automatic $\overline{\text{CAS}}$ to preamble calibration after POR.

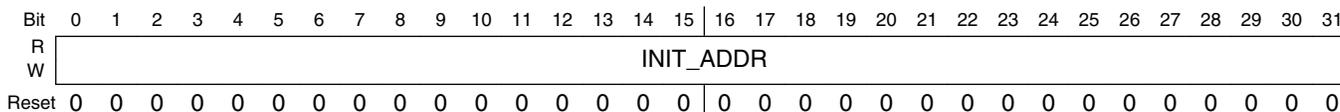
NOTE

After the skew adjustment, this address contains bad ECC data. This is not important at POR, as all of memory should be subsequently initialized if ECC is enabled (either by software or through the use of DDR_SDRAM_CFG_2[D_INIT]).

NOTE

If an $\overline{\text{HRESET}}$ has been issued after the DRAM is in self-refresh mode, however, memory is not initialized, so this address should be written to using a 32-byte transaction to avoid possible ECC errors if this address could later be accessed.

Address: 2000h base + 148h offset = 2148h



DDR_DDR_INIT_ADDR field descriptions

Field	Description
0–31 INIT_ADDR	Initialization address. Represents the address that is used for the data strobe to data skew adjustment and automatic CAS to preamble calibration at POR. This address is written to during the initialization sequence.

8.4.17 DDR training initialization extended address (DDR_DDR_INIT_EXT_ADDR)

The DDR SDRAM initialization extended address register provides the extended address that is used for the data strobe to data skew adjustment and automatic $\overline{\text{CAS}}$ to preamble calibration after POR.

Address: 2000h base + 14Ch offset = 214Ch

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R																	
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R																	
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

DDR_DDR_INIT_EXT_ADDR field descriptions

Field	Description
0 UIA	Use initialization address 0 Use the default address for training sequence as calculated by the controller. This is the first valid address in the first enabled chip select. 1 Use the initialization address programmed in DDR_INIT_ADDR and DDR_INIT_EXT_ADDR.
1–27 -	This field is reserved. Reserved
28–31 INIT_EXT_ADDR	Initialization extended address Represents the extended address that is used for the data strobe to data skew adjustment and automatic $\overline{\text{CAS}}$ to preamble calibration at POR. This extended address is written to during the initialization sequence.

8.4.18 DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4)

The DDR SDRAM timing configuration 4 register provides additional timing fields required to support DDR3 memories.

Address: 2000h base + 160h offset = 2160h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	RWT				WRT				RRT				WWT			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved														DLL_LOCK	
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_TIMING_CFG_4 field descriptions

Field	Description
0-3 RWT	<p>Read-to-write turnaround for same chip select</p> <p>Specifies how many cycles are added between a read to write turnaround for transactions to the same chip select. If a value of 0000 is chosen, then the DDR controller uses the value used for transactions to different chip selects, as defined in TIMING_CFG_0[RWT]. This field can be used to improve performance when operating in burst-chop mode by forcing transactions to the same chip select to use extra cycles, while transaction to different chip selects can utilize the tri-state time on the DRAM interface. Regardless of the value that is set in this field, the value defined by TIMING_CFG_0[RWT] also is met before issuing a write command.</p> <p>0000 Default 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks</p>
4-7 WRT	<p>Write-to-read turnaround for same chip select</p> <p>Specifies how many cycles are added between a write to read turnaround for transactions to the same chip select. If a value of 0000 is chosen, then the DDR controller uses the value used for transactions to different chip selects, as defined in TIMING_CFG_0[WRT]. This field can be used to improve performance</p>

Table continues on the next page...

DDR_TIMING_CFG_4 field descriptions (continued)

Field	Description
	<p>when operating in burst-chop mode by forcing transactions to the same chip select to use extra cycles, while transaction to different chip selects can utilize the tri-state time on the DRAM interface. Regardless of the value that is set in this field, the value defined by TIMING_CFG_0[WRT] also is met before issuing a read command.</p> <p>0000 Default 0001 1 clock 0010 2 clocks 0011 3 clocks 0100 4 clocks 0101 5 clocks 0110 6 clocks 0111 7 clocks 1000 8 clocks 1001 9 clocks 1010 10 clocks 1011 11 clocks 1100 12 clocks 1101 13 clocks 1110 14 clocks 1111 15 clocks</p>
<p>8–11 RRT</p>	<p>Read-to-read turnaround for same chip select</p> <p>Specifies how many cycles are added between reads to the same chip select. If a value of 0000 is chosen, then 2 cycles are required between read commands to the same chip select if 4-beat bursts are used (4 cycles are required if 8-beat bursts are used). Note that DDR3 does not support 4-beat bursts. However, this field may be used to add extra cycles when burst-chop mode is used, and the DDR controller must wait 4 cycles for read-to-read transactions to the same chip select.</p> <p>0000 BL/2 clocks 0001 BL/2 + 1 clock 0010 BL/2 + 2 clocks 0011 BL/2 + 3 clocks 0100 BL/2 + 4 clocks 0101 BL/2 + 5 clocks 0110 BL/2 + 6 clocks 0111 BL/2 + 7 clocks 1000 BL/2 + 8 clocks 1001 BL/2 + 9 clocks 1010 BL/2 + 10 clocks 1011 BL/2 + 11 clocks 1100 BL/2 + 12 clocks 1101 BL/2 + 13 clocks 1110 BL/2 + 14 clocks 1111 BL/2 + 15 clocks</p>
<p>12–15 WWT</p>	<p>Write-to-write turnaround for same chip select</p>

Table continues on the next page...

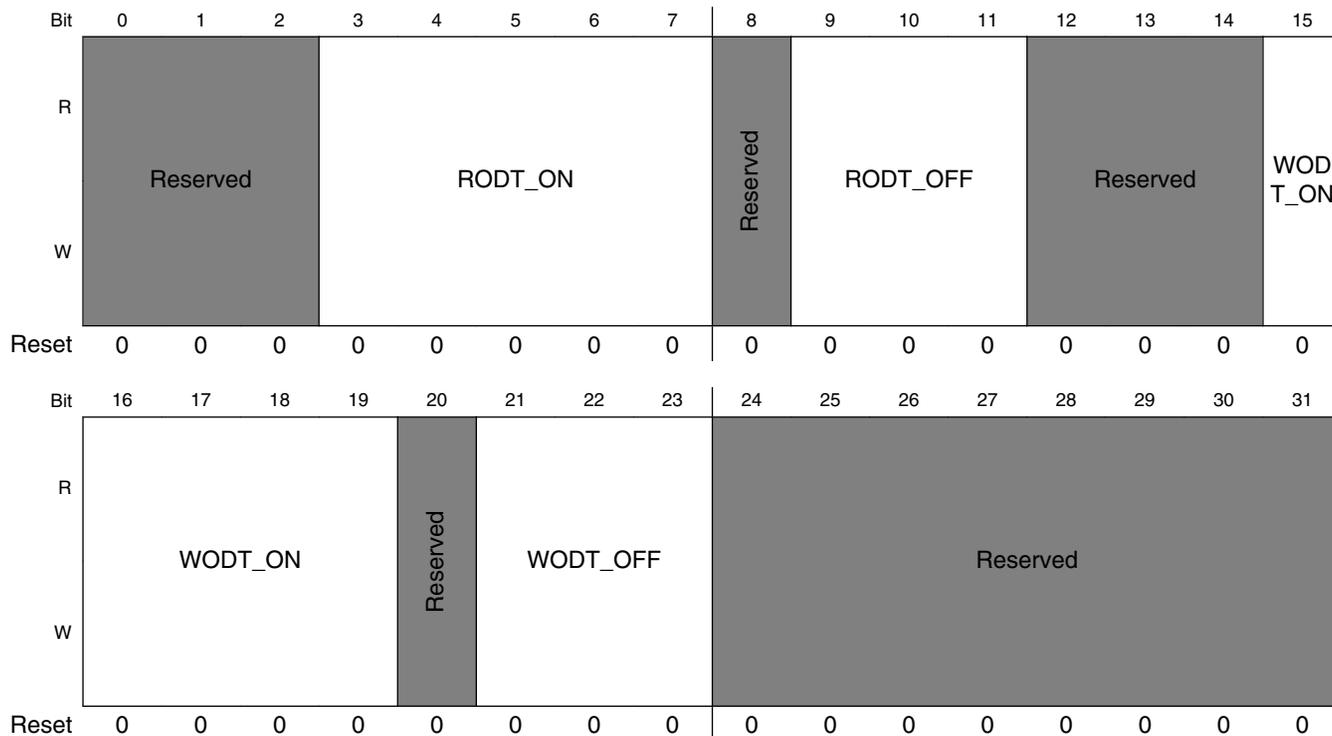
DDR_TIMING_CFG_4 field descriptions (continued)

Field	Description
	<p>Specifies how many cycles are added between writes to the same chip select. If a value of 0000 is chosen, then 2 cycles are required between write commands to the same chip select if 4-beat bursts are used (4 cycles are required if 8-beat bursts are used). Note that DDR3 does not support 4-beat bursts. However, this field may be used to add extra cycles when burst-chop mode is used, and the DDR controller must wait 4 cycles for write-to-write transactions to the same chip select.</p> <p>0000 BL/2 clocks 0001 BL/2 + 1 clock 0010 BL/2 + 2 clocks 0011 BL/2 + 3 clocks 0100 BL/2 + 4 clocks 0101 BL/2 + 5 clocks 0110 BL/2 + 6 clocks 0111 BL/2 + 7 clocks 1000 BL/2 + 8 clocks 1001 BL/2 + 9 clocks 1010 BL/2 + 10 clocks 1011 BL/2 + 11 clocks 1100 BL/2 + 12 clocks 1101 BL/2 + 13 clocks 1110 BL/2 + 14 clocks 1111 BL/2 + 15 clocks</p>
16–29 -	This field is reserved. Reserved
30–31 DLL_LOCK	<p>DDR SDRAM DLL Lock Time. This provides the number of cycles that it takes for the DRAMs DLL to lock at POR and after exiting self refresh. The controller waits the specified number of cycles before issuing any commands after exiting POR or self refresh.</p> <p>00 200 clocks 01 512 clocks 10 Reserved 11 Reserved</p>

8.4.19 DDR SDRAM timing configuration 5 (DDR_TIMING_CFG_5)

The DDR SDRAM timing configuration 5 register provides additional timing fields required to support DDR3 memories.

Address: 2000h base + 164h offset = 2164h



DDR_TIMING_CFG_5 field descriptions

Field	Description
0-2 -	This field is reserved. Reserved
3-7 RODT_ON	Read to ODT on Specifies the number of cycles that passes from when a read command is placed on the DRAM bus until the assertion of the relevant ODT signal(s). The default case (00000) provides a decode of RL - 3 cycles to support legacy of past products. RL is the read latency, derived from CAS latency + additive latency. If 2T/3T timing is used, one/two extra cycle(s) is/are automatically added to the value selected in this field. 00000 RL 3 clocks 00001 0 clocks 00010 1 clocks 00011 2 clocks ... 01111 14 clocks

Table continues on the next page...

DDR_TIMING_CFG_5 field descriptions (continued)

Field	Description
	10000 15 clocks 10001 16 clocks 10011 18 clocks ... 11111 30 clocks
8 -	This field is reserved. Reserved
9–11 RODT_OFF	Read to ODT off Specifies the number of cycles that the relevant ODT signal(s) remains asserted for each read transaction. The default case (000) leaves the ODT signal(s) asserted for 3 DRAM cycles. 000 3 clocks 001 1 clock 010 2 clocks 011 3 clocks 100 4 clocks 101 5 clocks 110 6 clocks 111 7 clocks
12–14 -	This field is reserved. Reserved
15–19 WODT_ON	Write to ODT On Specifies the number of cycles that passes from when a write command is placed on the DRAM bus until the assertion of the relevant ODT signal(s). The default case (00000) provides a decode of WL - 3 cycles to support legacy of past products. WL is the write latency, derived from Write Latency + Additive Latency. If 2T/3T timing is used, one/two extra cycle(s) is/are automatically added to the value selected in this field. 00000 WL - 3 clocks 00001 0 clocks 00010 1 clocks 00011 2 clocks ... 01111 14 clocks 10000 15 clocks 10011 18 clocks ... 11111 30 clocks
20 -	This field is reserved. Reserved
21–23 WODT_OFF	Write to ODT Off Specifies the number of cycles that the relevant ODT signal(s) remains asserted for each write transaction. The default case (000) leaves the ODT signal(s) asserted for 3 DRAM cycles. 000 3 clocks 001 1 clock 010 2 clocks

Table continues on the next page...

DDR_TIMING_CFG_5 field descriptions (continued)

Field	Description
011	3 clocks
100	4 clocks
101	5 clocks
110	6 clocks
111	7 clocks
24–31 -	This field is reserved. Reserved

8.4.20 DDR ZQ calibration control (DDR_DDR_ZQ_CNTL)

The DDR ZQ calibration control register provides the enable and controls required for ZQ calibration when using DDR3 SDRAM devices.

There is a limitation for various DRAM timing parameters when ZQ calibration is used. The following factors are involved in this limitation:

- DDR_ZQ_CNTL[ZQOPER]
- DDR_ZQ_CNTL[ZQCS]
- TIMING_CFG_1[PRETOACT]
- TIMING_CFG_1[REFREC]
- DDR_SDRAM_INTERVAL[REFINT]
- the number of chip selects enabled

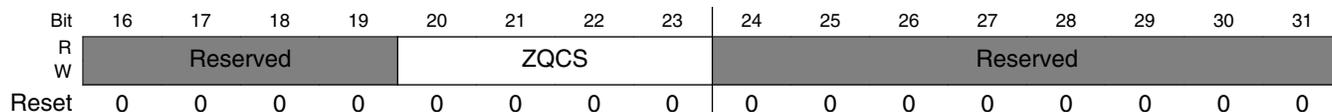
If the following condition is true:

$$\begin{aligned}
 & [((\text{DDR_ZQ_CNTL}[\text{ZQOPER}] + \text{DDR_ZQ_CNTL}[\text{ZQCS}]) \times (\# \text{ enabled chip selects})) + \\
 & \text{TIMING_CFG_1}[\text{PRETOACT}] + \\
 & \text{TIMING_CFG_1}[\text{REFREC}] + 2t_{\text{CK}}] > (\text{DDR_SDRAM_INTERVAL}[\text{REFINT}]),
 \end{aligned}$$

Then it is possible that one refresh is skipped when the controller is exiting self refresh. If this is an issue, then posted refreshes could be used to extend the refresh interval. Another alternative is to use the DDR_SDRAM_MD_CNTL register to force an extra refresh to each chip select after exiting self refresh mode. However, DDR3 timing parameters for most devices/frequencies do not allow for a refresh to be missed.

Address: 2000h base + 170h offset = 2170h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	ZQ_EN	Reserved			ZQINIT				Reserved			ZQOPER				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0



DDR_DDR_ZQ_CNTL field descriptions

Field	Description
0 ZQ_EN	ZQ Calibration Enable. This bit determines if ZQ calibration is used. This bit should only be set if DDR3 memory is used (DDR_SDRAM_CFG[SDRAM_TYPE] = 3'b111). 0 ZQ Calibration is not used. 1 ZQ Calibration is used. A ZQCL command is issued by the DDR controller after POR and anytime the DDR controller is exiting self refresh. A ZQCS command is issued every 32 refresh sequences to account for VT variations.
1-3 -	This field is reserved. Reserved
4-7 ZQINIT	POR ZQ Calibration Time (t_{ZQinit}). Determines the number of cycles that must be allowed for DRAM ZQ calibration at POR. Each chip select is calibrated separately, and this time must elapse after the ZQCL command is issued for each chip select before a separate command may be issued. 0000-0110 Reserved 0111 128 clocks 1000 256 clocks 1001 512 clocks 1010 1024 clocks 1011-1111 Reserved
8-11 -	This field is reserved. Reserved
12-15 ZQOPER	Normal Operation Full Calibration Time (t_{ZQoper}). Determines the number of cycles that must be allowed for DRAM ZQ calibration when exiting self refresh. Each chip select is calibrated separately, and this time must elapse after the ZQCL command is issued for each chip select before a separate command may be issued. 0000-0110 Reserved 0111 128 clocks 1000 256 clocks 1001 512 clocks 1010 1024 clocks 1011-1111 Reserved
16-19 -	This field is reserved. Reserved
20-23 ZQCS	Normal Operation Short Calibration Time (t_{ZQCS}). Determines the number of cycles that must be allowed for DRAM ZQ calibration during dynamic calibration which is issued every 32 refresh cycles. Each chip select is calibrated separately, and this time must elapse after the ZQCS command is issued for each chip select before a separate command may be issued. 0000 1 clocks 0001 2 clocks 0010 4 clocks 0011 8 clocks 0100 16 clocks

Table continues on the next page...

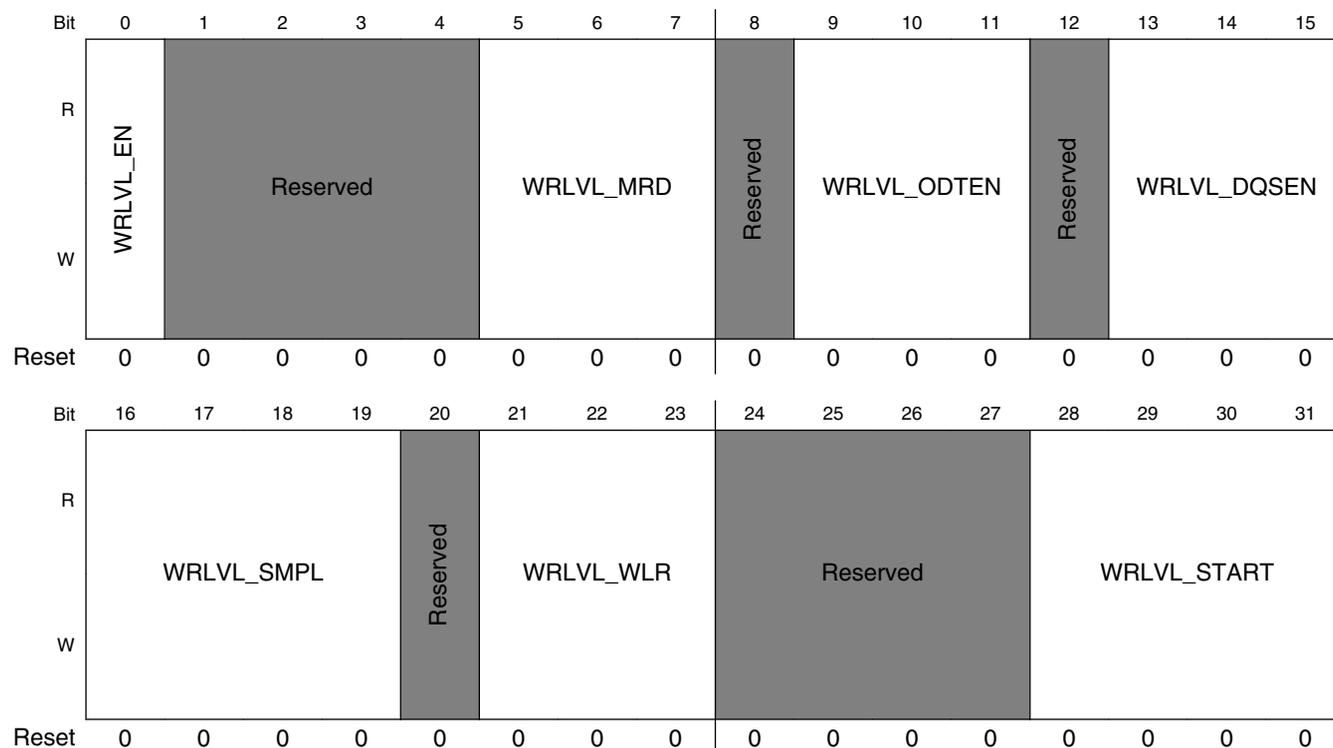
DDR_DDR_ZQ_CNTL field descriptions (continued)

Field	Description
0101	32 clocks
0110	64 clocks
0111	128 clocks
1000	256 clocks
1001	512 clocks
1010-1111	Reserved
24-31 -	This field is reserved. Reserved

8.4.21 DDR write leveling control (DDR_DDR_WRLVL_CNTL)

The DDR write leveling control register provides controls for write leveling, as it is supported for DDR3 memory devices.

Address: 2000h base + 174h offset = 2174h



DDR_DDR_WRLVL_CNTL field descriptions

Field	Description
0 WRLVL_EN	Write Leveling Enable

Table continues on the next page...

DDR_DDR_WRLVL_CNTL field descriptions (continued)

Field	Description
	This bit determines if write leveling is used. If this bit is set, then the DDR controller performs write leveling immediately after initializing the DRAM. This bit should only be set if DDR3 memory is used (DDR_SDRAM_CFG[SDRAM_TYPE] = 3'b111). 0 Write leveling is not used 1 Write leveling is used
1–4 -	This field is reserved. Reserved
5–7 WRLVL_MRD	First DQS pulse rising edge after margining mode is programmed (t_{WL_MRD}). Determines how many cycles to wait after margining mode has been programmed before the first DQS pulse may be issued. This field is only relevant when DDR_WRLVL_CNTL[WRLVL_EN] is set. 000 1 clocks 001 2 clocks 010 4 clocks 011 8 clocks 100 16 clocks 101 32 clocks 110 64 clocks 111 128 clocks
8 -	This field is reserved. Reserved
9–11 WRLVL_ODTEN	ODT delay after margining mode is programmed (t_{WL_ODTEN}). Determines how many cycles to wait after margining mode has been programmed until ODT may be asserted. This field is only relevant when DDR_WRLVL_CNTL[WRLVL_EN] is set. 000 1 clocks 001 2 clocks 010 4 clocks 011 8 clocks 100 16 clocks 101 32 clocks 110 64 clocks 111 128 clocks
12 -	This field is reserved. Reserved
13–15 WRLVL_DQSEN	DQS/ \overline{DQS} delay after margining mode is programmed (t_{WL_DQSEN}). Determines how many cycles to wait after margining mode has been programmed until DQS may be actively driven. This field is only relevant when DDR_WRLVL_CNTL[WRLVL_EN] is set. 000 1 clocks 001 2 clocks 010 4 clocks 011 8 clocks 100 16 clocks 101 32 clocks 110 64 clocks 111 128 clocks

Table continues on the next page...

DDR_DDR_WRLVL_CNTL field descriptions (continued)

Field	Description
16–19 WRLVL_SMPL	<p>Write leveling sample time</p> <p>Determines the number of cycles that must pass before the data signals are sampled after a DQS pulse during margining mode. This field should be programmed at least 6 cycles higher than t_{WLO} to allow enough time for propagation delay and sampling of the prime data bits. This field is only relevant when DDR_WRLVL_CNTL[WRLVL_EN] is set.</p> <p>0000 Reserved (if DDR_WRLVL_CNTL[WRLVL_EN] is set)</p> <p>0001 1 clocks</p> <p>0010 2 clocks</p> <p>0011 3 clocks</p> <p>0100 4 clocks</p> <p>0101 5 clocks</p> <p>0110 6 clocks</p> <p>0111 7 clocks</p> <p>1000 8 clocks</p> <p>1001 9 clocks</p> <p>1010 10 clocks</p> <p>1011 11 clocks</p> <p>1100 12 clocks</p> <p>1101 13 clocks</p> <p>1110 14 clocks</p> <p>1111 15 clocks</p>
20 -	<p>This field is reserved.</p> <p>Reserved</p>
21–23 WRLVL_WLR	<p>Write leveling repetition time</p> <p>Determines the number of cycles that must pass between DQS pulses during write leveling. This field is only relevant when DDR_WRLVL_CNTL[WRLVL_EN] is set.</p> <p>000 1 clocks</p> <p>001 2 clocks</p> <p>010 4 clocks</p> <p>011 8 clocks</p> <p>100 16 clocks</p> <p>101 32 clocks</p> <p>110 64 clocks</p> <p>111 128 clocks</p>
24–27 -	<p>This field is reserved.</p> <p>Reserved, should be cleared.</p>
28–31 WRLVL_START	<p>Write leveling start time. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled.</p> <p>0000 0 clock delay</p> <p>0001 1/8 clock delay</p> <p>0010 1/4 clock delay</p> <p>0011 3/8 clock delay</p> <p>0100 1/2 clock delay</p> <p>0101 5/8 clock delay</p> <p>0110 3/4 clock delay</p>

Table continues on the next page...

DDR_DDR_WRLVL_CNTL field descriptions (continued)

Field	Description
0111	7/8 clock delay
1000	1 clock delay
1001	9/8 clock delay
1010	5/4 clock delay
1011	11/8 clock delay
1100	3/2 clock delay
1101	13/8 clock delay
1110	7/4 clock delay
1111	15/8 clock delay

8.4.22 DDR Self Refresh Counter (DDR_DDR_SR_CNTR)

The DDR self-refresh counter register can be programmed to force the DDR controller to enter self-refresh after a predefined period of idle time.

Address: 2000h base + 17Ch offset = 217Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved													SR_IT				Reserved														
W	Reserved													SR_IT				Reserved														
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDR_SR_CNTR field descriptions

Field	Description
0–11 -	This field is reserved. Reserved
12–15 SR_IT	<p>Self Refresh Idle Threshold</p> <p>Defines the number of DRAM cycles that must pass while the DDR controller is idle before it will enter self refresh. Anytime a transaction is issued to the DDR controller, it will reset its internal counter. When a new transaction is received by the DDR controller, it will exit self refresh and reset its internal counter. If this field is zero, then the described power savings feature is disabled. In addition, if a non-zero value is programmed into this field, then the DDR controller will exit self refresh anytime a transaction is issued to the DDR controller, regardless of the reason self refresh was initially entered.</p> <p>If this field is set to a non-zero value, then DDR_SDRAM_CFG[SREN] must also be set.</p> <p>Patterns not shown are reserved.</p> <p>NOTE: The term 'clock' refers to the data rate rather than the clock frequency of the DDR interface.</p> <p>0000 Automatic self refresh entry disabled 0001 2^10 DRAM clocks 0010 2^12 DRAM clocks 0011 2^14 DRAM clocks 0100 2^16 DRAM clocks</p>

Table continues on the next page...

DDR_DDR_SR_CNTR field descriptions (continued)

Field	Description
0101	2 ¹⁸ DRAM clocks
0110	2 ²⁰ DRAM clocks
0111	2 ²² DRAM clocks
1000	2 ²⁴ DRAM clocks
1001	2 ²⁶ DRAM clocks
1010	2 ²⁸ DRAM clocks
1011	2 ³⁰ DRAM clocks
16–31 -	This field is reserved. Reserved

8.4.23 DDR Register Control Words 1 (DDR_DDR_SDRAM_RCW_1)

The DDR register control word 1 register should be programmed with the intended values of the register control words if DDR_SDRAM_CFG[RCW_EN] is set. Each 4-bit field represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during register control word writes.

Address: 2000h base + 180h offset = 2180h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDR_SDRAM_RCW_1 field descriptions

Field	Description
0–3 RCW0	Register Control Word 0. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 0.
4–7 RCW1	Register Control Word 1. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 1.
8–11 RCW2	Register Control Word 2. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 2.
12–15 RCW3	Register Control Word 3. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 3.
16–19 RCW4	Register Control Word 4. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 4.
20–23 RCW5	Register Control Word 5. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 5.
24–27 RCW6	Register Control Word 6. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 6.

Table continues on the next page...

DDR_DDR_SDRAM_RCW_1 field descriptions (continued)

Field	Description
28–31 RCW7	Register Control Word 7. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 7.

**8.4.24 DDR Register Control Words 2
(DDR_DDR_SDRAM_RCW_2)**

The DDR register control word 2 register should be programmed with the intended values of the register control words if DDR_SDRAM_CFG[RCW_EN] is set. Each 4-bit field represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during register control word writes.

Address: 2000h base + 184h offset = 2184h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDR_SDRAM_RCW_2 field descriptions

Field	Description
0–3 RCW8	Register Control Word 8. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 8.
4–7 RCW9	Register Control Word 9. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 9.
8–11 RCW10	Register Control Word 10. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 10.
12–15 RCW11	Register Control Word 11. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 11.
16–19 RCW12	Register Control Word 12. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 12.
20–23 RCW13	Register Control Word 13. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 13.
24–27 RCW14	Register Control Word 14. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 14.
28–31 RCW15	Register Control Word 15. Represents the value that is placed on MBA[1], MBA[0], MA[4], and MA[3] during writes to register control word 15.

8.4.25 DDR write leveling control 2 (DDR_DDR_WRLVL_CNTL_2)

The DDR write leveling control 2 register provides controls for write leveling, as it is supported for DDR3 memory devices. This register specifically defines the starting points for the individual data strobes.

Address: 2000h base + 190h offset = 2190h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W	Reserved			WRLVL_START_1				Reserved			WRLVL_START_2				Reserved			WRLVL_START_3				Reserved			WRLVL_START_4							
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDR_WRLVL_CNTL_2 field descriptions

Field	Description
0–2 -	This field is reserved. Reserved
3–7 WRLVL_START_1	Write leveling start time for DQS[1]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay 01001 9/8 clock delay 01010 5/4 clock delay 01011 11/8 clock delay 01100 3/2 clock delay 01101 13/8 clock delay 01110 7/4 clock delay 01111 15/8 clock delay 10000 2 clock delay 10001 17/8 clock delay 10010 9/4 clock delay 10011 19/8 clock delay 10100 5/2 clock delay 10101-11111 Reserved
8–10 -	This field is reserved. Reserved

Table continues on the next page...

DDR_DDR_WRLVL_CNTL_2 field descriptions (continued)

Field	Description
11–15 WRLVL_START_ 2	Write leveling start time for DQS[2]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay 01001 9/8 clock delay 01010 5/4 clock delay 01011 11/8 clock delay 01100 3/2 clock delay 01101 13/8 clock delay 01110 7/4 clock delay 01111 15/8 clock delay 10000 2 clock delay 10001 17/8 clock delay 10010 9/4 clock delay 10011 19/8 clock delay 10100 5/2 clock delay 10101-11111 Reserved
16–18 -	This field is reserved. Reserved
19–23 WRLVL_START_ 3	Write leveling start time for DQS[3]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay 01001 9/8 clock delay 01010 5/4 clock delay 01011 11/8 clock delay 01100 3/2 clock delay 01101 13/8 clock delay 01110 7/4 clock delay 01111 15/8 clock delay 10000 2 clock delay 10001 17/8 clock delay

Table continues on the next page...

DDR_DDR_WRLVL_CNTL_2 field descriptions (continued)

Field	Description
	10010 9/4 clock delay 10011 19/8 clock delay 10100 5/2 clock delay 10101-11111 Reserved
24–26 -	This field is reserved. Reserved
27–31 WRLVL_START_4	Write leveling start time for DQS[4]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay 01001 9/8 clock delay 01010 5/4 clock delay 01011 11/8 clock delay 01100 3/2 clock delay 01101 13/8 clock delay 01110 7/4 clock delay 01111 15/8 clock delay 10000 2 clock delay 10001 17/8 clock delay 10010 9/4 clock delay 10011 19/8 clock delay 10100 5/2 clock delay 10101-11111 Reserved

8.4.26 DDR write leveling control 3 (DDR_DDR_WRLVL_CNTL_3)

The DDR write leveling control 3 register provides controls for write leveling, as it is supported for DDR3 memory devices. This register specifically defines the starting points for the individual data strobes.

Address: 2000h base + 194h offset = 2194h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W	Reserved	WRLVL_START_5				Reserved	WRLVL_START_6				Reserved	WRLVL_START_7				Reserved	WRLVL_START_8															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

DDR_DDR_WRLVL_CNTL_3 field descriptions

Field	Description
0-2 -	This field is reserved. Reserved
3-7 WRLVL_START_ 5	Write leveling start time for DQS[5]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay 01001 9/8 clock delay 01010 5/4 clock delay 01011 11/8 clock delay 01100 3/2 clock delay 01101 13/8 clock delay 01110 7/4 clock delay 01111 15/8 clock delay 10000 2 clock delay 10001 17/8 clock delay 10010 9/4 clock delay 10011 19/8 clock delay 10100 5/2 clock delay 10101-11111 Reserved
8-10 -	This field is reserved. Reserved
11-15 WRLVL_START_ 6	Write leveling start time for DQS[6]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay 01001 9/8 clock delay 01010 5/4 clock delay 01011 11/8 clock delay 01100 3/2 clock delay 01101 13/8 clock delay 01110 7/4 clock delay 01111 15/8 clock delay

Table continues on the next page...

DDR_DDR_WRLVL_CNTL_3 field descriptions (continued)

Field	Description
	10000 2 clock delay 10001 17/8 clock delay 10010 9/4 clock delay 10011 19/8 clock delay 10100 5/2 clock delay 10101-11111 Reserved
16–18 -	This field is reserved. Reserved
19–23 WRLVL_START_ 7	Write leveling start time for DQS[7]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay 01001 9/8 clock delay 01010 5/4 clock delay 01011 11/8 clock delay 01100 3/2 clock delay 01101 13/8 clock delay 01110 7/4 clock delay 01111 15/8 clock delay 10000 2 clock delay 10001 17/8 clock delay 10010 9/4 clock delay 10011 19/8 clock delay 10100 5/2 clock delay 10101-11111 Reserved
24–26 -	This field is reserved. Reserved
27–31 WRLVL_START_ 8	Write leveling start time for DQS[8]. Determines the value to use for the DQS_ADJUST for the first sample when write leveling is enabled. 00000 Use value from DDR_WRLVL_CNTL[WRLVL_START] 00001 1/8 clock delay 00010 1/4 clock delay 00011 3/8 clock delay 00100 1/2 clock delay 00101 5/8 clock delay 00110 3/4 clock delay 00111 7/8 clock delay 01000 1 clock delay

Table continues on the next page...

DDR_DDR_WRLVL_CNTL_3 field descriptions (continued)

Field	Description
01001	9/8 clock delay
01010	5/4 clock delay
01011	11/8 clock delay
01100	3/2 clock delay
01101	13/8 clock delay
01110	7/4 clock delay
01111	15/8 clock delay
10000	2 clock delay
10001	17/8 clock delay
10010	9/4 clock delay
10011	19/8 clock delay
10100	5/2 clock delay
10101-11111	Reserved

8.4.27 DDR Debug Status Register 1 (DDR_DDRDSR_1)

The DDRDSR_1 register contains the DDR driver compensation input value and the current settings of the P and N FET impedance for MDICn, command/control, and data.

Address: 2000h base + B20h offset = 2B20h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	DDRDC		MDICPZ				MDICNZ			Reserved						
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	CPZ				CNZ				DPZ				DNZ			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDRDSR_1 field descriptions

Field	Description
0–1 DDRDC	DDR driver compensation input value
2–5 MDICPZ	Current setting of PFET driver MDIC impedance
6–9 MDICNZ	Current setting of NFET driver MDIC impedance
10–15 -	This field is reserved. Reserved, should be cleared.

Table continues on the next page...

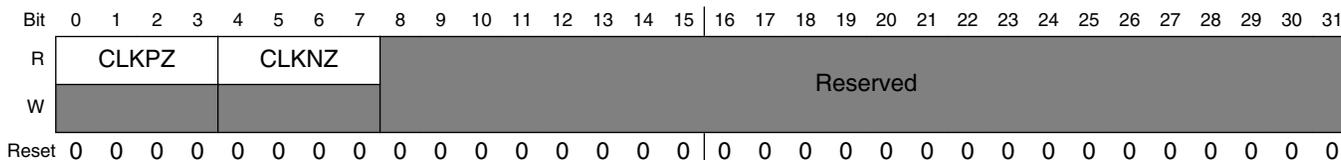
DDR_DDRDSR_1 field descriptions (continued)

Field	Description
16–19 CPZ	Current setting of PFET driver command impedance
20–23 CNZ	Current setting of NFET driver command impedance
24–27 DPZ	Current setting of PFET driver data impedance
28–31 DNZ	Current setting of NFET driver data impedance

8.4.28 DDR Debug Status Register 2 (DDR_DDRDSR_2)

The DDRDSR_2 register contains the current settings of the P and N FET impedance for the DDR drivers for clocks.

Address: 2000h base + B24h offset = 2B24h



DDR_DDRDSR_2 field descriptions

Field	Description
0–3 CLKPZ	Current setting of PFET driver clock impedance
4–7 CLKNZ	Current setting of NFET driver clock impedance
8–31 -	This field is reserved. Reserved

8.4.29 DDR Control Driver Register 1 (DDR_DDRCDR_1)

DDRCDR_1 sets the driver hardware compensation enable, the DDR MDIC driver P/N impedance, ODT termination value for IOs, driver software override enable for MDIC, driver software override enable for address/command, driver software override enable for data, the DDR address/command driver P/N impedance, and the DDR data driver P/N impedance.

The fields in DDRCDR_1, other than DDRCDR_1[ODT], are used to enable driver calibration with the MDIC[0:1] pins. This can be used to calibrate the DDR drivers to the external calibration resistors on the MDIC pins. However, this should only be used for full-strength driver applications.

Hardware DDR driver calibration is enabled by setting DDRCDR_1[DHC_EN].

NOTE

All driver calibration, whether by software or hardware, should be done before the DDR controller is enabled (before DDR_SDRAM_CFG[MEM_EN] is set).

Software can be used to calibrate the drivers instead of the automatic hardware calibration. If software calibration is used, the following steps should be taken:

1. Set DDRCDR_1[DSO_MDIC_EN] and ensure that DDRCDR_1[DHC_EN] is cleared
2. Set the highest impedance (value 0000) for DDRCDR_1[DSO_MDICPZ]
3. Set DDRCDR_1[DSO_MDIC_PZ_OE] to enable the output enable for MDIC[0]
4. After at least 4 cycles, read DDRDSR_1[0]. If the value is 0, then use the next lowest impedance, and read DDRDSR_1[0] again. Once a value of 1 is detected, then leave DDRCDR_1[DSO_MDICPZ] at the calibrated value
5. Clear DDRCDR_1[DSO_MDIC_PZ_OE]
6. After DDRCDR_1[DSO_MDICPZ] is calibrated, set a value of 0000 for DDRCDR_1[DSO_MDICNZ]
7. Set DDRCDR_1[DSO_MDIC_NZ_OE] to enable the output enable for MDIC[1]
8. After at least 4 cycles, read DDRDSR_1[1]. If the value is 1, then use the next lowest impedance, and read DDRDSR_1[1] again. Once a value of 0 is detected, then leave DDRCDR_1[DSO_MDICNZ] at the calibrated value
9. Clear DDRCDR_1[DSO_MDIC_NZ_OE]

Note that the legal impedance values (from lowest impedance to highest impedance) for DDR2 (1.8 V) are:

- 0000
- 0001
- 0011
- 0010
- 0110
- 0011
- 0111 (used for half-strength mode when driver calibration is not used)
- 0101
- 0100
- 1100

DDR memory map/register definition

- 1101
- 1111
- 1110
- 1010
- 1011 (default full-strength impedance)
- 1001
- 1000

Note that the legal impedance values (from highest impedance to lowest impedance) for DDR3 (1.5 V) are as follows:

- 0000
- 0001
- 0011
- 0010
- 0110
- 0111 (used for half-strength mode when driver calibration is not used)
- 0101
- 0100
- 1100
- 1101
- 1110
- 1010
- 1011
- 1001
- 1000 (default full-strength impedance)

Note that the drivers may either be calibrated to full-strength or half-strength.

Address: 2000h base + B28h offset = 2B28h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	DHC_EN	DSO_MDIC_EN	DSO_MDICPZ				DSO_MDICNZ				DSO_MDIC_PZ_OE	DSO_MDIC_NZ_OE	ODT		DSO_C_EN	DSO_D_EN
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	DSO_CPZ				DSO_CNZ				DSO_DPZ				DSO_DNZ			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_DDRCDR_1 field descriptions

Field	Description
0 DHC_EN	DDR driver hardware compensation enable
1 DSO_MDIC_EN	Driver software override enable for MDIC
2–5 DSO_MDICPZ	DDR driver software MDIC p-impedance override
6–9 DSO_MDICNZ	DDR driver software MDIC n-impedance override
10 DSO_MDIC_PZ_OE	Driver software override p-impedance output enable
11 DSO_MDIC_NZ_OE	Driver software override n-impedance output enable
12–13 ODT	<p>ODT termination value for IOs. This field is combined with DDRCDR_2[ODT] to determine the termination value. Below is the termination based on concatenating these two fields. Note that the order of concatenation is (from left to right) DDRCDR_1[ODT], DDRCDR_2[ODT]</p> <p>000 75 O 001 55 O 010 60 O 011 50 O 100 150 O 101 43 O 110 120 O 111 Reserved</p>
14 DSO_C_EN	Driver software override enable for address/command
15 DSO_D_EN	Driver software override enable for data
16–19 DSO_CPZ	DDR driver software command p-impedance override
20–23 DSO_CNZ	DDR driver software command n-impedance override
24–27 DSO_DPZ	Driver software data p-impedance override
28–31 DSO_DNZ	Driver software data n-impedance override

8.4.30 DDR Control Driver Register 2 (DDR_DDRCDR_2)

The DDRCDR_2 sets the driver software override enable for clocks and the DDR clocks driver P/N impedance.

Address: 2000h base + B2Ch offset = 2B2Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	DSO_CLK_EN	Reserved			DSO_CLKPZ				DSO_CLKNZ				Reserved			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															ODT
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

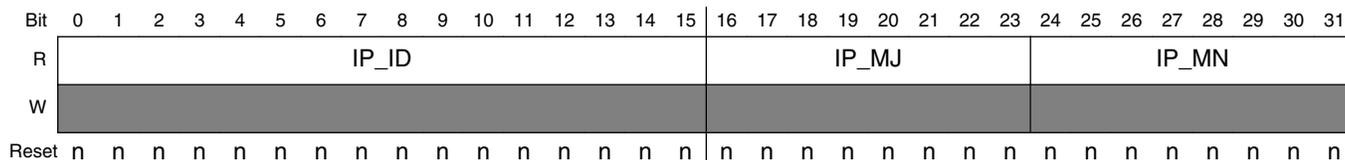
DDR_DDRCDR_2 field descriptions

Field	Description
0 DSO_CLK_EN	Driver software override enable for clocks
1-3 -	This field is reserved. Reserved
4-7 DSO_CLKPZ	Driver software clocks p-impedance override
8-11 DSO_CLKNZ	Driver software clocks n-impedance override
12-30 -	This field is reserved. Reserved
31 ODT	ODT termination value for IOs. This field is combined with DDRCDR_1[ODT] to determine the termination value. Below is the termination based on concatenating these two fields. Note that the order of concatenation is (from left to right) DDRCDR_1[ODT], DDRCDR_2[ODT] 000 75 O 001 55 O 010 60 O 011 50 O 100 150 O 101 43 O 110 120 O 111 Reserved

8.4.31 DDR IP block revision 1 (DDR_DDR_IP_REV1)

The DDR IP block revision 1 register provides read-only fields with the IP block ID, along with major and minor revision information.

Address: 2000h base + BF8h offset = 2BF8h



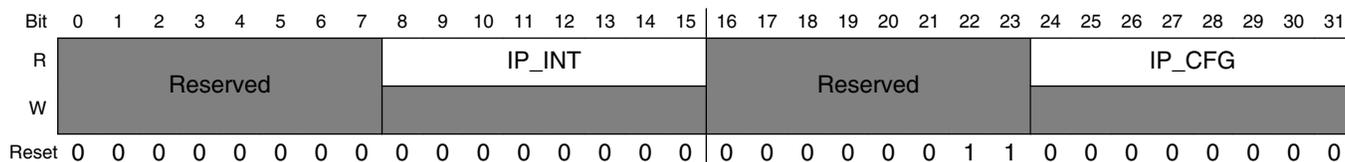
DDR_DDR_IP_REV1 field descriptions

Field	Description
0–15 IP_ID	IP block ID. For the DDR controller, this value is 0x0002.
16–23 IP_MJ	Major revision. This is currently set to 0x04.
24–31 IP_MN	Minor revision. This is currently set to 0x03.

8.4.32 DDR IP block revision 2 (DDR_DDR_IP_REV2)

The DDR IP block revision 2 register provides read-only fields with the IP block integration and configuration options.

Address: 2000h base + BFCh offset = 2BFCh



DDR_DDR_IP_REV2 field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–15 IP_INT	IP block integration options

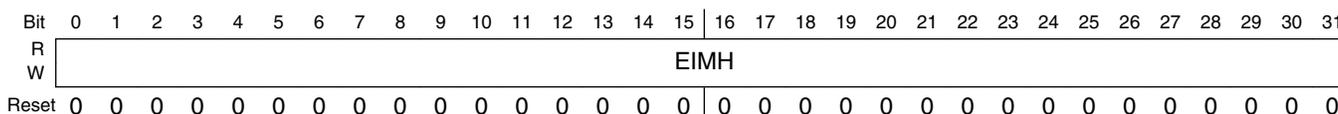
Table continues on the next page...

DDR_DDR_IP_REV2 field descriptions (continued)

Field	Description
16–23 -	This field is reserved. Reserved
24–31 IP_CFG	IP block configuration options

8.4.33 Memory data path error injection mask high (DDR_DATA_ERR_INJECT_HI)

Address: 2000h base + E00h offset = 2E00h

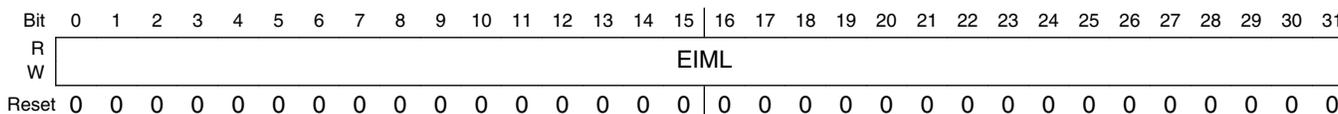


DDR_DATA_ERR_INJECT_HI field descriptions

Field	Description
0–31 EIMH	Error injection mask high data path Used to test ECC by forcing errors on the high word of the data path. Setting a bit causes the corresponding data path bit to be inverted on memory bus writes.

8.4.34 Memory data path error injection mask low (DDR_DATA_ERR_INJECT_LO)

Address: 2000h base + E04h offset = 2E04h



DDR_DATA_ERR_INJECT_LO field descriptions

Field	Description
0–31 EIML	Error injection mask low data path Used to test ECC by forcing errors on the low word of the data path. Setting a bit causes the corresponding data path bit to be inverted on memory bus writes.

8.4.35 Memory data path error injection mask ECC (DDR_ERR_INJECT)

The memory data path error injection mask ECC register sets the ECC mask, enables errors to be written to ECC memory, and allows the ECC byte to mirror the most significant data byte. In addition, a single address parity error may be injected through this register.

Address: 2000h base + E08h offset = 2E08h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R	Reserved																APIEN
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	Reserved						EMB	EIEN	EEIM								
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

DDR_ERR_INJECT field descriptions

Field	Description
0–14 -	This field is reserved. Reserved
15 APIEN	Address parity error injection enable. This bit is cleared by hardware after a single address parity error has been injected. 0 Address parity error injection disabled. 1 Address parity error injection enabled.
16–21 -	This field is reserved. Reserved
22 EMB	ECC mirror byte 0 Mirror byte functionality disabled. 1 Mirror the most significant data path byte onto the ECC byte.
23 EIEN	Error injection enable 0 Error injection disabled. 1 Error injection enabled. This applies to the data mask bits, the ECC mask bits, and the ECC mirror bit. Note that error injection should not be enabled until the memory controller has been enabled through DDR_SDRAM_CFG[MEM_EN].

Table continues on the next page...

DDR_ERR_INJECT field descriptions (continued)

Field	Description
24–31 EEIM	ECC error injection mask. Setting a mask bit causes the corresponding ECC bit to be inverted on memory bus writes.

8.4.36 Memory data path read capture high (DDR_CAPTURE_DATA_HI)

The memory data path read capture high register stores the high word of the read data path during error capture.

Address: 2000h base + E20h offset = 2E20h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_CAPTURE_DATA_HI field descriptions

Field	Description
0–31 ECHD	Error capture high data path. Captures the high word of the data path when errors are detected.

8.4.37 Memory data path read capture low (DDR_CAPTURE_DATA_LO)

The memory data path read capture low register stores the low word of the read data path during error capture.

Address: 2000h base + E24h offset = 2E24h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

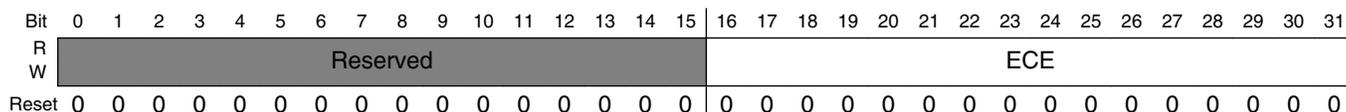
DDR_CAPTURE_DATA_LO field descriptions

Field	Description
0–31 ECLD	Error capture low data path. Captures the low word of the data path when errors are detected.

8.4.38 Memory data path read capture ECC (DDR_CAPTURE_ECC)

The memory data path read capture ECC register stores the ECC syndrome bits that were on the data bus when an error was detected.

Address: 2000h base + E28h offset = 2E28h



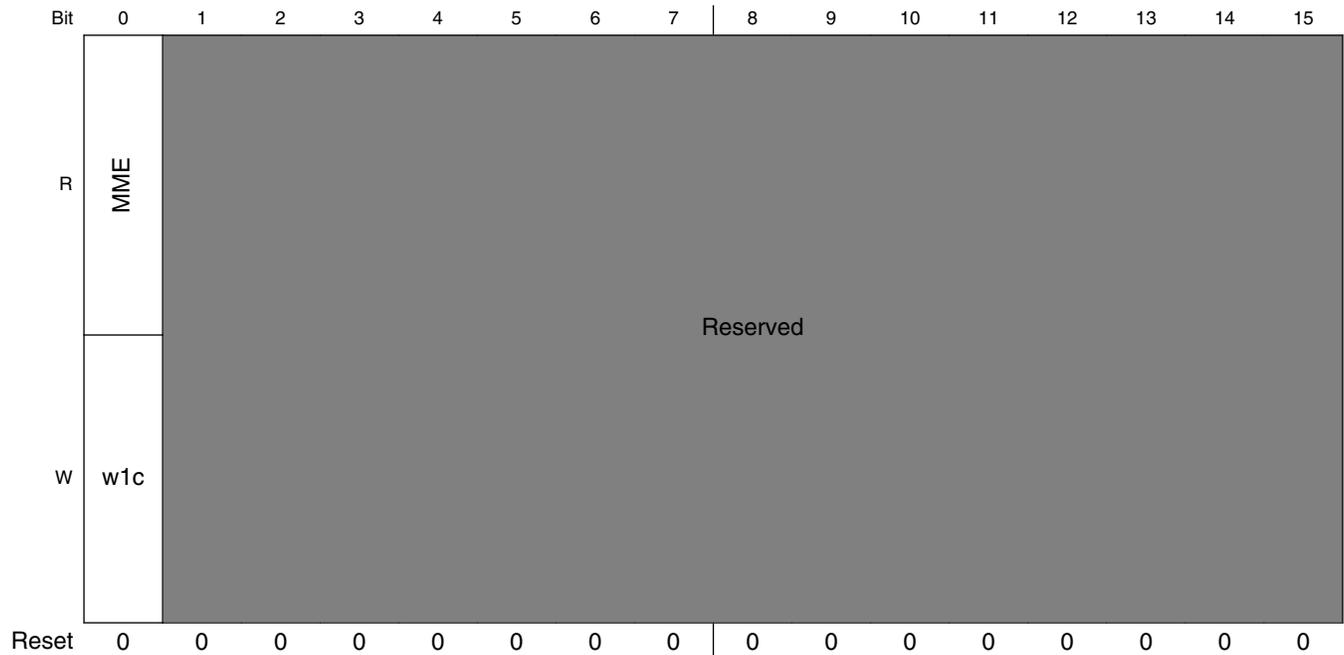
DDR_CAPTURE_ECC field descriptions

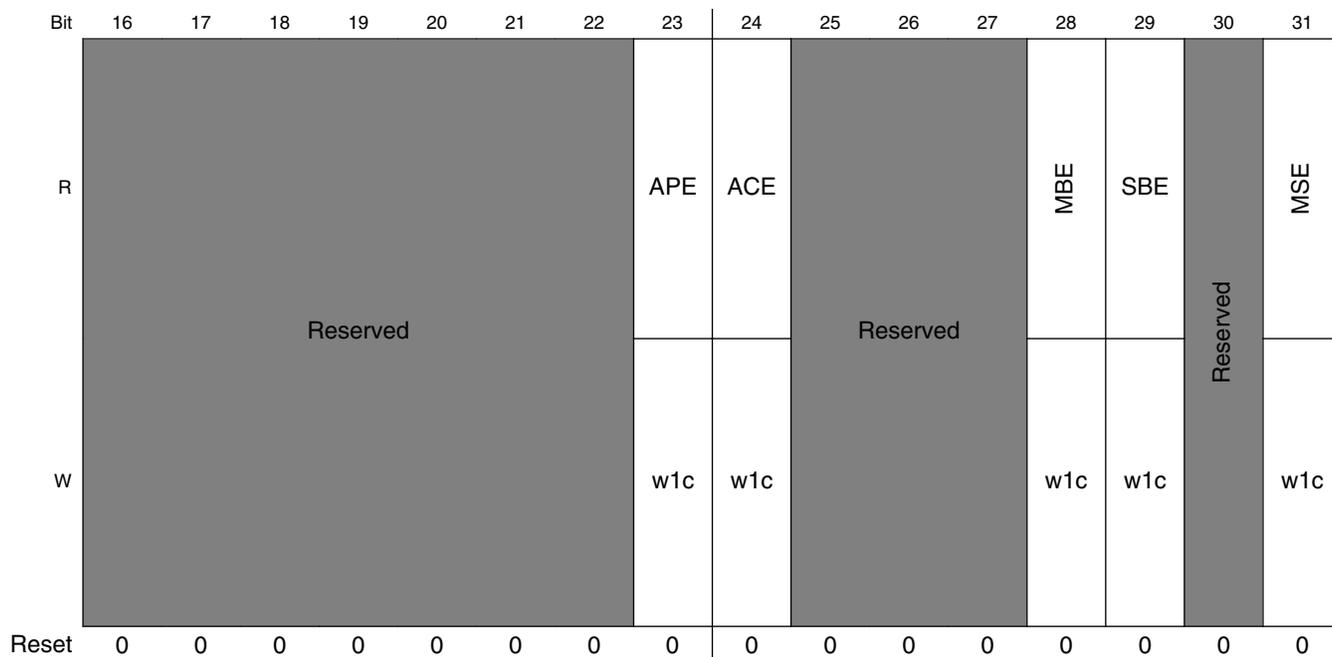
Field	Description
0–15 -	This field is reserved. Reserved
16–31 ECE	Error capture ECC. Captures the ECC bits on the data path whenever errors are detected. 16:23-8-bit ECC for the 32 bits in beats 0, 2, 4 and 6 in 32-bit bus mode ; should be ignored for 64-bit bus mode 24:31-8-bit ECC for the 32 bits in beats 1, 3, 5, and 7 in 32-bit bus mode ; should be used for every beat in 64-bit bus mode

8.4.39 Memory error detect (DDR_ERR_DETECT)

The memory error detect register stores the detection bits for multiple memory errors, single- and multiple-bit ECC errors, and memory select errors. It is a read/write register. A bit can be cleared by writing a one to the bit. System software can determine the type of memory error by examining the contents of this register. If an error is disabled with ERR_DISABLE, the corresponding error is never detected or captured in ERR_DETECT.

Address: 2000h base + E40h offset = 2E40h





DDR_ERR_DETECT field descriptions

Field	Description
0 MME	Multiple memory errors. This bit is cleared by software writing a 1. 0 Multiple memory errors of the same type were not detected. 1 Multiple memory errors of the same type were detected.
1–22 -	This field is reserved. Reserved
23 APE	Address parity error. This bit is cleared by software writing a 1. 0 An address parity error has not been detected. 1 An address parity error has been detected.
24 ACE	Automatic calibration error. This bit is cleared by software writing a 1. 0 An automatic calibration error has not been detected. 1 An automatic calibration error has been detected.
25–27 -	This field is reserved. Reserved
28 MBE	Multiple-bit error. This bit is cleared by software writing a 1. 0 A multiple-bit error has not been detected. 1 A multiple-bit error has been detected.
29 SBE	Single-bit ECC error. This bit is cleared by software writing a 1. 0 The number of single-bit ECC errors detected has not crossed the threshold set in ERR_SBE[SBET]. 1 The number of single-bit ECC errors detected crossed the threshold set in ERR_SBE[SBET].
30 -	This field is reserved. Reserved

Table continues on the next page...

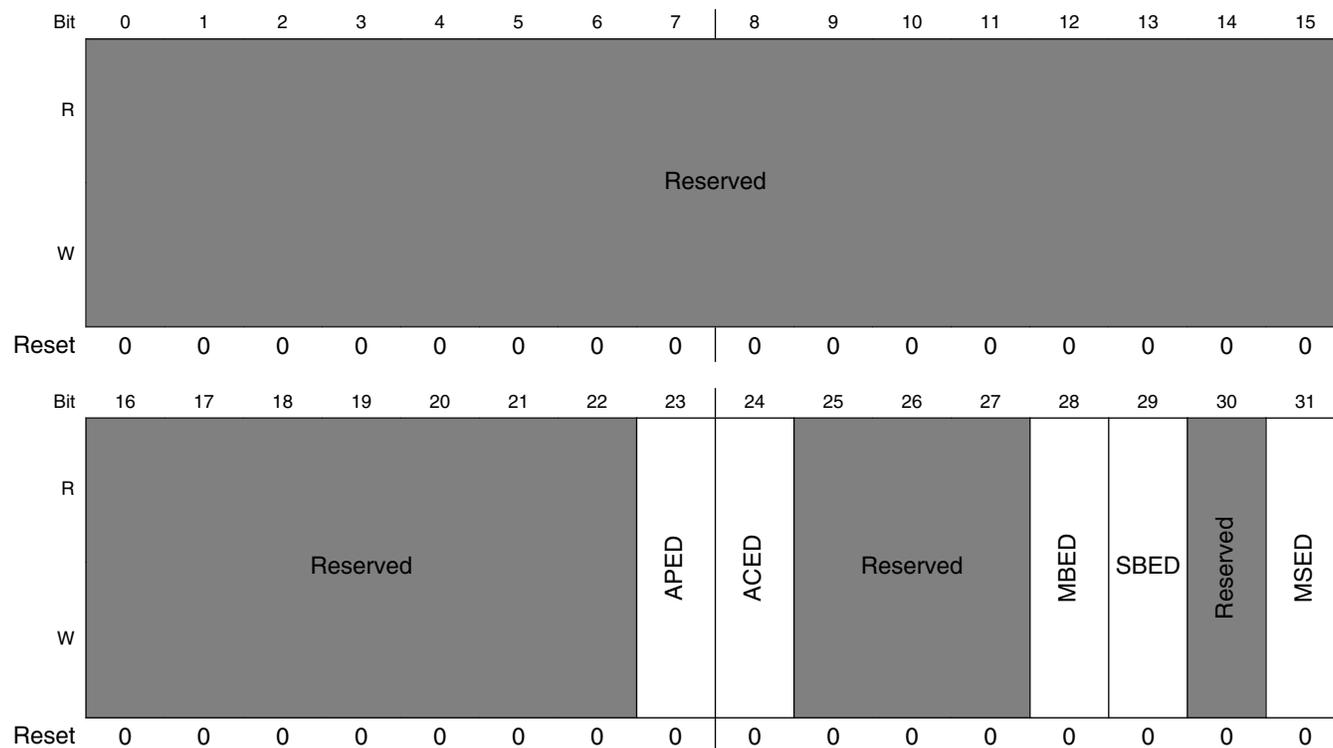
DDR_ERR_DETECT field descriptions (continued)

Field	Description
31 MSE	Memory select error. This bit is cleared by software writing a 1. 0 A memory select error has not been detected. 1 A memory select error has been detected.

8.4.40 Memory error disable (DDR_ERR_DISABLE)

The memory error disable register allows selective disabling of the DDR controller's error detection circuitry. Disabled errors are not detected or reported.

Address: 2000h base + E44h offset = 2E44h



DDR_ERR_DISABLE field descriptions

Field	Description
0–22 -	This field is reserved. Reserved
23 APED	Address parity error disable

Table continues on the next page...

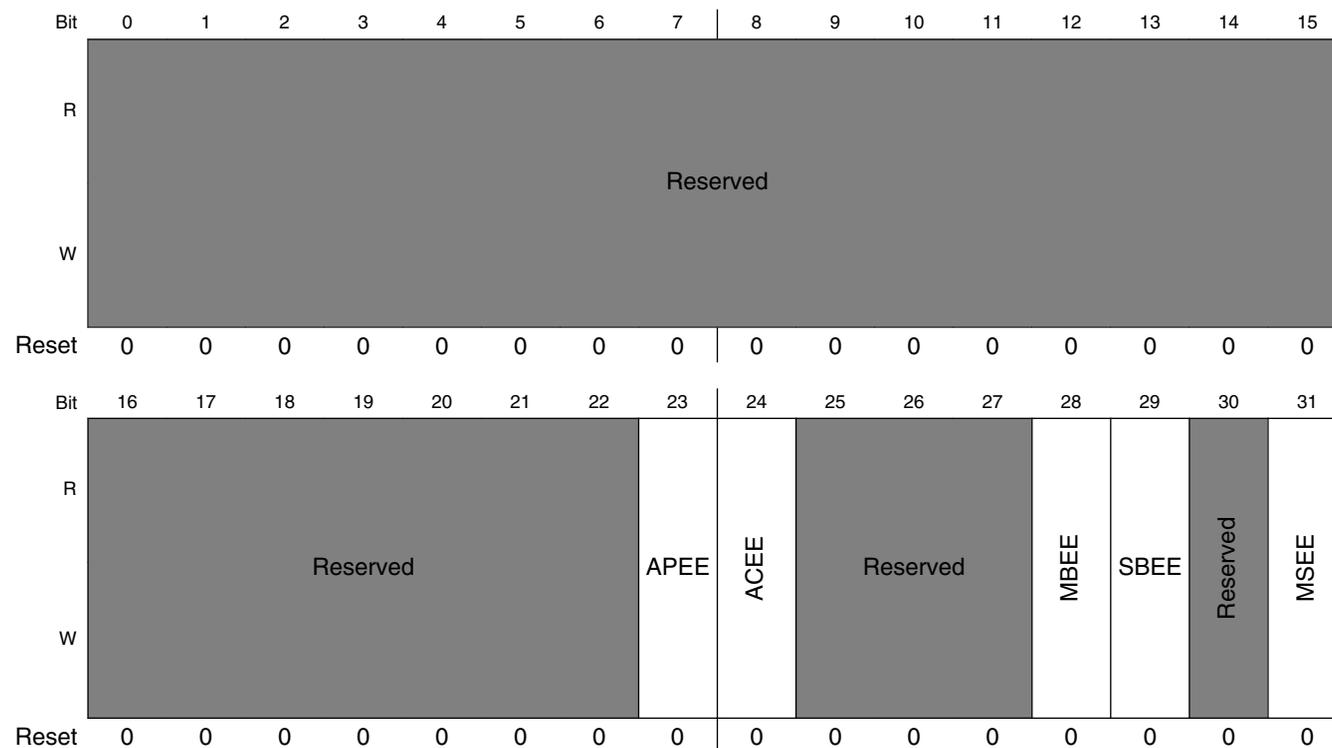
DDR_ERR_DISABLE field descriptions (continued)

Field	Description
	0 Address parity errors are detected if DDR_SDRAM_CFG_2[AP_EN] is set. They are reported if ERR_INT_EN[APEE] is set. 1 Address parity errors are not detected or reported.
24 ACED	Automatic calibration error disable 0 Automatic calibration errors are enabled. 1 Automatic calibration errors are disabled.
25–27 -	This field is reserved. Reserved
28 MBED	Multiple-bit ECC error disable 0 Multiple-bit ECC errors are detected if DDR_SDRAM_CFG[ECC_EN] is set. They are reported if ERR_INT_EN[MBEE] is set. Note that non-correctable read errors cause the assertion of <i>core_fault_in</i> , which causes the core to generate a machine check interrupt, unless it is disabled (by clearing HID1[RFXE]). If RFXE is zero and this error occurs, ERR_DISABLE[MBED] must be zero, and ERR_INT_EN[MBEE] and ECC_EN must be one to ensure that an interrupt is generated. 1 Multiple-bit ECC errors are not detected or reported.
29 SBED	Single-bit ECC error disable 0 Single-bit ECC errors are enabled. 1 Single-bit ECC errors are disabled.
30 -	This field is reserved. Reserved
31 MSED	Memory select error disable 0 Memory select errors are enabled. 1 Memory select errors are disabled.

8.4.41 Memory error interrupt enable (DDR_ERR_INT_EN)

The memory error interrupt enable register enables ECC interrupts or memory select error interrupts. When an enabled interrupt condition occurs, the internal $\overline{\text{int}}$ signal is asserted to the programmable interrupt controller (PIC).

Address: 2000h base + E48h offset = 2E48h



DDR_ERR_INT_EN field descriptions

Field	Description
0–22 -	This field is reserved. Reserved
23 APEE	Address parity error interrupt enable 0 Address parity errors cannot generate interrupts. 1 Address parity errors generate interrupts.
24 ACEE	Automatic calibration error interrupt enable 0 Automatic calibration errors cannot generate interrupts. 1 Automatic calibration errors generate interrupts.
25–27 -	This field is reserved. Reserved

Table continues on the next page...

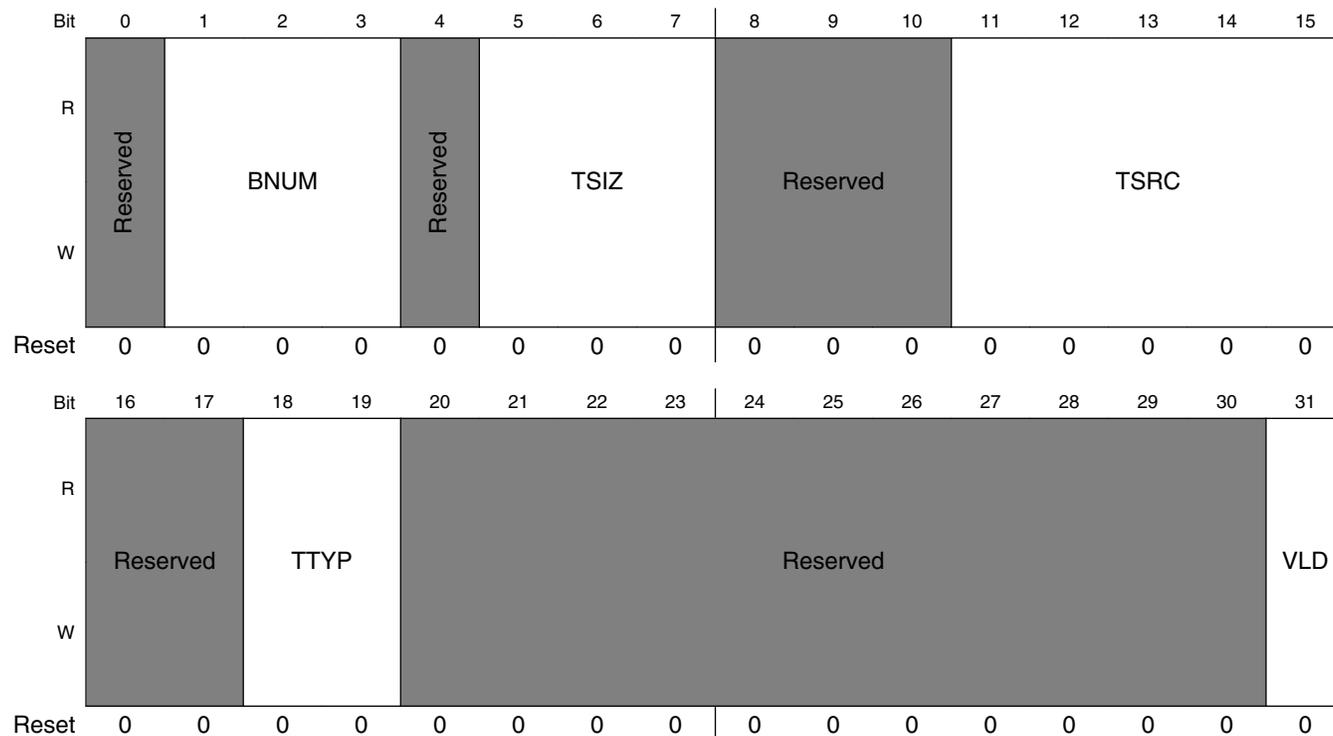
DDR_ERR_INT_EN field descriptions (continued)

Field	Description
28 MBEE	Multiple-bit ECC error interrupt enable. Note that non-correctable read errors may cause the assertion of <i>core_fault_in</i> , which causes the core to generate a machine check interrupt, unless it is disabled (by clearing HID1[RFXE]). If RFXE is zero and this error occurs, ERR_DISABLE[MBED] must be zero, and ERR_INT_EN[MBEE] and ECC_EN must be one to ensure that an interrupt is generated. 0 Multiple-bit ECC errors cannot generate interrupts. 1 Multiple-bit ECC errors generate interrupts.
29 SBEE	Single-bit ECC error interrupt enable 0 Single-bit ECC errors cannot generate interrupts. 1 Single-bit ECC errors generate interrupts.
30 -	This field is reserved. Reserved
31 MSEE	Memory select error interrupt enable 0 Memory select errors do not cause interrupts. 1 Memory select errors generate interrupts.

8.4.42 Memory error attributes capture (DDR_CAPTURE_ATTRIBUTES)

The memory error attributes capture register sets attributes for errors including type, size, source, and others.

Address: 2000h base + E4Ch offset = 2E4Ch



DDR_CAPTURE_ATTRIBUTES field descriptions

Field	Description
0 -	This field is reserved. Reserved
1-3 BNUM	Data beat number. Captures the double-word number for the detected error. Relevant only for ECC errors.
4 -	This field is reserved. Reserved
5-7 TSIZ	Transaction size for the error. Captures the transaction size in double words. Patterns not shown are reserved. 000 4 double words 001 1 double word 010 2 double words 011 3 double words

Table continues on the next page...

DDR_CAPTURE_ATTRIBUTES field descriptions (continued)

Field	Description
8–10 -	This field is reserved. Reserved
11–15 TSRC	Source ID. Specifies the source device mastering the transaction. 00000 PCI Express 3 00001 PCI Express 2 00010 PCI Express 1 00011 Reserved 00100 Reserved 00101 USB 00110 Reserved 00111 Security 01000 Reserved 01001 Reserved 01010 Boot sequencer 01011 eSDHC 01100 Serial RapidIO 1 01101 Serial RapidIO 2 01110 Reserved 01111 Reserved 10000 Processor 0 (instruction) 10001 Processor 0 (data) 10010 Processor 1 (instruction) 10011 Processor 1 (data) 10100 Reserved 10101 DMA 1 10110 DMA 2 10111 Reserved 11000 eTSEC 1 11001 eTSEC 2 11010 eTSEC 3 11011 Reserved 11100 RapidIO message unit 11101 RapidIO doorbell unit 11110 RapidIO port-write unit 11111 Reserved
16–17 -	This field is reserved. Reserved
18–19 TTYP	Transaction type for the error. 00 Reserved 01 Write 10 Read 11 Read-modify-write
20–30 -	This field is reserved. Reserved

Table continues on the next page...

DDR_CAPTURE_ATTRIBUTES field descriptions (continued)

Field	Description
31 VLD	Valid. Set as soon as valid information is captured in the error capture registers.

8.4.43 Memory error address capture (DDR_CAPTURE_ADDRESS)

The memory error address capture register holds the 32 lsbs of a transaction when a DDR ECC error is detected.

Address: 2000h base + E50h offset = 2E50h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	CADDR																															
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_CAPTURE_ADDRESS field descriptions

Field	Description
0–31 CADDR	Captured address. Captures the 32 lsbs of the transaction address when an error is detected.

8.4.44 Memory error extended address capture (DDR_CAPTURE_EXT_ADDRESS)

The memory error extended address capture register holds the four most significant transaction bits when an error is detected.

Address: 2000h base + E54h offset = 2E54h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																											CEADDR				
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

DDR_CAPTURE_EXT_ADDRESS field descriptions

Field	Description
0–27 -	This field is reserved. Reserved

Table continues on the next page...

DDR_CAPTURE_EXT_ADDRESS field descriptions (continued)

Field	Description
28–31 CEADDR	Captured extended address. Captures the 4 msbs of the transaction address when an error is detected

8.4.45 Single-Bit ECC memory error management (DDR_ERR_SBE)

The single-bit ECC memory error management register stores the threshold value for reporting single-bit errors and the number of single-bit errors counted since the last error report. When the counter field reaches the threshold, it wraps back to the reset value (0). If necessary, software must clear the counter after it has managed the error.

Address: 2000h base + E58h offset = 2E58h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DDR_ERR_SBE field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–15 SBET	Single-bit error threshold. Establishes the number of single-bit errors that must be detected before an error condition is reported.
16–23 -	This field is reserved. Reserved
24–31 SBEC	Single-bit error counter. Indicates the number of single-bit errors detected and corrected since the last error report. If single-bit error reporting is enabled, an error is reported and a regular or critical interrupt is generated when this value equals SBET. SBEC is automatically cleared when the threshold value is reached.

8.5 Functional description

The DDR SDRAM controller controls processor and I/O interactions with system memory.

It provides support for JEDEC-compliant DDR2 and DDR3 SDRAMs. The memory system allows a wide range of memory devices to be mapped to any arbitrary chip select, and support is provided for registered DIMMs and unbuffered DIMMs. However,

registered DIMMs cannot be mixed with unbuffered DIMMs. In addition, DDR3 DIMM module specifications allow for vendors to use mirrored DIMMs, where some address and bank address lines are mirrored on the DIMM.

The figure below is a high-level block diagram of the DDR memory controller. Requests are received from the internal mastering device and the address is decoded to generate the physical bank, logical bank, row, and column addresses. The transaction is compared with values in the row open table to determine if the address maps to an open page. If the transaction does not map to an open page, an active command is issued.

The memory interface supports the following configurations:

- as many as four physical banks of 64-/72-bit wide or 32-/40-bit wide memory
- bank sizes up to 8 Gbytes providing up to a maximum of 32 Gbytes of DDR main memory

Programmable parameters allow for a variety of memory organizations and timings. Optional error checking and correcting (ECC) protection is provided for the DDR SDRAM data bus. Using ECC, the DDR memory controller detects and corrects all single/double-bit errors within the data bus, and detects all errors within a nibble. The controller allows as many as 32 pages to be open simultaneously. The amount of time (in clock cycles) the pages remain open is programmable with `DDR_SDRAM_INTERVAL[BSTOPRE]`.

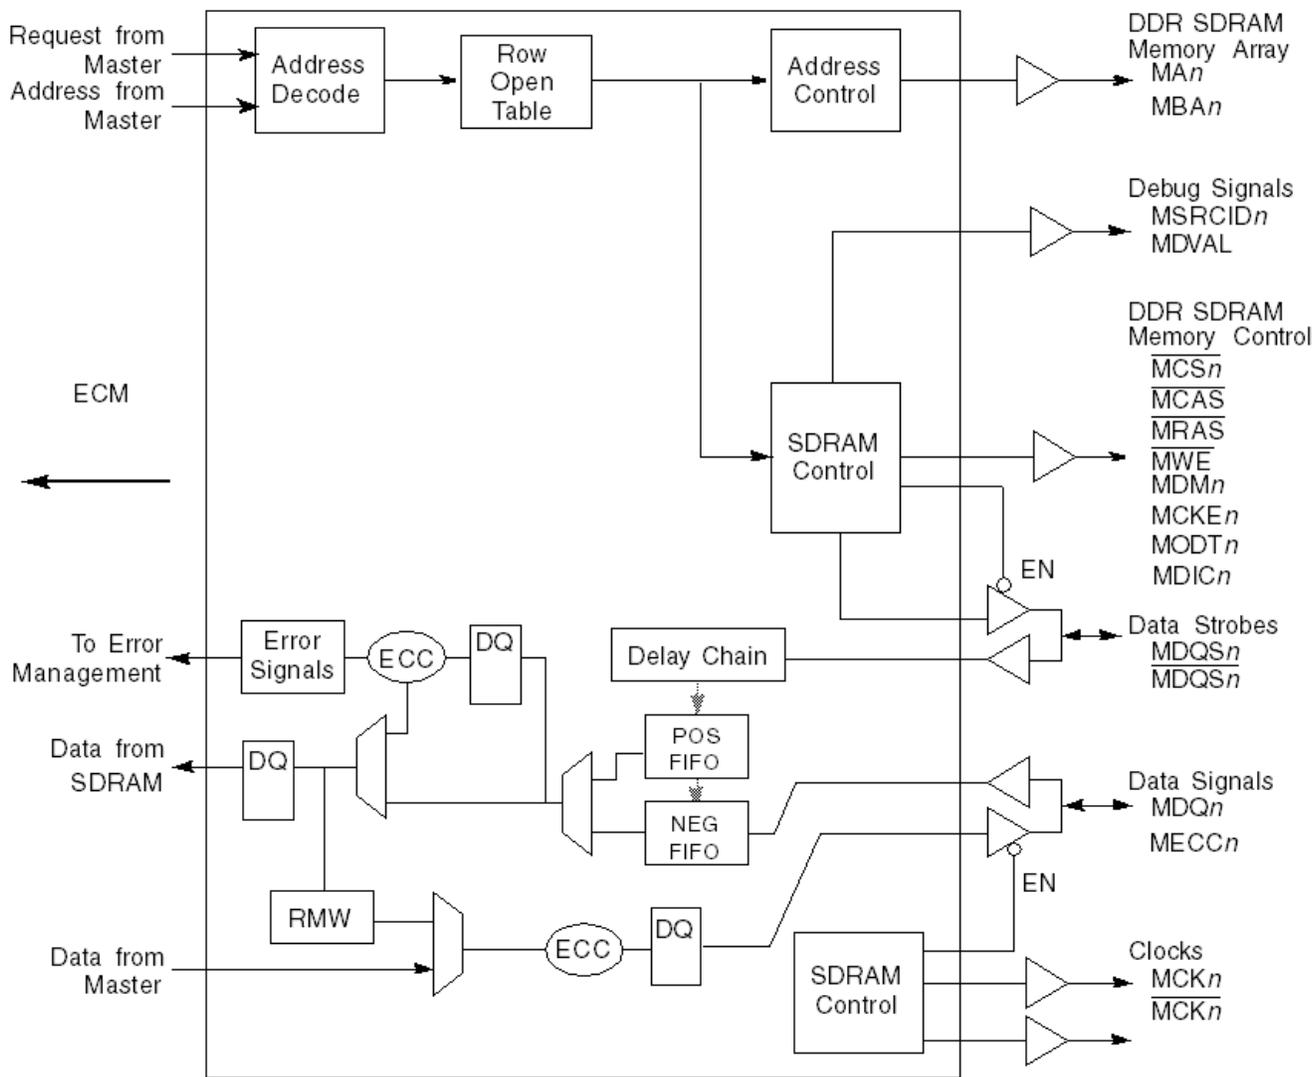


Figure 8-59. DDR memory controller block diagram

Read and write accesses to memory are burst oriented; accesses start at a selected location and continue for a programmed number of higher locations (4 or 8) in a programmed sequence. Accesses to closed pages start with the registration of an ACTIVE command followed by a READ or WRITE. (Accessing open pages does not require an ACTIVE command.) The address bits registered coincident with the activate command specifies the logical bank and row to be accessed. The address coincident with the READ or WRITE command specify the logical bank and starting column for the burst access.

The data interface is source synchronous, meaning whatever sources the data also provides a clocking signal to synchronize data reception. These bidirectional data strobes (MDQS[0:8]) are inputs to the controller during reads and outputs during writes. The

functional description

DDR SDRAM specification requires the data strobe signals to be centered within the data tenure during writes and to be offset by the controller to the center of the data tenure during reads. This delay is implemented in the controller for both reads and writes.

When ECC is enabled, 1 clock cycle is added to the read path to check ECC and correct single-bit errors. ECC generation does not add a cycle to the write path.

The address and command interface is also source synchronous, although 1/8 cycle adjustments are provided for adjusting the clock alignment.

This figure shows an example of DDR SDRAM configuration with four logical banks.

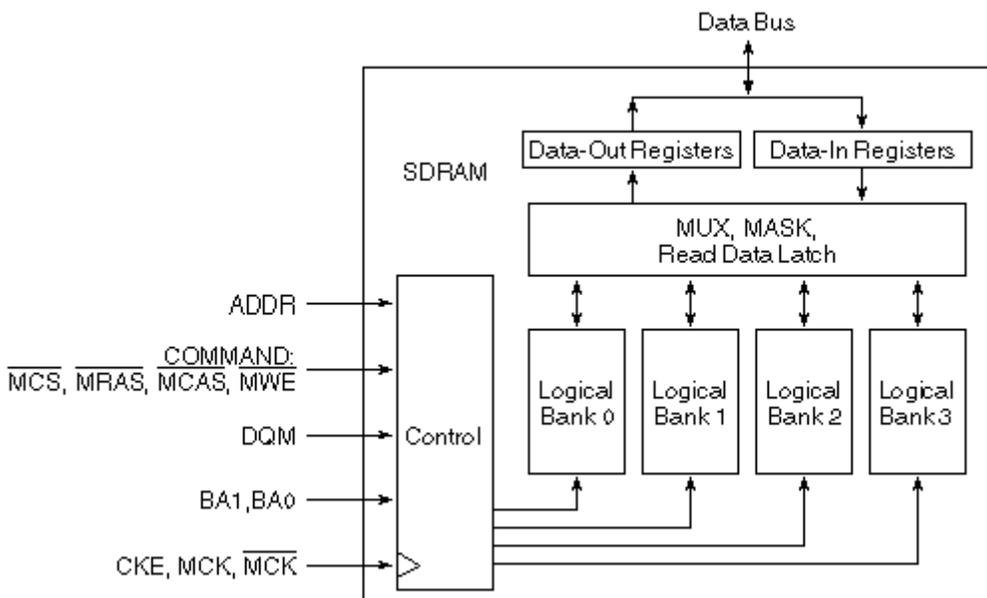


Figure 8-60. Typical dual data rate SDRAM internal organization

This figure shows some typical signal connections.

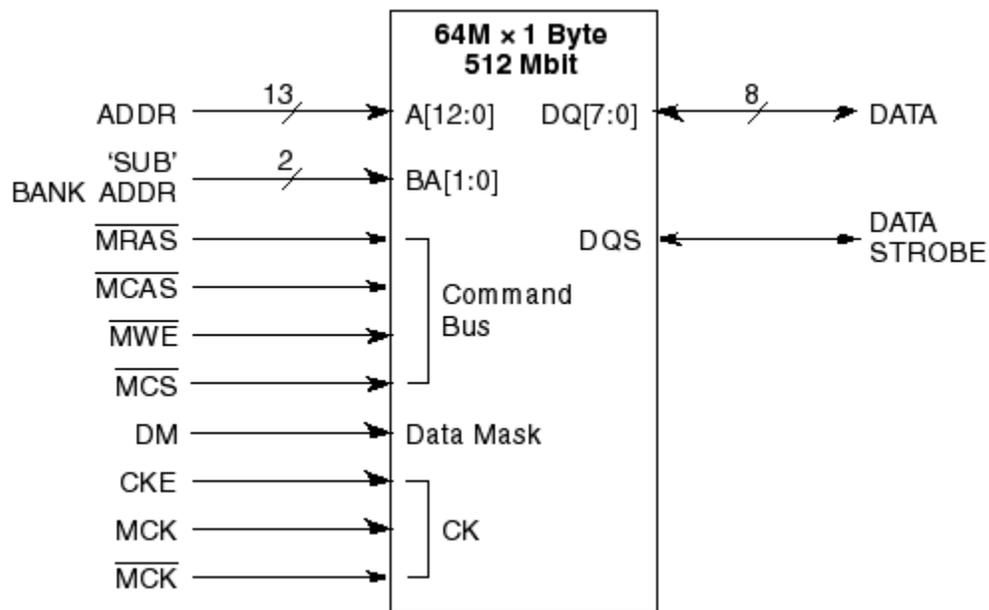
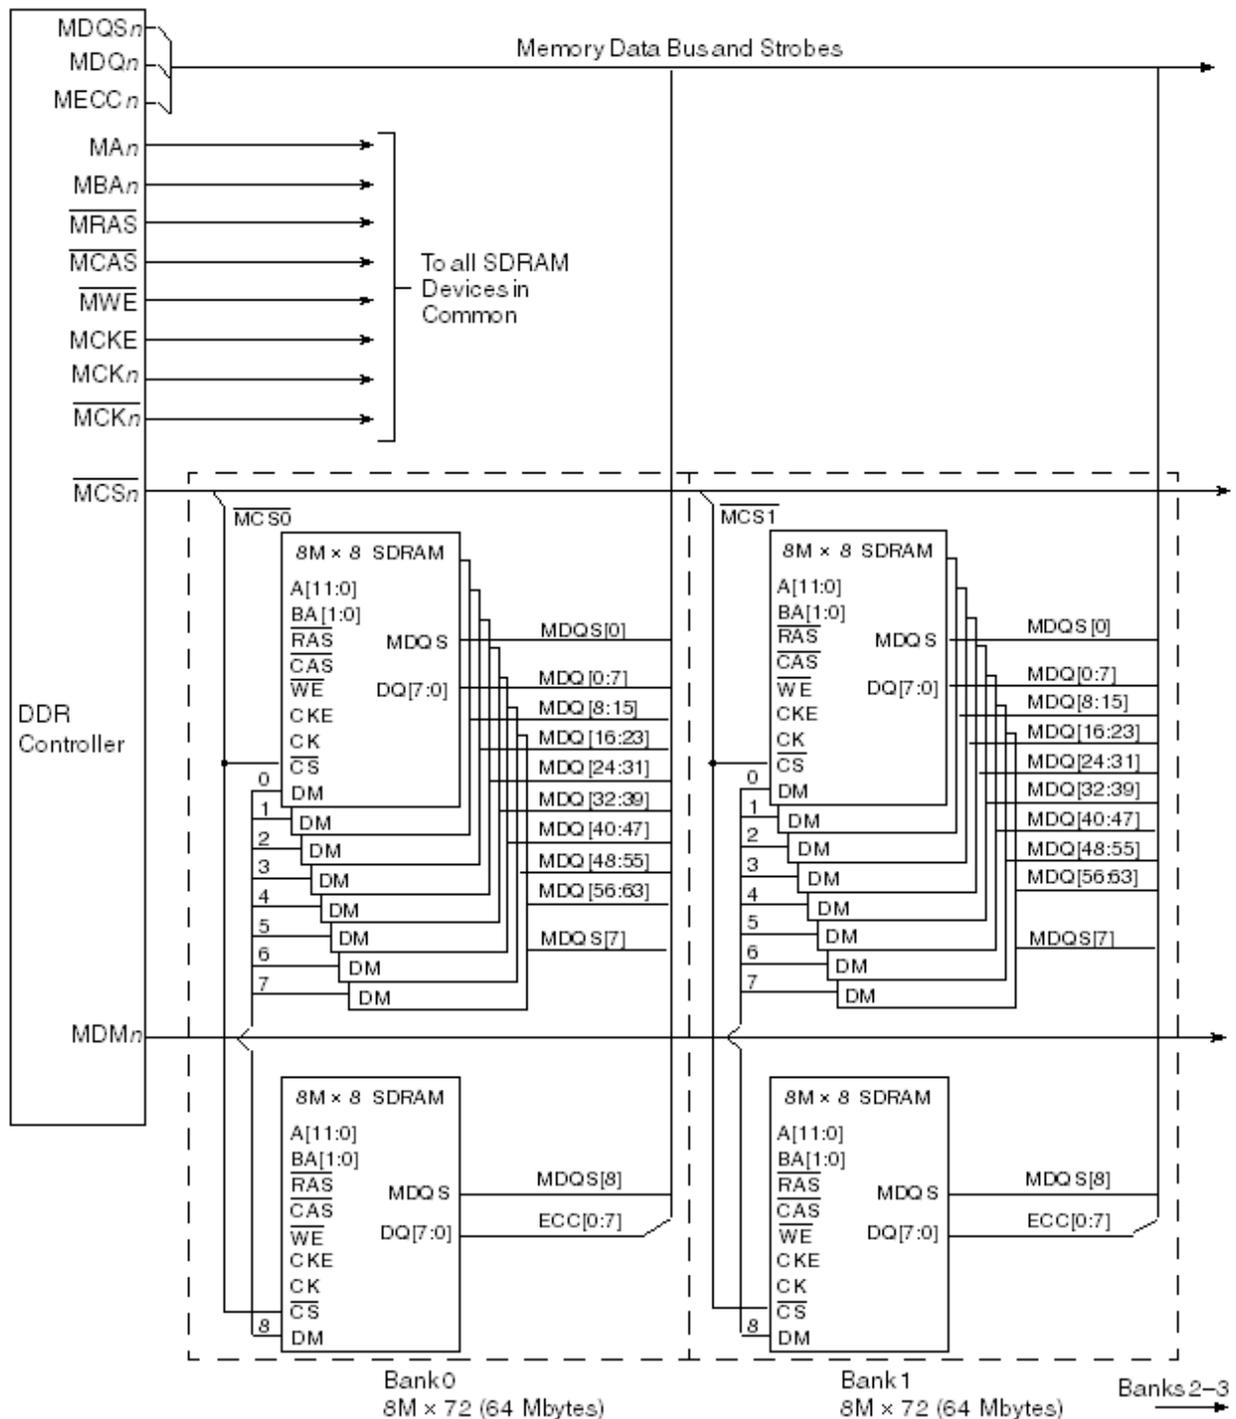


Figure 8-61. Typical DDR SDRAM interface signals

Figure 8-62 shows an example DDR SDRAM configuration with four physical banks each comprised of nine 8Mbyte x 8 DDR modules for a total of 256 Mbytes of system memory. One of the nine modules is used for the memory's ECC checking function. Certain address and control lines may require buffering. Analysis of the device's AC timing specifications, desired memory operating frequency, capacitive loads, and board routing loads can assist the system designer in deciding signal buffering requirements. The DDR memory controller drives 16 address pins, but in this example the DDR SDRAM devices use only 12 bits.



1. All signals are connected in common (in parallel) except for $\overline{MCS}[0:3]$, \overline{MCK}_n , $\overline{MDM}[0:8]$, and the data bus signals.
2. Each of the $\overline{MCS}[0:3]$ signals correspond with a separate physical bank of memory.
3. Buffering may be needed if large memory arrays are used.
4. \overline{MCK}_n may be apportioned among all memory devices. Complementary bus is not shown.

Figure 8-62. Example 256-Mbyte DDR SDRAM configuration with ECC

[Error management](#) explains how the DDR memory controller handles errors.

8.5.1 DDR SDRAM interface operation

The DDR memory controller supports many different DDR SDRAM configurations. SDRAMs with different sizes can be used in the same system.

Sixteen multiplexed address signals and three logical bank select signals support device densities from 64 Mbits to 8 Gbits. Four chip select (\overline{CS}) signals support up to two DIMMs of memory. The DDR SDRAM physical banks can be built from standard memory modules or directly-attached memory devices. The data path to individual physical banks is up to 64- or 32-bits wide, 72 or 40 bits with ECC. The DDR memory controller supports physical bank sizes from 16 Mbytes to 8 Gbytes. The physical banks can be constructed using $\times 8$, $\times 16$, or $\times 32$ memory devices. The memory densities supported are 64 Mbits, 128 Mbits, 256 Mbits, 512 Mbits, 1 Gbit, 2 Gbits, 4 Gbits and 8Gbits. Nine data mask (DM) signals provide byte selection for memory accesses.

NOTE

An 8-bit DDR SDRAM device has a DQM signal and eight data signals (DQ[0:7]). A 16-bit DDR SDRAM device has two DQM signals associated with specific halves of the 16 data signals (DQ[0:7] and DQ[8:15]).

When ECC is enabled, all memory accesses are performed on double-word boundaries (that is, all DQM signals are set simultaneously). However, when ECC is disabled, the memory system uses the DQM signals for byte lane selection.

The table below shows the DDR memory controller's relationships between data byte lane 0-7, MDM[0:7], MDQS[0:7], and MDQ[0:63] when DDR SDRAM memories are used with $\times 8$ or $\times 16$ devices.

Table 8-63. Byte lane to data relationship

Data byte lane	Data bus mask	Data bus strobe	Data bus 64-Bit Mode
0 (MSB)	MDM[0]	MDQS[0]	MDQ[0:7]
1	MDM[1]	MDQS[1]	MDQ[8:15]
2	MDM[2]	MDQS[2]	MDQ[16:23]
3	MDM[3]	MDQS[3]	MDQ[24:31]
4	MDM[4]	MDQS[4]	MDQ[32:39]
5	MDM[5]	MDQS[5]	MDQ[40:47]
6	MDM[6]	MDQS[6]	MDQ[48:55]
7 (LSB)	MDM[7]	MDQS[7]	MDQ[56:63]

NOTE

For a 32-bit bus mode data byte lanes 0-3 in the above table would be used.

8.5.1.1 Supported DDR SDRAM organizations

Although the DDR memory controller multiplexes row and column address bits onto 16 memory address signals and 3 logical bank select signals, a physical bank may be implemented with memory devices requiring fewer than 31 address bits.

The physical bank may be configured to provide from 12 to 16 row address bits, plus 2 or 3 logical bank-select bits and from 8-11 column address bits.

The following tables describe DDR SDRAM device configurations supported by the DDR memory controller.

NOTE

DDR SDRAM is limited to 30 total address bits.

Table 8-64. Supported DDR2 SDRAM device configurations

SDRAM device	Device configuration	Row x column x sub-bank bits	64 -Bit bank size	Four Banks of memory
256 Mbits	32 Mbits x 8	13 x 10 x 2	256 Mbytes	1 Gbyte
256 Mbits	16 Mbits x 16	13 x 9 x 2	128 Mbytes	512 Mbytes
512 Mbits	64 Mbits x 8	14 x 10 x 2	512 Mbytes	2 Gbytes
512 Mbits	32 Mbits x 16	13 x 10 x 2	256 Mbytes	1 Gbyte
1 Gbit	128 Mbits x 8	14 x 10 x 3	1 Gbyte	4 Gbytes
1 Gbit	64 Mbits x 16	13 x 10 x 3	512 Mbytes	2 Gbytes
2 Gbits	256 Mbits x 8	15 x 10 x 3	2 Gbytes	8 Gbytes
2 Gbits	128 Mbits x 16	14 x 10 x 3	1 Gbyte	4 Gbytes
4 Gbits	512 Mbits x 8	16 x 10 x 3	4 Gbytes	16 Gbytes
4 Gbits	256 Mbits x 16	15 x 10 x 3	2 Gbytes	8 Gbytes

Table 8-65. Supported DDR3 SDRAM device configurations

SDRAM device	Device configuration	Row x column x sub-bank bits	64-bit bank size	Four Banks of memory
512 Mbits	64 Mbits x 8	13 x 10 x 3	512 Mbytes	2 Gbytes

Table continues on the next page...

Table 8-65. Supported DDR3 SDRAM device configurations (continued)

SDRAM device	Device configuration	Row x column x sub-bank bits	64-bit bank size	Four Banks of memory
512 Mbits	32 Mbits x 16	12 x 10 x 2	256 Mbytes	1 Gbyte
1 Gbits	128 Mbits x 8	14 x 10 x 3	1 Gbyte	4 Gbytes
1 Gbits	64 Mbits x 16	13 x 10 x 3	512 Mbytes	2 Gbytes
2 Gbits	256 Mbits x 8	15 x 10 x 3	2 Gbytes	8 Gbytes
2 Gbits	128 Mbits x 16	14 x 10 x 3	1 Gbyte	4 Gbytes
4 Gbits	512 Mbits x 8	16 x 10 x 3	4 Gbytes	16 Gbytes
4 Gbits	256 Mbits x 16	15 x 10 x 3	2 Gbytes	8 Gbytes
8 Gbits	1 Gbit x 8	16 x 11 x 3	8 Gbytes	32 Gbytes
8 Gbits	512 Mbits x 16	16 x 10 x 3	4 Gbytes	16 Gbytes

If a transaction request is issued to the DDR memory controller and the address does not lie within any of the programmed address ranges for an enabled chip select, a memory select error is flagged. Errors are described in detail in [Error management](#).

Using a memory-polling algorithm at power-on reset or by querying the JEDEC serial presence detect capability of memory modules, system firmware uses the memory-boundary registers to configure the DDR memory controller to map the size of each bank in memory. The memory controller uses its bank map to assert the appropriate MCS_Bn signal for memory accesses according to the provided bank starting and ending addresses. The memory banks are not required to be mapped to a contiguous address space.

8.5.2 DDR SDRAM address multiplexing

The following tables show the address bit encodings for each DDR SDRAM configuration.

functional description

The address presented at the memory controller signals MA[15:0] use MA[15] as the msb and MA[0] as the lsb. Also, MA[10] is used as the auto-precharge bit in DDR2/DDR3 modes for reads and writes, so the column address can never use MA[10].

Table 8-66. DDR2/DDR3 Address multiplexing for 64-bit data bus with interleaving and partial array self refresh disabled

Row x Col	msb	Address from Core Master																														lsb			
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33		34	35	
16 x 10 x 3	M R A S_ B	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																		
	M B A																	2	1	0															
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0					
15 x 10 x 3	M R A S_ B		14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																		
	M B A																	2	1	0															
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0					
14 x 10 x 3	M R A S_ B			13	12	11	10	9	8	7	6	5	4	3	2	1	0																		
	M B A																	2	1	0															
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0					

Table continues on the next page...

Table 8-66. DDR2/DDR3 Address multiplexing for 64-bit data bus with interleaving and partial array self refresh disabled (continued)

Row x Col	m s b	Address from Core Master																																Is b	
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35		
14 x 10 x 2	M R A S_ B				13	12	11	10	9	8	7	6	5	4	3	2	1	0																	
	M B A																		1	0															
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0					
13 x 10 x 3	M R A S_ B				12	11	10	9	8	7	6	5	4	3	2	1	0																		
	M B A																		2	1	0														
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0					
13 x 10 x 2	M R A S_ B				12	11	10	9	8	7	6	5	4	3	2	1	0																		
	M B A																			1	0														
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0					

Table continues on the next page...

Table 8-67. DDR2/DDR3 Address multiplexing for 32-bit data bus with interleaving and partial array self-refresh disabled (continued)

Row x Col	ms b	Address from Core Master																																Isb					
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34-35							
14 x 10 x 2	MRA S_B						1	1	11	1	9	8	7	6	5	4	3	2	1	0																			
	MBA																																						
	MCA S_B																																						
13 x 10 x 3	MRA S_B						1	1	10	9	8	7	6	5	4	3	2	1	0																				
	MBA																																						
	MCA S_B																																						
13 x 10 x 2	MRA S_B						1	11	1	9	8	7	6	5	4	3	2	1	0																				
	MBA																																						
	MCA S_B																																						
13 x 9 x 2	MRA S_B																																						
	MBA																																						
	MCA S_B																																						

Table 8-68. DDR2/DDR3 Address multiplexing for 16-bit data bus

Row x Col	ms b	Address from Core Master																																Isb						
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35							
15 x 10 x 3	M R A S _ B					14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																				
	M B A																																							
	M C A S _ B																																							

Table continues on the next page...

Table 8-68. DDR2/DDR3 Address multiplexing for 16-bit data bus (continued)

Row x Col	msb	Address from Core Master																																lsb	
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35		
14 x 10 x 3	M R A S _ B						13	12	11	10	9	8	7	6	5	4	3	2	1	0															
	M B A																				2	1	0												
	M C A S _ B																						9	8	7	6	5	4	3	2	1	0			
14 x 10 x 2	M R A S _ B						13	12	11	10	9	8	7	6	5	4	3	2	1	0															
	M B A																					1	0												
	M C A S _ B																						9	8	7	6	5	4	3	2	1	0			
13 x 10 x 3	M R A S _ B						12	11	10	9	8	7	6	5	4	3	2	1	0																
	M B A																					2	1	0											
	M C A S _ B																						9	8	7	6	5	4	3	2	1	0			

Table continues on the next page...

Table 8-68. DDR2/DDR3 Address multiplexing for 16-bit data bus (continued)

Row x Col	msb	Address from Core Master																												lsb					
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		32	33	34	35	
13 x 10 x 2	M R A S _ B									12	11	10	9	8	7	6	5	4	3	2	1	0													
	M B A																						1	0											
	M C A S _ B																								9	8	7	6	5	4	3	2	1	0	
13 x 9 x 2	M R A S _ B									12	11	10	9	8	7	6	5	4	3	2	1	0													
	M B A																							1	0										
	M C A S _ B																								8	7	6	5	4	3	2	1	0		

Chip select interleaving is supported for the memory controller, and is programmed in DDR_SDRAM_CFG[BA_INTLV_CTL]. Interleaving is supported between chip selects 0 and 1 or chip selects 2 and 3. In addition, interleaving between all four chip selects can be enabled. When interleaving is enabled, the chip selects being interleaved must use the same size of memory. If two chip selects are interleaved, then 1 extra bit in the address decode is used for the interleaving to determine which chip select to access. If four chip selects are interleaved, then two extra bits are required in the address decode.

The following tables show examples of interleaving between chip selects.

Table 8-69. Example of address multiplexing for 64-bit data bus interleaving between two banks with partial array self refresh disabled

Row x Col	msb	Address from core master																																lsb				
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35					
14 x 10 x 3	M R A S_ B		13	12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																					
	M B A																	2	1	0																		
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0								
14 x 10 x 2	M R A S_ B			13	12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																				
	M B A																	1	0																			
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0								
13 x 10 x 3	M R A S_ B			12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																					
	M B A																	2	1	0																		
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0								

Table continues on the next page...

Table 8-69. Example of address multiplexing for 64-bit data bus interleaving between two banks with partial array self refresh disabled (continued)

Row x Col	m s b	Address from core master																																Is b				
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35					
13 x 10 x 2	M R A S_ B				12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																				
	M B A																		1	0																		
	M C A S_ B																					9	8	7	6	5	4	3	2	1	0							

Table 8-70. Example of address multiplexing for 32-bit data bus interleaving between two banks with partial array self refresh disabled

Row x Col	m s b	Address from Core Master																																Is b				
		5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35						
14 x 10 x 3	M R A S_ B		13	12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																					
	M B A																	2	1	0																		
	M C A S_ B																				9	8	7	6	5	4	3	2	1	0								

Table continues on the next page...

Table 8-70. Example of address multiplexing for 32-bit data bus interleaving between two banks with partial array self refresh disabled (continued)

Row x Col	m s b	Address from Core Master																												Is b					
		5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32		33	34	35		
14 x 10 x 2	M R A S_ B			13	12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																	
	M B A																		1	0															
	M C A S_ B																						9	8	7	6	5	4	3	2	1	0			
13 x 10 x 3	M R A S_ B			12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																		
	M B A																		2	1	0														
	M C A S_ B																						9	8	7	6	5	4	3	2	1	0			
13 x 10 x 2	M R A S_ B			12	11	10	9	8	7	6	5	4	3	2	1	0	C S S E L																		
	M B A																		1	0															
	M C A S_ B																						9	8	7	6	5	4	3	2	1	0			

Table 8-71. Example of address multiplexing for 64-bit data bus interleaving between four banks with partial array self refresh disabled

Row x Col	m s b	Address from Core Master																																Is b			
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35				
14 x 10 x 3	M R A S_ B	13	12	11	10	9	8	7	6	5	4	3	2	1	0	CS SEL																					
	M B A																2	1	0																		
	M C A S_ B																			9	8	7	6	5	4	3	2	1	0								
14 x 10 x 2	M R A S_ B		13	12	11	10	9	8	7	6	5	4	3	2	1	0	CS SEL																				
	M B A																1	0																			
	M C A S_ B																			9	8	7	6	5	4	3	2	1	0								
13 x 10 x 3	M R A S_ B		12	11	10	9	8	7	6	5	4	3	2	1	0	CS SEL																					
	M B A																2	1	0																		
	M C A S_ B																			9	8	7	6	5	4	3	2	1	0								

Table continues on the next page...

Table 8-71. Example of address multiplexing for 64-bit data bus interleaving between four banks with partial array self refresh disabled (continued)

Row x Col	m s b	Address from Core Master																														Is b				
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33		34	35		
13 x 10 x 2	M R A S_ B			12	11	10	9	8	7	6	5	4	3	2	1	0	CS SEL																			
	M B A																	1	0																	
	M C A S_ B																			9	8	7	6	5	4	3	2	1	0							

Partial array self-refresh (PASR) can be enabled for any chip select using the CS_n_CONFIG_2[PASR_CFG] fields. If PASR is enabled for a given chip select, then the sub-bank and row decode is swapped, and the sub-bank is decoded as the most significant portion of the DRAM address, as shown in the table below. If chip select interleaving and PASR are enabled for a chip select, then the interleaved chip select bit is placed immediately to the left of the column decode, as shown in [Table 8-73](#).

Table 8-72. Address multiplexing with partial array self refresh Enabled

Row x Col	m s b	Address from Core Master																																Is b	
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34-35			
16 x 10 x 3	M R A S - B					15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0														
	M B A		2	1	0																														
	M C A S - B																					9	8	7	6	5	4	3	2	1	0				
15 x 10 x 3	M R A S - B						14	13	12	11	10	9	8	7	6	5	4	3	2	1	0														
	M B A			2	1	0																													
	M C A S - B																					9	8	7	6	5	4	3	2	1	0				
14 x 10 x 3	M R A S - B							13	12	11	10	9	8	7	6	5	4	3	2	1	0														
	M B A				2	1	0																												
	M C A S - B																					9	8	7	6	5	4	3	2	1	0				

Table continues on the next page...

Table 8-72. Address multiplexing with partial array self refresh Enabled (continued)

Row x Col	m s b	Address from Core Master																																Is b	
		4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34-35			
14 x 10 x 2	M R A S - B									13	12	11	10	9	8	7	6	5	4	3	2	1	0												
	M B A					1	0																												
	M C A S - B																							9	8	7	6	5	4	3	2	1	0		
13 x 10 x 3	M R A S - B									12	11	10	9	8	7	6	5	4	3	2	1	0													
	M B A					2	1	0																											
	M C A S - B																							9	8	7	6	5	4	3	2	1	0		
13 x 10 x 2	M R A S - B									12	11	10	9	8	7	6	5	4	3	2	1	0													
	M B A							1	0																										
	M C A S - B																							9	8	7	6	5	4	3	2	1	0		

Table continues on the next page...

Table 8-73. Example of address multiplexing for 64-bit data bus interleaving between two banks with partial array self refresh enabled (continued)

Row x Col	msb		Address from Core Master																								lsb					
	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29		30	31	32	33-35	
13 x 10	MRAS					1	1	1	9	8	7	6	5	4	3	2	1	0	CS SE L													
x 2	MBA			1	0																											
	MCAS																				9	8	7	6	5	4	3	2	1	0		

8.5.3 JEDEC standard DDR SDRAM interface commands

The following section describes the commands and timings the controller uses when operating in DDR3 or DDR2 modes.

All read or write accesses to DDR SDRAM are performed by the DDR memory controller using JEDEC standard DDR SDRAM interface commands. The SDRAM device samples command and address inputs on rising edges of the memory clock; data is sampled using both the rising and falling edges of DQS. Data read from the DDR SDRAM is also sampled on both edges of DQS.

The following DDR SDRAM interface commands (summarized in the table below) are provided by the DDR controller. All actions for these commands are described from the perspective of the SDRAM device.

- Row activate-Latches row address and initiates memory read of that row. Row data is latched in SDRAM sense amplifiers and must be restored by a precharge command before another row activate occurs.
- Precharge-Restores data from the sense amplifiers to the appropriate row. Also initializes the sense amplifiers in preparation for reading another row in the memory array (performing another activate command). Precharge must occur after read or write, if the row address changes on the next open page mode access.
- Read-Latches column address and transfers data from the selected sense amplifier to the output buffer as determined by the column address. During each succeeding clock edge, additional data is driven without additional read commands. The amount of data transferred is determined by the burst size.
- Write-Latches column address and transfers data from the data pins to the selected sense amplifier as determined by the column address. During each succeeding clock edge, additional data is transferred to the sense amplifiers from the data pins without

additional write commands. The amount of data transferred is determined by the data masks and the burst size.

- Refresh (similar to MCAS_B before MRAS_B)-Causes a row to be read in all logical banks (JEDEC SDRAM) as determined by the refresh row address counter. This refresh row address counter is internal to the SDRAM. After being read, the row is automatically rewritten in the memory array. All logical banks must be in a precharged state before executing a refresh. The memory controller also supports posted refreshes, where several refreshes may be executed at once, and the refresh interval may be extended.
- Mode register set (for configuration)-Allows setting of DDR SDRAM options. These options are: MCAS_B latency, additive latency (for DDR2/DDR3), write recovery (for DDR2/DDR3), burst type, and burst length. MCAS_B latency may be chosen as provided by the preferred SDRAM. Burst type is always sequential. This memory controller supports a burst length of 4. A burst length of 8 is supported for DDR3 memory only. For DDR2 in 32-bit bus mode, all 32-byte burst accesses from the platform are split into two 16-byte (that is, 4-beat) accesses to the SDRAMs in the memory controller. The mode register set command is performed by the DDR memory controller during system initialization. Parameters such as mode register data, MCAS_B latency, burst length, and burst type, are set by software in DDR_SDRAM_MODE[SDMODE] and transferred to the SDRAM array by the DDR memory controller after DDR_SDRAM_CFG[MEM_EN] is set. If DDR_SDRAM_CFG[BI] is set to bypass the automatic initialization, then the MODE registers can be configured through software through use of the DDR_SDRAM_MD_CNTL register.
- Self refresh (for long periods of standby)-Used when the device is in standby for very long periods of time. Automatically generates internal refresh cycles to keep the data in all memory banks refreshed. Before execution of this command, the DDR controller places all logical banks in a precharged state.

Table 8-74. DDR SDRAM command table

Operation	CKE Prev.	CKE Current	MCS _B	MRAS _B	MCAS _B	MWE _B	MBA	MA10	MA
Activate	H	H	L	L	H	H	Logical bank select	Row	Row
Precharge select logical bank	H	H	L	L	H	L	Logical bank select	L	X
Precharge all logical banks	H	H	L	L	H	L	X	H	X
Read	H	H	L	H	L	H	Logical bank select	L	Column
Read with auto- precharge	H	H	L	H	L	H	Logical bank select	H	Column
Write	H	H	L	H	L	L	Logical bank select	L	Column

Table continues on the next page...

Table 8-74. DDR SDRAM command table (continued)

Operation	CKE Prev.	CKE Current	MCS _B	MRAS _B	MCAS _B	MWE _B	MBA	MA10	MA
Write with auto-precharge	H	H	L	H	L	L	Logical bank select	H	Column
Mode register set	H	H	L	L	L	L	Opcode	Opcode	Opcode and mode
Auto refresh	H	H	L	L	L	H	X	X	X
Self refresh	H	L	L	L	L	H	X	X	X

8.5.4 DDR SDRAM interface timing

The DDR memory controller supports four- (or eight-) beat bursts to SDRAM.

For single-beat reads, the DDR memory controller performs a four- (or eight-) beat burst read, but ignores the last three (or seven) beats. Single-beat writes are performed by masking the last three (or seven) beats of the four- (or eight-) beat burst using the data mask MDM[0:8]. If ECC is disabled, writes smaller than double words are performed by appropriately activating the data mask. If ECC is enabled, the controller performs a read-modify write.

NOTE

If a second read or write is pending, reads shorter than four beats are not terminated early even if some data is irrelevant.

To accommodate available memory technologies across a wide spectrum of operating frequencies, the DDR memory controller allows the setting of the intervals defined in the following table with granularity of one memory clock cycle.

Table 8-75. DDR SDRAM interface timing intervals

Timing intervals	Definition
ACTTOACT	The number of clock cycles from a bank-activate command until another bank-activate command within a physical bank. This interval is listed in the AC specifications of the SDRAM as t_{RRD} .
ACTTOPRE	The number of clock cycles from an activate command until a precharge command is allowed. This interval is listed in the AC specifications of the SDRAM as t_{RAS} .
ACTTORW	The number of clock cycles from an activate command until a read or write command is allowed. This interval is listed in the AC specifications of the SDRAM as t_{RCD} .
BSTOPRE	The number of clock cycles to maintain a page open after an access. The page open duration counter is reloaded with BSTOPRE each time the page is accessed (including page hits). When the counter expires, the open page is closed with an SDRAM precharge bank command as soon as possible.
CASLAT	Used in conjunction with additive latency to obtain the READ latency. The number of clock cycles between the registration of a READ command by the SDRAM and the availability of the first piece of output data. If a READ command is registered at clock edge n , and the read latency is m clocks, the data is available nominally coincident with clock edge $n + m$.

Table continues on the next page...

Table 8-75. DDR SDRAM interface timing intervals (continued)

Timing intervals	Definition
PRETOACT	The number of clock cycles from a precharge command until an activate or a refresh command is allowed. This interval is listed in the AC specifications of the SDRAM as t_{RP} .
REFINT	Refresh interval. Represents the number of memory bus clock cycles between refresh cycles. Depending on <code>DDR_SDRAM_CFG_2[NUM_PR]</code> , some number of rows are refreshed in each SDRAM bank during each refresh cycle. The value of REFINT depends on the specific SDRAMs used and the frequency of the interface as t_{RP} .
REFREC	The number of clock cycles from the refresh command until an activate command is allowed. This can be calculated by referring to the AC specification of the SDRAM device. The AC specification indicates a maximum refresh-to-activate interval in nanoseconds.
WR_DATA_DELAY	Provides different options for the timing between a write command and the write data strobe. This allows write data to be sent later than the nominal time to meet the SDRAM timing requirement between the registration of a write command and the reception of a data strobe associated with the write command. The specification dictates that the data strobe may not be received earlier than 75% of a cycle, or later than 125% of a cycle, from the registration of a write command. This parameter is not defined in the SDRAM specification. It is implementation-specific, defined for the DDR memory controller in <code>TIMING_CFG_2</code> .
WRREC	The number of clock cycles from the last beat of a write until a precharge command is allowed. This interval, write recovery time, is listed in the AC specifications of the SDRAM as t_{WR} .
WRTORD	Last write pair to read command. Controls the number of clock cycles from the last write data pair to the subsequent read command to the same bank as t_{WTR} .

The value of the above parameters (in whole clock cycles) must be set by boot code at system start-up (as described in [DDR SDRAM timing configuration 0 \(DDR_TIMING_CFG_0\)](#), [DDR SDRAM timing configuration 1 \(DDR_TIMING_CFG_1\)](#), [DDR SDRAM timing configuration 2 \(DDR_TIMING_CFG_2\)](#), and [DDR SDRAM timing configuration 3 \(DDR_TIMING_CFG_3\)](#)) and be kept in the DDR memory controller configuration register space.

The following figures show SDRAM timing for various types of accesses. System software is responsible (at reset) for optimally configuring SDRAM timing parameters. The programmable timing parameters apply to both read and write timing configuration. The configuration process must be completed and the DDR SDRAM initialized before any accesses to SDRAM are attempted.

See [Figure 8-63](#) for a single-beat read operation, [Figure 8-64](#) for a single-beat write operation, and [Figure 8-65](#) for a double-word write operation. Note that all signal transitions occur on the rising edge of the memory bus clock and that single-beat read operations are identical to burst-reads. These figures assume the `CLK_ADJUST` is set to 1/2 DRAM cycle, an additive latency of 0 DRAM cycles is used, and the write latency is 1 DRAM cycle.

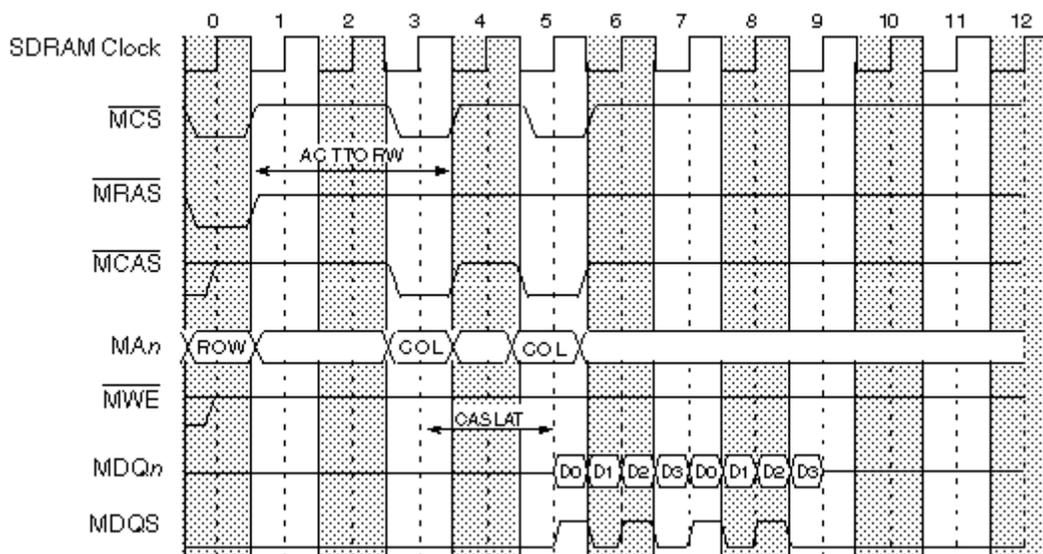


Figure 8-63. DDR SDRAM burst read timing- $ACTTORW = 3$, \overline{MCAS} Latency = 2

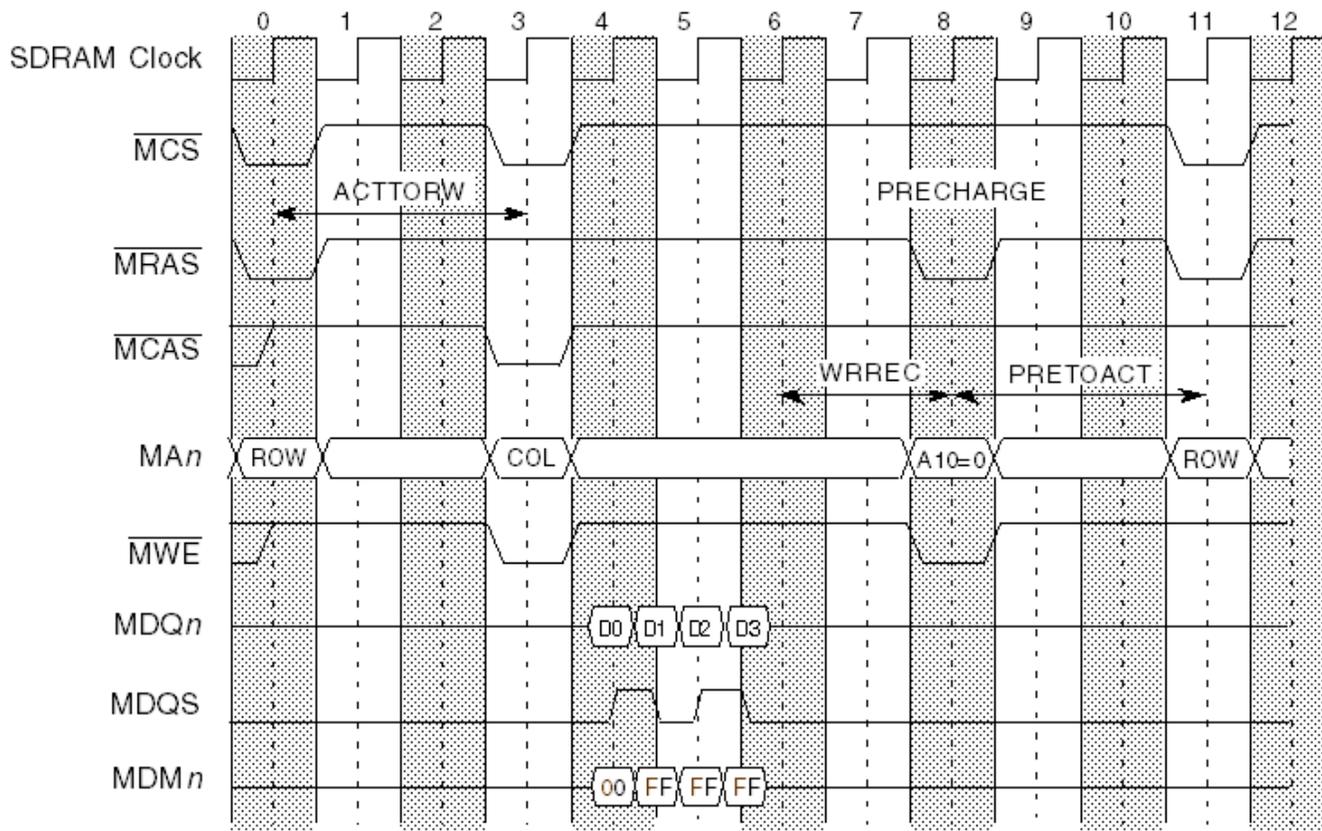


Figure 8-64. DDR SDRAM single-beat (double word) write timing- $ACTTORW = 3$

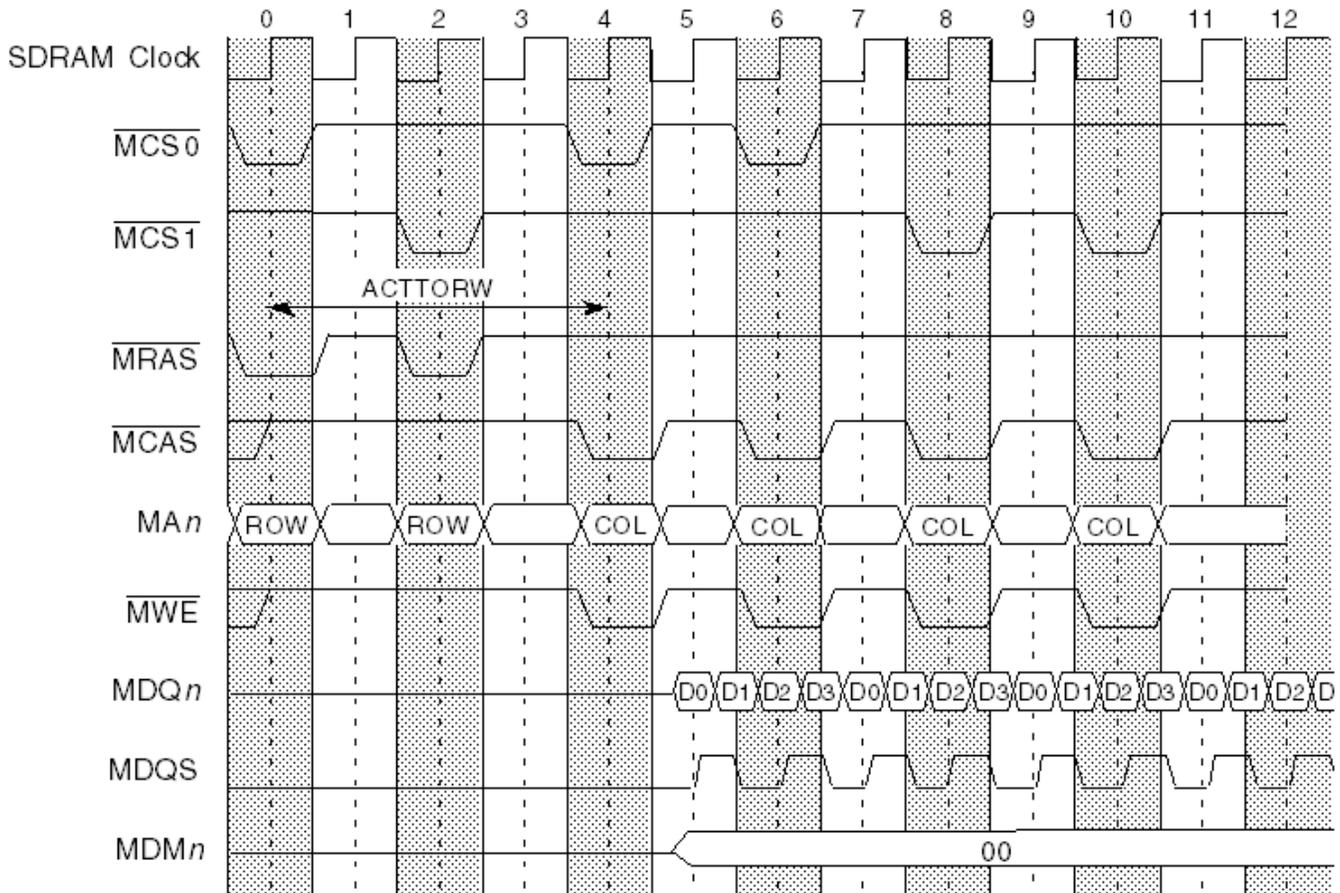


Figure 8-65. DDR SDRAM 4-beat burst write timing- $ACTTORW = 4$

8.5.5 DDR SDRAM registered DIMM mode

To reduce loading, registered DIMMs latch the DDR SDRAM control signals internally before using them to access the array.

Setting `DDR_SDRAM_CFG[RD_EN]` compensates for this delay on the DIMMs' control bus by delaying the data and data mask writes (on SDRAM buses) by an extra SDRAM clock cycle.

NOTE

Application system board must assert the reset signal on DDR memory devices until software is able to program the DDR memory controller configuration registers, and must deassert the reset signal on DDR memory devices before `DDR_SDRAM_CFG[MEM_EN]` is set. This ensures that the DDR memory devices are held in reset until a stable clock is provided and, further, that a stable clock is provided before memory devices are released from reset.

The figure below shows the registered DDR SDRAM DIMM burst write timing.

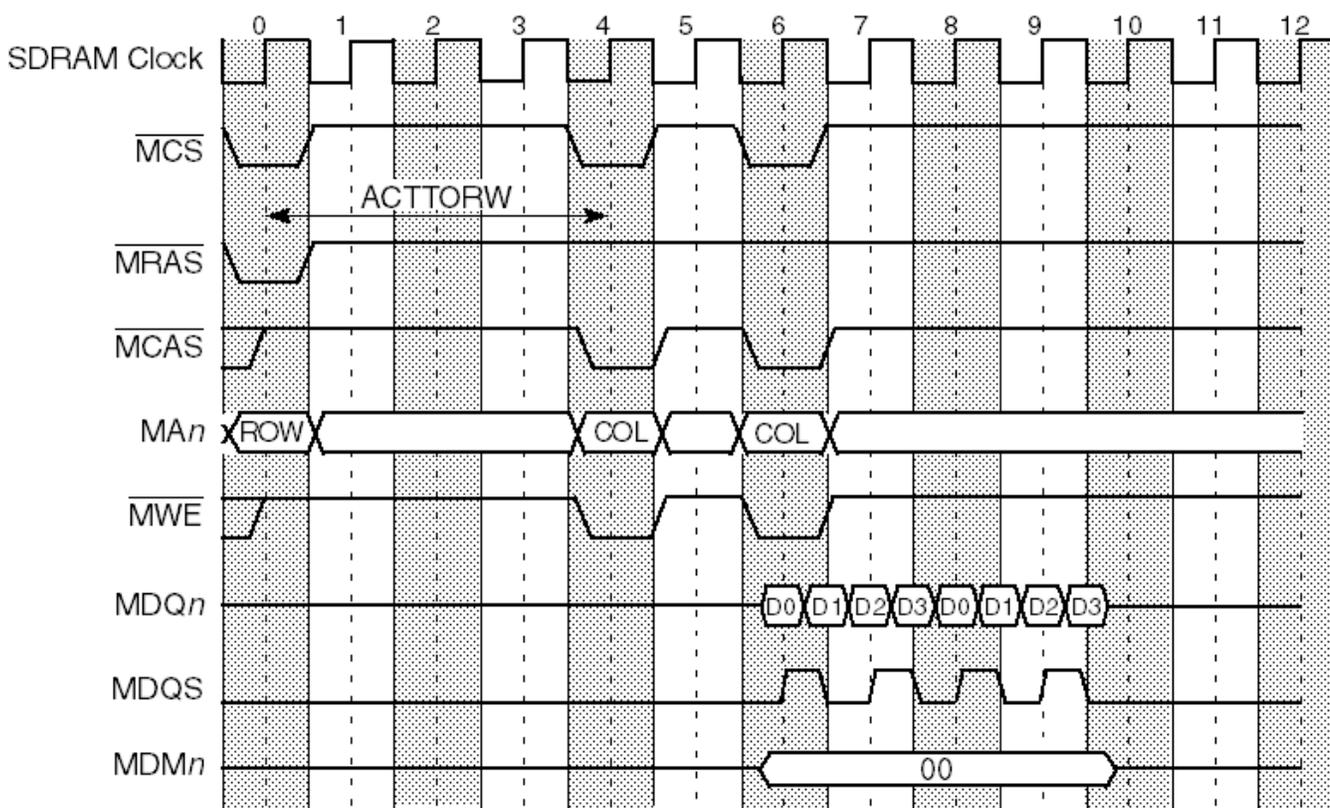


Figure 8-66. Registered DDR SDRAM DIMM burst write timing

8.5.6 DDR SDRAM write timing adjustments

The DDR memory controller facilitates system design flexibility by providing a write timing adjustment parameter, write data delay, (TIMING_CFG_2[WR_DATA_DELAY]) for data and DQS.

The DDR SDRAM specification requires DQS be received no sooner than 75% of an SDRAM clock period-and no later than 125% of a clock period-from the capturing clock edge of the command/address at the SDRAM. The WR_DATA_DELAY parameter may be used to meet this timing requirement for a variety of system configurations, ranging from a system with one DIMM to a fully populated system with two DIMMs.

TIMING_CFG_2[WR_DATA_DELAY] specifies how much to delay the launching of DQS and data from the first clock edge occurring. The delay increment step sizes are in 1/4 SDRAM clock periods starting with the default value of 0.

The figure below shows the use of the WR_DATA_DELAY parameter.

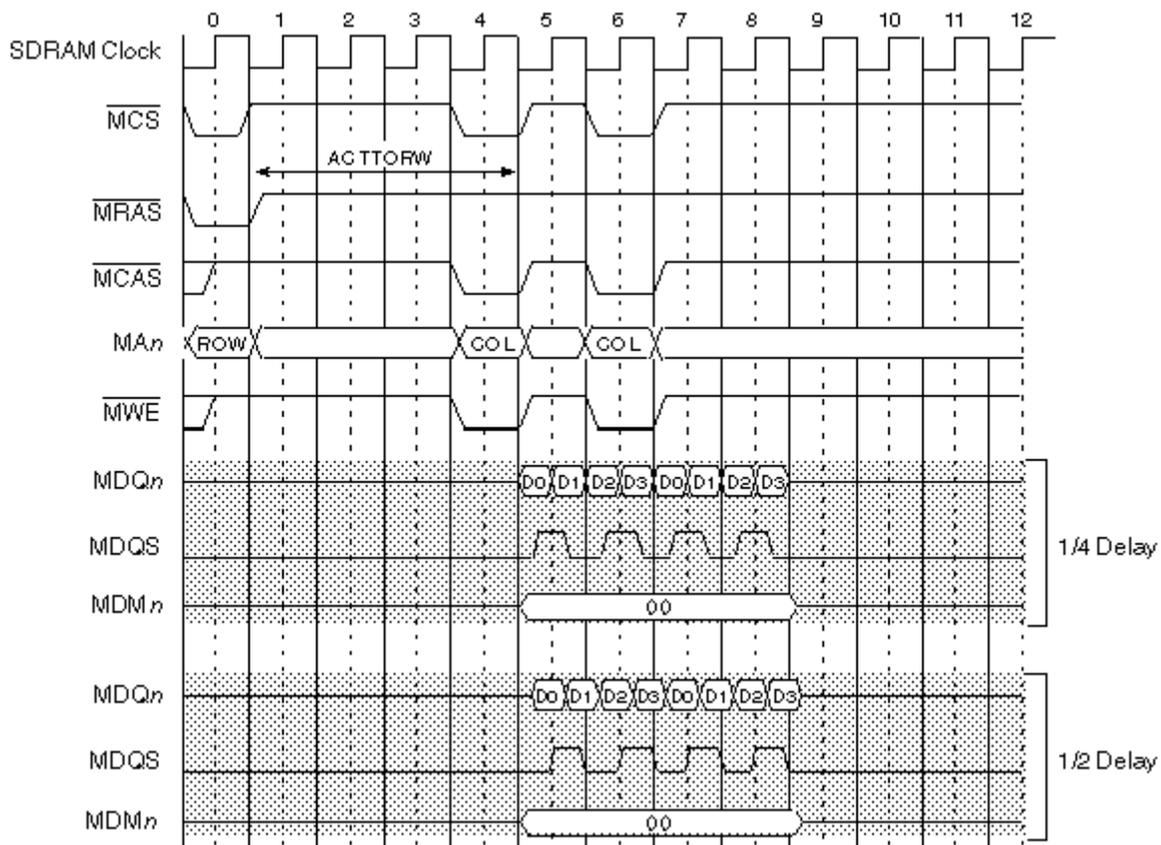


Figure 8-67. Write timing adjustments example for write latency = 1

8.5.7 DDR SDRAM refresh

The DDR memory controller supports auto-refresh and self-refresh.

Auto-refresh is used during normal operation and is controlled by the DDR_SDRAM_INTERVAL[REFINT] value; self-refresh is used only when the DDR memory controller is set to enter a sleep power management state. The REFINT value, which represents the number of memory bus clock cycles between refresh cycles, must allow for possible outstanding transactions to complete before a refresh request is sent to the memory after the REFINT value is reached. If a memory transaction is in progress when the refresh interval is reached, the refresh cycle waits for the transaction to complete. In the worst case, the refresh cycle must wait the number of bus clock cycles required by the longest programmed access.

When a refresh cycle is required, the DDR memory controller does the following:

1. Completes all current memory requests.
2. Closes all open pages with a PRECHARGE-ALL command to each DDR SDRAM bank with an open page (as indicated by the row open table).

- Issues one or more auto-refresh commands to each DDR SDRAM bank (as identified by its chip select) to refresh one row in each logical bank of the selected physical bank.

The auto-refresh commands are staggered across the four possible banks to reduce the system's instantaneous power requirements. Three sets of auto refresh commands are issued on consecutive cycles when the memory is fully populated with two DIMMs. The initial PRECHARGE-ALL commands are also staggered in three groups for convenience. It is important to note that when entering self-refresh mode, only one refresh command is issued simultaneously to all physical banks. For this entire refresh sequence, no cycle optimization occurs for the usual case where fewer than four banks are installed. After the refresh sequence completes, any pending memory request is initiated after an inactive period specified by TIMING_CFG_1 [REFREC] and TIMING_CFG_3[EXT_REFREC]. In addition, posted refreshes are supported to allow the refresh interval to be set to a larger value.

8.5.7.1 DDR SDRAM refresh timing

Refresh timing for the DDR SDRAM is controlled by the programmable timing parameter TIMING_CFG_1 [REFREC], which specifies the number of memory bus clock cycles from the refresh command until a logical bank activate command is allowed.

The DDR memory controller implements bank staggering for refreshes, as shown in the figure below (TIMING_CFG_1 [REFREC] = 10 in this example).

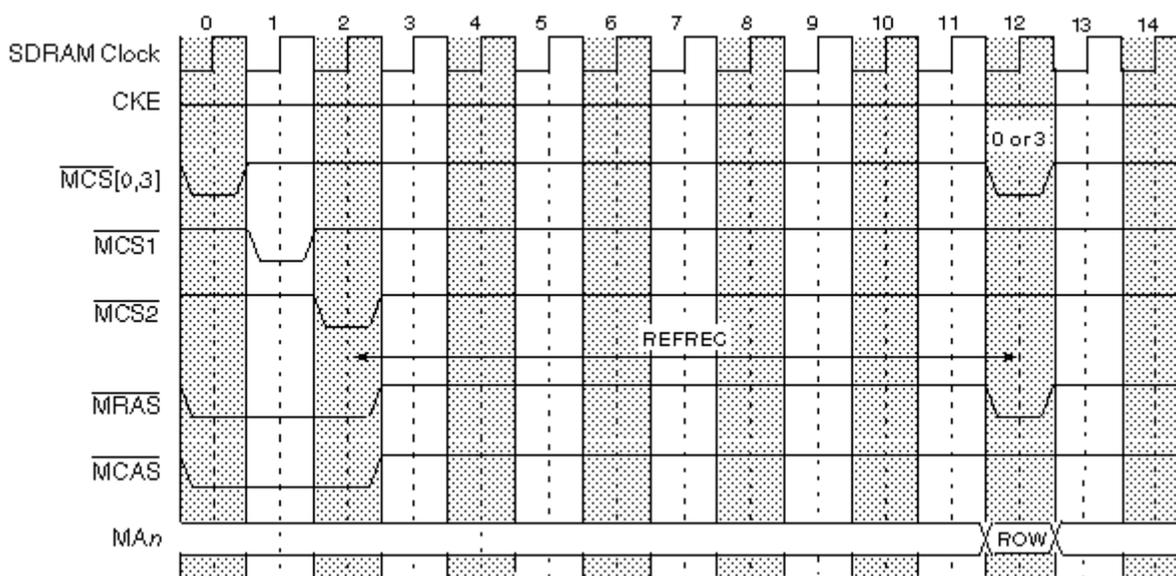


Figure 8-68. DDR SDRAM bank staggered auto refresh timing

System software is responsible for optimal configuration of TIMING_CFG_1 [REFREC] and TIMING_CFG_3[EXT_REFREC] at reset. Configuration must be completed before DDR SDRAM accesses are attempted.

8.5.7.2 DDR SDRAM refresh and power-saving modes

In full-on mode, the DDR memory controller supplies the normal auto refresh to SDRAM.

In sleep mode, the DDR memory controller can be configured to take advantage of self-refreshing SDRAMs or to provide no refresh support. Self-refresh support is enabled with the SREN memory control parameter.

The table below summarizes the refresh types available in each power-saving mode.

Table 8-76. DDR SDRAM power-saving modes refresh configuration

Power saving mode	Refresh type	SREN
Sleep	Self	1
	None	-

Note that in the absence of refresh support, system software must preserve DDR SDRAM data (such as by copying the data to disk) before entering the power-saving mode.

The dynamic power-saving mode uses the CKE DDR SDRAM pin to dynamically power down when there is no system memory activity. The CKE pin is negated when both of the following conditions are met:

- No memory refreshes are scheduled
- No memory accesses are scheduled

CKE is reasserted when a new access or refresh is scheduled or the dynamic power mode is disabled. This mode is controlled with DDR_SDRAM_CFG[DYN_PWR_MGMT].

Dynamic power management mode offers tight control of the memory system's power consumption by trading power for performance through the use of CKE. Powering up the DDR SDRAM when a new memory reference is scheduled causes an access latency penalty, depending on whether active or precharge powerdown is used, along with the settings of TIMING_CFG_0[ACT_PD_EXIT] and TIMING_CFG_0[PRE_PD_EXIT]. For example, a penalty of 1 cycle is shown in the following figure.

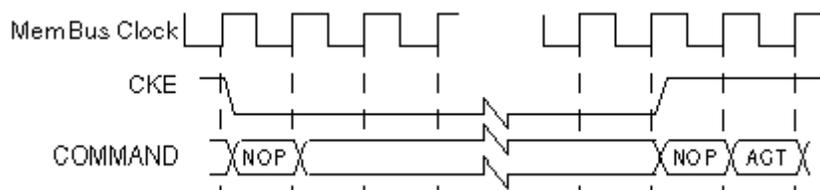


Figure 8-69. DDR SDRAM power-down mode

8.5.7.2.1 Self-refresh in sleep mode

The entry and exit timing for self-refreshing SDRAMs is shown in the following figures.

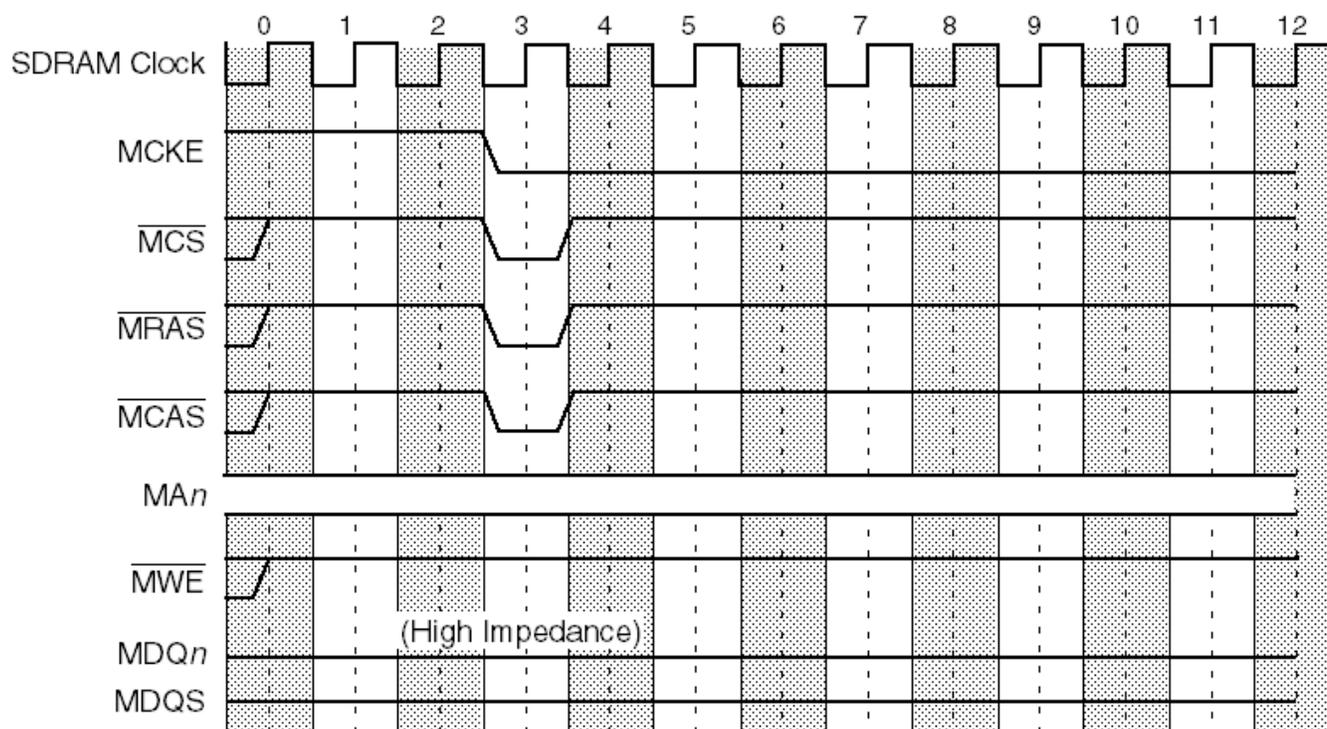


Figure 8-70. DDR SDRAM self-refresh entry timing

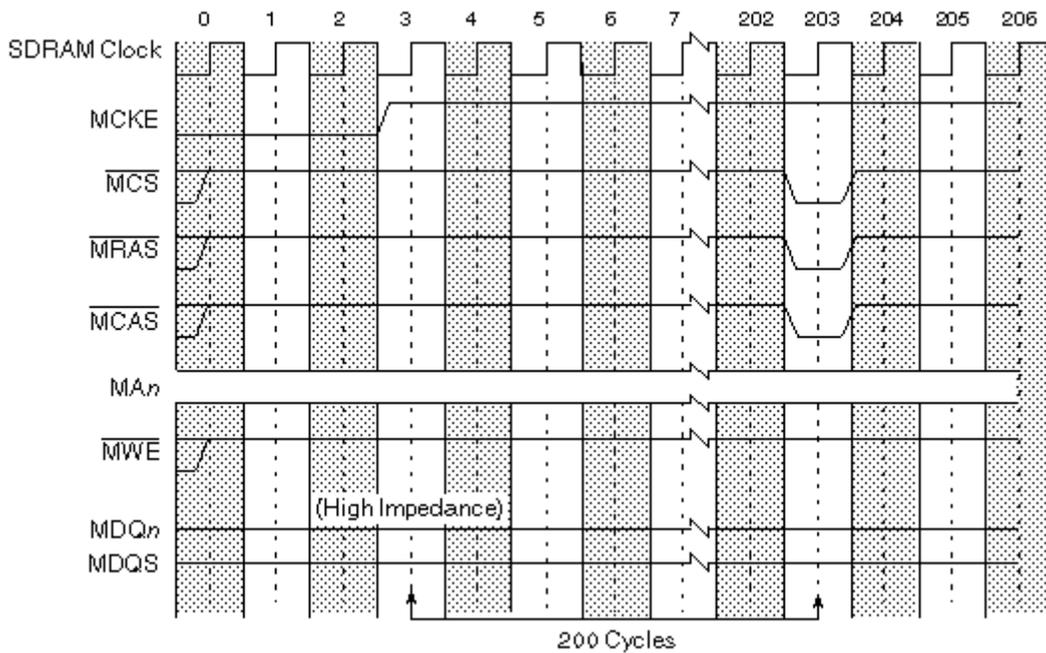


Figure 8-71. DDR2 SDRAM self-refresh exit timing

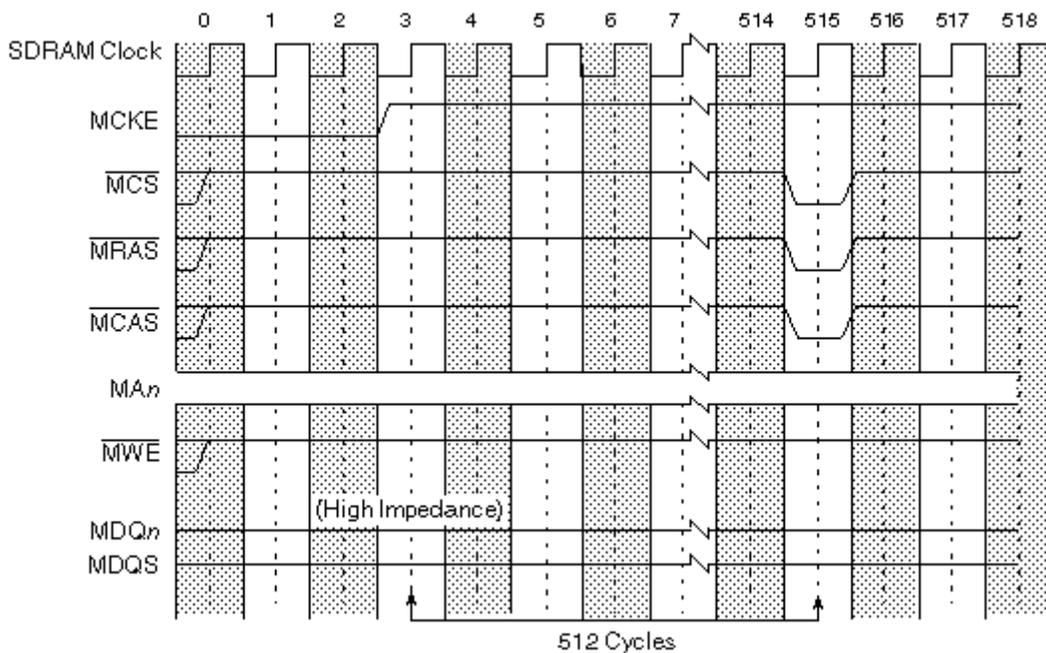


Figure 8-72. DDR3 SDRAM self-refresh exit timing

NOTE

Deep sleep mode is not supported for DDR3 registered DIMMS.

8.5.8 DDR data beat ordering

Transfers to and from memory are always performed in four- or eight-beat bursts (four beats = 32 bytes when a 64-bit bus is used).

For transfer sizes other than four or eight beats, the data transfers are still operated as four- or eight-beat bursts. If ECC is enabled and either the access is not double-word aligned or the size is not a multiple of a double word, a full read-modify-write is performed for a write to SDRAM. If ECC is disabled or both the access is double-word aligned with a size that is a multiple of a double word, the data masks MDM[0:8] (MDM[0:3] and MDM[8] for 32-bit bus) can be used to prevent the writing of unwanted data to SDRAM. The DDR memory controller also uses data masks to prevent all unintended full double words from writing to SDRAM. For example, if a write transaction is desired with a size of one double word (8 bytes), then the second, third, and fourth beats of data are not written to DRAM, as the width of the data bus is 64 bits.

The table below lists the data beat sequencing to and from the DDR SDRAM and the data queues for each of the possible transfer sizes with each of the possible starting double-word offsets. All underlined double-word offsets are valid for the transaction.

Table 8-77. Memory controller-data beat ordering

Transfer Size	Starting double-word offset	Double-word sequence ¹ to/from DRAM and queues
1 double word	0	<u>0</u> - 1 - 2 - 3
	1	<u>1</u> - 2 - 3 - 0
	2	<u>2</u> - 3 - 0 - 1
	3	<u>3</u> - 0 - 1 - 2
2 double words	0	<u>0 - 1</u> - 2 - 3
	1	<u>1 - 2</u> - 3 - 0
	2	<u>2 - 3</u> - 0 - 1
3 double words	0	<u>0 - 1 - 2</u> - 3
	1	<u>1 - 2 - 3</u> - 0

1. All underlined double-word offsets are valid for the transaction. All writes are aligned to double-word 0 for DDR3 memories.

8.5.9 Page mode and logical bank retention

The DDR memory controller supports an open/closed page mode with an allowable open page for each logical bank of DRAM used.

In closed page mode for DDR SDRAMs, the DDR memory controller uses the SDRAM auto-precharge feature, which allows the controller to indicate that the page must be automatically closed by the DDR SDRAM after the READ or WRITE access. This is performed using MA[10] of the address during the COMMAND phase of the access to enable auto-precharge. Auto-precharge is non-persistent in that it is either enabled or disabled for each individual READ or WRITE command. It can, however, be enabled or disabled separately for each chip select.

When the DDR memory controller operates in open page mode, it retains the currently active SDRAM page by not issuing a precharge command. The page remains open until one of the following conditions occurs:

- Refresh interval is met.
- The user-programmable DDR_SDRAM_INTERVAL[BSTOPRE] value is exceeded.
- There is a logical bank row collision with another transaction that must be issued.

Page mode can dramatically reduce access latencies for page hits. Depending on the memory system design and timing parameters, using page mode can save two to three clock cycles for subsequent burst accesses that hit in an active page. Also, better performance can be obtained using more banks, especially in systems which use many different channels. Page mode is disabled by clearing DDR_SDRAM_INTERVAL[BSTOPRE] or setting CS_n_CONFIG[AP_nEN].

8.5.10 Error checking and correcting (ECC)

The DDR memory controller supports error checking and correcting (ECC) for the data path between the core master and system memory.

The memory detects all double-bit errors, detects all multi-bit errors within a nibble, and corrects all single-bit errors. Other errors may be detected, but are not guaranteed to be corrected or detected. Multi-bit errors are always reported when error reporting is enabled. When a single-bit error occurs, the single-bit error counter register is incremented, and its value compared to the single-bit error trigger register. An error is reported when these values are equal. The single-bit error registers can be programmed such that minor memory faults are corrected and ignored, but a catastrophic memory failure generates an interrupt.

For writes that are smaller than 64 bits, the DDR memory controller performs a double-word read from system memory of the address for the write (checking for errors), and merges the write data with the data read from memory. Then, a new ECC code is generated for the merged double word. The data and ECC code is then written to

memory. If a multi-bit error is detected on the read, the transaction completes the read-modify-write to keep the DDR memory controller from hanging. However, the corrupt data is masked on the write, so the original contents in SDRAM remain unchanged.

The syndrome encodings for the ECC code are shown in the following tables.

In 32-bit mode, [Table 8-78](#) is split into 2 halves. The first half, consisting of rows 0-31, is used to calculate the ECC bits for the first 32 data bits of any 64-bit granule of data. This always applies to the odd data beats on the DDR data bus. The second half of the table, consisting of rows 32-63, is used to calculate the ECC bits for the second 32 bits of any 64-bit granule of data. This always applies to the even data beats on the DDR data bus.

Table 8-78. DDR SDRAM ECC syndrome encoding

Data bit	Syndrome bit								Data bit	Syndrome bit									
	0	1	2	3	4	5	6			7	0	1	2	3	4	5	6	7	
0	•	•						•		32			•	•					•
1	•		•						•		33			•		•			•
2	•			•					•		34	•		•		•			
3	•				•				•		35		•	•		•			
4	•	•				•					36			•	•		•		
5	•		•			•					37			•		•	•		
6	•			•		•					38	•		•		•	•		•
7	•				•	•					39		•	•		•	•		•
8	•	•					•				40			•	•			•	
9	•		•					•			41			•		•		•	
10	•			•				•			42	•		•		•		•	•
11	•				•			•			43		•	•		•		•	•
12	•	•				•	•	•			44			•	•		•	•	•
13	•		•			•	•	•			45			•		•	•	•	•
14	•			•		•	•	•			46	•		•		•	•	•	
15	•				•	•	•	•			47		•	•		•	•	•	
16		•	•					•			48		•			•	•		
17		•		•					•		49			•			•	•	
18		•			•				•		50				•		•	•	
19	•	•			•						51	•					•	•	
20		•	•			•					52		•				•		•
21		•		•		•					53			•			•		•
22		•			•	•					54				•		•		•
23	•	•			•	•		•			55	•					•		•
24		•	•				•				56		•					•	•
25		•		•			•				57			•				•	•
26		•			•		•				58				•			•	•

Table continues on the next page...

Table 8-78. DDR SDRAM ECC syndrome encoding (continued)

Data bit	Syndrome bit								Data bit	Syndrome bit								
	0	1	2	3	4	5	6	7		0	1	2	3	4	5	6	7	
27	•	•			•		•	•		59	•						•	•
28		•	•			•	•	•		60				•	•		•	
29		•		•		•	•	•		61	•			•	•		•	•
30		•			•	•	•	•		62		•		•	•		•	•
31	•	•			•	•	•			63			•	•	•		•	•

Table 8-79. DDR SDRAM ECC syndrome encoding (check bits)

Check bit	Syndrome bit							
	0	1	2	3	4	5	6	7
0	•							
1		•						
2			•					
3				•				
4					•			
5						•		
6							•	
7								•

8.5.11 Error management

The DDR memory controller detects four different kinds of errors: training, single-bit, multi-bit, and memory select errors.

The following discussion assumes all the relevant error detection, correction, and reporting functions are enabled as described in [Memory error interrupt enable \(DDR_ERR_INT_EN\)](#), [Memory error disable \(DDR_ERR_DISABLE\)](#), and [Memory error detect \(DDR_ERR_DETECT\)](#).

Single-bit errors are counted and reported based on the ERR_SBE value. When a single-bit error is detected, the DDR memory controller does the following:

- Corrects the data
- Increments the single-bit error counter ERR_SBE[SBEC]
- Generates an interrupt if the counter value ERR_SBE[SBEC] equals the programmable threshold ERR_SBE[SBET]
- Completes the transaction normally

If a multi-bit error is detected for a read, the DDR memory controller logs the error and generates the interrupt, and transfer error acknowledge (TEA) is asserted internally on the CSB bus (if enabled, as described in [Memory error disable \(DDR_ERR_DISABLE\)](#)). Another error the DDR memory controller detects is a memory select error, which causes the DDR memory controller to log the error and generate an interrupt (if enabled, as described in [Memory error detect \(DDR_ERR_DETECT\)](#)). This error is detected if the address from the memory request does not fall into any of the enabled, programmed chip select address ranges. For all memory select errors, the DDR memory controller does not issue any transactions onto the pins after the first read has returned data strobes. If the DDR memory controller is not using sample points, then a dummy transaction is issued to DDR SDRAM with the first enabled chip select. In this case, the source port on the pins is forced to 0x1F to show the transaction is not real. The table below shows the errors with their descriptions. The final error the memory controller detects is the automatic calibration error. This error is set if the memory controller detects an error during its training sequence.

Table 8-80. Memory controller errors

Category	Error	Descriptions	Action	Detect Register
Notification	Single-bit ECC threshold	The number of ECC errors has reached the threshold specified in the ERR_SBE.	Regular or critical interrupt (if enabled)	The error control register only logs read versus write, not full type
Access Error	Multi-bit ECC error	A multi-bit ECC error is detected during a read, or read-modify-write memory operation.	Machine check for e500 core initiated reads (if enabled), or regular/critical interrupt (if enabled)	
	Memory select error	Read, or write, address does not fall within the address range of any of the memory banks.		

8.6 Initialization/application information

System software must configure the DDR memory controller, using a memory polling algorithm at system start-up, to correctly map the size of each bank in memory.

Then, the DDR memory controller uses its bank map to assert the appropriate $\overline{\text{MCSn}}$ signal for memory accesses according to the provided bank depths. System software must also configure the DDR memory controller at system start-up to appropriately multiplex the row and column address bits for each bank. Refer to row-address configuration in [Chip select n configuration \(DDR_CS_n_CONFIG\)](#). Address multiplexing occurs according to these configuration bits.

At system reset, initialization software (boot code) must set up the programmable parameters in the memory interface configuration registers. See [DDR memory map/register definition](#) for more detailed descriptions of the configuration registers. These parameters are shown in the following table.

Table 8-81. Memory interface configuration register initialization parameters

Name	Description	Parameter		Section/ page
CS _n _BNDS	Chip select memory bounds	SAn EAn		Chip select n memory bounds (DDR_CS _n _BNDS)
CS _n _CONFIG	Chip select configuration	CS _n _EN AP _n _EN ODT_RD_CFG ODT_WR_CFG	BA_BITS_CS _n ROW_BITS_CS _n COL_BITS_CS _n	Chip select n configuration (DDR_CS _n _CONFIG)
CS _n _CONFIG_2	Chip select configuration 2	PASR_CFG		Chip select n configuration 2 (DDR_CS _n _CONFIG_2)
TIMING_CFG_3	Extended timing parameters for fields in TIMING_CFG_1	EXT_REFREC EXT_ACTTOPRE EXT_CASLAT CNTL_ADJ		DDR SDRAM timing configuration 3 (DDR_TIMING_CFG_3)
TIMING_CFG_0	Timing configuration	RWT WRT RRT WWT	ACT_PD_EXIT PRE_PD_EXIT ODT_PD_EXIT MRS_CYC	DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)
TIMING_CFG_1	Timing configuration	PRETOACT ACTTOPRE ACTTORW CASLAT	REFREC WRREC ACTTOACT WRTORD	DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)
TIMING_CFG_2	Timing configuration	ADD_LAT CPO WR_LAT RD_TO_PRE	WR_DATA_DELAY CKE_PLS FOUR_ACT	DDR SDRAM timing configuration 2 (DDR_TIMING_CFG_2)

Table continues on the next page...

Table 8-81. Memory interface configuration register initialization parameters (continued)

Name	Description	Parameter		Section/ page
DDR_SDRAM_CFG	Control configuration	SREN ECC_EN RD_EN SDRAM_TYPE DYN_PWR 8_BE DBW MEM_HALT	2T_EN 3T_EN BA_INTLV_CTL x32_EN HSE BI MEM_EN PCHB8	DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG)
DDR_SDRAM_CFG_2	Control configuration	SR_IE DLL_RST_DIS DQS_CFG ODT_CFG FRC_SR	NUM_PR OBC_CFG AP_EN D_INIT RCW_EN MD_EN	DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2)
DDR_SDRAM_MODE	Mode configuration	ESDMODE SDMODE		DDR SDRAM mode configuration (DDR_DDR_SDRAM_MODE)
DDR_SDRAM_MODE_2	Mode configuration	ESDMODE2 ESDMODE3		DDR SDRAM mode configuration 2 (DDR_DDR_SDRAM_MODE_2)
DDR_SDRAM_INTERVAL	Interval configuration	REFINT BSTOPRE		DDR SDRAM interval configuration (DDR_DDR_SDRAM_INTERVAL)
DDR_DATA_INIT	Data initialization configuration register	INIT_VALUE		DDR SDRAM data initialization (DDR_DDR_DATA_INIT)
DDR_SDRAM_CLK_CNTL	Clock adjust	CLK_ADJUST		DDR SDRAM clock control (DDR_DDR_SDRAM_CLK_CNTL)

Table continues on the next page...

Table 8-81. Memory interface configuration register initialization parameters (continued)

Name	Description	Parameter	Section/ page
DDR_INIT_ADDR	Initialization address	INIT_ADDR	DDR training initialization address (DDR_DDR_INIT_ADDR)
TIMING_CFG_4	Timing configuration	RWT WRT RRT WWT DLL_LOCK	DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4)
TIMING_CFG_5	Timing configuration	RODT_ON RODT_OFF WODT_ON WODT_OFF	DDR SDRAM timing configuration 5 (DDR_TIMING_CFG_5)
DDR_ZQ_CNTL	ZQ calibration control	ZQ_EN ZQINIT ZQOPER ZQCS	DDR ZQ calibration control (DDR_DDR_ZQ_CNTL)
DDR_WRLVL_CNTL	Write leveling control	WRLVL_EN WRLVL_MRD WRLVL_ODTEN WRLVL_DQSEN WRLVL_SMPL WRLVL_WLR WRLVL_START	DDR write leveling control (DDR_DDR_WRLVL_CNTL)
DDR_WRLVL_CNTL_2	Write leveling control	WRLVL_START_1 WRLVL_START_2 WRLVL_START_3 WRLVL_START_4	DDR write leveling control 2 (DDR_DDR_WRLVL_CNTL_2)
DDR_WRLVL_CNTL_3	Write leveling control	WRLVL_START_5 WRLVL_START_6 WRLVL_START_7 WRLVL_START_8	DDR write leveling control 3 (DDR_DDR_WRLVL_CNTL_3)
DDR_SR_CNTR	Self refresh control	SR_IT	DDR Self Refresh Counter (DDR_DDR_SR_CNTR)

Table continues on the next page...

Table 8-81. Memory interface configuration register initialization parameters (continued)

Name	Description	Parameter		Section/ page
DDR_SDRAM_RCW_1	Register control words configuration	RCW0 RCW1 RCW2 RCW3 RCW4 RCW5 RCW6 RCW7		DDR Register Control Words 1 (DDR_DDR_SDRAM_RCW_1)
DDR_SDRAM_RCW_2	Register control words configuration	RCW8 RCW9 RCW10 RCW11 RCW12 RCW13 RCW14 RCW15		DDR Register Control Words 2 (DDR_DDR_SDRAM_RCW_2)
DDRCDR_1	Driver control	DHC_EN ODT DSO_C_EN DSO_D_EN	DSO_CPZ DSO_CNZ DSO_DPZ DSO_DNZ	DDR Control Driver Register 1 (DDR_DDRCDR_1)
DDRCDR_2	Driver control	DSO_CLK_EN DSO_CLKPZ DSO_CLKNZ ODT		DDR Control Driver Register 2 (DDR_DDRCDR_2)
DDR_INIT_EXT_ADDR	Extended initialization address	UIA INIT_EXT_ADDR		DDR training initialization extended address (DDR_DDR_INIT_EXT_ADDR)

8.6.1 Programming differences between memory types

Depending on the memory type used, certain fields must be programmed differently.

The table below illustrates the differences in certain fields for different memory types. Note that this table does not list all fields that must be programmed.

Table 8-82. Programming differences between memory types

Parameter	Description	Differences
AP _n _EN	Chip Select <i>n</i> Auto precharge enable	DDR2/DDR3 Can be used to place chip select <i>n</i> in auto precharge mode See also: Chip select <i>n</i> configuration (DDR_CS_n_CONFIG)
ODT_RD_CFG	Chip Select ODT read configuration	DDR2/DDR3 Can be enabled to assert ODT if desired. This could be set differently depending on system topology. However, systems with only 1 chip select will typically not use ODT when issuing reads to the memory. See also: Chip select <i>n</i> configuration (DDR_CS_n_CONFIG)
ODT_WR_CFG	Chip Select ODT write configuration	DDR2/DDR3 Can be enabled to assert ODT if desired. This could be set differently depending on system topology. However, ODT is typically set to assert for the chip select that is getting written to (value would be set to 001). See also: Chip select <i>n</i> configuration (DDR_CS_n_CONFIG)
ODT_PD_EXIT	ODT powerdown exit	DDR2 Should be set according to the DDR2 specifications for the memory used. The JEDEC parameter this applies to is <i>t</i> _{AXPD} . DDR3 Should be set to 0001 for DDR3. The powerdown times (<i>t</i> _{XP} and <i>t</i> _{XPDLL}) required for DDR3 are controlled via TIMING_CFG_0[ACT_PD_EXIT] and TIMING_CFG_0[PRE_PD_EXIT]. See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)
PRETOACT	Precharge to activate timing	DDR2/DDR3 Should be set according to the specifications for the memory used (<i>t</i> _{RP}) See also: DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)
ACTTOPRE	Activate to precharge timing	DDR2/DDR3 Should be set, along with the Extended Activate to Precharge Timing, according to the specifications for the memory used (<i>t</i> _{RAS}) See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)
ACTTORW	Activate to read/write timing	DDR2/DDR3 Should be set according to the specifications for the memory used (<i>t</i> _{RCD}) See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)
CASLAT	CAS latency	DDR2/DDR3 Should be set, along with the Extended CAS Latency, to the desired CAS latency See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)

Table continues on the next page...

Table 8-82. Programming differences between memory types (continued)

Parameter	Description	Differences
REFREC	Refresh recovery	<p>DDR2/DDR3</p> <p>Should be set, along with the Extended Refresh Recovery, to the specifications for the memory used (T_{RFC})</p> <p>See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)</p>
WRREC	Write Recovery	<p>DDR2</p> <p>Should be set according to the specifications for the memory used (t_{WR})</p> <p>DDR3</p> <p>Should be set according to the specifications for the memory used (t_{WR}). If DDR_SDRAM_CFG_2[OBC_CFG] is set, then this should be programmed to $t_{WR} + 2$ DRAM cycles.</p> <p>See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)</p>
ACTTOACT	Activate A to Activate B	<p>DDR2/DDR3</p> <p>Should be set according to the specifications for the memory used (t_{RRD})</p> <p>See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)</p>
WRTORD	Write to read timing	<p>DDR2</p> <p>Should be set according to the specifications for the memory used (t_{WTR})</p> <p>DDR3</p> <p>Should be set according to the specifications for the memory used (t_{WTR}). If DDR_SDRAM_CFG_2[OBC_CFG] is set, then this should be programmed to $t_{WTR} + 2$ DRAM cycles.</p> <p>See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)</p>
ADD_LAT	Additive latency	<p>DDR2/DDR3</p> <p>Should be set to the desired additive latency. This must be set to a value less than TIMING_CFG_1[ACTTORW]</p> <p>See also: DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)</p>
WR_LAT	Write latency	<p>DDR2</p> <p>Should be set to CAS latency - 1 cycle. For example, if the CAS latency is 5 cycles, then this field should be set to 100 (4 cycles).</p> <p>DDR3</p> <p>Should be set to the desired write latency. Note that DDR3 SDRAMs do not necessarily require the write latency to equal the CAS latency minus 1 cycle.</p> <p>See also: DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)</p>

Table continues on the next page...

Table 8-82. Programming differences between memory types (continued)

Parameter	Description	Differences
RD_TO_PRE	Read to precharge timing	<p>DDR2</p> <p>Should be set according to the specifications for the memory used (t_{RTP}). Time between read and precharge for non-zero value of additive latency (AL) is a minimum of $AL + t_{RTP}$ cycles.</p> <p>DDR3</p> <p>Should be set according to the specifications for the memory used (t_{RTP}). Time between read and precharge for non-zero value of additive latency (AL) is a minimum of $AL + t_{RTP}$ cycles. If <code>DDR_SDRAM_CFG_2[OBC_CFG]</code> is set, then this should be programmed to $t_{RTP} + 2$ DRAM cycles.</p> <p>See also: DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)</p>
CKE_PLS	Minimum CKE pulse width	<p>DDR2/DDR3</p> <p>Should be set according to the specifications for the memory used (t_{CKE})</p> <p>See also: DDR SDRAM timing configuration 1 (DDR_TIMING_CFG_1)</p>
FOUR_ACT	Four activate window	<p>DDR2</p> <p>Should be set according to the specifications for the memory used (t_{FAW}). Only applies to eight logical banks.</p> <p>DDR3</p> <p>Should be set according to the specifications for the memory used (t_{FAW}).</p> <p>See also: DDR SDRAM timing configuration 0 (DDR_TIMING_CFG_0)</p>
RD_EN	Registered DIMM enable	<p>DDR2/DDR3</p> <p>If registered DIMMs are used, then this field should be set to 1</p> <p>See also: DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG)</p>
8_BE	8-beat burst enable	<p>DDR2</p> <p>Should be set to 0</p> <p>DDR3</p> <p>If a 64-bit bus is used, this should be set to 0. For 32-bit, this may be set to 1 or 0. If this is set to 0, then other requirements in <code>TIMING_CFG_4</code> are needed to ensure t_{CCD} is met.</p> <p>See also: DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG)</p>
2T_EN	2T timing enable	<p>DDR2/DDR3</p> <p>In heavily loaded systems, this can be set to 1 to gain extra timing margin on the interface at the cost of address/command bandwidth.</p> <p>See also: DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG)</p>

Table continues on the next page...

Table 8-82. Programming differences between memory types (continued)

Parameter	Description	Differences
DLL_RST_DIS	DLL reset disable	<p>DDR2</p> <p>Should typically be set to 0, unless it is desired to bypass the DLL reset when exiting self refresh.</p> <p>DDR3</p> <p>Should be set to 1</p> <p>See also: DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2)</p>
DQS_CFG	DQS configuration	<p>DDR2</p> <p>Should be set to 01</p> <p>DDR3</p> <p>Should be set to 01</p> <p>See also: DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2)</p>
ODT_CFG	ODT configuration	<p>DDR2/DDR3</p> <p>Can be set for termination at the IOs according to system topology. Typically, if ODT is enabled, then the internal IOs should be set up for termination only during reads to DRAM.</p> <p>See also: DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2)</p>
OBC_CFG	On-the-fly burst chop configuration	<p>DDR2</p> <p>Should be set to 0</p> <p>DDR3</p> <p>Can be set to 1 if on-the-fly burst chop is used. This is expected to give the best performance in DDR3 mode. This feature can only be used if a 64-bit data bus is used.</p> <p>See also: DDR SDRAM control configuration 2 (DDR_DDR_SDRAM_CFG_2)</p>
RWT	Read-to-write turnaround for same chip select (in TIMING_CFG_4)	<p>DDR2</p> <p>Should typically be set to 0000</p> <p>DDR3</p> <p>This can be used to force a longer read-to-write turnaround time when accessing the same chip select. This is useful for burst chop mode, as there are some timing requirements to the same chip select that still must be met.</p> <p>See also: DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4)</p>

Table continues on the next page...

Table 8-82. Programming differences between memory types (continued)

Parameter	Description	Differences
WRT	Write-to-read turnaround for same chip select (in TIMING_CFG_4)	<p>DDR2</p> <p>Should typically be set to 0000</p> <p>DDR3</p> <p>This could be used to force a certain turnaround time between a write and read to the same chip select. This is useful for burst chop mode. However, it is expected that TIMING_CFG_1[WRTORD] is programmed appropriately such that TIMING_CFG_4[WRT] can be set to 0000.</p> <p>See also: DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4)</p>
RRT	Read-to-read turnaround for same chip select (in TIMING_CFG_4)	<p>DDR2</p> <p>Should typically be set to 0000</p> <p>DDR3</p> <p>Should typically be set to 0010 in burst chop mode (on-the-fly or fixed) or set to 0000 in fixed 8-beat burst mode.</p> <p>See also: DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4)</p>
WWT	Write-to-write turnaround for same chip select (in TIMING_CFG_4)	<p>DDR2</p> <p>Should typically be set to 0000</p> <p>DDR3</p> <p>Should typically be set to 0010 in burst chop mode (on-the-fly or fixed) or set to 0000 in fixed 8-beat burst mode.</p> <p>See also: DDR SDRAM timing configuration 4 (DDR_TIMING_CFG_4)</p>
ZQ_EN	ZQ Calibration enable	<p>DDR2</p> <p>Should be set to 0</p> <p>DDR3</p> <p>Should be set to 1. The other fields in DDR_ZQ_CNTL should also be programmed appropriately based on the DRAM specifications.</p> <p>See also: DDR ZQ calibration control (DDR_DDR_ZQ_CNTL)</p>
WRLVL_EN	Write leveling enable	<p>DDR2</p> <p>Should be set to 0</p> <p>DDR3</p> <p>Can be set to 1 if write leveling is desired. Otherwise the value used in TIMING_CFG_2[WR_DATA_DELAY] is used to shift all bytes during writes to DRAM. If write leveling is used, all other fields in DDR_WRLVL_CNTL should be programmed appropriately based on the DRAM specifications.</p> <p>See also: DDR write leveling control (DDR_DDR_WRLVL_CNTL)</p>
BSTOPR	Burst To precharge interval	<p>DDR2/DDR3</p> <p>Can be set to any value, depending on the application. Auto precharge can be enabled by setting this field to all 0s.</p> <p>See also: DDR SDRAM interval configuration (DDR_DDR_SDRAM_INTERVAL)</p>

8.6.2 DDR SDRAM initialization sequence

After configuration of all parameters is complete, system software must set `DDR_SDRAM_CFG[MEM_EN]` to enable the memory interface.

Note that 200 μ s (500 μ s for DDR3) must elapse after DRAM clocks are stable (`DDR_SDRAM_CLK_CNTL[CLK_ADJUST]` is set and any chip select is enabled) before `MEM_EN` can be set, so a delay loop in the initialization code may be necessary if software is enabling the memory controller. If `DDR_SDRAM_CFG[BI]` is not set, the DDR memory controller conducts an automatic initialization sequence to the memory, which follows the memory specifications. If the bypass initialization mode is used, then software can initialize the memory through the `DDR_SDRAM_MD_CNTL` register.

Chapter 9

Programmable Interrupt Controller (PIC)

This chapter describes the programmable interrupt controller (PIC) interrupt protocol, various types of interrupt sources controlled by the PIC, and the PIC registers with some programming guidelines.

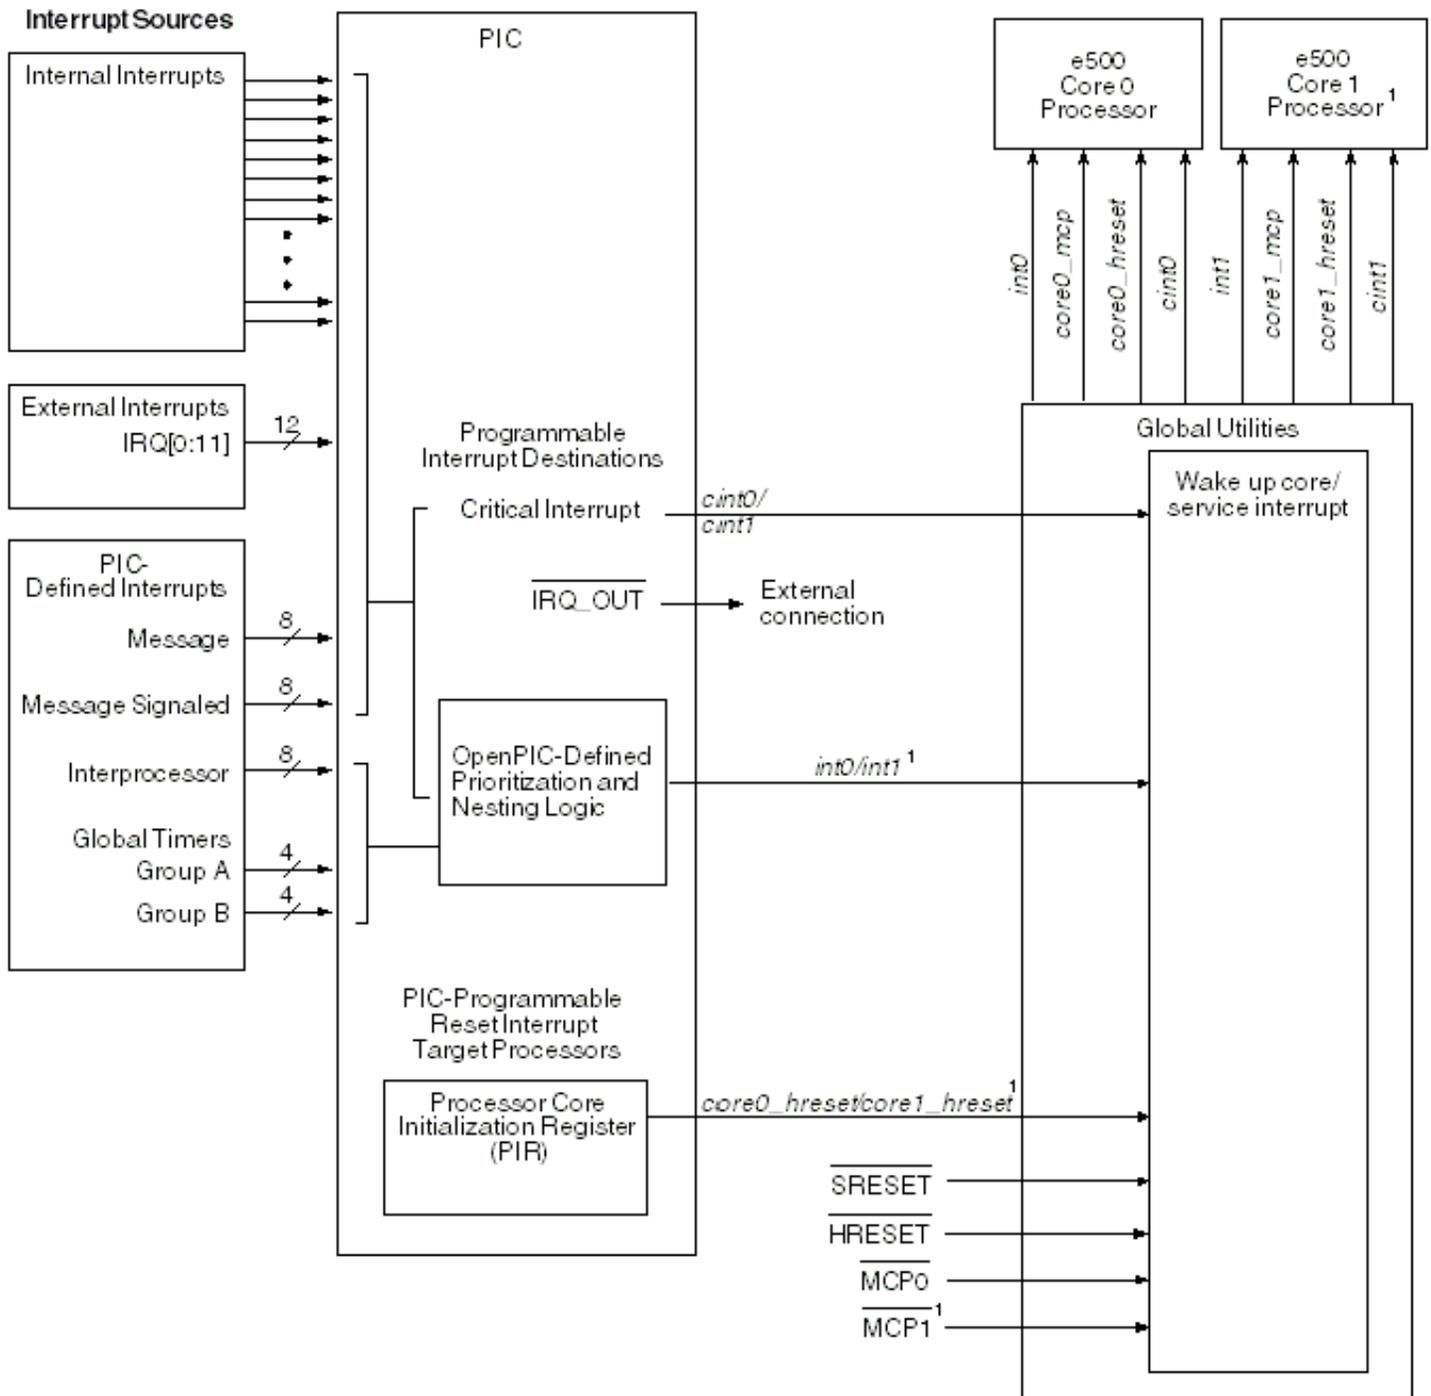
9.1 Introduction

This chapter describes the programmable interrupt controller (PIC) interrupt protocol, various types of interrupt sources controlled by the PIC, and the PIC registers with some programming guidelines.

The PIC conforms to the OpenPIC architecture. The interrupt controller provides multiprocessor interrupt management, and is responsible for receiving hardware-generated interrupts from different sources (both internal and external), prioritizing them, delivering them to the appropriate destination for servicing.

9.1.1 Overview

The figure below is a block diagram showing the relationship of the various functional blocks and how the signals external to the PIC are connected to other blocks on the device, including the cores.



¹ Not supported in single-processor implementations.

Figure 9-1. Interrupt sources block diagram features

The PIC has the following features:

- Support for the following interrupt sources:

- External-Off-chip signals, IRQ[0:11]
- Internal-These are on-chip sources from peripheral logic within the integrated device signaling error conditions that need to be addressed by software.
- Interrupts generated from within the PIC itself, which are as follows:
 - Global timers A and B internal to the PIC
 - Two groups of four global 32-bit timers clocked with the CCB clock or the RTC input. Timers within each group can be concatenated to time longer durations.
 - Interprocessor interrupts (IPI)-Intended for communication between different processor cores on the same device. (Can be used for self-interrupt in single-core devices.)
 - Message registers-From within the PIC. Triggered on register write, cleared on read. Used for interprocessor communication.
 - Shared message signaled registers-From within the PIC. Triggered on register write, cleared on read. Used for cross-program communication.
 - Eight 32-bit message interrupt channels.
- Three types of programmable interrupt outputs:
 - External interrupt (*int0* and *int1*). Any of the PIC interrupt sources can be programmed to direct interrupt requests to *int0* and *int1*. Handling of such interrupt requests follows the OpenPIC specification, which guarantees that the highest priority interrupt supersedes lower priority interrupts. [Interrupts routed to *int*](#), describes how the PIC logic handles these interrupts.
 - Critical interrupt (*cint0* and *cint1*). Connected to the respective e500 core's critical interrupt input.
 - IRQ_OUT_B .

[Interrupts routed to *cint* or IRQ_OUT_B](#), describes how the PIC logic supports this interrupt.

- Programming model compliant with the OpenPIC architecture.
 - Message, interprocessor and global timer interrupts. (Note that the interprocessor and global timer interrupts can only be routed to *int*.)
 - The following OpenPIC-defined features support only interrupts routed to the *int* signal:
 - Fully-nested interrupt delivery, guaranteeing that the interrupt source with the highest priority is given precedence over lower priority interrupts, including any that are in service.
 - 16 programmable interrupt priority levels
 - Support for identifying and handling spurious interrupts
- Support for two processors

- Interrupts can be routed to processor core 0 or 1
- Multi-cast delivery mode for interprocessor and global timer interrupts allowing these interrupts to be routed to either core 0 or 1, or both cores
- Processor core initialization control
- Programmable resetting of the PIC through the global configuration register
- Support for connection of external interrupt controller device such as an 8259 programmable interrupt controller. In 8259 mode, an interrupt causes assertion of a local (that is, internal to the integrated device) interrupt output signal IRQ_OUT_B.
- Pass-through mode (PIC disabled) in which the PIC directs interrupts off-chip for external servicing. See [Pass-through mode \(GCR\[M\] = 0\)](#).

9.1.2 The PIC in multiple-processor implementations

In a multiprocessor implementation, the PIC is replicated for each core, and where necessary, the duplicated resources are indicated with a 0 or a 1.

For example, *int0* identifies the *int* signal to processor 0 and *int1* identifies the *int* signal to processor 1. Other resources associated with the cores are identified by the number. For example, setting PIR[P1] triggers a reset of core 1, and setting PIR[P0] triggers a reset of core 0.

However, where the distinction is not necessary, none is made and such resources are referred to generically. For example, the *int* signal refers to the behavior of either *int0* or *int1*.

In a multiprocessor system, it is possible for the PIC to direct interrupts to the other processor. This functionality is supported by certain multiprocessor aspects to the programming model. For example, an interrupt source in processor 0 can be programmed to target *int1* by setting the P1 bit in its destination register, xIDR_n.

Note that although they are part of the programming model, such resources are reserved in a single-processor device.

9.1.3 Interrupts to the e500 processor core

The external interrupt signal, *int*, is the main interrupt output from the PIC to the processor core.

The interrupt sources can also specify the critical interrupt output, *cint0* or *cint1*, if the corresponding xIDR_n[CI0] or xIDR_n[CI1] is set.

The PIC also defines the PIR, described in [Processor core initialization register \(PIC_PIR\)](#) which can be used to reset the core. Processor core interrupts generated by the PIC are described in the table below.

Table 9-1. e500 processor interrupts generated outside the core-types and sources

Core interrupt type	Signaled by (input to core)	Sources
PIC-Programmable interrupts		
External interrupt	<i>int</i>	Generated by the PIC, as described in Interrupt sources .
Critical interrupt	<i>cint</i>	Generated by the PIC, as described in Interrupt sources .
Other interrupts generated outside the core		
Machine check	<i>coren_mcp</i>	<ul style="list-style-type: none"> • MCP • SRESET • Assertion of <i>coren_mcp</i> by global utilities block
Unconditional debug event	<i>coren_ude</i>	UDE. Asserting UDE generates an unconditional debug exception type debug interrupt and sets a bit in the debug status register, DBSR[UDE], as described in the <i>e500 Core Family Reference Manual</i> .
Reset	<i>coren_hreset</i>	<ul style="list-style-type: none"> • HRESET assertion (and negation) • <i>coren_hreset_req</i>. Output from core-caused by writing to the core DBCR0[RST], as described in the <i>e500 Reference Manual</i>. This condition is additionally qualified with MSR[DE] and DBCR0[IDM] bits. Note that assertion of this signal causes a hard reset of the core only. • <i>coren_reset</i>. Output from PIC. See Processor core initialization register (PIC_PIR).

9.1.4 Modes of operation

Mixed or pass-through mode of operation is chosen by setting or clearing GCR[M].

This is described in [Global configuration register \(PIC_GCR\)](#).

9.1.4.1 Mixed mode (GCR[M] = 1)

In mixed mode, external and internal interrupts are delivered using the normal priority and delivery mechanisms detailed in [Flow of interrupt control](#).

9.1.4.2 Pass-through mode (GCR[M] = 0)

The PIC provides a mechanism to support alternate external interrupt controllers such as the PC/AT-compatible 8259 interrupt controller architecture.

After a hard reset, the PIC defaults to pass-through mode, in which active-high interrupts from external source IRQ0 are passed directly to core 0 as shown in the figure below; all other external interrupt signals are ignored. Thus, the interrupt signal from an external interrupt controller can be connected to IRQ0 and cause direct interrupts to the processor core 0. The PIC does not perform a vector fetch from an external interrupt controller.

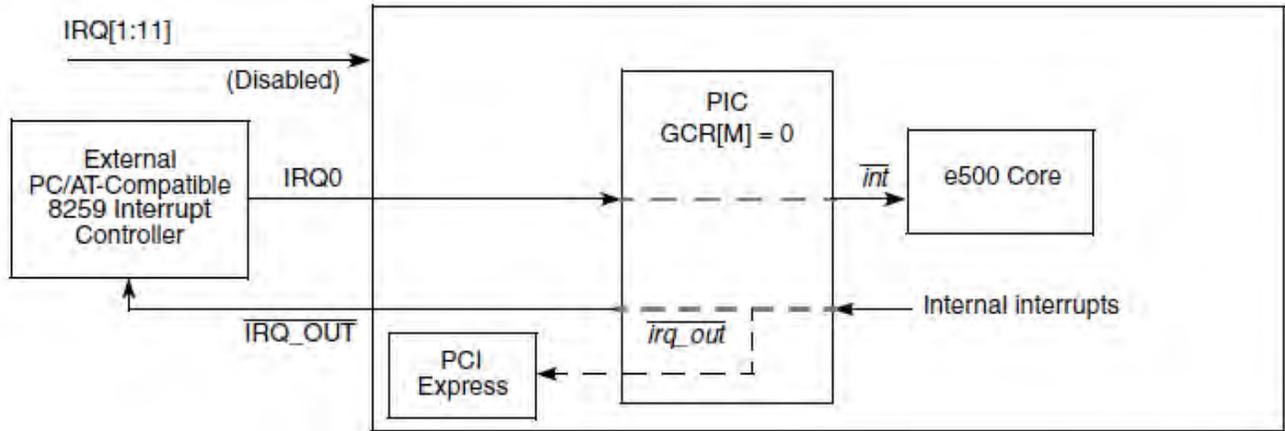


Figure 9-2. Pass-through mode example

When pass-through mode is enabled, the internally-generated interrupts shown in [Table 9-2](#) are not forwarded to core 0. Instead, the PIC passes the raw interrupts from the internal sources to IRQ_OUT_B . Note that when the PCI Express controller is configured as an endpoint (EP) device, the *irq_out* signal from the PIC may be used to automatically generate an outbound PCI Express MSI transaction toward the remote interrupt controller resource on the root complex (RC). See [Hardware MSI generation](#).

Note that in pass-through mode, interrupts generated within the PIC (global timers, interprocessor, and message register interrupts) are disabled. If internal or PIC-generated interrupts must be reported internally to the processor, mixed mode must be used.

It is required that in pass-through mode the internal and external interrupt targets should be *int*, which is by default.

9.1.5 Interrupt sources

The PIC can receive separate interrupts from the following sources:

- External-Off-chip signals, IRQ[0:11]
- Internal-On-chip sources from peripheral logic within the integrated device. See [Table 9-2](#).
- Global timers A and B internal to the PIC

- Interprocessor interrupts (IPI)-Intended for communication between different e500 processor cores on the same device. (Can be used for self-interrupt in single-core implementations.)
- Message registers-From within the PIC. Triggered on register write, cleared on read. Used for interprocessor communication.
- Shared message signaled registers-From within the PIC. Triggered on register write, cleared on read. Used for cross-program communication.

9.1.5.1 Interrupt routing-mixed mode

When an interrupt request is delivered to the PIC, the corresponding interrupt destination register is checked to determine where the request should be routed, as follows:

- If $xIDRn[EP] = 1$ (and all other destination bits are zero), the interrupt is routed off-chip to the external `IRQ_OUT_B` signal. If the PCI Express controller is in EP mode (and the corresponding $xIDRn[EP] = 1$), the controller automatically generates a PCI Express MSI transaction. See [Hardware MSI generation](#).
- If either, but not both, $xIDRn[CI0]$ or $xIDRn[CI1]$ is set (and all other destination bits are zero), the interrupt is routed to `cint0` or `cint1`.
- If $xIDRn[P0]$ is set (and all other destination bits are zero) the interrupt is routed to `int0`. Setting $xIDRn[P1]$ likewise routes the interrupt to `int1`. In this case, the interrupt is latched by the interrupt pending register (IPR) and the interrupt flow is as described in [Flow of interrupt control](#). Note that multicasting interrupts (global timer and interprocessor interrupts) can set both P0 and P1; other interrupt sources cannot.

9.1.5.2 Interrupt destinations

Following its reset (by default), the PIC directs all timer, shared message signaled, and interrupts from external and internal sources to `int` output (connected to the `int` signal of the e500 processor core).

Interprocessor and global timers interrupts can be programmed to be routed to either core's `int` signal or to both cores (multi-casting).

All other interrupts have more destination options, but only one destination can be chosen for a single non-multi-casting interrupt. Instead of being routed to `int`, these interrupts can be routed to the core through `cint` or can be directed to `IRQ_OUT_B`. These options are selected by writing to the EP or CI fields in the appropriate destination register.

Note that EP and CI are only supported for external and internal interrupts that does not support multi-casting.

9.1.5.3 Internal interrupt sources

The table below shows the assignments of the internal interrupt sources and how they are mapped to the registers that control them.

Only the internal interrupts used are listed; that is, the numbers are not consecutive.

Table 9-2. Internal interrupt assignments

Internal Interrupt Number	Interrupt Source
0	L2 cache
1	ECM
2	DDR DRAM controller
3	eLBC controller
4	DMA 1 channel 1
5	DMA 1 channel 2
6	DMA 1 channel 3
7	DMA 1 channel 4
8	PCI Express Port 3
9	PCI Express Port 2
10	PCI Express Port 1
11	Reserved
12	USB
13	eTSEC 1 transmit
14	eTSEC 1 receive
15	eTSEC 3 transmit
16	eTSEC 3 receive
17	eTSEC 3 error
18	eTSEC 1 error
19	eTSEC 2 transmit
20	eTSEC 2 receive
21-23	Reserved
24	eTSEC 2 error
25	Reserved
26	DUART
27	I ² C controllers
28	Performance monitor
29	Security interrupt 1
30	Reserved
31	GPIO
32	SRIO1/SRIO2 error/write-port unit
33	SRIO outbound doorbell 1

Table continues on the next page...

Table 9-2. Internal interrupt assignments (continued)

Internal Interrupt Number	Interrupt Source
34	SRIO inbound doorbell 1
35	Reserved
36	Reserved
37	SRIO outbound message unit1
38	SRIO inbound message unit 1
39	SRIO outbound message unit 2
40	SRIO inbound message unit 2
41	Reserved
42	Security interrupt 2
43	SPI
44-51	Reserved
52	IEEE 1588 TSEC1
53	IEEE 1588 TSEC 2
54	IEEE 1588 TSEC 3
55	Reserved
56	eSDHC
57-59	Reserved
60	DMA 2 channel 1
61	DMA 2 channel 2
62	DMA 2 channel 3
63	DMA 2 channel 4

9.2 PIC external signal descriptions

The following sections describe the PIC signals.

9.2.1 Signal overview

The PIC external interface signals are described in [Table 9-3](#).

There are distinct external interrupt request input signals and an interrupt request output signal `IRQ_OUT_B`.

As [Table 9-3](#) shows, the IRQ inputs are also used for delivering `INTx` signals for the PCI Express root complexes.

9.2.2 Detailed signal descriptions

Table 9-3. Interrupt signals-detailed signal descriptions

Signal	I/O	Description
IRQ[0:11]	I	Interrupt request 0-11. The polarity and sense of each of these signals is programmable. All of these inputs can be driven asynchronously. NOTE: Some interrupt request signals IRQ _n may share PIC external interrupt registers with PCI Express INTx signaling. See PCI Express INTx/IRQ_n sharing .
		<p>State Meaning</p> <p>Asserted-When an external interrupt signal is asserted (according to the programmed polarity), the PIC checks its priority and the interrupt is conditionally passed to the processor designated in the corresponding destination register. In pass-through mode, only interrupts detected on IRQ0 are passed directly to core 0.</p> <p>Negated-There is no incoming interrupt from that source.</p>
		<p>Timing</p> <p>Assertion-All of these inputs can be asserted asynchronously.</p> <p>Negation-Interrupts programmed as level-sensitive must remain asserted until serviced. Timing requirements for edge-sensitive interrupts can be found in the Hardware Specifications.</p>
IRQ_OUT_B	O	Interrupt request out. When the PIC is programmed in pass-through mode, this output reflects the raw interrupts generated by on-chip sources. See Modes of operation .
		<p>State Meaning</p> <p>Asserted-At least one interrupt is currently being signaled to the external system.</p> <p>Negated-Indicates no interrupt source currently routed to IRQ_OUT_B .</p>
		<p>Timing</p> <p>Because external interrupts are asynchronous with respect to the system clock, both assertion and negation of IRQ_OUT_B occurs asynchronously with respect to the interrupt source. All timing given here is approximate.</p> <p>Assertion-Internal interrupt source: 2 CCB clock cycles after interrupt occurs.</p> <p>External interrupt source: 4 cycles after interrupt occurs.</p> <p>Message interrupts: 2 cycles after write to message register.</p> <p>Negation-Follows interrupt source negation with the following delay:</p> <p>Internal interrupt: 2 CCB clock cycles</p> <p>External interrupt: 4 cycles.</p> <p>Message interrupts: 2 cycles after message register cleared.</p>
MCP0_B and MCP1_B	I	Machine check processor <i>n</i> . Assertion causes a machine check interrupt to the corresponding core. Note that if the core is not configured to process machine check interrupts (MSR[ME] = 0), assertion of MCP _n _B causes a core checkstop condition. Note that internal sources for the internal <i>coren_mcp</i> can also cause a machine check interrupt to the processor core as described in Machine check summary register (GUTS_MCPSUMR) and Table 9-1 .
		<p>State Meaning</p> <p>Asserted-Integrated logic should direct the corresponding core to take a machine check interrupt or enter the checkstop state as directed by the MSR.</p> <p>Negated-Machine check handling is not being requested by the external system.</p>
		<p>Timing</p> <p>Assertion-May occur at any time, asynchronous to any clock.</p> <p>Negation-Because MCP_B <i>n</i> is edge-triggered, it can be negated one clock after its assertion.</p>

9.3 PIC memory map/register definition

The PIC programmable register map occupies 256 Kbytes of memory-mapped space. Reading undefined portions of the memory map returns all zeros; writing has no effect. All PIC registers are 32 bits wide and, although located on 128-bit address boundaries, should be accessed only as 32-bit quantities.

The PIC address offset map is divided into three areas:

- 0xnn4_0000-0xnn4_FFF0-Global registers
- 0xnn5_0000-0xnn5_FFF0-Interrupt source configuration registers
- 0xnn6_0000- 0xnn7_FFF0 -Per-CPU registers

PIC memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
4_0000	Block revision register 1 (PIC_BRR1)	32	R	0040_0301h	9.3.1/407
4_0010	Block revision register 2 (PIC_BRR2)	32	R	0000_0001h	9.3.2/407
4_0040	Interprocessor n dispatch register (PIC_IPIDR0)	32	W	0000_0000h	9.3.3/408
4_0050	Interprocessor n dispatch register (PIC_IPIDR1)	32	W	0000_0000h	9.3.3/408
4_0060	Interprocessor n dispatch register (PIC_IPIDR2)	32	W	0000_0000h	9.3.3/408
4_0070	Interprocessor n dispatch register (PIC_IPIDR3)	32	W	0000_0000h	9.3.3/408
4_0080	Current task priority register (PIC_CTPR)	32	R/W	0000_000Fh	9.3.4/409
4_0090	Who am I register (PIC_WHOAMI)	32	R	See section	9.3.5/410
4_00A0	Interrupt acknowledge register (PIC_IACK)	32	R	0000_0000h	9.3.6/410
4_00B0	End of interrupt register (PIC_EOI)	32	W	0000_0000h	9.3.7/411
4_1000	Feature reporting register (PIC_FRR)	32	R	See section	9.3.8/412
4_1020	Global configuration register (PIC_GCR)	32	R/W	0000_0000h	9.3.9/413
4_1080	Vendor identification register (PIC_VIR)	32	R	0000_0000h	9.3.10/414
4_1090	Processor core initialization register (PIC_PIR)	32	R/W	0000_0000h	9.3.11/414
4_10A0	Interprocessor interrupt n vector/priority register (PIC_IPIVPR0)	32	R/W	8000_0000h	9.3.12/415
4_10B0	Interprocessor interrupt n vector/priority register (PIC_IPIVPR1)	32	R/W	8000_0000h	9.3.12/415
4_10C0	Interprocessor interrupt n vector/priority register (PIC_IPIVPR2)	32	R/W	8000_0000h	9.3.12/415
4_10D0	Interprocessor interrupt n vector/priority register (PIC_IPIVPR3)	32	R/W	8000_0000h	9.3.12/415
4_10E0	Spurious vector register (PIC_SVR)	32	R/W	0000_FFFFh	9.3.13/416
4_10F0	Timer frequency reporting register group X (PIC_TFRR)	32	R/W	0000_0000h	9.3.14/416
4_1100	Global timer n current count register group A (PIC_GTCCRA0)	32	R	0000_0000h	9.3.15/417

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
4_1110	Global timer n base count register group A (PIC_GTBCRA0)	32	R/W	8000_0000h	9.3.16/418
4_1120	Global timer n vector/priority register group A (PIC_GTVPRA0)	32	R/W	8000_0000h	9.3.17/419
4_1130	Global timer n destination register group A (PIC_GTDRA0)	32	R/W	0000_0001h	9.3.18/420
4_1140	Global timer n current count register group A (PIC_GTCCRA1)	32	R	0000_0000h	9.3.15/417
4_1150	Global timer n base count register group A (PIC_GTBCRA1)	32	R/W	8000_0000h	9.3.16/418
4_1160	Global timer n vector/priority register group A (PIC_GTVPRA1)	32	R/W	8000_0000h	9.3.17/419
4_1170	Global timer n destination register group A (PIC_GTDRA1)	32	R/W	0000_0001h	9.3.18/420
4_1180	Global timer n current count register group A (PIC_GTCCRA2)	32	R	0000_0000h	9.3.15/417
4_1190	Global timer n base count register group A (PIC_GTBCRA2)	32	R/W	8000_0000h	9.3.16/418
4_11A0	Global timer n vector/priority register group A (PIC_GTVPRA2)	32	R/W	8000_0000h	9.3.17/419
4_11B0	Global timer n destination register group A (PIC_GTDRA2)	32	R/W	0000_0001h	9.3.18/420
4_11C0	Global timer n current count register group A (PIC_GTCCRA3)	32	R	0000_0000h	9.3.15/417
4_11D0	Global timer n base count register group A (PIC_GTBCRA3)	32	R/W	8000_0000h	9.3.16/418
4_11E0	Global timer n vector/priority register group A (PIC_GTVPRA3)	32	R/W	8000_0000h	9.3.17/419
4_11F0	Global timer n destination register group A (PIC_GTDRA3)	32	R/W	0000_0001h	9.3.18/420
4_1300	Timer control register group n (PIC_TCRA)	32	R/W	0000_0000h	9.3.19/420
4_1308	External interrupt summary register (PIC_ERQSR)	32	R	0000_0000h	9.3.20/423
4_1310	IRQ_OUT_B summary register 0 (PIC_IRQSR0)	32	R	0000_0000h	9.3.21/423
4_1320	IRQ_OUT_B summary register 1 (PIC_IRQSR1)	32	R	0000_0000h	9.3.22/424
4_1324	IRQ_OUT_B summary register 2 (PIC_IRQSR2)	32	R	0000_0000h	9.3.23/424
4_1330	Critical interrupt summary register 0 (PIC_CISR0)	32	R	0000_0000h	9.3.24/425
4_1340	Critical interrupt summary register 1 (PIC_CISR1)	32	R	0000_0000h	9.3.25/425
4_1344	Critical interrupt summary register 2 (PIC_CISR2)	32	R	0000_0000h	9.3.26/426
4_1350	Performance monitor n mask register 0 (PIC_PM0MR0)	32	R/W	FFFF_FFFFh	9.3.27/426
4_1360	Performance monitor n mask register 1 (PIC_PM0MR1)	32	R/W	FFFF_FFFFh	9.3.28/427
4_1364	Performance monitor n mask register 2 (PIC_PM0MR2)	32	R/W	FFFF_FFFFh	9.3.29/427
4_1370	Performance monitor n mask register 0 (PIC_PM1MR0)	32	R/W	FFFF_FFFFh	9.3.27/426
4_1380	Performance monitor n mask register 1 (PIC_PM1MR1)	32	R/W	FFFF_FFFFh	9.3.28/427
4_1384	Performance monitor n mask register 2 (PIC_PM1MR2)	32	R/W	FFFF_FFFFh	9.3.29/427
4_1390	Performance monitor n mask register 0 (PIC_PM2MR0)	32	R/W	FFFF_FFFFh	9.3.27/426
4_13A0	Performance monitor n mask register 1 (PIC_PM2MR1)	32	R/W	FFFF_FFFFh	9.3.28/427
4_13A4	Performance monitor n mask register 2 (PIC_PM2MR2)	32	R/W	FFFF_FFFFh	9.3.29/427
4_13B0	Performance monitor n mask register 0 (PIC_PM3MR0)	32	R/W	FFFF_FFFFh	9.3.27/426

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
4_13C0	Performance monitor n mask register 1 (PIC_PM3MR1)	32	R/W	FFFF_FFFFh	9.3.28/427
4_13C4	Performance monitor n mask register 2 (PIC_PM3MR2)	32	R/W	FFFF_FFFFh	9.3.29/427
4_1400	Message register n (PIC_MSGR0)	32	R/W	0000_0000h	9.3.30/428
4_1410	Message register n (PIC_MSGR1)	32	R/W	0000_0000h	9.3.30/428
4_1420	Message register n (PIC_MSGR2)	32	R/W	0000_0000h	9.3.30/428
4_1430	Message register n (PIC_MSGR3)	32	R/W	0000_0000h	9.3.30/428
4_1500	Message enable register (PIC_MER)	32	R/W	0000_0000h	9.3.31/428
4_1510	Message status register (PIC_MSR)	32	R/W	0000_0000h	9.3.32/429
4_1600	Shared message signaled interrupt register n (PIC_MSIR0)	32	R	0000_0000h	9.3.33/429
4_1610	Shared message signaled interrupt register n (PIC_MSIR1)	32	R	0000_0000h	9.3.33/429
4_1620	Shared message signaled interrupt register n (PIC_MSIR2)	32	R	0000_0000h	9.3.33/429
4_1630	Shared message signaled interrupt register n (PIC_MSIR3)	32	R	0000_0000h	9.3.33/429
4_1640	Shared message signaled interrupt register n (PIC_MSIR4)	32	R	0000_0000h	9.3.33/429
4_1650	Shared message signaled interrupt register n (PIC_MSIR5)	32	R	0000_0000h	9.3.33/429
4_1660	Shared message signaled interrupt register n (PIC_MSIR6)	32	R	0000_0000h	9.3.33/429
4_1670	Shared message signaled interrupt register n (PIC_MSIR7)	32	R	0000_0000h	9.3.33/429
4_1720	Shared message signaled interrupt status register (PIC_MSISR)	32	R	0000_0000h	9.3.34/430
4_1740	Shared message signaled interrupt index register (PIC_MSIIIR)	32	W	0000_0000h	9.3.35/431
4_20F0	Timer frequency reporting register group X (PIC_TFRRB)	32	R/W	0000_0000h	9.3.14/416
4_2100	Global timer n current count register group B (PIC_GTCCRB0)	32	R	0000_0000h	9.3.36/432
4_2110	Global timer n base count register group B (PIC_GTBCRB0)	32	R/W	8000_0000h	9.3.37/433
4_2120	Global timer n vector/priority register group B (PIC_GTVPRB0)	32	R/W	8000_0000h	9.3.38/434
4_2130	Global timer n destination register group B (PIC_GTDRB0)	32	R/W	0000_0001h	9.3.39/435
4_2140	Global timer n current count register group B (PIC_GTCCRB1)	32	R	0000_0000h	9.3.36/432
4_2150	Global timer n base count register group B (PIC_GTBCRB1)	32	R/W	8000_0000h	9.3.37/433
4_2160	Global timer n vector/priority register group B (PIC_GTVPRB1)	32	R/W	8000_0000h	9.3.38/434
4_2170	Global timer n destination register group B (PIC_GTDRB1)	32	R/W	0000_0001h	9.3.39/435
4_2180	Global timer n current count register group B (PIC_GTCCRB2)	32	R	0000_0000h	9.3.36/432
4_2190	Global timer n base count register group B (PIC_GTBCRB2)	32	R/W	8000_0000h	9.3.37/433
4_21A0	Global timer n vector/priority register group B (PIC_GTVPRB2)	32	R/W	8000_0000h	9.3.38/434
4_21B0	Global timer n destination register group B (PIC_GTDRB2)	32	R/W	0000_0001h	9.3.39/435

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
4_21C0	Global timer n current count register group B (PIC_GTCCRB3)	32	R	0000_0000h	9.3.36/432
4_21D0	Global timer n base count register group B (PIC_GTBCRB3)	32	R/W	8000_0000h	9.3.37/433
4_21E0	Global timer n vector/priority register group B (PIC_GTVPRB3)	32	R/W	8000_0000h	9.3.38/434
4_21F0	Global timer n destination register group B (PIC_GTDRB3)	32	R/W	0000_0001h	9.3.39/435
4_2300	Timer control register group n (PIC_TCRB)	32	R/W	0000_0000h	9.3.19/420
4_2400	Message register n (PIC_MSGRa4)	32	R/W	0000_0000h	9.3.40/435
4_2410	Message register n (PIC_MSGRa5)	32	R/W	0000_0000h	9.3.40/435
4_2420	Message register n (PIC_MSGRa6)	32	R/W	0000_0000h	9.3.40/435
4_2430	Message register n (PIC_MSGRa7)	32	R/W	0000_0000h	9.3.40/435
4_2500	Message enable register (PIC_MERa)	32	R/W	0000_0000h	9.3.41/436
4_2510	Message status register (PIC_MSRa)	32	R/W	0000_0000h	9.3.42/437
5_0000	External interrupt n (IRQn) vector/priority register (PIC_EIVPR0)	32	R/W	8000_0000h	9.3.43/437
5_0010	External interrupt n (IRQn) destination register (PIC_EIDR0)	32	R/W	0000_0001h	9.3.44/439
5_0020	External interrupt n (IRQn) vector/priority register (PIC_EIVPR1)	32	R/W	8000_0000h	9.3.43/437
5_0030	External interrupt n (IRQn) destination register (PIC_EIDR1)	32	R/W	0000_0001h	9.3.44/439
5_0040	External interrupt n (IRQn) vector/priority register (PIC_EIVPR2)	32	R/W	8000_0000h	9.3.43/437
5_0050	External interrupt n (IRQn) destination register (PIC_EIDR2)	32	R/W	0000_0001h	9.3.44/439
5_0060	External interrupt n (IRQn) vector/priority register (PIC_EIVPR3)	32	R/W	8000_0000h	9.3.43/437
5_0070	External interrupt n (IRQn) destination register (PIC_EIDR3)	32	R/W	0000_0001h	9.3.44/439
5_0080	External interrupt n (IRQn) vector/priority register (PIC_EIVPR4)	32	R/W	8000_0000h	9.3.43/437
5_0090	External interrupt n (IRQn) destination register (PIC_EIDR4)	32	R/W	0000_0001h	9.3.44/439
5_00A0	External interrupt n (IRQn) vector/priority register (PIC_EIVPR5)	32	R/W	8000_0000h	9.3.43/437
5_00B0	External interrupt n (IRQn) destination register (PIC_EIDR5)	32	R/W	0000_0001h	9.3.44/439
5_00C0	External interrupt n (IRQn) vector/priority register (PIC_EIVPR6)	32	R/W	8000_0000h	9.3.43/437
5_00D0	External interrupt n (IRQn) destination register (PIC_EIDR6)	32	R/W	0000_0001h	9.3.44/439
5_00E0	External interrupt n (IRQn) vector/priority register (PIC_EIVPR7)	32	R/W	8000_0000h	9.3.43/437
5_00F0	External interrupt n (IRQn) destination register (PIC_EIDR7)	32	R/W	0000_0001h	9.3.44/439
5_0100	External interrupt n (IRQn) vector/priority register (PIC_EIVPR8)	32	R/W	8000_0000h	9.3.43/437
5_0110	External interrupt n (IRQn) destination register (PIC_EIDR8)	32	R/W	0000_0001h	9.3.44/439

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5_0120	External interrupt n (IRQn) vector/priority register (PIC_EIVPR9)	32	R/W	8000_0000h	9.3.43/437
5_0130	External interrupt n (IRQn) destination register (PIC_EIDR9)	32	R/W	0000_0001h	9.3.44/439
5_0140	External interrupt n (IRQn) vector/priority register (PIC_EIVPR10)	32	R/W	8000_0000h	9.3.43/437
5_0150	External interrupt n (IRQn) destination register (PIC_EIDR10)	32	R/W	0000_0001h	9.3.44/439
5_0160	External interrupt n (IRQn) vector/priority register (PIC_EIVPR11)	32	R/W	8000_0000h	9.3.43/437
5_0170	External interrupt n (IRQn) destination register (PIC_EIDR11)	32	R/W	0000_0001h	9.3.44/439
5_0200	Internal interrupt n vector/priority register (PIC_IIVPR0)	32	R/W	8080_0000h	9.3.45/440
5_0210	Internal interrupt n destination register (PIC_IIDR0)	32	R/W	0000_0001h	9.3.46/441
5_0220	Internal interrupt n vector/priority register (PIC_IIVPR1)	32	R/W	8080_0000h	9.3.45/440
5_0230	Internal interrupt n destination register (PIC_IIDR1)	32	R/W	0000_0001h	9.3.46/441
5_0240	Internal interrupt n vector/priority register (PIC_IIVPR2)	32	R/W	8080_0000h	9.3.45/440
5_0250	Internal interrupt n destination register (PIC_IIDR2)	32	R/W	0000_0001h	9.3.46/441
5_0260	Internal interrupt n vector/priority register (PIC_IIVPR3)	32	R/W	8080_0000h	9.3.45/440
5_0270	Internal interrupt n destination register (PIC_IIDR3)	32	R/W	0000_0001h	9.3.46/441
5_0280	Internal interrupt n vector/priority register (PIC_IIVPR4)	32	R/W	8080_0000h	9.3.45/440
5_0290	Internal interrupt n destination register (PIC_IIDR4)	32	R/W	0000_0001h	9.3.46/441
5_02A0	Internal interrupt n vector/priority register (PIC_IIVPR5)	32	R/W	8080_0000h	9.3.45/440
5_02B0	Internal interrupt n destination register (PIC_IIDR5)	32	R/W	0000_0001h	9.3.46/441
5_02C0	Internal interrupt n vector/priority register (PIC_IIVPR6)	32	R/W	8080_0000h	9.3.45/440
5_02D0	Internal interrupt n destination register (PIC_IIDR6)	32	R/W	0000_0001h	9.3.46/441
5_02E0	Internal interrupt n vector/priority register (PIC_IIVPR7)	32	R/W	8080_0000h	9.3.45/440
5_02F0	Internal interrupt n destination register (PIC_IIDR7)	32	R/W	0000_0001h	9.3.46/441
5_0300	Internal interrupt n vector/priority register (PIC_IIVPR8)	32	R/W	8080_0000h	9.3.45/440
5_0310	Internal interrupt n destination register (PIC_IIDR8)	32	R/W	0000_0001h	9.3.46/441
5_0320	Internal interrupt n vector/priority register (PIC_IIVPR9)	32	R/W	8080_0000h	9.3.45/440
5_0330	Internal interrupt n destination register (PIC_IIDR9)	32	R/W	0000_0001h	9.3.46/441
5_0340	Internal interrupt n vector/priority register (PIC_IIVPR10)	32	R/W	8080_0000h	9.3.45/440
5_0350	Internal interrupt n destination register (PIC_IIDR10)	32	R/W	0000_0001h	9.3.46/441
5_0360	Internal interrupt n vector/priority register (PIC_IIVPR11)	32	R/W	8080_0000h	9.3.45/440
5_0370	Internal interrupt n destination register (PIC_IIDR11)	32	R/W	0000_0001h	9.3.46/441
5_0380	Internal interrupt n vector/priority register (PIC_IIVPR12)	32	R/W	8080_0000h	9.3.45/440
5_0390	Internal interrupt n destination register (PIC_IIDR12)	32	R/W	0000_0001h	9.3.46/441
5_03A0	Internal interrupt n vector/priority register (PIC_IIVPR13)	32	R/W	8080_0000h	9.3.45/440
5_03B0	Internal interrupt n destination register (PIC_IIDR13)	32	R/W	0000_0001h	9.3.46/441

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5_03C0	Internal interrupt n vector/priority register (PIC_IIVPR14)	32	R/W	8080_0000h	9.3.45/440
5_03D0	Internal interrupt n destination register (PIC_IIDR14)	32	R/W	0000_0001h	9.3.46/441
5_03E0	Internal interrupt n vector/priority register (PIC_IIVPR15)	32	R/W	8080_0000h	9.3.45/440
5_03F0	Internal interrupt n destination register (PIC_IIDR15)	32	R/W	0000_0001h	9.3.46/441
5_0400	Internal interrupt n vector/priority register (PIC_IIVPR16)	32	R/W	8080_0000h	9.3.45/440
5_0410	Internal interrupt n destination register (PIC_IIDR16)	32	R/W	0000_0001h	9.3.46/441
5_0420	Internal interrupt n vector/priority register (PIC_IIVPR17)	32	R/W	8080_0000h	9.3.45/440
5_0430	Internal interrupt n destination register (PIC_IIDR17)	32	R/W	0000_0001h	9.3.46/441
5_0440	Internal interrupt n vector/priority register (PIC_IIVPR18)	32	R/W	8080_0000h	9.3.45/440
5_0450	Internal interrupt n destination register (PIC_IIDR18)	32	R/W	0000_0001h	9.3.46/441
5_0460	Internal interrupt n vector/priority register (PIC_IIVPR19)	32	R/W	8080_0000h	9.3.45/440
5_0470	Internal interrupt n destination register (PIC_IIDR19)	32	R/W	0000_0001h	9.3.46/441
5_0480	Internal interrupt n vector/priority register (PIC_IIVPR20)	32	R/W	8080_0000h	9.3.45/440
5_0490	Internal interrupt n destination register (PIC_IIDR20)	32	R/W	0000_0001h	9.3.46/441
5_04A0	Internal interrupt n vector/priority register (PIC_IIVPR21)	32	R/W	8080_0000h	9.3.45/440
5_04B0	Internal interrupt n destination register (PIC_IIDR21)	32	R/W	0000_0001h	9.3.46/441
5_04C0	Internal interrupt n vector/priority register (PIC_IIVPR22)	32	R/W	8080_0000h	9.3.45/440
5_04D0	Internal interrupt n destination register (PIC_IIDR22)	32	R/W	0000_0001h	9.3.46/441
5_04E0	Internal interrupt n vector/priority register (PIC_IIVPR23)	32	R/W	8080_0000h	9.3.45/440
5_04F0	Internal interrupt n destination register (PIC_IIDR23)	32	R/W	0000_0001h	9.3.46/441
5_0500	Internal interrupt n vector/priority register (PIC_IIVPR24)	32	R/W	8080_0000h	9.3.45/440
5_0510	Internal interrupt n destination register (PIC_IIDR24)	32	R/W	0000_0001h	9.3.46/441
5_0520	Internal interrupt n vector/priority register (PIC_IIVPR25)	32	R/W	8080_0000h	9.3.45/440
5_0530	Internal interrupt n destination register (PIC_IIDR25)	32	R/W	0000_0001h	9.3.46/441
5_0540	Internal interrupt n vector/priority register (PIC_IIVPR26)	32	R/W	8080_0000h	9.3.45/440
5_0550	Internal interrupt n destination register (PIC_IIDR26)	32	R/W	0000_0001h	9.3.46/441
5_0560	Internal interrupt n vector/priority register (PIC_IIVPR27)	32	R/W	8080_0000h	9.3.45/440
5_0570	Internal interrupt n destination register (PIC_IIDR27)	32	R/W	0000_0001h	9.3.46/441
5_0580	Internal interrupt n vector/priority register (PIC_IIVPR28)	32	R/W	8080_0000h	9.3.45/440
5_0590	Internal interrupt n destination register (PIC_IIDR28)	32	R/W	0000_0001h	9.3.46/441
5_05A0	Internal interrupt n vector/priority register (PIC_IIVPR29)	32	R/W	8080_0000h	9.3.45/440
5_05B0	Internal interrupt n destination register (PIC_IIDR29)	32	R/W	0000_0001h	9.3.46/441
5_05C0	Internal interrupt n vector/priority register (PIC_IIVPR30)	32	R/W	8080_0000h	9.3.45/440
5_05D0	Internal interrupt n destination register (PIC_IIDR30)	32	R/W	0000_0001h	9.3.46/441
5_05E0	Internal interrupt n vector/priority register (PIC_IIVPR31)	32	R/W	8080_0000h	9.3.45/440
5_05F0	Internal interrupt n destination register (PIC_IIDR31)	32	R/W	0000_0001h	9.3.46/441
5_0600	Internal interrupt n vector/priority register (PIC_IIVPR32)	32	R/W	8080_0000h	9.3.45/440
5_0610	Internal interrupt n destination register (PIC_IIDR32)	32	R/W	0000_0001h	9.3.46/441

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5_0620	Internal interrupt n vector/priority register (PIC_IIVPR33)	32	R/W	8080_0000h	9.3.45/440
5_0630	Internal interrupt n destination register (PIC_IIDR33)	32	R/W	0000_0001h	9.3.46/441
5_0640	Internal interrupt n vector/priority register (PIC_IIVPR34)	32	R/W	8080_0000h	9.3.45/440
5_0650	Internal interrupt n destination register (PIC_IIDR34)	32	R/W	0000_0001h	9.3.46/441
5_0660	Internal interrupt n vector/priority register (PIC_IIVPR35)	32	R/W	8080_0000h	9.3.45/440
5_0670	Internal interrupt n destination register (PIC_IIDR35)	32	R/W	0000_0001h	9.3.46/441
5_0680	Internal interrupt n vector/priority register (PIC_IIVPR36)	32	R/W	8080_0000h	9.3.45/440
5_0690	Internal interrupt n destination register (PIC_IIDR36)	32	R/W	0000_0001h	9.3.46/441
5_06A0	Internal interrupt n vector/priority register (PIC_IIVPR37)	32	R/W	8080_0000h	9.3.45/440
5_06B0	Internal interrupt n destination register (PIC_IIDR37)	32	R/W	0000_0001h	9.3.46/441
5_06C0	Internal interrupt n vector/priority register (PIC_IIVPR38)	32	R/W	8080_0000h	9.3.45/440
5_06D0	Internal interrupt n destination register (PIC_IIDR38)	32	R/W	0000_0001h	9.3.46/441
5_06E0	Internal interrupt n vector/priority register (PIC_IIVPR39)	32	R/W	8080_0000h	9.3.45/440
5_06F0	Internal interrupt n destination register (PIC_IIDR39)	32	R/W	0000_0001h	9.3.46/441
5_0700	Internal interrupt n vector/priority register (PIC_IIVPR40)	32	R/W	8080_0000h	9.3.45/440
5_0710	Internal interrupt n destination register (PIC_IIDR40)	32	R/W	0000_0001h	9.3.46/441
5_0720	Internal interrupt n vector/priority register (PIC_IIVPR41)	32	R/W	8080_0000h	9.3.45/440
5_0730	Internal interrupt n destination register (PIC_IIDR41)	32	R/W	0000_0001h	9.3.46/441
5_0740	Internal interrupt n vector/priority register (PIC_IIVPR42)	32	R/W	8080_0000h	9.3.45/440
5_0750	Internal interrupt n destination register (PIC_IIDR42)	32	R/W	0000_0001h	9.3.46/441
5_0760	Internal interrupt n vector/priority register (PIC_IIVPR43)	32	R/W	8080_0000h	9.3.45/440
5_0770	Internal interrupt n destination register (PIC_IIDR43)	32	R/W	0000_0001h	9.3.46/441
5_0780	Internal interrupt n vector/priority register (PIC_IIVPR44)	32	R/W	8080_0000h	9.3.45/440
5_0790	Internal interrupt n destination register (PIC_IIDR44)	32	R/W	0000_0001h	9.3.46/441
5_07A0	Internal interrupt n vector/priority register (PIC_IIVPR45)	32	R/W	8080_0000h	9.3.45/440
5_07B0	Internal interrupt n destination register (PIC_IIDR45)	32	R/W	0000_0001h	9.3.46/441
5_07C0	Internal interrupt n vector/priority register (PIC_IIVPR46)	32	R/W	8080_0000h	9.3.45/440
5_07D0	Internal interrupt n destination register (PIC_IIDR46)	32	R/W	0000_0001h	9.3.46/441
5_07E0	Internal interrupt n vector/priority register (PIC_IIVPR47)	32	R/W	8080_0000h	9.3.45/440
5_07F0	Internal interrupt n destination register (PIC_IIDR47)	32	R/W	0000_0001h	9.3.46/441
5_0800	Internal interrupt n vector/priority register (PIC_IIVPR48)	32	R/W	8080_0000h	9.3.45/440
5_0810	Internal interrupt n destination register (PIC_IIDR48)	32	R/W	0000_0001h	9.3.46/441
5_0820	Internal interrupt n vector/priority register (PIC_IIVPR49)	32	R/W	8080_0000h	9.3.45/440
5_0830	Internal interrupt n destination register (PIC_IIDR49)	32	R/W	0000_0001h	9.3.46/441
5_0840	Internal interrupt n vector/priority register (PIC_IIVPR50)	32	R/W	8080_0000h	9.3.45/440
5_0850	Internal interrupt n destination register (PIC_IIDR50)	32	R/W	0000_0001h	9.3.46/441
5_0860	Internal interrupt n vector/priority register (PIC_IIVPR51)	32	R/W	8080_0000h	9.3.45/440
5_0870	Internal interrupt n destination register (PIC_IIDR51)	32	R/W	0000_0001h	9.3.46/441

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5_0880	Internal interrupt n vector/priority register (PIC_IIVPR52)	32	R/W	8080_0000h	9.3.45/440
5_0890	Internal interrupt n destination register (PIC_IIDR52)	32	R/W	0000_0001h	9.3.46/441
5_08A0	Internal interrupt n vector/priority register (PIC_IIVPR53)	32	R/W	8080_0000h	9.3.45/440
5_08B0	Internal interrupt n destination register (PIC_IIDR53)	32	R/W	0000_0001h	9.3.46/441
5_08C0	Internal interrupt n vector/priority register (PIC_IIVPR54)	32	R/W	8080_0000h	9.3.45/440
5_08D0	Internal interrupt n destination register (PIC_IIDR54)	32	R/W	0000_0001h	9.3.46/441
5_08E0	Internal interrupt n vector/priority register (PIC_IIVPR55)	32	R/W	8080_0000h	9.3.45/440
5_08F0	Internal interrupt n destination register (PIC_IIDR55)	32	R/W	0000_0001h	9.3.46/441
5_0900	Internal interrupt n vector/priority register (PIC_IIVPR56)	32	R/W	8080_0000h	9.3.45/440
5_0910	Internal interrupt n destination register (PIC_IIDR56)	32	R/W	0000_0001h	9.3.46/441
5_0920	Internal interrupt n vector/priority register (PIC_IIVPR57)	32	R/W	8080_0000h	9.3.45/440
5_0930	Internal interrupt n destination register (PIC_IIDR57)	32	R/W	0000_0001h	9.3.46/441
5_0940	Internal interrupt n vector/priority register (PIC_IIVPR58)	32	R/W	8080_0000h	9.3.45/440
5_0950	Internal interrupt n destination register (PIC_IIDR58)	32	R/W	0000_0001h	9.3.46/441
5_0960	Internal interrupt n vector/priority register (PIC_IIVPR59)	32	R/W	8080_0000h	9.3.45/440
5_0970	Internal interrupt n destination register (PIC_IIDR59)	32	R/W	0000_0001h	9.3.46/441
5_0980	Internal interrupt n vector/priority register (PIC_IIVPR60)	32	R/W	8080_0000h	9.3.45/440
5_0990	Internal interrupt n destination register (PIC_IIDR60)	32	R/W	0000_0001h	9.3.46/441
5_09A0	Internal interrupt n vector/priority register (PIC_IIVPR61)	32	R/W	8080_0000h	9.3.45/440
5_09B0	Internal interrupt n destination register (PIC_IIDR61)	32	R/W	0000_0001h	9.3.46/441
5_09C0	Internal interrupt n vector/priority register (PIC_IIVPR62)	32	R/W	8080_0000h	9.3.45/440
5_09D0	Internal interrupt n destination register (PIC_IIDR62)	32	R/W	0000_0001h	9.3.46/441
5_09E0	Internal interrupt n vector/priority register (PIC_IIVPR63)	32	R/W	8080_0000h	9.3.45/440
5_09F0	Internal interrupt n destination register (PIC_IIDR63)	32	R/W	0000_0001h	9.3.46/441
5_1600	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR0)	32	R/W	8000_0000h	9.3.47/442
5_1610	Messaging interrupt n (MSGn) destination register (PIC_MIDR0)	32	R/W	0000_0001h	9.3.48/443
5_1620	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR1)	32	R/W	8000_0000h	9.3.47/442
5_1630	Messaging interrupt n (MSGn) destination register (PIC_MIDR1)	32	R/W	0000_0001h	9.3.48/443
5_1640	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR2)	32	R/W	8000_0000h	9.3.47/442
5_1650	Messaging interrupt n (MSGn) destination register (PIC_MIDR2)	32	R/W	0000_0001h	9.3.48/443
5_1660	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR3)	32	R/W	8000_0000h	9.3.47/442
5_1670	Messaging interrupt n (MSGn) destination register (PIC_MIDR3)	32	R/W	0000_0001h	9.3.48/443

Table continues on the next page...

PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5_1680	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR4)	32	R/W	8000_0000h	9.3.47/442
5_1690	Messaging interrupt n (MSGn) destination register (PIC_MIDR4)	32	R/W	0000_0001h	9.3.48/443
5_16A0	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR5)	32	R/W	8000_0000h	9.3.47/442
5_16B0	Messaging interrupt n (MSGn) destination register (PIC_MIDR5)	32	R/W	0000_0001h	9.3.48/443
5_16C0	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR6)	32	R/W	8000_0000h	9.3.47/442
5_16D0	Messaging interrupt n (MSGn) destination register (PIC_MIDR6)	32	R/W	0000_0001h	9.3.48/443
5_16E0	Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPR7)	32	R/W	8000_0000h	9.3.47/442
5_16F0	Messaging interrupt n (MSGn) destination register (PIC_MIDR7)	32	R/W	0000_0001h	9.3.48/443
5_1C00	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR0)	32	R/W	8000_0000h	9.3.49/444
5_1C10	Shared message signaled interrupt destination register n (PIC_MSIDR0)	32	R/W	See section	9.3.50/445
5_1C20	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR1)	32	R/W	8000_0000h	9.3.49/444
5_1C30	Shared message signaled interrupt destination register n (PIC_MSIDR1)	32	R/W	See section	9.3.50/445
5_1C40	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR2)	32	R/W	8000_0000h	9.3.49/444
5_1C50	Shared message signaled interrupt destination register n (PIC_MSIDR2)	32	R/W	See section	9.3.50/445
5_1C60	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR3)	32	R/W	8000_0000h	9.3.49/444
5_1C70	Shared message signaled interrupt destination register n (PIC_MSIDR3)	32	R/W	See section	9.3.50/445
5_1C80	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR4)	32	R/W	8000_0000h	9.3.49/444
5_1C90	Shared message signaled interrupt destination register n (PIC_MSIDR4)	32	R/W	See section	9.3.50/445
5_1CA0	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR5)	32	R/W	8000_0000h	9.3.49/444
5_1CB0	Shared message signaled interrupt destination register n (PIC_MSIDR5)	32	R/W	See section	9.3.50/445
5_1CC0	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR6)	32	R/W	8000_0000h	9.3.49/444
5_1CD0	Shared message signaled interrupt destination register n (PIC_MSIDR6)	32	R/W	See section	9.3.50/445

Table continues on the next page...

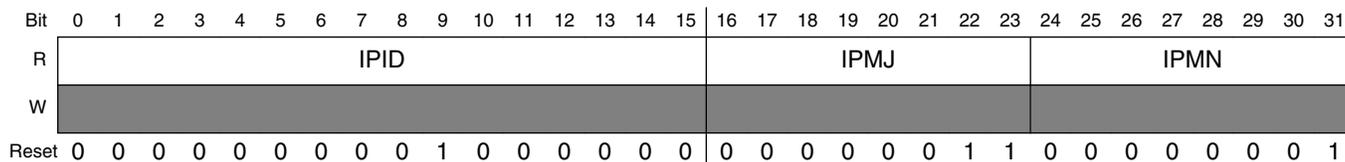
PIC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5_1CE0	Shared message signaled interrupt vector/priority register n (PIC_MSIVPR7)	32	R/W	8000_0000h	9.3.49/444
5_1CF0	Shared message signaled interrupt destination register n (PIC_MSIDR7)	32	R/W	See section	9.3.50/445
6_0040	Processor core 0 interprocessor n dispatch register (PIC_IPIDR_CPU00)	32	W	0000_0000h	9.3.51/446
6_0050	Processor core 0 interprocessor n dispatch register (PIC_IPIDR_CPU01)	32	W	0000_0000h	9.3.51/446
6_0060	Processor core 0 interprocessor n dispatch register (PIC_IPIDR_CPU02)	32	W	0000_0000h	9.3.51/446
6_0070	Processor core 0 interprocessor n dispatch register (PIC_IPIDR_CPU03)	32	W	0000_0000h	9.3.51/446
6_0080	Processor core current task priority register 0 Processor core (PIC_CTPR_CPU0)	32	R/W	0000_000Fh	9.3.52/446
6_0090	Processor core 0 who am I register (PIC_WHOAMI_CPU0)	32	R	See section	9.3.53/447
6_00A0	Processor core 0 interrupt acknowledge register (PIC_IACK_CPU0)	32	R	0000_0000h	9.3.54/448
6_00B0	Processor core 0 end of interrupt register (PIC_EOI_CPU0)	32	W	0000_0000h	9.3.55/449
6_1040	Processor core 1 interprocessor n dispatch register (PIC_IPIDR_CPU10)	32	W	0000_0000h	9.3.56/450
6_1050	Processor core 1 interprocessor n dispatch register (PIC_IPIDR_CPU11)	32	W	0000_0000h	9.3.56/450
6_1060	Processor core 1 interprocessor n dispatch register (PIC_IPIDR_CPU12)	32	W	0000_0000h	9.3.56/450
6_1070	Processor core 1 interprocessor n dispatch register (PIC_IPIDR_CPU13)	32	W	0000_0000h	9.3.56/450
6_1080	Processor core 1 current task priority register (PIC_CTPR_CPU1)	32	R/W	0000_000Fh	9.3.57/450
6_1090	Processor core 1 who am I register (PIC_WHOAMI_CPU1)	32	R	See section	9.3.58/451
6_10A0	Processor core 1 interrupt acknowledge register (PIC_IACK_CPU1)	32	R	0000_0000h	9.3.59/452
6_10B0	Processor core 1 end of interrupt register (PIC_EOI_CPU1)	32	W	0000_0000h	9.3.60/453

9.3.1 Block revision register 1 (PIC_BRR1)

BRR1 provides information about the PIC IP block.

Address: 4_0000h base + 0h offset = 4_0000h



PIC_BRR1 field descriptions

Field	Description
0–15 IPID	IP block ID.
16–23 IPMJ	The major revision of the IP block.
24–31 IPMN	The minor revision of the IP block.

9.3.2 Block revision register 2 (PIC_BRR2)

BRR2 provides information about the IP block integration option and IP block configuration options.

Address: 4_0000h base + 10h offset = 4_0010h



PIC_BRR2 field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–15 IPINTO	IP block integration options

Table continues on the next page...

PIC_BRR2 field descriptions (continued)

Field	Description
16–23 -	This field is reserved. Reserved
24–31 IPCFGO	IP block configuration options

9.3.3 Interprocessor n dispatch register (PIC_IPIDRn)

The figure below shows the four IPIDRs, one for each interprocessor interrupt channel. Writing to an IPIDR with a bit set causes a self interrupt for a single-core device. Because external bus masters can write to these registers, this feature can serve as a doorbell type interrupt.

Address: 4_0000h base + 40h offset + (16d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_IPIDRn field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1. Specifies if processor core 1 receives the interrupt. This interrupt is multicasting, so both P0 and P1 can be set. NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive the interrupt 1 Directs the interrupt to processor core 1
31 P0	Processor core 0 . Determines if processor core 0 receives the interrupt. 0 Processor core 0 does not receive the interrupt. 1 Directs the interrupt to processor core 0.

9.3.4 Current task priority register (PIC_CTPR)

There is one CTPR per processor core on this device.

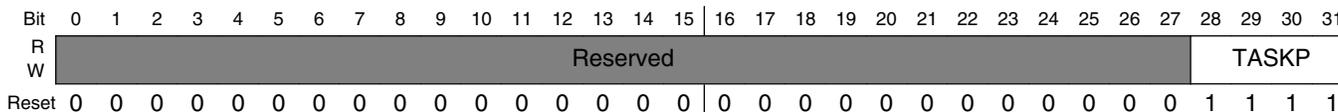
NOTE

CTPR has meaning only for interrupts routed to *int* .

Software must write the priority of the current processor core task in the CTPR for each core. The PIC uses this value for comparison with the priority of incoming interrupts. Given several concurrent incoming interrupts, the highest priority interrupt is asserted to that core if the following apply:

- The interrupt is not masked.
- The priority of the interrupt is higher than the values in the corresponding CTPR[TASKP] and ISR.

Address: 4_0000h base + 80h offset = 4_0080h



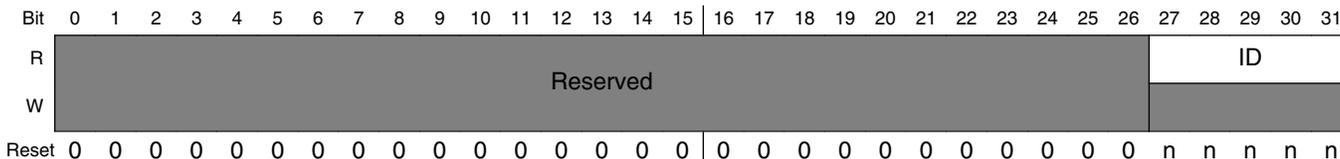
PIC_CTPR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 TASKP	Task priority. Indicates the threshold that individual interrupt priorities must exceed for the interrupt request to be serviced. 0000-1111 xVPRn[PRIORITY] must exceed this value for the interrupt request to be serviced. Note the following special cases: 0000 Lowest priority. All interrupts except those whose priority are 0 can be serviced. 1111 Highest priority. No interrupts are signaled to that processor core. Hardware selects this value on a device hard reset or when the corresponding PIR[Pn] is set.

9.3.5 Who am I register (PIC_WHOAMI)

The processor core WHOAMI *n* register can be read by a processor core to determine its physical connection to the PIC. The value returned when reading this register may be used to determine the value for the destination masks used for dispatching interrupts.

Address: 4_0000h base + 90h offset = 4_0090h



PIC_WHOAMI field descriptions

Field	Description
0–26 -	This field is reserved. Reserved
27–31 ID	Returns the ID of the processor core reading this register. Does not always return zero. 0_0000 Processor core 0 0_0001 Processor core 1. (Value not supported in single-processor implementations.) 1_1111 All other devices

9.3.6 Interrupt acknowledge register (PIC_IACK)

NOTE

IACK has meaning only for interrupts routed to *int* and should not be accessed for interrupts routed to *cint* or *IRQ_OUT_B* .

In systems based on processors built on Power Architecture technology, the interrupt acknowledge function occurs as an explicit read operation to a memory-mapped interrupt acknowledge register (IACK). Each processor core has an IACK register assigned to it. Reading IACK returns the interrupt vector corresponding to the highest priority pending interrupt. Reading IACK also has the following side effects:

- The associated field in the corresponding interrupt pending register (IPR) is cleared for edge-sensitive interrupts. See [Interrupts routed to int](#) .
- The corresponding in-service register (ISR) is updated.
- The corresponding *int* output signal from the PIC is negated.

Reading IACK when no interrupt is pending returns the spurious vector value, as described in [Spurious vector register \(PIC_SVR\)](#).

Address: 4_0000h base + A0h offset = 4_00A0h



PIC_IACK field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 VECTOR	Interrupt vector. Vector of the highest pending interrupt

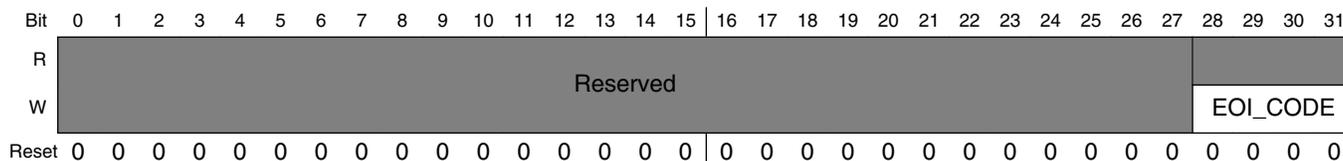
9.3.7 End of interrupt register (PIC_EOI)

NOTE

EOI has meaning only for interrupts routed to *int* and should not be accessed for interrupts routed to *cint* or *IRQ_OUT_B*.

Each core is assigned an EOI register. Writing to EOI signals the end of processing for the highest-priority interrupt (routed to *int*) currently in service. It also updates the corresponding ISR *n* by retiring the highest priority interrupt. Data values written to EOI are ignored, and zero is assumed.

Address: 4_0000h base + B0h offset = 4_00B0h



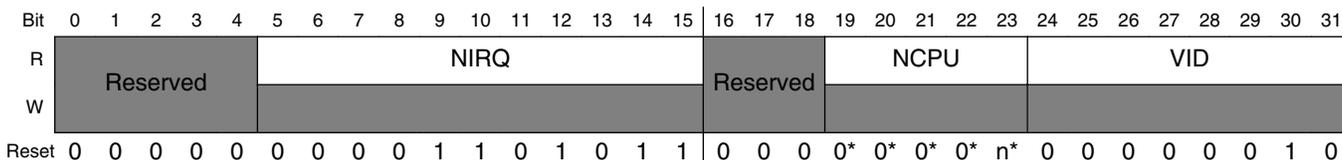
PIC_EOI field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EOI_CODE	0000 (write only)

9.3.8 Feature reporting register (PIC_FRR)

FRR provides information about interrupt and e500 processor core configurations. It also informs the programming environment of the controller version.

Address: 4_0000h base + 1000h offset = 4_1000h



* Notes:

- NCPU field: See description for NCPU for reset value.

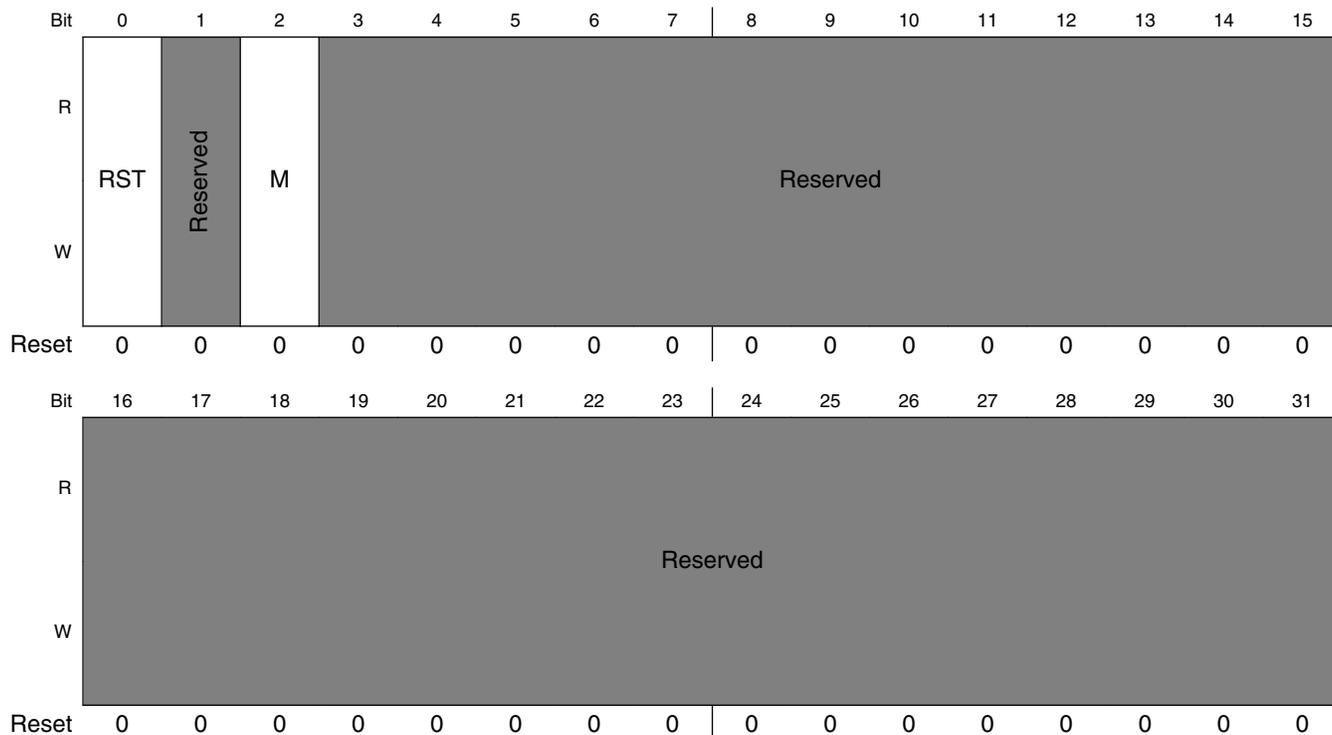
PIC_FRR field descriptions

Field	Description
0–4 -	This field is reserved. Reserved
5–15 NIRQ	Number of interrupts. Holds the binary value of the number of the highest interrupt source supported minus one. The value is 107 (0x6B or 0b000_0110_1011). A zero in this field corresponds to one source.
16–18 -	This field is reserved. Reserved, should be cleared
19–23 NCPU	Number of CPUs. The number of the highest physical CPUs (or e500 processor cores) supported minus one. 00000 Single core-core0 00001 Two cores-core0 and core1
24–31 VID	Version ID. Reports the OpenPIC specification revision level supported by this interrupt controller implementation. The VID field's value of two (0x02) corresponds to revision 1.2 which is the revision level currently supported.

9.3.9 Global configuration register (PIC_GCR)

GCR controls the PIC's operating mode, and allows software to reset the PIC.

Address: 4_0000h base + 1020h offset = 4_1020h



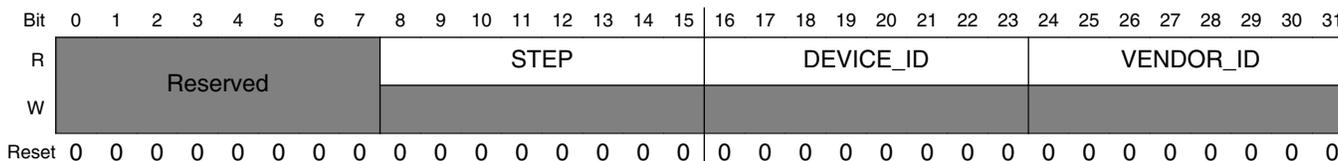
PIC_GCR field descriptions

Field	Description
0 RST	Reset. Setting RST forces the PIC to be reset. Cleared automatically when the reset sequence is complete. See Resetting the PIC for more information.
1 -	This field is reserved. Reserved
2 M	Mode. PIC operating mode. Modes of operation provides details about these modes. 0 Pass-through mode. On-chip PIC is disabled and interrupts detected on IRQ0 are passed directly to core 0. 1 Mixed mode. Interrupts are handled by the normal priority and delivery mechanisms of the PIC.
3–31 -	This field is reserved. Reserved

9.3.10 Vendor identification register (PIC_VIR)

VIR is defined by the OpenPIC specifications and is provided for compliance. The zero value for VIR[VENDORID] indicates a generic OpenPIC-compliant device, which makes the other VIR fields meaningless.

Address: 4_0000h base + 1080h offset = 4_1080h



PIC_VIR field descriptions

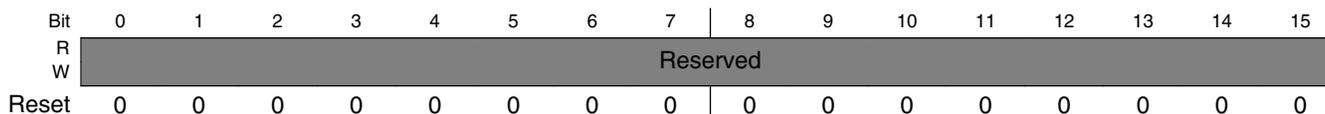
Field	Description
0–7 -	This field is reserved. Reserved
8–15 STEP	Stepping. Indicates the silicon revision for this device. Has no meaning if VENDOR_ID value is zero.
16–23 DEVICE_ID	Device identification. Vendor-specified identifier for this device. Has no meaning if VENDOR_ID is zero.
24–31 VENDOR_ID	Vendor identification. Specifies the manufacturer of this part. A value of zero implies a generic OpenPIC-compliant device.

9.3.11 Processor core initialization register (PIC_PIR)

PIR provides a way for software to generate a core reset. Setting P1 or P0 causes the respective *core 0_hreset* or *core1_hreset* signal to assert. However, if one core is used to reset another one, the core being reset can effectively be held off indefinitely from issuing its initial boot vector fetch to the platform by leaving its appropriate PIR[Px] bit asserted. Clearing it releases the core to fetch.

Note that although the OpenPIC architecture was defined to support up to 32 processing cores, only fields corresponding to the number of cores on the device are implemented.

Address: 4_0000h base + 1090h offset = 4_1090h



Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved														P1	P0
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_PIR field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1 reset. Setting this bit causes the PIC to assert the <i>core1_hreset</i> signal. Reserved in single-processor implementations.
31 P0	Processor core 0 reset. Setting this bit causes the PIC to assert the <i>core0_hreset</i> signal.

9.3.12 Interprocessor interrupt n vector/priority register (PIC_IPIVPR_n)

IPIVPRs contain the interrupt vector and priority fields for the four interprocessor interrupt channels. There is one vector/priority register per channel. The VECTOR and PRIORITY values should not be changed while IPIVPR *n* [A] is set. See [Flow of interrupt control](#) for information on IPR and ISR.

Address: 4_0000h base + 10A0h offset + (16d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	MSK	A	Reserved										PRIORITY			
W																
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	VECTOR															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_IPIVPR_n field descriptions

Field	Description
0 MSK	Mask. Mask interrupts to <i>int</i> from this source. 0 An interrupt request is generated if the corresponding IPR bit is set. 1 Further interrupts from this source are disabled.
1 A	Activity. Indicates an interrupt has been requested or is in service. The VECTOR and PRIORITY values should not be changed while this bit is set.

Table continues on the next page...

PIC_IPIVPR_n field descriptions (continued)

Field	Description
	0 No current interrupt activity associated with this source. 1 The interrupt field for this source is set in the IPR or ISR.
2–11 -	This field is reserved. Reserved
12–15 PRIORITY	Priority. Specifies the interrupt priority. The lowest priority is 0 and the highest priority is 15. A priority level of 0 inhibits signaling of this interrupt to the core. Affects only interrupts routed to int.
16–31 VECTOR	Vector (Affects only interrupts routed to int). Contains the value returned when IACK is read and this interrupt resides in the corresponding interrupt request register (IRR) for that core, as shown in Figure 9-331 .

9.3.13 Spurious vector register (PIC_SVR)

SVR contains the 16-bit vector returned to the e500 processor core when the corresponding IACK register is read for a spurious interrupt.

Address: 4_0000h base + 10E0h offset = 4_10E0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

PIC_SVR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 VECTOR	Spurious interrupt vector. Value returned when IACK is read during a spurious vector fetch. Spurious vector generation gives information about the conditions that may cause a spurious vector fetch.

9.3.14 Timer frequency reporting register group X (PIC_TFRR_n)

The TFRRs are written by software to report the clocking frequency of the PIC timers. Note that although TFRRs are read/write, the PIC ignores the register values.

Address: 4_0000h base + 10F0h offset + (4096d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_TFRRn field descriptions

Field	Description
0–31 FREQ	Timer frequency (in ticks/second (Hz)). Used to communicate the frequency of the global timers' clock source, (either the CCB clock or the frequency of the RTC signal), to user software. TFRRx is set only by software for later use by other applications and its value in no way affects the operating frequency of the global timers. The timers operate at a ratio of this clock frequency, as set by TCRx[CLKR]. See Timer control register group n (PIC_TCRn) .

9.3.15 Global timer n current count register group A (PIC_GTCCRAn)

The GTCCRs contain the current count for each of the four PIC timers in each of the two groups.

Address: 4_0000h base + 1100h offset + (64d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	TOG	COUNT														
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	COUNT															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_GTCCRAn field descriptions

Field	Description
0 TOG	Toggle. Toggles when the current count decrements to zero. Cleared when GTBCR xn [CI] goes from 1 to 0.
1–31 COUNT	Current count. Decrementd while GTBCR xn [CI] is zero. When the timer count reaches zero, an interrupt is generated (provided it is not masked), the toggle bit is inverted, and the count is reloaded. For non-cascaded timers, the reload value is the contents of the corresponding GTBCR xn . Cascaded timers are reloaded with either all ones, or the GTBCR xn contents, depending on the value of TCRn[ROVR]. See Timer control register group n (PIC_TCRn) , for more details.

9.3.16 Global timer n base count register group A (PIC_GTBCRAn)

The GTBCRs contain the base counts for each of the four PIC timers in each of the two groups. This value is reloaded into the corresponding GTCCR *xn* when the current count reaches zero. Note that when zero is written to the base count field, (and GTCCR *xn* [CI] = 0), the timer generates an interrupt on every timer cycle.

Address: 4_0000h base + 1110h offset + (64d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	CI	BASE_CNT														
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W	BASE_CNT															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

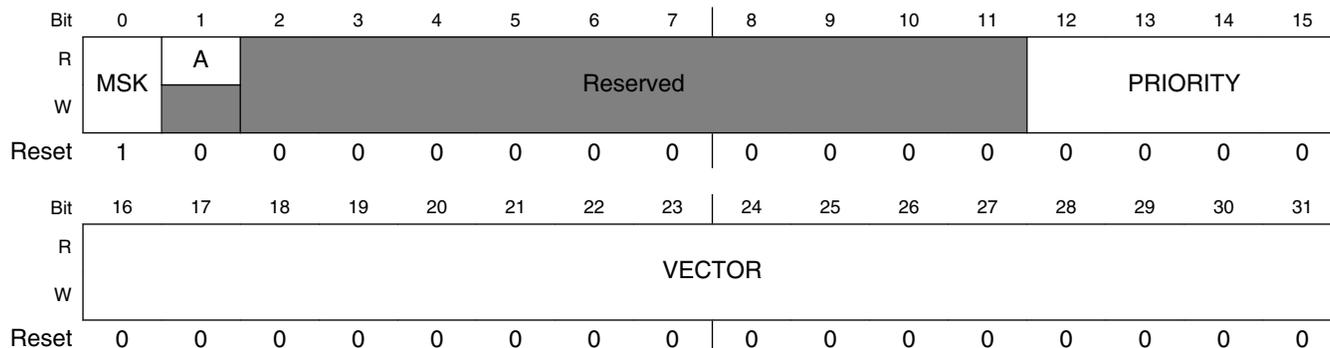
PIC_GTBCRAn field descriptions

Field	Description
0 CI	Count inhibit. Always set following reset 0 Counting enabled 1 Counting inhibited
1–31 BASE_CNT	Base count. When CI transitions from 1 to 0, this value is copied into the corresponding GTCCR <i>xn</i> and the toggle bit is cleared. If CI is already cleared (counting is in progress), the base count is copied to the GTCCR <i>xn</i> at the next zero crossing of the current count.

9.3.17 Global timer n vector/priority register group A (PIC_GTVPRAn)

The GTVPRs contain the interrupt vector and the interrupt priority values for the timers. They also contain the mask and activity fields for all the timers. See [Flow of interrupt control](#) for information on IPR and ISR.

Address: 4_0000h base + 1120h offset + (64d × i), where i=0d to 3d



PIC_GTVPRAn field descriptions

Field	Description
0 MSK	Mask. Mask interrupts to <i>int</i> from this source. 0 An interrupt request is generated if the corresponding IPR bit is set. 1 Further interrupts from this source are disabled.
1 A	Activity. Indicates an interrupt has been requested or is in service. The VECTOR and PRIORITY values should not be changed while this bit is set. 0 No current interrupt activity associated with this source. 1 The interrupt field for this source is set in the IPR or ISR.
2–11 -	This field is reserved. Reserved
12–15 PRIORITY	Priority. Specifies the interrupt priority. The lowest priority is 0 and the highest priority is 15. A priority level of 0 inhibits signaling of this interrupt to the core. Affects only interrupts routed to <i>int</i> .
16–31 VECTOR	Vector (Affects only interrupts routed to <i>int</i>). Contains the value returned when IACK is read and this interrupt resides in the corresponding interrupt request register (IRR) for that core, as shown in Figure 9-331 .

9.3.18 Global timer n destination register group A (PIC_GTDRAn)

The GTDR *xn* registers control the destination (core) to which each timer's interrupt is directed. Note that GTDR *xn* bits can be set independent of each other and that either P1 or P0 or both can be set for this type of interrupt .

Address: 4_0000h base + 1130h offset + (64d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved																	
W	Reserved																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved															P1	P0	
W	Reserved															P1	P0	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	1	

PIC_GTDRAn field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1. This interrupt is multicasting, so both P0 and P1 can be set. NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive this interrupt 1 Directs the timer interrupt to processor core 1
31 P0	Processor core 0 . Default destination after PIC is reset. Both P0 and P1 can be set. 0 Processor core 0 does not receive this interrupt. 1 Directs the timer interrupt to processor core 0 .

9.3.19 Timer control register group n (PIC_TCRn)

The TCR *n* registers provide various configuration options such as count frequency and roll-over behavior for the timers.

There are two choices for the clock source for the timers: a selectable frequency ratio from the CCB bus clock, or the RTC signal. TCRs can be cascaded to create timers larger than the default 31-bit global timers. Timer cascade fields allow configuration of up to two 63-bit timers, one 95-bit timer, or one 127-bit timer (within each group).

With one exception mentioned below, the value reloaded into a timer is determined by its roll-over control field, TCR *n* [ROVR]. Setting TCR *n* [ROVR] causes its GTCCR *xn* to roll over to all ones when the count reaches zero. This is equivalent to reloading the count register with 0xFFFF_FFFF instead of its base count value. Clearing a timer's associated ROVR bit ensures the timer always reloads with its base count value.

When timers are cascaded, the last (most significant) counter in the cascade also affects their roll-over behavior. Cascaded timers always reload their base count when the most significant counter has decremented to zero, regardless of the TCR *n* [ROVR] settings.

For example, timers 0-2 can be cascaded to generate one interrupt per hour. Given a CCB clock frequency of 333 MHz, letting the timer clock frequency default to 1/8th the system clock, (TCR *n* [CLKR] = 0 sets a clock ratio of 8), provides a basic input of 41.625 MHz to timer 0. Setting timer 0 to count 41,625,000 (0x27B_25A8) timer clock cycles generates one output per second. Setting both timers 1 and 2 to 59, and cascading all three timers, generates one interrupt every hour from timer 2.

Table 9-48. Parameters for Hourly Interrupt Timer Cascade Example

System Clock	Clock Ratio	Timer Clock	Timer 0 Count	Timer 1 Count	Timer 2 Count
333 MHz	1 / 8	41.625 MHz	41.625 x 10 ⁶ (0x027B_25A8)	59 ¹ (0x0000_0036)	59 (0x0000_0036)

- 1. Counting down from 59 through 0 requires 60 ticks.
- 1. Counting down from 59 through 0 requires 60 ticks.
- 1. Counting down from 59 through 0 requires 60 ticks.

$$(41.625 \times 10^6 \text{ ticks/sec}) \times (60 \text{ sec/min}) \times (60 \text{ min/hr}) = \text{total ticks/hr generating 1 interrupt/hr}$$

Figure 9-47. Example Calculation for Cascaded Timers

Address: 4_0000h base + 1300h offset + (4096d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved				ROVR				Reserved				RTM			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved				CLKR				Reserved				CASC			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_TCR*n* field descriptions

Field	Description
0-4 -	This field is reserved. Reserved
5-7 ROVR	Roll-over control for cascaded timers only. Specifies behavior when count reaches zero by identifying the source of the reload value. Cascaded timers are always reloaded with their base count value when the

Table continues on the next page...

PIC_TCR_n field descriptions (continued)

Field	Description
	<p>more significant timer in the cascade (the upstream timer) is zero. Bits 5-7 correspond to timers 2-0. Note that global timer 3 always reloads with its GTBCR <i>xn</i> .</p> <p>0 The timer does not roll over. When the count reaches zero, GTCCR <i>xn</i> is reloaded with the GTBCR <i>xn</i> value.</p> <p>1 Timer rolls over at zero to all ones. (When the count reaches zero, GTCCR <i>xn</i> is reloaded with 0xFFFF_FFFF.)</p> <p>000 All timers reload with base count.</p> <p>001 Timers 1 and 2 reload with base count, timer 0 rolls over (reloads with 0xFFFF_FFFF).</p> <p>010 Timers 0 and 2 reload with base count, timer 1 rolls over (reloads with 0xFFFF_FFFF).</p> <p>011 Timer 2 reloads with base count, timers 0 and 1 roll over (reload with 0xFFFF_FFFF).</p> <p>100 Timers 0 and 1 reload with base count, timer 2 rolls over (reloads with 0xFFFF_FFFF).</p> <p>101 Timer 1 reloads with base count, timers 0 and 2 roll over (reload with 0xFFFF_FFFF).</p> <p>110 Timer 0 reloads with base count, timers 1 and 2 roll over (reload with 0xFFFF_FFFF).</p> <p>111 Timers 0, 1, and 2 roll over (reload with 0xFFFF_FFFF).</p>
8–14 -	This field is reserved. Reserved
15 RTM	<p>Real time mode. Specifies the clock source for the PIC timers.</p> <p>0 Timer clock frequency is a ratio of the frequency of the CCB clock as determined by the CLKR field. This is the default value.</p> <p>1 The RTC signal is used to clock the PIC timers. If this bit is set, the CLKR field has no meaning.</p>
16–21 -	This field is reserved. Reserved
22–23 CLKR	<p>Clock ratio. Specifies the ratio of the timer frequency to the CCB clock. The following are supported:</p> <p>00 Default. Divide by 8</p> <p>01 Divide by 16</p> <p>10 Divide by 32</p> <p>11 Divide by 64</p>
24–28 -	This field is reserved. Reserved
29–31 CASC	<p>Cascade timers. Specifies the output of particular global timers as input to others.</p> <p>000 Default. Timers not cascaded</p> <p>001 Cascade timers 0 and 1</p> <p>010 Cascade timers 1 and 2</p> <p>011 Cascade timers 0, 1, and 2</p> <p>100 Cascade timers 2 and 3</p> <p>101 Cascade timers 0 and 1; timers 2 and 3</p> <p>110 Cascade timers 1, 2, and 3</p> <p>111 Cascade timers 0, 1, 2, and 3</p>

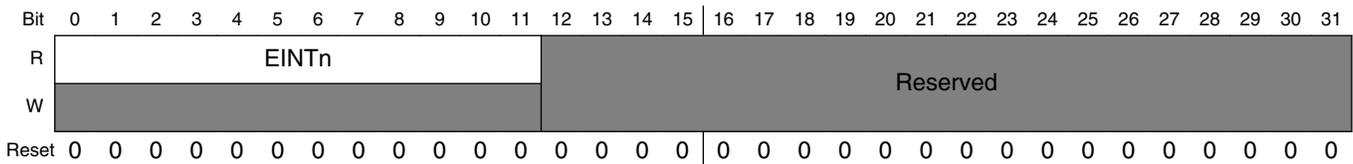
1. Counting down from 59 through 0 requires 60 ticks.

9.3.20 External interrupt summary register (PIC_ERQSR)

NOTE

ERQSR fields report only the current logic level of IRQ0-IRQ11 pins. These fields were designed to work with level-sensitive interrupts; values returned for edge-sensitive interrupts may be unreliable.

Address: 4_0000h base + 1308h offset = 4_1308h

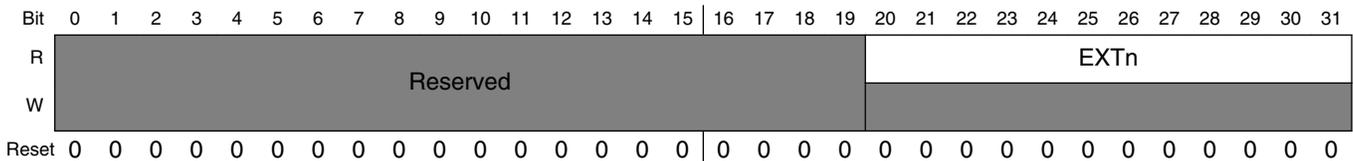


PIC_ERQSR field descriptions

Field	Description
0–11 EINTn	External interrupts signal 0-11 status. Bit 0 represents EINT0. Bit 11 represents EINT 11. 0 The corresponding external interrupt signal is at a low logic level . 1 The corresponding external interrupt signal is at a high logic level .
12–31 -	This field is reserved. Reserved

9.3.21 IRQ_OUT_B summary register 0 (PIC_IRQSR0)

Address: 4_0000h base + 1310h offset = 4_1310h



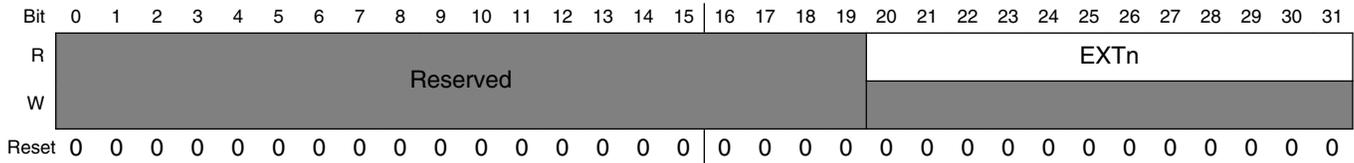
PIC_IRQSR0 field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 EXTn	External interrupts 0-11 . Each bit corresponds to a unique interrupt according to the following: Bit Interrupt 20 IRQ0

Table continues on the next page...

9.3.24 Critical interrupt summary register 0 (PIC_CISR0)

Address: 4_0000h base + 1330h offset = 4_1330h

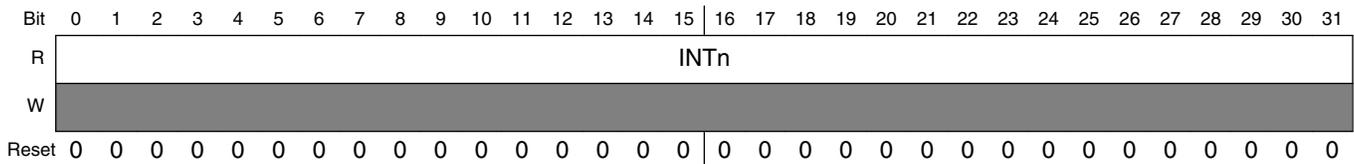


PIC_CISR0 field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 EXTn	External interrupts 0-11. Bit 20 represents IRQ0. Bit 31 represents IRQ11. 0 The corresponding interrupt is not active . 1 The corresponding interrupt is active

9.3.25 Critical interrupt summary register 1 (PIC_CISR1)

Address: 4_0000h base + 1340h offset = 4_1340h

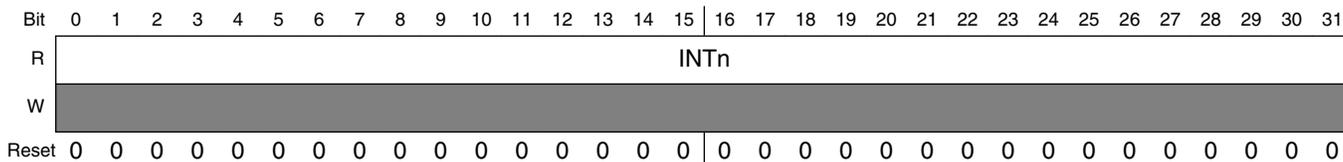


PIC_CISR1 field descriptions

Field	Description
0–31 INTn	Internal interrupts 0-31. Bit 0 represents INT0. Bit 31 represents INT31. 0 Corresponding interrupt is not active or not routed to cint 0 or cint1 . 1 The corresponding interrupt is active and is routed to the cint 0 or cint1 (if the corresponding x IDRn[CI] is set).

9.3.26 Critical interrupt summary register 2 (PIC_CISR2)

Address: 4_0000h base + 1344h offset = 4_1344h



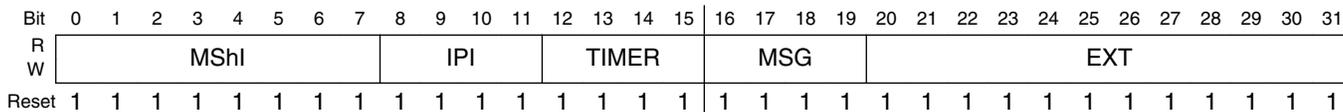
PIC_CISR2 field descriptions

Field	Description
0–31 INTn	Internal interrupts 32-63. Bit 0 represents INT32. Bit 31 represents INT63. 0 Corresponding interrupt is not active or not routed to cint 0 or cint1 . 1 The corresponding interrupt is active and is routed to the cint 0 or cint1 , if the corresponding x IDRn[CI] is set.

9.3.27 Performance monitor n mask register 0 (PIC_PMnMR0)

Each PM n MR0 register is matched with a PM n MR1 and a PM n MR2 register. Because each unreserved bit in the 96-bit vector (PM n MR0/1/2) specifies a different interrupt, only one bit in the 96-bit vector can be unmasked at a time. Unmasking more than 1 bit per set is considered a programming error and results in unpredictable behavior.

Address: 4_0000h base + 1350h offset + (32d × i), where i=0d to 3d



PIC_PMnMR0 field descriptions

Field	Description
0–7 MShI	Shared message signaled interrupts 0-7 0 The corresponding interrupt source generates a performance monitor event when the interrupt occurs. 1 The corresponding interrupt does not generate a performance monitor event.
8–11 IPI	Interprocessor interrupts 0-3 0 The corresponding interrupt source generates a performance monitor event when the interrupt occurs. 1 The corresponding interrupt does not generate a performance monitor event.

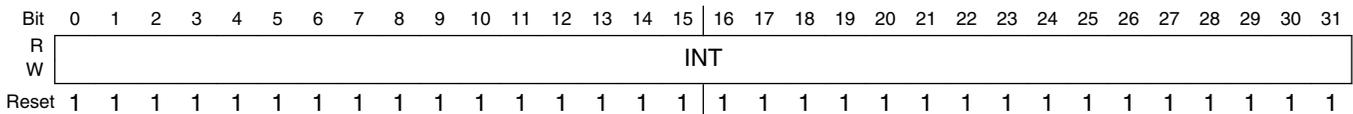
Table continues on the next page...

PIC_PMnMR0 field descriptions (continued)

Field	Description
12–15 TIMER	<p>Timer interrupts 0-3 (Group A and Group B: Each bit represents an OR of the event for the correspondingly numbered timer in Group A and that in Group B).</p> <p>0 The corresponding interrupt source generates a performance monitor event when the interrupt occurs. 1 The corresponding interrupt does not generate a performance monitor event.</p>
16–19 MSG	<p>Message interrupts 0-7</p> <p>Bit 0 is used for MSG0 and MSG4 Bit 1 is used for MSG1 and MSG5 Bit 2 is used for MSG2 and MSG6 Bit 3 is used for MSG3 and MSG7</p> <p>0 The corresponding interrupt source generates a performance monitor event when the interrupt occurs. 1 The corresponding interrupt does not generate a performance monitor event.</p>
20–31 EXT	<p>External interrupts IRQ[0:11]</p> <p>0 The corresponding interrupt source generates a performance monitor event when the interrupt occurs. 1 The corresponding interrupt does not generate a performance monitor event.</p>

9.3.28 Performance monitor n mask register 1 (PIC_PMnMR1)

Address: 4_0000h base + 1360h offset + (32d × i), where i=0d to 3d

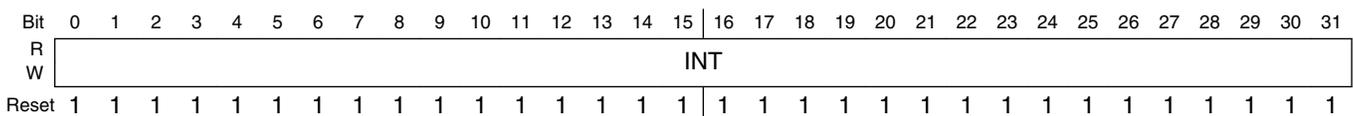


PIC_PMnMR1 field descriptions

Field	Description
0–31 INT	<p>Internal interrupts 0-31</p> <p>0 The corresponding interrupt source generates a performance monitor event when the interrupt occurs. 1 The corresponding interrupt does not generate a performance monitor event.</p>

9.3.29 Performance monitor n mask register 2 (PIC_PMnMR2)

Address: 4_0000h base + 1364h offset + (32d × i), where i=0d to 3d



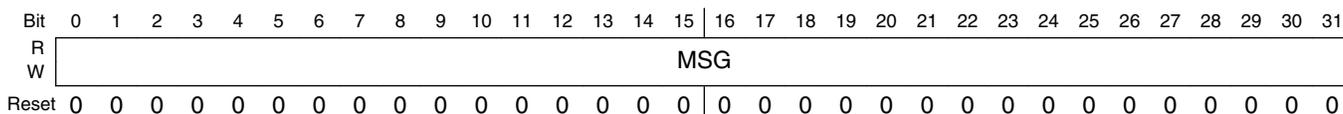
PIC_PMnMR2 field descriptions

Field	Description
0–31 INT	Internal interrupts 32-63 0 The corresponding interrupt source generates a performance monitor event when the interrupt occurs. 1 The corresponding interrupt does not generate a performance monitor event.

9.3.30 Message register n (PIC_MSGRn)

The message registers (MSGR0-MSGR 7) can contain a 32-bit message. For (MSGR4-MSGR7), see [Message register n \(PIC_MSGRn\)](#)

Address: 4_0000h base + 1400h offset + (16d × i), where i=0d to 3d



PIC_MSGRn field descriptions

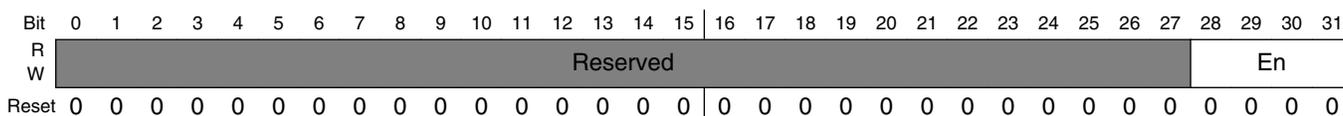
Field	Description
0–31 MSG	Message. Contains the 32-bit message data.

9.3.31 Message enable register (PIC_MER)

The MER contains the enable bits for each message register. The enable bit must be set to enable interrupt generation when the corresponding message register is written.

When bits in MER are set to mask message interrupts, an interrupt is not generated if the message register is written while it is masked in MER and the MER bit is then cleared. To mask the interrupt without loss, set MIVPR *n* [MSK] (See [Messaging interrupt n \(MSGn\) vector/priority register \(PIC_MIVPRn\)](#)). MER should be set to 0x0000_000F at reset and then left unchanged during normal operation.

Address: 4_0000h base + 1500h offset = 4_1500h



PIC_MER field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 En	Enable 3-enable 0 . Used to enable interrupt generation for MSGR _n (where $n = 0- 3$). 0 Interrupt generation for MSGR _n disabled. 1 Interrupt generation for MSGR _n enabled.

9.3.32 Message status register (PIC_MSR)

The message status register (MSR) contains status bits for each message register. A status bit is set when the corresponding messaging interrupt is active. Writing a 1 to a status bit clears the corresponding message interrupt and the status bit.

Address: 4_0000h base + 1510h offset = 4_1510h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																Sn															
W	Reserved																Sn															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_MSR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 Sn	Status 3-status 0 . Reports status of messaging interrupt n . Writing a 1 clears this field. 0 Messaging interrupt n is not active. 1 Messaging interrupt n is active.

9.3.33 Shared message signaled interrupt register n (PIC_MSIR_n)

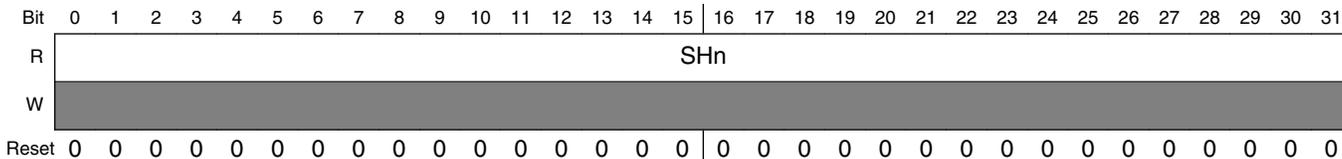
The eight MSIRs indicate which of the up to 32 interrupt sources sharing the message register have pending interrupts. These registers are cleared when read. A write to these registers has no effect.

NOTE

After soft reset, read the MSIRs to ensure that any interrupts previously pending are cleared.

PIC memory map/register definition

Address: 4_0000h base + 1600h offset + (16d × i), where i=0d to 7d



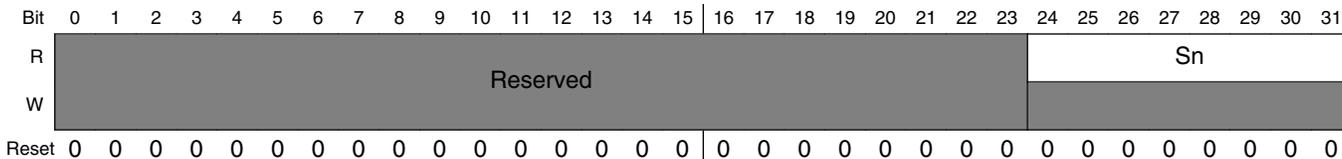
PIC_MSIRn field descriptions

Field	Description
0–31 SHn	Message sharer n has a pending interrupt.

9.3.34 Shared message signaled interrupt status register (PIC_MSISR)

MSISR contains the status bits for the shared message signaled interrupts. A status bit is set when the corresponding MSIR has an active interrupt. The status bit is 0 if all the corresponding shared interrupt sources are cleared for that MSIR.

Address: 4_0000h base + 1720h offset = 4_1720h



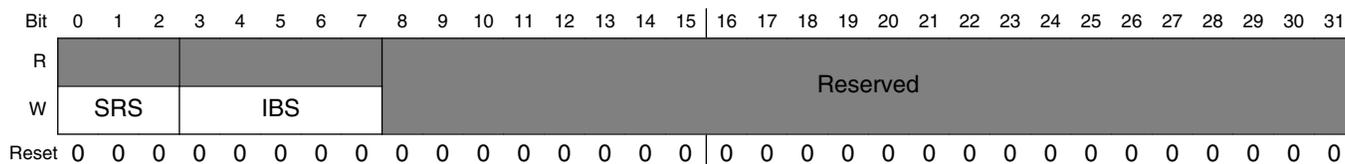
PIC_MSISR field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24–31 Sn	Status n. 0 MSIRn is not active. 1 MSIRn has an active interrupt.

9.3.35 Shared message signaled interrupt index register (PIC_MSIIR)

MSIIR provides the mechanism for setting an interrupt in the MSIRs. When MSIIR is written, MSIIR[SRS] selects the register in which an interrupt bit is to be set; MSIIR[IBS] selects the shared interrupt field in the selected MSIR register to be set. MSIIR is primarily intended to support PCI Express MSIs.

Address: 4_0000h base + 1740h offset = 4_1740h



PIC_MSIIR field descriptions

Field	Description
0–2 SRS	Shared interrupt register select. Selects the MSIR to be written Example settings: 000 MSIR 0 001 MSIR 1 010 MSIR 2 ... 111 MSIR 7
3–7 IBS	Interrupt bit select-Selects the bit to set in the MSIR Example settings: 00000 Set field SH0 (bit 31) 00001 Set field SH1 (bit 30) 00010 Set field SH2 (bit 29) ... 11111 Set field SH31 (bit 0)
8–31 -	This field is reserved. Reserved

9.3.36 Global timer n current count register group B (PIC_GTCCRBn)

The GTCCRs contain the current count for each of the four PIC timers in each of the two groups.

Address: 4_0000h base + 2100h offset + (64d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	TOG	COUNT																
W	[Shaded]																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	COUNT																	
W	[Shaded]																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0

PIC_GTCCRBn field descriptions

Field	Description
0 TOG	Toggle. Toggles when the current count decrements to zero. Cleared when GTBCR xn [CI] goes from 1 to 0.
1–31 COUNT	Current count. Decrementing while GTBCR xn [CI] is zero. When the timer count reaches zero, an interrupt is generated (provided it is not masked), the toggle bit is inverted, and the count is reloaded. For non-cascaded timers, the reload value is the contents of the corresponding GTBCR xn . Cascaded timers are reloaded with either all ones, or the GTBCR xn contents, depending on the value of TCRn[ROVR]. See Timer control register group n (PIC_TCRn) for more details.

9.3.37 Global timer n base count register group B (PIC_GTBCRBn)

The GTBCRs contain the base counts for each of the four PIC timers in each of the two groups. This value is reloaded into the corresponding GTCCR *xn* when the current count reaches zero. Note that when zero is written to the base count field, (and GTCCR *xn* [CI] = 0), the timer generates an interrupt on every timer cycle.

Address: 4_0000h base + 2110h offset + (64d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R																		
W	CI	BASE_CNT																
Reset	1	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R																		
W	BASE_CNT																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0

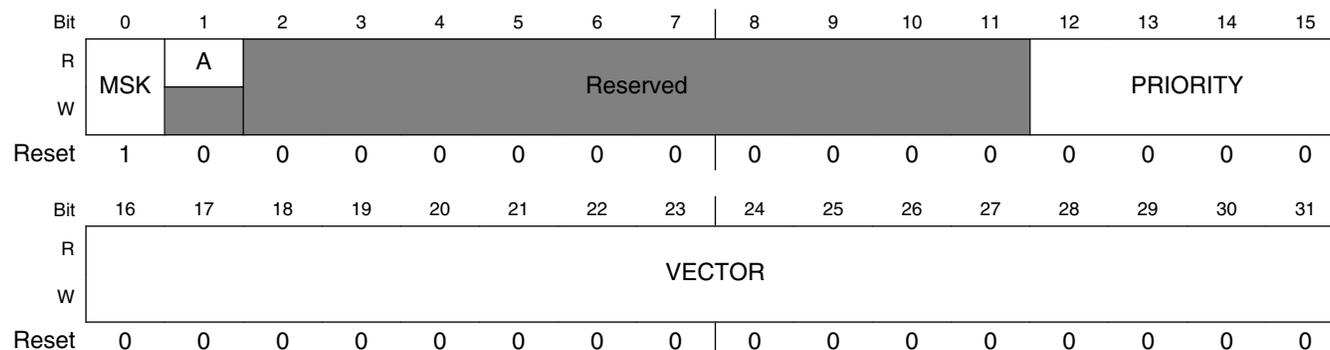
PIC_GTBCRBn field descriptions

Field	Description
0 CI	Count inhibit. Always set following reset 0 Counting enabled 1 Counting inhibited
1–31 BASE_CNT	Base count. When CI transitions from 1 to 0, this value is copied into the corresponding GTCCR <i>xn</i> and the toggle bit is cleared. If CI is already cleared (counting is in progress), the base count is copied to the GTCCR <i>xn</i> at the next zero crossing of the current count.

9.3.38 Global timer n vector/priority register group B (PIC_GTVPRBn)

The GTVPRs contain the interrupt vector and the interrupt priority values for the timers. They also contain the mask and activity fields for all the timers. See [Flow of interrupt control](#) for information on IPR and ISR.

Address: 4_0000h base + 2120h offset + (64d × i), where i=0d to 3d



PIC_GTVPRBn field descriptions

Field	Description
0 MSK	Mask. Mask interrupts to <i>int</i> from this source. 0 An interrupt request is generated if the corresponding IPR bit is set. 1 Further interrupts from this source are disabled.
1 A	Activity. Indicates an interrupt has been requested or is in service. The VECTOR and PRIORITY values should not be changed while this bit is set. 0 No current interrupt activity associated with this source. 1 The interrupt field for this source is set in the IPR or ISR.
2–11 -	This field is reserved. Reserved
12–15 PRIORITY	Priority. Specifies the interrupt priority. The lowest priority is 0 and the highest priority is 15. A priority level of 0 inhibits signaling of this interrupt to the core. Affects only interrupts routed to <i>int</i> .
16–31 VECTOR	Vector (Affects only interrupts routed to <i>int</i>). Contains the value returned when IACK is read and this interrupt resides in the corresponding interrupt request register (IRR) for that core, as shown in Figure 9-331 .

9.3.39 Global timer n destination register group B (PIC_GTDRBn)

The GTDR *xn* registers control the destination (core) to which each timer's interrupt is directed. Note that GTDR *xn* bits can be set independently of each other and that either P1 or P0 or both can be set for this type of interrupt .

Address: 4_0000h base + 2130h offset + (64d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved																	
W	Reserved																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved																P1	P0
W	Reserved																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	1	

PIC_GTDRBn field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1. This interrupt is multicasting, so both P0 and P1 can be set. NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive this interrupt 1 Directs the timer interrupt to processor core 1
31 P0	Processor core 0 . Default destination after PIC is reset. Both P0 and P1 can be set. 0 Processor core 0 does not receive this interrupt. 1 Directs the timer interrupt to processor core 0 .

9.3.40 Message register n (PIC_MSGRan)

The message registers (MSGR0-MSGR 7) can contain a 32-bit message. For (MSGR0-MSGR3), see [Message register n \(PIC_MSGRn\)](#)

Address: 4_0000h base + 2400h offset + (16d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	MSG																																	
W	MSG																																	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_MSGRn field descriptions

Field	Description
0–31 MSG	Message. Contains the 32-bit message data.

9.3.41 Message enable register (PIC_MERa)

The MER contains the enable bits for each message register. The enable bit must be set to enable interrupt generation when the corresponding message register is written.

When bits in MER are set to mask message interrupts, an interrupt is not generated if the message register is written while it is masked in MER and the MER bit is then cleared.

To mask the interrupt without loss, set MIVPR *n* [MSK]. (See [Messaging interrupt *n* \(MSGn\) vector/priority register \(PIC_MIVPRn\)](#) .) MER should be set to 0x0000_000F at reset and then left unchanged during normal operation.

Address: 4_0000h base + 2500h offset = 4_2500h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																			En												
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

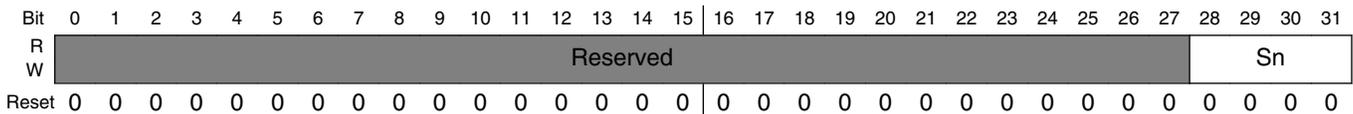
PIC_MERa field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 En	Enable 3-enable 0 . Used to enable interrupt generation for MSGRn (where <i>n</i> = 0-3). 0 Interrupt generation for MSGRn disabled. 1 Interrupt generation for MSGRn enabled.

9.3.42 Message status register (PIC_MSRA)

The message status register (MSR) contains status bits for each message register. A status bit is set when the corresponding messaging interrupt is active. Writing a 1 to a status bit clears the corresponding message interrupt and the status bit.

Address: 4_0000h base + 2510h offset = 4_2510h



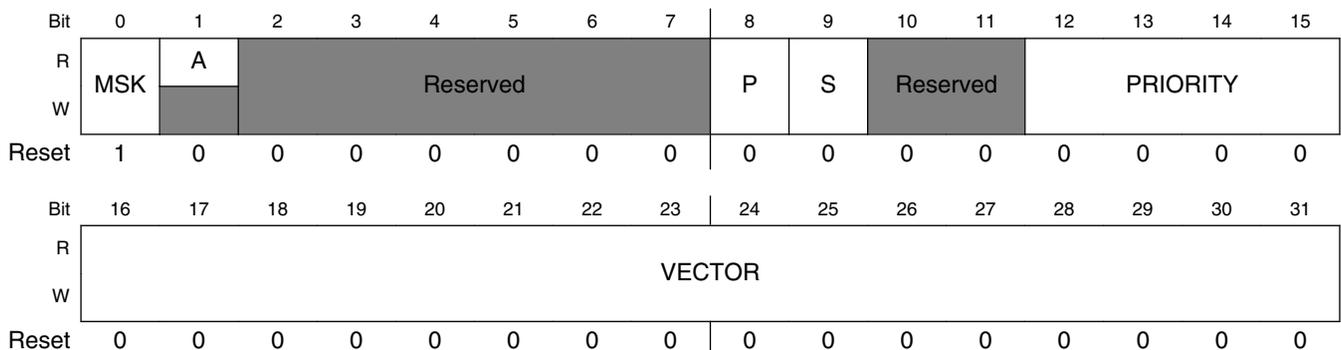
PIC_MSRA field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 Sn	Status 3-status 0 . Reports status of messaging interrupt <i>n</i> . Writing a 1 clears this field. 0 Messaging interrupt <i>n</i> is not active. 1 Messaging interrupt <i>n</i> is active.

9.3.43 External interrupt n (IRQn) vector/priority register (PIC_EIVPRn)

The EIVPRs contain polarity and sense fields for the external interrupts, that is, those caused by the assertion of any of IRQ[0:11]. See [Flow of interrupt control](#) for information on IPR and ISR.

Address: 4_0000h base + 1_0000h offset + (32d × i), where i=0d to 11d



PIC_EIVPR_n field descriptions

Field	Description
0 MSK	Mask. Mask interrupts from this source. MSK affects only interrupts routed to int. 0 An interrupt request is generated if the corresponding IPR bit is set. 1 Further interrupts from this source are disabled.
1 A	Activity. Indicates an interrupt has been requested or is in service. The VECTOR and PRIORITY values should not be changed while this bit is set. Affects only interrupts routed to int. 0 No current interrupt activity associated with this source. 1 The interrupt field for this source is set in the IPR or ISR.
2–7 -	This field is reserved. Reserved
8 P	Polarity. Specifies the polarity for the external interrupt. NOTE: If any of IRQ[0:11] are used to receive INT _x signals from one of the PCI Express ports as a root complex, the polarity must be set to be active-high. 0 Polarity is active-low or negative edge-triggered. 1 Polarity is active-high or positive edge-triggered.
9 S	Sense. Specifies the sense for external interrupts. NOTE: If an IRQ <i>n</i> signal is used to receive INT <i>x</i> signals from one of the PCI Express ports as a root complex, S must be set to be level-sensitive. 0 The external interrupt is edge sensitive. 1 The external interrupt is level sensitive. This setting must be used to direct the interrupt to IRQ_OUT_B or cint.
10–11 -	This field is reserved. Reserved
12–15 PRIORITY	Priority. Specifies the interrupt priority. The lowest priority is 0 and the highest priority is 15. A priority level of 0 inhibits signaling of this interrupt to the core. Affects only interrupts routed to int.
16–31 VECTOR	Vector (Affects only interrupts routed to int). Contains the value returned when IACK is read and this interrupt resides in the corresponding interrupt request register (IRR) for that core, as shown in Figure 9-331 .

9.3.44 External interrupt n (IRQn) destination register (PIC_EIDRn)

The EIDRs control the destination of external interrupts caused by the assertion of any of IRQ[0:11]. Only one destination bit may be set; otherwise, behavior is undefined.

Address: 4_0000h base + 1_0010h offset + (32d × i), where i=0d to 11d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	EP	CI0	CI1	Reserved												
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved														P1	P0
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

PIC_EIDRn field descriptions

Field	Description
0 EP	External signal. Allows interrupt to be serviced externally. EP should be set only for level-sensitive external interrupts (EIVPRn[S]= 1). Setting for edge-sensitive does not provide reliable interrupt response. 0 Interrupt is not routed to IRQ_OUT_B . 1 Interrupt is routed to IRQ_OUT_B for external service.
1 CI0	Critical interrupt 0. Cin fields should be set only for level-sensitive external interrupts (EIVPRn[S]= 1). Setting them for edge-sensitive does not provide reliable interrupt response. 0 Processor core 0 does not receive this interrupt. 1 Directs the external interrupt to processor core 0 by causing the <i>cint0</i> output signal from the PIC to assert. See Interrupts to the e500 processor core .
2 CI1	Critical interrupt 1. Cin fields should be set only for level-sensitive external interrupts (EIVPRn[S]= 1). Setting them for edge-sensitive does not provide reliable interrupt response. Reserved in single-processor implementations. 0 Processor core 1 does not receive this interrupt. 1 Directs the external interrupt to processor core 1 by causing the <i>cint1</i> output signal from the PIC to assert. See Interrupts to the e500 processor core .
3–29 -	This field is reserved. Reserved
30 P1	Processor core 1. Indicates whether processor core 1 receives the interrupt through <i>int</i> . NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive this interrupt. 1 Directs the interrupt to processor core 1 through the assertion of <i>int1</i> .
31 P0	Processor core 0 . Indicates whether processor core 0 receives the interrupt. The default destination is for processor core 0 to receive this external interrupt after the PIC is reset.

Table continues on the next page...

PIC_EIDR_n field descriptions (continued)

Field	Description
0	Processor core 0 does not receive this interrupt.
1	Directs the interrupt to processor core 0 through the assertion of <i>int0</i> .

9.3.45 Internal interrupt n vector/priority register (PIC_IIVPR_n)

The IIVPRs have the same fields and format as the GTVPRs, except that they apply to the internal interrupt sources listed in [Table 9-3](#). These interrupts are all level-sensitive. See [Flow of interrupt control](#) for information on IPR and ISR.

NOTE

Because all internal interrupts are active-high, clearing the polarity field, IIVPR *n* [P], disables that interrupt. Care should be taken to ensure this field is set during initialization and that it is not inadvertently corrupted when loading or reloading IIVPRs with priority, mask, or vector data.

Address: 4_0000h base + 1_0200h offset + (32d × i), where i=0d to 63d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	MSK	A	Reserved							P	Reserved			PRIORITY		
Reset	1	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	VECTOR															
W	VECTOR															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PIC_IIVPR_n field descriptions

Field	Description
0 MSK	Mask. Mask interrupts from this source. MSK affects only interrupts routed to int. 0 An interrupt request is generated if the corresponding IPR bit is set. 1 Further interrupts from this source are disabled.
1 A	Activity. Indicates an interrupt has been requested or is in service. The VECTOR and PRIORITY values should not be changed while this bit is set. Affects only interrupts routed to int. 0 No current interrupt activity associated with this source. 1 The interrupt field for this source is set in the IPR or ISR.
2–7 -	This field is reserved. Reserved

Table continues on the next page...

PIC_IIVPR_n field descriptions (continued)

Field	Description
8 P	Polarity. Specifies the polarity for the internal interrupt. Note: Because all internal interrupts are active-high, clearing this bit disables the interrupt. 0 Interrupt polarity is active-low. This value disables the interrupt. 1 Interrupt polarity is active-high.
9–11 -	This field is reserved. Reserved
12–15 PRIORITY	Priority. Specifies the interrupt priority. The lowest priority is 0 and the highest priority is 15. A priority level of 0 inhibits signaling of this interrupt to the core. Affects only interrupts routed to int.
16–31 VECTOR	Vector (Affects only interrupts routed to int). Contains the value returned when IACK is read and this interrupt resides in the corresponding interrupt request register (IRR) for that core, as shown in Figure 9-331 .

9.3.46 Internal interrupt n destination register (PIC_IIDR_n)

The IIDRs have the same fields and format as EIVDRs, except that they apply to the internal interrupt sources listed in [Table 9-3](#). Only one destination bit may be set; otherwise, behavior is undefined.

Address: 4_0000h base + 1_0210h offset + (32d × i), where i=0d to 63d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	EP	CI0	CI1	Reserved												
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved														P1	P0
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

PIC_IIDR_n field descriptions

Field	Description
0 EP	External signal. Allows internal interrupt to be serviced externally. 0 Interrupt is not routed to IRQ_OUT_B. 1 Interrupt is routed to IRQ_OUT_B for external service.
1 CI0	Critical interrupt 0. See Interrupts to the e500 processor core for more information. 0 Processor core 0 does not receive this interrupt. 1 Directs the internal interrupt to processor core 0 by causing the <i>cint0</i> output signal from the PIC to assert.
2 CI1	Critical interrupt 1. See Interrupts to the e500 processor core , for more information. Reserved in single-processor implementations.

Table continues on the next page...

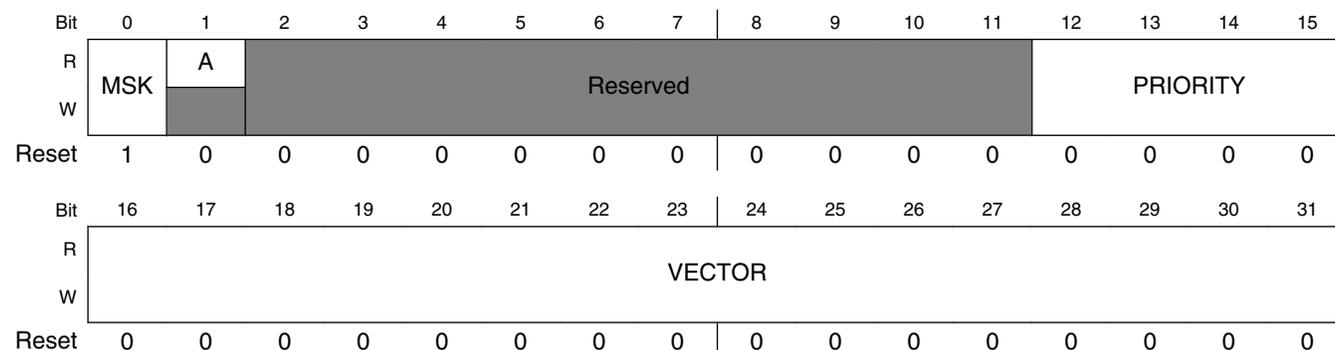
PIC_IIDRn field descriptions (continued)

Field	Description
	0 Processor core 1 does not receive this interrupt. 1 Directs the internal interrupt to processor core 1 by causing the <i>cint1</i> output signal from the PIC to assert.
3–29 -	This field is reserved. Reserved
30 P1	Processor core 1. Indicates whether processor core 1 receives the interrupt through <i>int</i> . NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive this interrupt. 1 Directs the interrupt to processor core 1 through the assertion of <i>int1</i> .
31 P0	Processor core 0 . Indicates whether processor core 0 receives the interrupt. The default destination is for processor core 0 to receive this external interrupt after the PIC is reset. 0 Processor core 0 does not receive this interrupt. 1 Directs the interrupt to processor core 0 through the assertion of <i>int0</i> .

9.3.47 Messaging interrupt n (MSGn) vector/priority register (PIC_MIVPRn)

The MIVPRs have the same fields and format as the GTVPRs, except they apply to messaging interrupts. See [Flow of interrupt control](#) for information on IPR and ISR.

Address: 4_0000h base + 1_1600h offset + (32d × i), where i=0d to 7d



PIC_MIVPRn field descriptions

Field	Description
0 MSK	Mask. Mask interrupts from this source. MSK affects only interrupts routed to int. 0 An interrupt request is generated if the corresponding IPR bit is set. 1 Further interrupts from this source are disabled.

Table continues on the next page...

PIC_MIVPR_n field descriptions (continued)

Field	Description
1 A	Activity. Indicates an interrupt has been requested or is in service. The VECTOR and PRIORITY values should not be changed while this bit is set. Affects only interrupts routed to int. 0 No current interrupt activity associated with this source. 1 The interrupt field for this source is set in the IPR or ISR.
2–11 -	This field is reserved. Reserved
12–15 PRIORITY	Priority. Specifies the interrupt priority. The lowest priority is 0 and the highest priority is 15. A priority level of 0 inhibits signaling of this interrupt to the core. Affects only interrupts routed to int.
16–31 VECTOR	Vector (Affects only interrupts routed to int). Contains value returned when IACK is read and this interrupt resides in the corresponding interrupt request register (IRR) for that core, as shown in Figure 9-331 .

9.3.48 Messaging interrupt n (MSG_n) destination register (PIC_MIDR_n)

The messaging interrupt destination registers (MIDRs) control the destination for the messaging interrupts. Only one destination bit may be set; otherwise, behavior is undefined.

Address: 4_0000h base + 1_1610h offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved														P1	P0
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

PIC_MIDR_n field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1. Indicates whether processor core 1 receives the interrupt through <i>int</i> . NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive this interrupt. 1 Directs the interrupt to processor core 1 through the assertion of <i>int1</i> .
31 P0	Processor core 0. Indicates whether processor core 0 receives the interrupt. The default destination is for processor core 0 to receive this external interrupt after the PIC is reset.

Table continues on the next page...

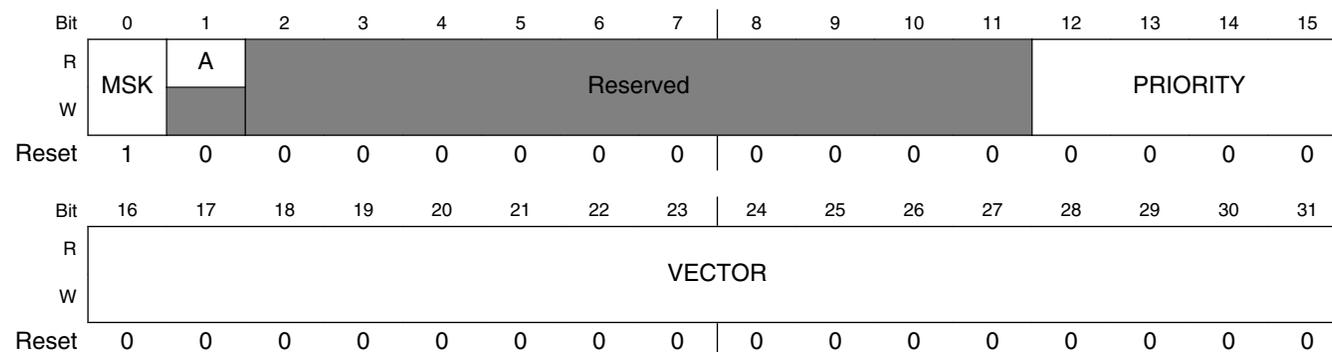
PIC_MIDR_n field descriptions (continued)

Field	Description
0	Processor core 0 does not receive this interrupt.
1	Directs the interrupt to processor core 0 through the assertion of <i>int0</i> .

9.3.49 Shared message signaled interrupt vector/priority register n (PIC_MSIVPR_n)

The MSIVPRs have the same fields and format as the GTVPRs. See [Flow of interrupt control](#) for information on IPR and ISR.

Address: 4_0000h base + 1_1C00h offset + (32d × i), where i=0d to 7d



PIC_MSIVPR_n field descriptions

Field	Description
0 MSK	Mask. Mask interrupts from this source. MSK affects only interrupts routed to int. 0 An interrupt request is generated if the corresponding IPR bit is set. 1 Further interrupts from this source are disabled.
1 A	Activity. Indicates an interrupt has been requested or is in service. The VECTOR and PRIORITY values should not be changed while this bit is set. Affects only interrupts routed to int. 0 No current interrupt activity associated with this source. 1 The interrupt field for this source is set in the IPR or ISR.
2–11 -	This field is reserved. Reserved
12–15 PRIORITY	Priority. Specifies the interrupt priority. The lowest priority is 0 and the highest priority is 15. A priority level of 0 inhibits signaling of this interrupt to the core. Affects only interrupts routed to int.
16–31 VECTOR	Vector (Affects only interrupts routed to int). Contains the value returned when IACK is read and this interrupt resides in the corresponding interrupt request register (IRR) for that core, as shown in Figure 9-331 .

9.3.50 Shared message signaled interrupt destination register n (PIC_MSIDRn)

The MSIDRs contain the destination bits for the shared message signaled interrupt. A shared message signaled interrupt can be directed to one of the processors by setting the appropriate bit in the shared message signaled interrupt destination register. Only one of the bits corresponding to destination processors may be set. The behavior if more than one bit is set is not defined. This register also contains the external pin field that can be programmed to direct the interrupt to the external pin .

Address: 4_0000h base + 1_1C10h offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved														P1	P0
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	*	1

* Notes:

- P1 field: Reserved in single-processor implementations.

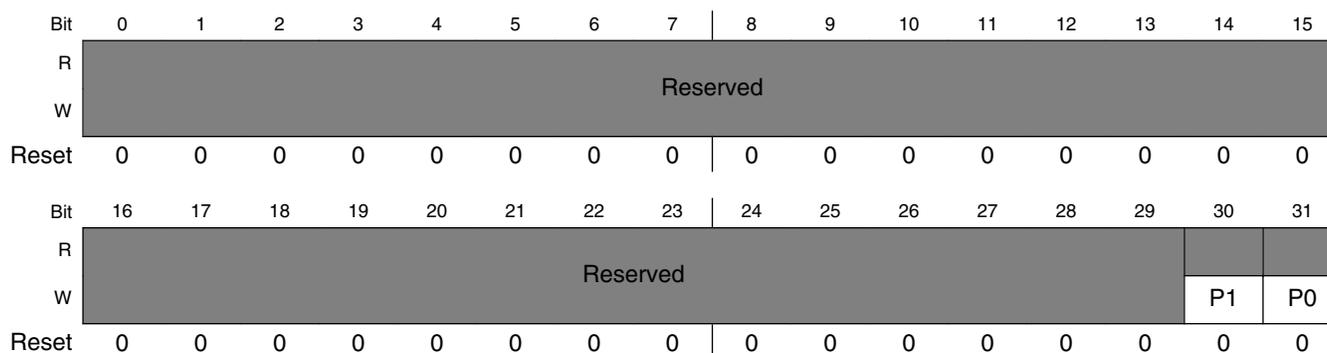
PIC_MSIDRn field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1. Indicates whether processor core 1 receives the interrupt through <i>int1</i> . NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive this interrupt. 1 Directs the interrupt to processor core 1 through the assertion of <i>int1</i> .
31 P0	Processor core 0 . Indicates whether processor core 0 receives the interrupt. The default destination is for processor core 0 to receive this shared message signaled interrupt after the PIC is reset. 0 Processor core 0 does not receive this interrupt. 1 Directs the interrupt to processor core 0 through the assertion of <i>int0</i> .

9.3.51 Processor core 0 interprocessor n dispatch register (PIC_IPIDR_CPU0n)

The figure below shows the four IPIDRs, one for each interprocessor interrupt channel. Writing to an IPIDR with a bit set causes a self interrupt for a single-core device. Because external bus masters can write to these registers, this feature can serve as a doorbell type interrupt.

Address: 4_0000h base + 2_0040h offset + (16d × i), where i=0d to 3d



PIC_IPIDR_CPU0n field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1. Specifies if processor core 1 receives the interrupt. This interrupt is multicasting, so both P0 and P1 can be set. NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive the interrupt 1 Directs the interrupt to processor core 1
31 P0	Processor core 0 . Determines if processor core 0 receives the interrupt. 0 Processor core 0 does not receive the interrupt. 1 Directs the interrupt to processor core 0.

9.3.52 Processor core current task priority register 0 Processor core (PIC_CTPR_CPU0)

There is one CTPR per processor core on this device.

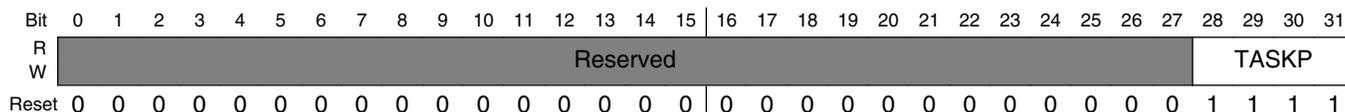
NOTE

CTPR has meaning only for interrupts routed to *int* .

Software must write the priority of the current processor core task in the CTPR for each core. The PIC uses this value for comparison with the priority of incoming interrupts. Given several concurrent incoming interrupts, the highest priority interrupt is asserted to that core if the following apply:

- The interrupt is not masked.
- The priority of the interrupt is higher than the values in the corresponding CTPR[TASKP] and ISR.

Address: 4_0000h base + 2_0080h offset = 6_0080h



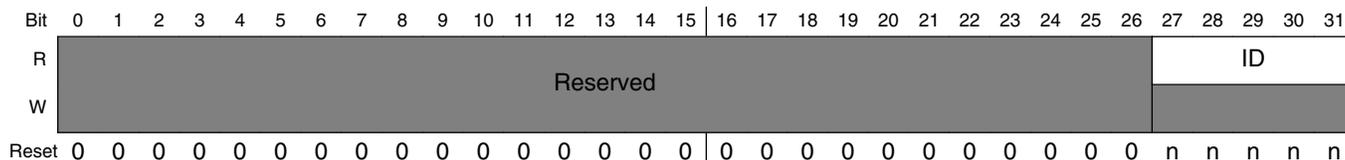
PIC_CTPR_CPU0 field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 TASKP	Task priority. Indicates the threshold that individual interrupt priorities must exceed for the interrupt request to be serviced. 0000-1111 xVPRn[PRIORITY] must exceed this value for the interrupt request to be serviced. Note the following special cases: 0000 Lowest priority. All interrupts except those whose priority are 0 can be serviced. 1111 Highest priority. No interrupts are signaled to that processor core. Hardware selects this value on a device hard reset or when the corresponding PIR[Pn] is set.

9.3.53 Processor core 0 who am I register (PIC_WHOAMI_CPU0)

The processor core WHOAMI *n* register can be read by a processor core to determine its physical connection to the PIC. The value returned when reading this register may be used to determine the value for the destination masks used for dispatching interrupts.

Address: 4_0000h base + 2_0090h offset = 6_0090h



PIC_WHOAMI_CPU0 field descriptions

Field	Description
0–26 -	This field is reserved. Reserved
27–31 ID	Returns the ID of the processor core reading this register. Does not always return zero. 0_0000 Processor core 0 0_0001 Processor core 1. (Value not supported in single-processor implementations.) 1_1111 All other devices

9.3.54 Processor core 0 interrupt acknowledge register (PIC_IACK_CPU0)

NOTE

IACK has meaning only for interrupts routed to *int* and should not be accessed for interrupts routed to *cint* or *IRQ_OUT_B* .

In systems based on processors built on Power Architecture technology, the interrupt acknowledge function occurs as an explicit read operation to a memory-mapped interrupt acknowledge register (IACK). Each processor core has an IACK register assigned to it. Reading IACK returns the interrupt vector corresponding to the highest priority pending interrupt. Reading IACK also has the following side effects:

- The associated field in the corresponding interrupt pending register (IPR) is cleared for edge-sensitive interrupts. See [Interrupts routed to int](#) .
- The corresponding in-service register (ISR) is updated.
- The corresponding *int* output signal from the PIC is negated.

Reading IACK when no interrupt is pending returns the spurious vector value, as described in [Spurious vector register \(PIC_SVR\)](#) .

Address: 4_0000h base + 2_00A0h offset = 6_00A0h



PIC_IACK_CPU0 field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 VECTOR	Interrupt vector. Vector of the highest pending interrupt

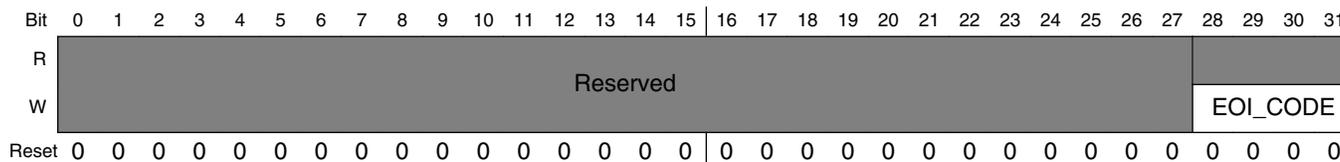
9.3.55 Processor core 0 end of interrupt register (PIC_EOI_CPU0)

NOTE

EOI has meaning only for interrupts routed to *int* and should not be accessed for interrupts routed to *cint* or *IRQ_OUT_B* .

Each core is assigned an EOI register. Writing to EOI signals the end of processing for the highest-priority interrupt (routed to *int*) currently in service. It also updates the corresponding ISR *n* by retiring the highest priority interrupt. Data values written to EOI are ignored, and zero is assumed.

Address: 4_0000h base + 2_00B0h offset = 6_00B0h



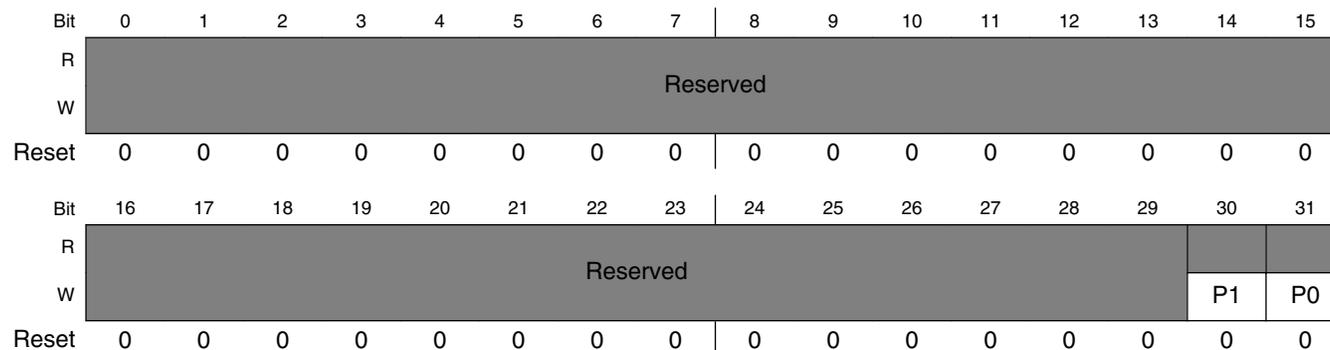
PIC_EOI_CPU0 field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EOI_CODE	0000 (write only)

9.3.56 Processor core 1 interprocessor n dispatch register (PIC_IPIDR_CPU1n)

The figure below shows the four IPIDRs, one for each interprocessor interrupt channel. Writing to an IPIDR with a bit set causes a self interrupt for a single-core device. Because external bus masters can write to these registers, this feature can serve as a doorbell type interrupt.

Address: 4_0000h base + 2_1040h offset + (16d × i), where i=0d to 3d



PIC_IPIDR_CPU1n field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 P1	Processor core 1. Specifies if processor core 1 receives the interrupt. This interrupt is multicasting, so both P0 and P1 can be set. NOTE: Reserved in single-processor implementations. 0 Processor core 1 does not receive the interrupt 1 Directs the interrupt to processor core 1
31 P0	Processor core 0 . Determines if processor core 0 receives the interrupt. 0 Processor core 0 does not receive the interrupt. 1 Directs the interrupt to processor core 0.

9.3.57 Processor core 1 current task priority register (PIC_CTPR_CPU1)

There is one CTPR per processor core on this device.

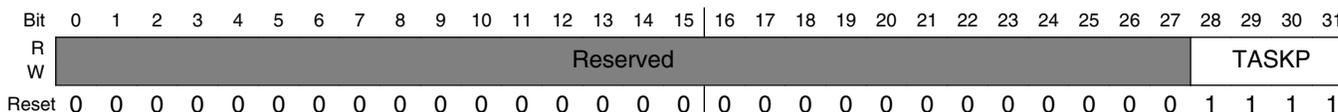
NOTE

CTPR has meaning only for interrupts routed to *int* .

Software must write the priority of the current processor core task in the CTPR for each core. The PIC uses this value for comparison with the priority of incoming interrupts. Given several concurrent incoming interrupts, the highest priority interrupt is asserted to that core if the following apply:

- The interrupt is not masked.
- The priority of the interrupt is higher than the values in the corresponding CTPR[TASKP] and ISR.

Address: 4_0000h base + 2_1080h offset = 6_1080h



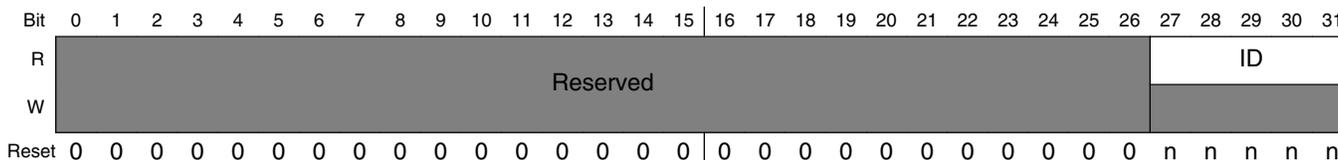
PIC_CTPR_CPU1 field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 TASKP	Task priority. Indicates the threshold that individual interrupt priorities must exceed for the interrupt request to be serviced. 0000-1111 xVPRn[PRIORITY] must exceed this value for the interrupt request to be serviced. Note the following special cases: 0000 Lowest priority. All interrupts except those whose priority are 0 can be serviced. 1111 Highest priority. No interrupts are signaled to that processor core. Hardware selects this value on a device hard reset or when the corresponding PIR[Pn] is set.

9.3.58 Processor core 1 who am I register (PIC_WHOAMI_CPU1)

The processor core WHOAMI *n* register can be read by a processor core to determine its physical connection to the PIC. The value returned when reading this register may be used to determine the value for the destination masks used for dispatching interrupts.

Address: 4_0000h base + 2_1090h offset = 6_1090h



PIC_WHOAMI_CPU1 field descriptions

Field	Description
0–26 -	This field is reserved. Reserved
27–31 ID	Returns the ID of the processor core reading this register. Does not always return zero. 0_0000 Processor core 0 0_0001 Processor core 1. (Value not supported in single-processor implementations.) 1_1111 All other devices

9.3.59 Processor core 1 interrupt acknowledge register (PIC_IACK_CPU1)

NOTE

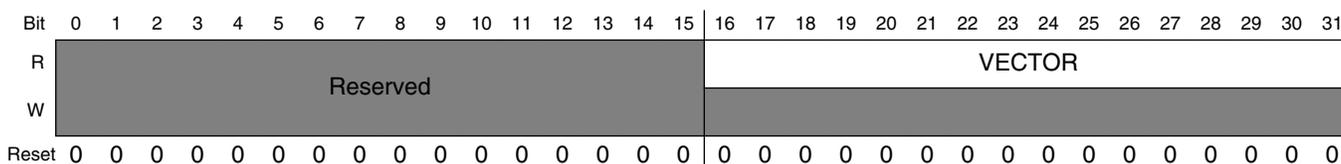
IACK has meaning only for interrupts routed to *int* and should not be accessed for interrupts routed to *cint* or *IRQ_OUT_B* .

In systems based on processors built on Power Architecture technology, the interrupt acknowledge function occurs as an explicit read operation to a memory-mapped interrupt acknowledge register (IACK). Each processor core has an IACK register assigned to it. Reading IACK returns the interrupt vector corresponding to the highest priority pending interrupt. Reading IACK also has the following side effects:

- The associated field in the corresponding interrupt pending register (IPR) is cleared for edge-sensitive interrupts. See [Interrupts routed to int](#) .
- The corresponding in-service register (ISR) is updated.
- The corresponding *int* output signal from the PIC is negated.

Reading IACK when no interrupt is pending returns the spurious vector value, as described in [Spurious vector register \(PIC_SVR\)](#) .

Address: 4_0000h base + 2_10A0h offset = 6_10A0h



PIC_IACK_CPU1 field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 VECTOR	Interrupt vector. Vector of the highest pending interrupt

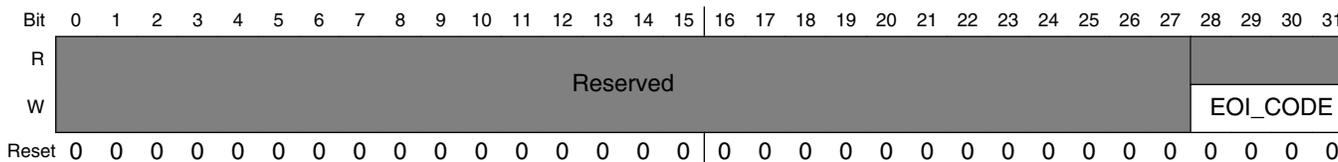
9.3.60 Processor core 1 end of interrupt register (PIC_EOI_CPU1)

NOTE

EOI has meaning only for interrupts routed to *int* and should not be accessed for interrupts routed to *cint* or *IRQ_OUT_B* .

Each core is assigned an EOI register. Writing to EOI signals the end of processing for the highest-priority interrupt (routed to *int*) currently in service. It also updates the corresponding ISR *n* by retiring the highest priority interrupt. Data values written to EOI are ignored, and zero is assumed.

Address: 4_0000h base + 2_10B0h offset = 6_10B0h



PIC_EOI_CPU1 field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EOI_CODE	0000 (write only)

9.4 Functional description

This section is a functional description of the PIC.

9.4.1 Programming model considerations

9.4.1.1 Global registers

Although most PIC registers have one address, some are replicated for each e500 processor core in a multiprocessor device.

For such registers, each core accesses its separate registers using the same address, the address decoding being sensitive to the processor core ID. A copy of the per-CPU registers is available to each e500 processor core at the same physical address, that is, in a private access address space that acts like an alias to a processor's own copy of the per-CPU registers. As shown in [Figure 9-330](#), the ID of the core initiating the read/write transaction determines which processor's per-CPU registers to access. For more information, see [Per-CPU \(private access\) registers](#).

NOTE

Register fields designated as write-1-to-clear are cleared only by writing ones to them. Writing zeros to them has no effect.

9.4.1.2 Global timer registers

The two independent groups of global timer registers, group A and group B, are identical in their functionality, except that they appear at different locations within the PIC register map.

Note that each of the four timers within an x group have four individual configuration registers ($GTCCR_{xn}$, $GTBCR_{xn}$, $GTVPR_{xn}$, and $GTDR_{xn}$), but they are only shown once in this section. These two groups of timers cannot be cascaded together.

9.4.1.3 IRQ_OUT_B and critical interrupt summary registers

The summary registers indicate the specific interrupt sources routed to the `IRQ_OUT_B` or `cint0/cint1` PIC outputs.

Summary register bits are cleared when the corresponding interrupt that caused a bit to be set is negated.

Note that only level-sensitive interrupts can be routed to `IRQ_OUT_B` or `cint0` and `cint1`.

The `IRQ_OUT_B` summary registers, shown in [IRQ_OUT_B summary register 0 \(PIC_IRQSR0\)](#) through [IRQ_OUT_B summary register 2 \(PIC_IRQSR2\)](#) contain one bit for each interrupt source that can be routed to `IRQ_OUT_B`. The corresponding bit is set if the interrupt is active and is routed to `IRQ_OUT_B` (that is, if the corresponding $xIDRn[EP]$ is set).

The critical interrupt summary registers, shown in [Critical interrupt summary register 0 \(PIC_CISR0\)](#) through [Critical interrupt summary register 2 \(PIC_CISR2\)](#), contain one bit for each interrupt source that can be designated as a critical interrupt. The corresponding bit is set if the interrupt is active and is routed to either the `cint0` or `cint1` outputs of the PIC (if $xIDRn[CIn] = 1$ in its corresponding destination register).

9.4.1.4 Performance monitor mask registers (PMMRs)

The twelve performance monitor mask registers consist of four sets of three 32-bit registers, PM_nMR0 , PM_nMR1 , and PM_nMR2 .

Each set can be configured to select one interrupt source (interprocessor, timer, message, shared message signaled, external, or internal) to generate a performance monitor event. The performance monitor can be configured to track this event in the performance monitor local control registers. See [Performance monitor local control register \$A_n\$ \(PERFMON_PMLCA \$n\$ \)](#) and [Performance monitor local control register \$B_n\$ \(PERFMON_PMLCB \$n\$ \)](#).

9.4.1.5 Message registers

The following registers support the message register interrupts:

- [Message register \$n\$ \(PIC_MSGR \$n\$ \)](#)
- [Message enable register \(PIC_MER\)](#)
- [Message status register \(PIC_MSR\)](#)
- [Messaging interrupt \$n\$ \(MSG \$n\$ \) vector/priority register \(PIC_MIVPR \$n\$ \)](#)
- [Messaging interrupt \$n\$ \(MSG \$n\$ \) destination register \(PIC_MIDR \$n\$ \)](#)

Writing to one of the eight message registers (MSGR0-MSGR7) causes a messaging interrupt as directed by the other message registers listed above. Reading a message register clears the messaging interrupt. Note that a messaging interrupt can also be cleared by writing a one to the corresponding status field of the PIC message status register (MSR), shown in [Message status register \(PIC_MSR\)](#).

9.4.1.6 Shared message signaled registers

This section contains description the shared message signaled interrupt registers (MSIRs).

The shared message signaled interrupt structure allows programs to interrupt each other by simply writing to these shared memory-mapped registers in the PIC.

Each of the eight MSIRs can be thought of as collecting interrupts from 32 different memory-mapped writes that can cause interrupts.

Table 9-332. Vector/priority register summary (continued)

Internal (Internal interrupt n vector/ priority register (PIC_IIVPRn))	R	MSK	A	-	P	-	PRIORIT Y	VECTOR	
	W								
External (External interrupt n (IRQn) vector/priority register (PIC_EIVPRn))	R	MSK	A	-	P	S	-	PRIORIT Y	VECTOR
	W								

Note the following:

- The MSK, A, PRIORITY, and VECTOR fields have meaning only for interrupts routed to the *int* signal.
- The polarity field, P, is provided to indicate whether the signals from the corresponding source are active high or low.
- The sense field, S, is provided to allow external interrupt sources to be configured as level-sensitive so they can be routed to either *cint* or *IRQ_OUT_B*.

9.4.1.8 Per-CPU (private access) registers

The OpenPIC programming model supports multiprocessor systems of up to 32 separate processors.

As such, the OpenPIC interface specification provides for coordinating both the requesting and servicing of interrupts among several processor cores within a single integrated device. To comply with the OpenPIC specification, the PIC incorporates several of these multiprocessor capabilities.

NOTE

Note that these registers are meaningful only for interrupts routed to *int*.

The registers in the table below are called per-CPU registers because they are duplicated for each core in a multi-core device. The OpenPIC interface specifies that a copy of these registers be available to each core at the same physical address by using the ID of the processor core that initiates the transaction to determine the set of per-CPU registers to access.

Table 9-333. Per-CPU registers-private access address offsets

Register name	Offset
Interprocessor 0 dispatch register (IPIDR0)	0x0040
Interprocessor 1 dispatch register (IPIDR1)	0x0050
Interprocessor 2 dispatch register (IPIDR2)	0x0060
Interprocessor 3 dispatch register (IPIDR3)	0x0070

Table continues on the next page...

Table 9-333. Per-CPU registers-private access address offsets (continued)

Register name	Offset
Current task priority register (CTPR)	0x0080
Who am I register (WHOAMI0)	0x0090
Interrupt acknowledge register (IACK)	0x00A0
End of interrupt register (EOI)	0x00B0

These addresses, shown in [Table 9-333](#), appear in the memory map at the same offset for every processor in what is called the private access space. This duplication allows user code to execute correctly in an multiprocessor environment without needing to know which core it is running on. On a single-core device, each register has two addresses, one in the normal address space and one in the private access space. It is included on even single-core devices to simplify the porting of such code.

[Figure 9-330](#) shows how the duplicated registers are addressed in a four-core device. Note that when accessing a register normally, each core sources a different address. However, when accessing the same register using the per-CPU address space, each core sources the same address.

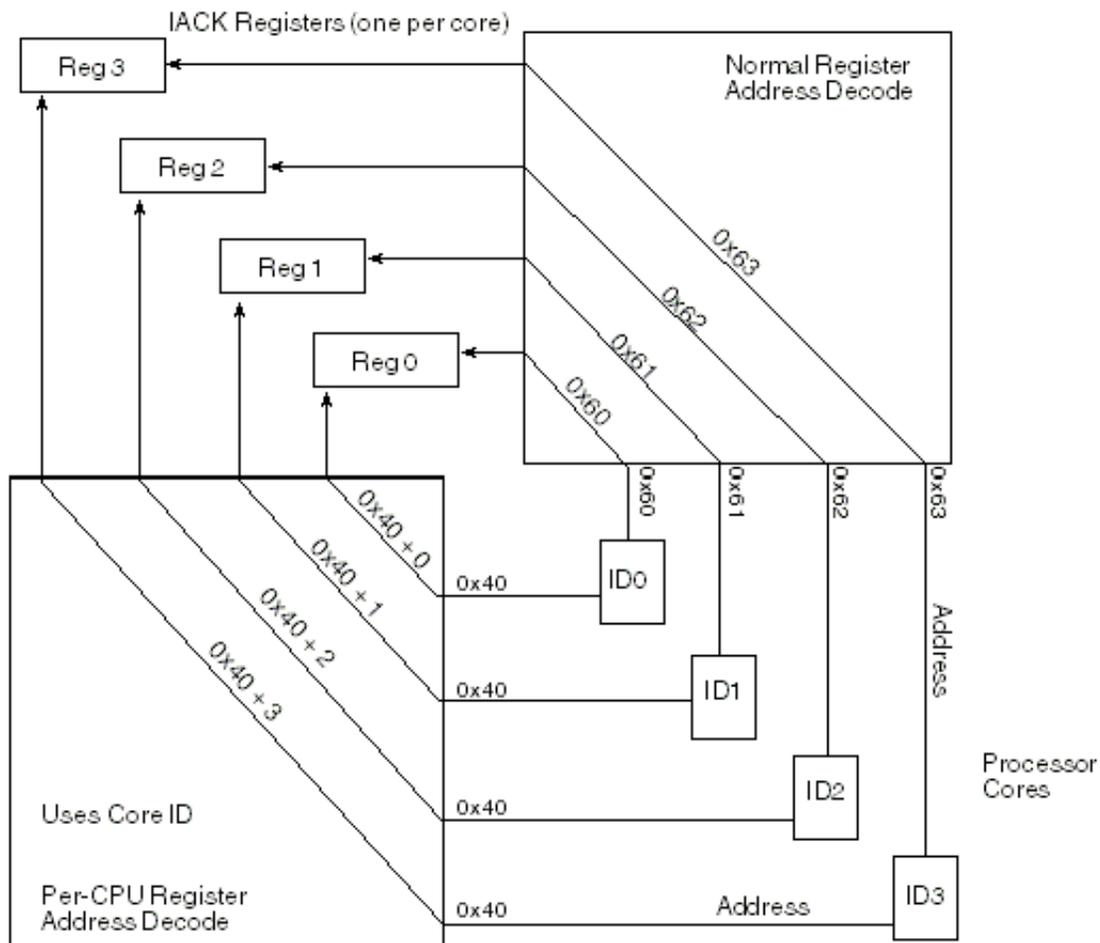


Figure 9-330. Per-CPU register address decoding in a four-core device

9.4.2 Flow of interrupt control

Figure 9-331 shows the flow of interrupts directed by the PIC to the *int*, *cint*, and IRQ_OUT_B outputs. Note that this diagram describes a conceptual model of an PIC on a single processor. This logic is replicated for each implemented processor. This conceptual diagram does not fully represent all internal circuitry of the implementation.

This figure focusses especially on the OpenPIC-defined logic and shows how the PIC controls interrupt requests that target the *int* signal. The flow in Figure 9-331 is from the bottom to the top, and shows at the bottom how the destination register associated with each source determines the path.

9.4.2.1 Interrupts routed to *cint* or IRQ_OUT_B

Interrupt requests routed to *cint* or IRQ_OUT_B bypass the logic that is dedicated to interrupt sources that compete for *int*.

That is, if $xIDRn[CI_n]$ or $xIDRn[EP] = 1$, corresponding $xIVPR$ field settings have no hardware effects; however, an interrupt handler may be able to make use of some of those fields.

cint signals are connected to the respective core's critical interrupt input.

NOTE

Because interrupt sources routed to *cint* or IRQ_OUT_B must be level sensitive, EIVPR[S] should be set. See [External interrupt n \(IRQn\) vector/priority register \(PIC_EIVPRn\)](#).

Because these interrupts bypass the OpenPIC logic, it is especially important that handlers do not read IACK. Doing so causes a spurious interrupt. Likewise, they should not write EOI.

9.4.2.2 Interrupts routed to *int*

As shown in the figure below, the PIC receives interrupt requests from external and internal sources and from within the PIC itself.

As the figure shows, all of these interrupt sources can be routed to *int*; the global timer and timer processor interrupts can be directed only to *int*.

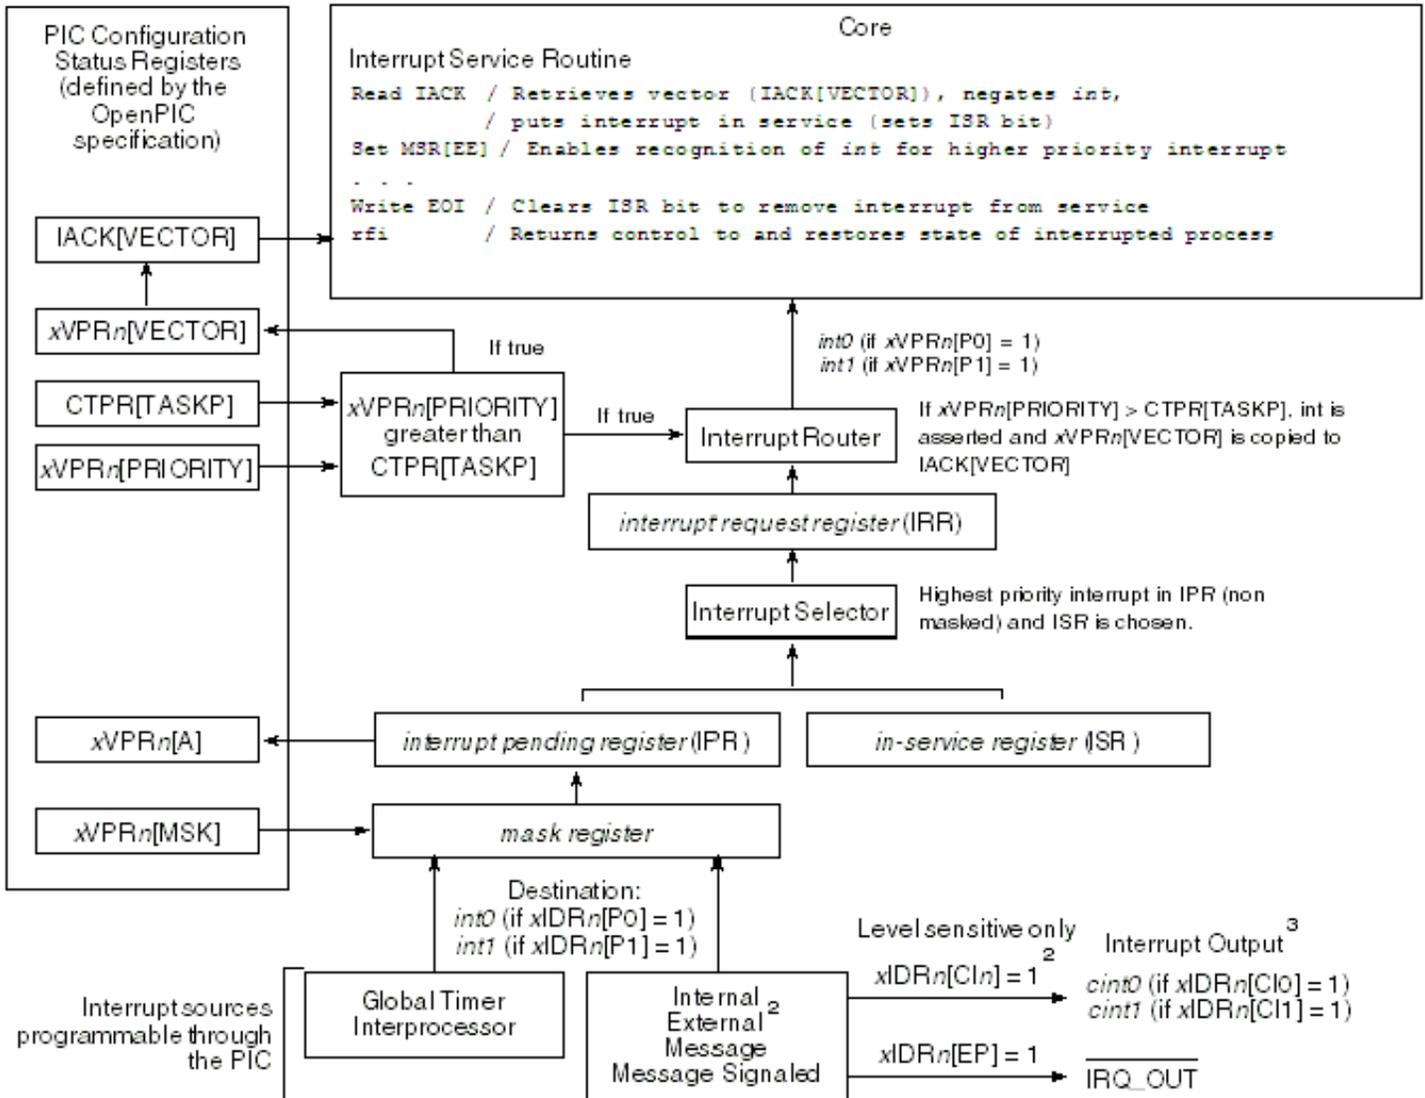
The sources' mask bits ($xVPRn[MSK]$) are tracked in the internal mask register. If a source's MSK bit is set, the mask register prevents the PIC from asserting *int* on its behalf.

Unmasked interrupt requests are qualified and latched in the interrupt pending register (IPR), an internal interrupt summary register with a bit for each source. If an interrupt request is multi-cast, a bit is set in the IPR for each targeted processor. Although the IPR cannot be read by software, when an IPR bit is set, the corresponding source's activity bit ($xVPRn[A]$) is automatically set.

The interrupt selector monitors the IPR and the in-service register (ISR), which tracks previously taken interrupts that were superseded by a higher-priority interrupt before the interrupt handler finished. The interrupt selector recognizes the highest priority unmasked interrupt request and latches it into the interrupt request register (IRR). The source's vector ($xVPRn[VECTOR]$) is copied to IACK[VECTOR], which the interrupt handler retrieves by reading IACK.

If the priority ($xVPR_n[PRIORITY]$) of an interrupt latched in the IRR is higher than the value in the target processor's $CTPR[TASKP]$, the interrupt router asserts the external interrupt signal (int), causing that processor core to vector to its external interrupt handler.

Note that like IPR, ISR and IRR cannot be read by software.



¹ If $cint$ or IRQ_OUT is the destination, $EIVPR_n[S]$ must be set to configure the source as level sensitive.
² If multiple destination register bits are set, PIC behavior is undefined.
³ Although setting Cl_n directs the interrupt request to the critical interrupt output ($cint0/cint1$), integrated logic may connect this signal to a different interrupt input to the core.

Figure 9-331. PIC interrupt processing flow diagram for each core (n)

The interrupt handler must acknowledge the interrupt by explicitly reading the corresponding IACK registers, described in [Processor core 0 interrupt acknowledge register](#) and [Processor core 0 interrupt acknowledge register](#). The PIC interprets this read

as an interrupt acknowledge (IACK) cycle. The IACK cycle not only returns the source's vector, it also negates the *int* signal to the processor (making it possible for a higher priority interrupt to assert *int*) and sets the source's bit in the ISR, indicating that this interrupt has been put in service. An interrupt remains in service from the time until the corresponding end-of-interrupt register (EOI) is written, generating what the PIC considers an EOI signal.

9.4.2.2.1 Interrupt source priority

Each interrupt source routed to *int* is assigned a value through its $xVPRn$ [PRIORITY] field.

Priority values range from 0 to 15, where 15 is the highest. Interrupts are delivered only when the priority of the source is greater than the destination processor's CTPR[TASKP]. Therefore, setting $xVPRn$ [PRIORITY] to zero inhibits that interrupt. Likewise, setting TASKP to 15 prevents the PIC from delivering interrupts to that core through the *int* signal. Note that this is the reset value, preventing the PIC from asserting *int* before the PIC is configured.

The PIC services simultaneous interrupts occurring with the same priority according to the following order:

1. MSG0-MSG 7
2. MSI0-MSI7
3. IPI0-IPI3
4. Group A timer0-timer3
5. Group B timer0-timer3
6. IRQ[0:11]/PCIINT x
7. Internal0-internal 63

For example, if MSG0, MSG2, and IPI0 are all assigned the same priority and receive simultaneous interrupts, they are serviced in the following order:

1. MSG0
2. MSG2
3. IPI0

9.4.2.2.2 Interrupt acknowledge

When the PIC causes *int* to be asserted, the external interrupt service routine acknowledges the request by reading that core's IACK register, which at this point holds the 16-bit vector value for the interrupt source that generated the request.

This is the value programmed in that source's $xVPR_n[VECTOR]$. Reading IACK has the following effects:

- The *int* signal for that core is negated, making it possible for another interrupt source to signal an external interrupt to the core, and more particularly, allowing the PIC to signal a higher-priority interrupt, as described in [Nesting of interrupts](#).
- The source that caused that resource is represented in the internal in-service register (ISR).

The interrupt is then considered to be in service. It remains so until the processor core performs a write to the corresponding EOI. Writing to EOI is referred to as an EOI cycle.

9.4.2.2.3 Spurious vector generation

Under certain circumstances, the PIC has no valid vector to return to a processor core during an interrupt acknowledge cycle.

In these cases, the spurious vector from the spurious vector register is returned. The following cases cause a spurious vector fetch:

- *int* is asserted in response to an externally or internally-sourced interrupt which is activated with level-sensitive logic, and the asserted level is negated before the interrupt is acknowledged.
- *int* is asserted for an interrupt source that is later masked (using the mask bit in the vector/priority register corresponding to that source) before the interrupt is acknowledged.
- *int* is asserted for an interrupt source that is later masked by an increase in the task priority level before the interrupt is acknowledged.
- An interrupt acknowledge cycle is performed by the processor core in spite of the fact that the *int* signal has not been asserted by the PIC.

In all cases, a spurious vector is returned only if no pending interrupt has sufficient priority to signal an interrupt, otherwise, the vector for that interrupt source is returned.

NOTE

EOI should not be written in response to a spurious vector. Otherwise, a previously accepted interrupt might be cleared unintentionally.

9.4.2.2.4 Nesting of interrupts

While an interrupt is being handled, if an interrupt request arrives with a higher $xVPR_n[PRIORITY]$ value, the interrupt being serviced is superseded.

As described in [Interrupts routed to *int*](#), the PIC asserts *int*, and the newer, higher priority interrupt is handled. This happens even if software, as part of its interrupt service routine, updates the corresponding CTPR with a lower value.

Thus, although several interrupts can be in service simultaneously (and tracked by the ISR), the highest priority interrupt by that processor is always the one actively handled. When the interrupt routine completes, it performs a write EOI cycle, a side effect of which is to take the current highest priority interrupt out of service (removes it from the ISR). At this point, the interrupt selector chooses the new highest priority interrupt request, and, assuming CTPR[TASKP] has not been updated to a value higher than the new interrupt, the PIC asserts *int* on its behalf.

The next write EOI cycle takes the current highest priority interrupt out of service. An interrupt with lower priority than those in service is not started until all higher priority interrupts complete even if its priority is greater than the CTPR value.

9.4.3 Interprocessor interrupts

Processors 0 and 1 can generate interprocessor interrupts that target either or both processors.

A self interrupt occurs when a core dispatches an interprocessor interrupt event to itself.

Interrupts are initiated by writing either or both of the PO_n bits in an interprocessor interrupt dispatch register (IPIDR0-IPIDR3) of one of the four IPI channels. If subsequent interprocessor interrupts from a given channel to a given target processor are initiated before the first is acknowledged, only one interrupt is generated.

9.4.4 Message interrupts

The eight MSGRs described in [Message register n \(PIC_MSGR \$n\$ \)](#), can be used to send 32-bit messages to one or more processors.

A messaging interrupt is generated by writing an MSGR if the corresponding MER bit is set and the interrupt is not masked. Reading a MSGR or writing a 1 to the status bit clears the interrupt.

9.4.5 Shared message signaled interrupts

There are eight shared MSIRs, [Shared message signaled interrupt register n \(PIC_MSIR \$n\$ \)](#), that indicate which of the interrupt sources sharing the MSI register have pending interrupts.

Up to 32 sources can share any individual MSI register. A shared message signaled interrupt is generated by writing to Shared Message Signaled Interrupt Index Register (MSIIR) fields SRS and IBS. This register is primarily intended to support inbound PCI Express message signaled interrupts (MSIs) when the PCI Express controller is configured as a root complex (RC).

MSIIR[SRS] selects the associated MSIR and MSIIR[IBS] selects the interrupt flag/bit in that register that is to be set. The corresponding interrupt needs to be unmasked for the interrupt to occur. A read to an MSIR clears the all of its flags.

9.4.6 PCI Express INTx/IRQ n sharing

Whenever the PCI Express controller is in root complex mode and it receives an inbound INTx asserted or negated message transaction, it asserts or negates an equivalent internal INTx signal to the PIC. This INTx virtual-wire interrupt signaling mechanism replaces the PCI standard sideband interrupts (INTA, INTB, INTC, and INTD) that historically were connected to the IRQ n external interrupt inputs.

The internal INTx signals from the PCI Express controller are logically combined with the interrupt request (IRQ n) signals so that they share the same OpenPIC external interrupt controlled by the associated EIVPR n and EIDR n registers.

Table 9-334. PCI Express INTx/IRQ n sharing

PCI Express number	INTx	IRQ n
PCI Express 1	INTA	IRQ0
	INTB	IRQ1
	INTC	IRQ2
	INTD	IRQ3
PCI Express 2	INTA	IRQ4
	INTB	IRQ5
	INTC	IRQ6
	INTD	IRQ7
PCI Express 3	INTA	IRQ8
	INTB	IRQ9
	INTC	IRQ10
	INTD	IRQ11

Note that IRQ[7:11] are also muxed with GPIO[0:4]. When using one of these shared GPIO/IRQ pins as GPIO, the associated external IRQ pin must be masked in the PIC (by default, external IRQ pins are masked in PIC). When using one of these as an external IRQ pin, it cannot be used as a GPIO pin. The default GPIO input direction should be maintained and ignore the state of GPIO data register (GPDAT).

In general, these signals should be considered mutually exclusive. If a PCI Express INTx signal is being used, the PIC must be configured so that external interrupts are level sensitive ($EIVPRn[S] = 1$). If an IRQ n signal is being used as edge-triggered ($EIVPRn[S] = 0$), the system must not allow inbound PCI Express INTx transactions.

Note that it is possible to share IRQ n and INTx if the external interrupt is level sensitive; however, if an interrupt occurs, the interrupt service routine must poll both the external sources connected to the IRQ n input and the PCI Express INTx sources to determine from which path the external interrupt came. In any case, IRQ n should be pulled to the negated state as determined by the associated polarity setting in $EIVPRn[P]$.

9.4.7 Global timers

There are appropriate clock prescalers and synchronizers to provide a time base for the internal PIC timers.

These 8 timers are organized as 2 groups of 4 timers each. The timers can be individually programmed to generate a processor core interrupt when they count down to zero and can be used to generate regular periodic interrupts. Each timer has the following four configuration and control registers:

- Global timer current count register (GTCCR xn)
- Global timer base count register (GTBCR xn)
- Global timer vector-priority register (GTVPR xn)
- Global timer destination register (GTDR xn)

The timer frequency should be written to the TFRR xn , described in [Timer frequency reporting register group X \(PIC_TFRR \$n\$ \)](#).

Timer interrupts are all edge-triggered interrupts. If a timer period expires while a previous interrupt from the same source is pending or in service, the subsequent interrupt is lost.

The timer control register (TCR) provides users with the ability to create timers larger than the 31-bit global timers. The timer frequency can also be changed by setting the appropriate TCR fields, as described in [Timer control register group n \(PIC_TCR \$n\$ \)](#).

9.4.8 Resets

This section describes the behavior of the PIC at reset and the PIC's ability to initiate processor resets.

9.4.9 Resetting the PIC

The PIC is reset by a device power-on reset (POR) or by software that sets GCR[RST], either of which causes the following:

- All pending and in-service interrupts are cleared.
- All interrupt mask bits are set.
- Polarity, sense, external signal, critical interrupt, and activity fields are reset to default values.
- PIR, TFRR, TCR, MER, MSR, and MSGR0-MSGR 7 are cleared.
- MSG and timer destination fields are set.
- The interprocessor dispatch registers are cleared.
- All timer base count values are reset to zero with count inhibited.
- CTPR[TASKP] is reset to 0x000F, disabling delivery of interrupts that target *int*.
- The spurious interrupt vector resets to 0xFFFF.
- The PM_nMRs are reset to 0xFFFF.
- The PIC defaults to the pass-through mode (GCR[M] = 0).
- All other registers remain at their pre-reset programmed values.

GCR[RST] is automatically cleared when the reset sequence is complete.

9.4.9.1 Processor core initialization

A software reset can be routed to either of the cores by writing to the processor core initialization register (PIR).

This causes the assertion of the corresponding *core_hreset* output signal from the PIC. When this occurs, the corresponding CTPR also gets written to 0x000F to prevent delivery of any interrupts to *int*.

9.5 Initialization/application information

This section contains initialization and application information for the PIC.

9.5.1 Programming guidelines

The following subsections contain information about programming PIC registers.

9.5.1.1 PIC registers

Most PIC control and status registers are readable and return the last value written.

The exceptions to this rule are as follows:

- Interprocessor dispatch and EOI registers, which return zeros on reads.
- Activity bits (A) of the vector/priority registers reflect the status of the corresponding interrupt source.
- IACK, which returns the vector of the highest priority pending interrupt or the spurious vector (SVR[VECTOR]) if none is pending.
- Reserved fields always return 0.

When the PIC is in mixed mode (GCR[M] = 1), the following guidelines are recommended:

- All PIC registers must be located in a cache-inhibited, guarded area (configured through the core's MMU).
- The PIC portion of the address map must be set up appropriately.

In addition, the following initialization sequence is recommended:

1. Write the vector, priority, and polarity values in each interrupt's vector/priority register, leaving their MSK (mask) bit set. This is required only if interrupts are used.
2. Clear CTPR (CTPR = all zeros).
3. Program the PIC to mixed mode by setting GCR[M].
4. Clear the MSK bit in the vector/priority registers to be used.
5. Perform a software loop to clear all pending interrupts:
 - Load counter with FRR[NIRQ].
 - While counter > 0, read IACK and write EOI to guarantee all the IPR and ISR bits are cleared.
6. Set the processor core CTPR values to the desired values.
7. Read the MSIRs to clear any pending message signaled interrupts that may have been pending before a soft reset.
8. Set MER to 0x0000_000F. See [Message enable register \(PIC_MER\)](#), for more information.

Depending on the interrupt system configuration, the PIC may generate spurious interrupts to clear interrupts latched during power-up. A spurious or non-spurious vector is returned for an interrupt acknowledge cycle in this case. See the programming note below for the non-spurious case.

NOTE

Because the default polarity/sense for external interrupts is edge-sensitive, and edge-sensitive interrupts are not cleared until they are acknowledged, it is possible for the PIC to store spurious edges detected during power-up as pending external interrupts. If software permanently configures an external interrupt source to be edge-sensitive, it may receive the vector for the interrupt source and not a spurious interrupt vector when software clears the mask bit. This can occur once for any edge-sensitive interrupt when its mask bit is first cleared and the PIC is in mixed mode.

To avoid a false interrupt for this case, software can clear the IPR of these spurious edge detections by first configuring the polarity/sense of external interrupt sources to be level-sensitive: high-level if the input is a positive-edge source and low-level if it is a negative-edge source (while the mask bit remains set). After this is complete, configuring the external interrupt source as edge-sensitive does not cause a false interrupt.

9.5.1.2 Changing interrupt source configuration

To change the vector, priority, polarity, sense, or destination of an active (unmasked) interrupt source, the following steps should be taken:

1. Mask the source using the mask (MSK) bit in the vector/priority register.
2. Wait for the activity (A) bit for that source to be cleared.
3. Make the desired changes.
4. Unmask the source.

Note that changing the destination from *int* to *cint* or *IRQ_OUT_B* makes the A, MSK, and PRIORITY fields meaningless.



Chapter 10

I2C Interfaces

This chapter describes the dual inter-integrated circuit (I2C) bus modules implemented on this device.

10.1 Overview

This chapter describes the dual inter-integrated circuit (I²C) bus modules implemented on this device and covers the following topics:

- [Introduction to I²C](#)
- [I²C external signal descriptions](#)
- [I2C memory map/register definition](#)
- [Functional description](#)
- [Initialization/application information](#)

10.2 Introduction to I²C

This section presents the following topics:

- [What is the I²C module?](#)
- [I²C module block diagram](#)
- [Features](#)
- [Advantages of the I²C bus](#)
- [Modes of operation](#)
- [I²C-specific conditions](#)

10.2.1 What is the I²C module?

The I²C module is a two-wire-serial data (SDA) and serial clock (SCL)-bidirectional serial bus that provides a simple efficient method of data exchange between this device and other devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs.

10.2.2 I²C module block diagram

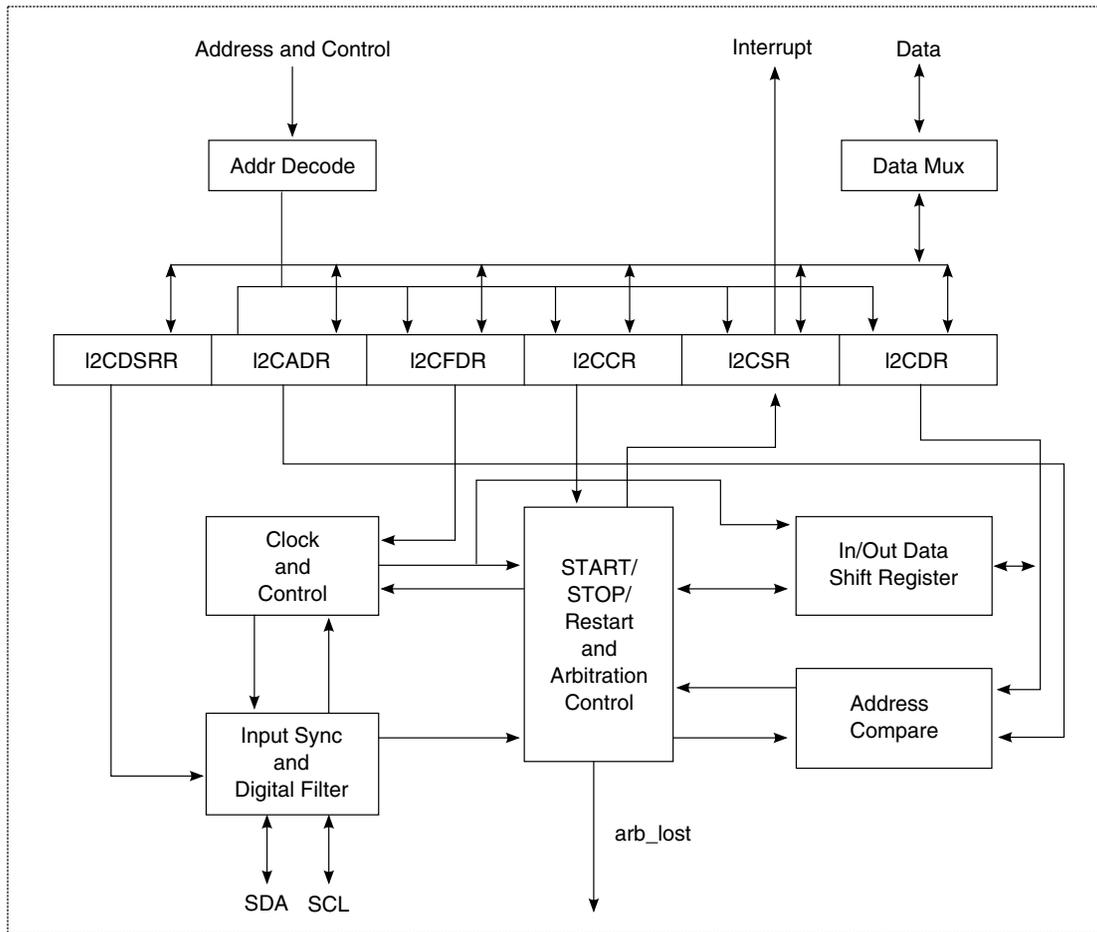


Figure 10-1. I²C block diagram

10.2.3 Features

The I²C interface includes the following features:

- Two-wire interface
- Multiple-master operation

- Arbitration lost interrupt with automatic mode switching from master to slave
- Calling address identification interrupt
- START and STOP signal generation/detection
- Acknowledge bit generation/detection
- Bus busy detection
- Software-programmable clock frequency
- Software-selectable acknowledge bit
- On-chip filtering for spikes on the bus

10.2.4 Advantages of the I²C bus

The two-wire I²C bus minimizes interconnections between devices.

The synchronous, multiple-master I²C bus allows the connection of additional devices to the bus for expansion and system development. The bus includes collision detection and arbitration that prevent data corruption if two or more masters attempt to control the bus simultaneously.

10.2.5 Modes of operation

The I²C units on this device can operate in one of the following modes:

- Master mode-The I²C is the driver of the SDA line. It cannot use its own slave address as a calling address. The I²C cannot be a master and a slave simultaneously.
- Slave mode-The I²C is not the driver of the SDA line. The module must be enabled before a START condition from a non-I²C master is detected.
- Interrupt-driven byte-to-byte data transfer-When successful slave addressing is achieved (and SCL returns to zero), the data transfer can proceed on a byte-to-byte basis in the direction specified by the R/W_B bit sent by the calling master. Each byte of data must be followed by an acknowledge bit, which is signaled from the receiving device. Several bytes can be transferred during a data transfer session.
- Boot sequencer mode-This mode can be used to initialize the configuration registers in the device after the I²C1 module is initialized. Note that the device powers up with boot sequencer mode disabled as a default, but this mode can be selected with the `cfg_boot_seq[0:1]` power-on reset (POR) configuration signals that are located on the LGPL3 and LGPL5 signals.

10.2.6 I²C-specific conditions

The following three conditions are defined for the I²C interface:

- **START condition**-This condition denotes the beginning of a new data transfer (each data transfer contains several bytes of data) and awakens all slaves.
- **Repeated START condition**-A START condition that is generated without a STOP condition to terminate the previous transfer.
- **STOP condition**-The master can terminate the transfer by generating a STOP condition to free the bus.

10.3 I²C external signal descriptions

The following sections give an overview of signals and provide detailed signal descriptions.

10.3.1 Signal overview

The I²C interface uses the SDA and SCL signals, described in the table below for data transfer.

Note that the signal patterns driven on SDA represent address, data, or read/write information at different stages of the protocol.

Table 10-1. I²C interface signal descriptions

Signal name	Idle state	I/O	State meaning
Serial clock (IICn_SCL)	HIGH	I	When the I ² C module is idle or acts as a slave, SCL defaults as an input. The unit uses SCL to synchronize incoming data on SDA. The bus is assumed to be busy when SCL is detected low.
		O	As a master, the I ² C module drives SCL along with SDA when transmitting. As a slave, the I ² C module drives SCL low for data pacing.
Serial data (IICn_SDA)	HIGH	I	When the I ² C module is idle or in a receiving mode, SDA defaults as an input. The unit receives data from other I ² C devices on SDA. The bus is assumed to be busy when SDA is detected low.
		O	When writing as a master or slave, the I ² C module drives data on SDA synchronous to SCL.

10.3.2 Detailed signal descriptions

SDA and SCL, described in the table below, serve as a communication interconnect with other devices. All devices connected to these two signals must have open-drain or open-collector outputs. The logic AND function is performed on both of these signals with external pull-up resistors. Refer to the device hardware specifications for the electrical characteristics of these signals.

Table 10-2. I²C interface signal-detailed signal descriptions

Signal	I/O	Description
IICn_SCL	I/O	Serial clock. Performs as an input when the device is programmed as an I ² C slave. SCL also performs as an output when the device is programmed as an I ² C master.
	O	As outputs for the bidirectional serial clock, these signals operate as described below.
		State meaning
	I	As inputs for the bidirectional serial clock, these signals operate as described below.
State Meaning		Asserted/Negated-The I ² C unit uses this signal to synchronize incoming data on SDA. The bus is assumed to be busy when this signal is detected low.
IICn_SDA	I/O	Serial data. Performs as an input when the device is in a receiving mode. SDA also performs as an output signal when the device is transmitting (as an I ² C master or a slave).
	O	As outputs for the bidirectional serial data, these signals operate as described below.
		State meaning
	I	As inputs for the bidirectional serial data, these signals operate as described below.
State meaning		Asserted/Negated-Used to receive data from other devices. The bus is assumed to be busy when SDA is detected low.

10.4 I2C memory map/register definition

The following table lists the I²C-specific registers and their offsets. It lists the offset, name, and a cross-reference to the complete description of each register. Note that the full register address is comprised of CCSRBAR together with the block base address and offset listed in the table below. The offsets to the memory map table are defined for both I²C interfaces. That is, I²C1 starts at address offset 0x000, and I²C2 starts at address offset 0x100. Note that the registers are the same for I²C2 except that the offsets change from 0x0nn to 0x1nn.

All I²C registers are one byte wide. Reads and writes to these registers must be byte-wide operations.

NOTE

Reserved bits should always be written with the value they returned when read. That is, the register should be programmed by reading the value, modifying appropriate fields, and writing back the value. The return value of the reserved fields should not be assumed, even though the reserved fields return zero.

This note does not apply to the I2C data register (I2CDR).

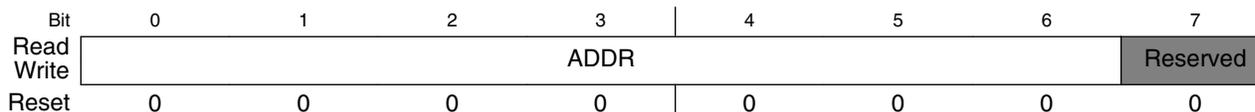
I2C memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
3000	I2C address register (I2C1_I2CADR)	8	R/W	00h	10.4.1/476
3004	I2C frequency divider register (I2C1_I2CFDR)	8	R/W	00h	10.4.2/477
3008	I2C control register (I2C1_I2CCR)	8	R/W	00h	10.4.3/479
300C	I2C status register (I2C1_I2CSR)	8	R/W	81h	10.4.4/480
3010	I2C data register (I2C1_I2CDR)	8	R/W	00h	10.4.5/481
3014	I2C digital filter sampling rate register (I2C1_I2CDFSRR)	8	R/W	10h	10.4.6/482
3100	I2C address register (I2C2_I2CADR)	8	R/W	00h	10.4.1/476
3104	I2C frequency divider register (I2C2_I2CFDR)	8	R/W	00h	10.4.2/477
3108	I2C control register (I2C2_I2CCR)	8	R/W	00h	10.4.3/479
310C	I2C status register (I2C2_I2CSR)	8	R/W	81h	10.4.4/480
3110	I2C data register (I2C2_I2CDR)	8	R/W	00h	10.4.5/481
3114	I2C digital filter sampling rate register (I2C2_I2CDFSRR)	8	R/W	10h	10.4.6/482

10.4.1 I2C address register (I2Cx_I2CADR)

The figure below shows the I2CADR register, which contains the address to which the I²C interface responds when addressed as a slave. Note that this is not the address that is sent on the bus during the address-calling cycle when the I²C module is in master mode.

Address: Base address + 0h offset



I2Cx_I2CADDR field descriptions

Field	Description
0–6 ADDR	Slave address. Contains the specific slave address that is used by the I ² C interface. Note that the default mode of the I ² C interface is slave mode for an address match. Note that an address match is one of the conditions that can cause I2CSR[MIF] to be set, signaling an interrupt pending condition.
7 -	This field is reserved.

10.4.2 I2C frequency divider register (I2Cx_I2CFDR)

Refer to application note AN2919, "Determining the I2C Frequency Divider Ratio for SCL," for additional guidance regarding the proper use of I2CFDR and I2CDFSRR.

Address: Base address + 4h offset

Bit	0	1	2	3	4	5	6	7
Read	Reserved			FDR				
Write	Reserved			FDR				
Reset	0	0	0	0	0	0	0	0

I2Cx_I2CFDR field descriptions

Field	Description																																
0–1 -	This field is reserved.																																
2–7 FDR	<p>Frequency divider ratio. Used to prescale the clock for bit rate selection. The serial bit clock frequency of SCL is equal to one half the platform (CCB) clock divided by the designated divider . Note that the frequency divider value can be changed at any point in a program. The serial bit clock frequency divider selections are described as follows:</p> <p>FDRDivider (Decimal)</p> <table border="0"> <tr><td>0x00</td><td>384</td></tr> <tr><td>0x01</td><td>416</td></tr> <tr><td>0x02</td><td>480</td></tr> <tr><td>0x03</td><td>576</td></tr> <tr><td>0x04</td><td>640</td></tr> <tr><td>0x05</td><td>704</td></tr> <tr><td>0x06</td><td>832</td></tr> <tr><td>0x07</td><td>1024</td></tr> <tr><td>0x08</td><td>1152</td></tr> <tr><td>0x09</td><td>1280</td></tr> <tr><td>0x0A</td><td>1536</td></tr> <tr><td>0x0B</td><td>1920</td></tr> <tr><td>0x0C</td><td>2304</td></tr> <tr><td>0x0D</td><td>2560</td></tr> <tr><td>0x0E</td><td>3072</td></tr> <tr><td>0x0F</td><td>3840</td></tr> </table>	0x00	384	0x01	416	0x02	480	0x03	576	0x04	640	0x05	704	0x06	832	0x07	1024	0x08	1152	0x09	1280	0x0A	1536	0x0B	1920	0x0C	2304	0x0D	2560	0x0E	3072	0x0F	3840
0x00	384																																
0x01	416																																
0x02	480																																
0x03	576																																
0x04	640																																
0x05	704																																
0x06	832																																
0x07	1024																																
0x08	1152																																
0x09	1280																																
0x0A	1536																																
0x0B	1920																																
0x0C	2304																																
0x0D	2560																																
0x0E	3072																																
0x0F	3840																																

Table continues on the next page...

I2Cx_I2CFDR field descriptions (continued)

Field	Description
0x10	4608
0x11	5120
0x12	6144
0x13	7680
0x14	9216
0x15	10240
0x16	12288
0x17	15360
0x18	18432
0x19	20480
0x1A	24576
0x1B	30720
0x1C	36864
0x1D	40960
0x1E	49152
0x1F	61440
0x20	256
0x21	288
0x22	320
0x23	352
0x24	384
0x25	448
0x26	512
0x27	576
0x28	640
0x29	768
0x2A	896
0x2B	1024
0x2C	1280
0x2D	1536
0x2E	1792
0x2F	2048
0x30	2560
0x31	3072
0x32	3584
0x33	4096
0x34	5120
0x35	6144
0x36	7168
0x37	8192
0x38	10240
0x39	12288
0x3A	14336
0x3B	16384
0x3C	20480
0x3D	24576
0x3E	28672
0x3F	32768

10.4.3 I2C control register (I2Cx_I2CCR)

Address: Base address + 8h offset

Bit	0	1	2	3	4	5	6	7
Read	MEN	MIEN	MSTA	MTX	TXAK		Reserved	BCST
Write						RSTA		
Reset	0	0	0	0	0	0	0	0

I2Cx_I2CCR field descriptions

Field	Description
0 MEN	<p>Module enable. This bit controls the software reset of the I²C module.</p> <p>0 The module is reset and disabled. When low, the interface is held in reset but the registers can still be accessed.</p> <p>1 The I²C module is enabled. This bit must be set before any other control register bits have any effect. All I²C registers for slave receive or master START can be initialized before setting this bit.</p>
1 MIEN	<p>Module interrupt enable</p> <p>0 Interrupts from the I²C module are disabled. This does not clear any pending interrupt conditions.</p> <p>1 Interrupts from the I²C module are enabled. An interrupt occurs provided I2CSR[MIF] is also set.</p>
2 MSTA	<p>Master/slave mode START</p> <p>0 When this bit is changed from one to zero, a STOP condition is generated and the mode changes from master to slave.</p> <p>1 Cleared without generating a STOP condition when the master loses arbitration. When this bit is changed from zero to one, a START condition is generated on the bus, and master mode is selected.</p>
3 MTX	<p>Transmit/receive mode select. This bit selects the direction of the master and slave transfers. When configured as a slave, this bit should be set by software according to I2CSR[SRW]. In master mode, the bit should be set according to the type of transfer required. Therefore, for address cycles, this bit is always high. The MTX bit is cleared when the master loses arbitration.</p> <p>0 Receive mode</p> <p>1 Transmit mode</p>
4 TXAK	<p>Transfer acknowledge. This bit specifies the value driven onto the SDA line during acknowledge cycles for both master and slave receivers. The value of this bit only applies when the I²C module is configured as a receiver, not a transmitter. It also does not apply to address cycles; when the device is addressed as a slave, an acknowledge is always sent.</p> <p>0 An acknowledge signal (low value on SDA) is sent out to the bus at the 9th clock after receiving one byte of data.</p> <p>1 No acknowledge signal response (high value on SDA) is sent.</p>
5 RSTA	<p>Repeated START. Setting this bit always generates a repeated START condition on the bus, provides the device with the current bus master. Attempting a repeated START at the wrong time (or if the bus is owned by another master), results in loss of arbitration. Note that this bit is not readable, which means if a read is performed to I2CCR[RSTA], a zero value is returned.</p>

Table continues on the next page...

I²Cx_I2CCR field descriptions (continued)

Field	Description
	0 No START condition is generated 1 Generates repeated START condition
6 -	This field is reserved.
7 BCST	Broadcast 0 Disables the broadcast accept capability 1 Enables the I ² C to accept broadcast messages at address zero

10.4.4 I²C status register (I²Cx_I2CSR)

The I²C status register is read only with the exception of the MIF and MAL bits, which can be cleared by software. The MCF and RXAK bits are set at reset; all other I2CSR bits are cleared on reset.

Address: Base address + Ch offset

Bit	0	1	2	3	4	5	6	7
Read	MCF	MAAS	MBB	MAL	BCSTM	SRW	MIF	RXAK
Write								
Reset	1	0	0	0	0	0	0	1

I²Cx_I2CSR field descriptions

Field	Description
0 MCF	Data transfer. When one byte of data is transferred, the bit is cleared. It is set by the falling edge of the 9th clock of a byte transfer. 0 Byte transfer in progress. MCF is cleared under either of the following conditions: When I2CDR is read in receive mode, When I2CDR is written in transmit mode 1 Byte transfer is completed
1 MAAS	Addressed as a slave. When the value in I2CDR matches with the calling address, this bit is set. The processor is interrupted, if I2CCR[MIE] is set. Next, the processor must check the SRW bit and set I2CCR[MTX] accordingly. Writing to the I2CCR automatically clears this bit. 0 Not addressed as a slave 1 Addressed as a slave
2 MBB	Bus busy. Indicates the status of the bus. When a START condition is detected, MBB is set. If a STOP condition is detected, it is cleared. 0 I ² C bus is idle 1 I ² C bus is busy
3 MAL	Arbitration lost. Automatically set when the arbitration procedure is lost. Note that the device does not automatically retry a failed transfer attempt.

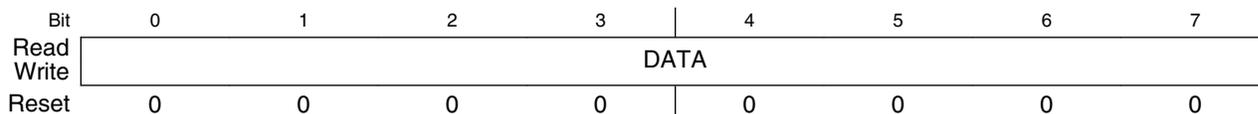
Table continues on the next page...

I2Cx_I2CSR field descriptions (continued)

Field	Description
	0 Arbitration is not lost. Can only be cleared by software 1 Arbitration is lost
4 BCSTM	Broadcast match 0 There has not been a broadcast match. 1 The calling address matches with the broadcast address instead of the programmed slave address. This also sets if this I ² C drives an address of all 0s and broadcast mode is enabled.
5 SRW	Slave read/write. When MAAS is set, SRW indicates the value of the R/W command bit of the calling address, which is sent from the master. 0 Slave receive, master writing to slave 1 Slave transmit, master reading from slave. This bit is valid only when both of the following conditions are true: A complete transfer occurred and no other transfers have been initiated. The I ² C interface is configured as a slave and has an address match. By checking this bit, the processor can select slave transmit/receive mode according to the command of the master.
6 MIF	Module interrupt. The MIF bit is set when an interrupt is pending, causing a processor interrupt request (provided I2CCR[MIEN] is set). The interrupts for I ² C1 and I ² C2 are combined into one interrupt, which is sourced by the dual I ² C controller. 0 No interrupt is pending. Can be cleared only by software. 1 Interrupt is pending. MIF is set when one of the following events occurs: One byte of data is transferred (set at the falling edge of the 9th clock). The value in I2CADR matches with the calling address in slave-receive mode. Arbitration is lost.
7 RXAK	Received acknowledge. The value of SDA during the reception of acknowledge bit of a bus cycle. If the received acknowledge bit (RXAK) is low, it indicates that an acknowledge signal has been received after the completion of eight bits of data transmission on the bus. If RXAK is high, it means no acknowledge signal has been detected at the 9th clock. 0 Acknowledge received 1 No acknowledge received

10.4.5 I2C data register (I2Cx_I2CDR)

Address: Base address + 10h offset



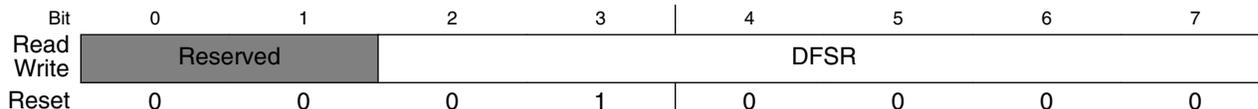
I2Cx_I2CDR field descriptions

Field	Description
0–7 DATA	Transmission starts when an address and the R/W bit are written to the data register and the I ² C interface performs as the master. A data transfer is initiated when data is written to the I2CDR. The most significant bit is sent first in both cases. In master receive mode, reading the data register allows the read to occur, but also allows the I ² C module to receive the next byte of data on the I2C interface. In slave mode, the same function is available after it is addressed. Note that in both master receive and slave receive modes, the very first read is always a dummy read.

10.4.6 I2C digital filter sampling rate register (I2Cx_I2CDFSRR)

The digital filter sampling rate register (I2CDFSRR) is shown in the figure below. Refer to application note AN2919, "Determining the I2C Frequency Divider Ratio for SCL," for additional guidance regarding the proper use of I2CFDR and I2CDFSRR.

Address: Base address + 14h offset



I2Cx_I2CDFSRR field descriptions

Field	Description
0–1 -	This field is reserved.
2–7 DFSRR	Digital filter sampling rate. To assist in filtering out signal noise, the sample rate is programmed. This field is used to prescale the frequency at which the digital filter takes samples from the I ² C bus. The resulting sampling rate is calculated by dividing one half the platform (CCBclock) frequency by the non-zero value of DFSRR.

10.5 Functional description

The I²C unit always performs as a slave receiver as a default, unless explicitly programmed to be a master or slave transmitter. After the boot sequencer has completed (when powered up in boot sequencer mode), the I²C interface performs as a slave receiver.

Note that the boot sequencer only functions from the I²C1 interface; the I²C2 interface cannot be used for this purpose.

10.5.1 Transaction protocol

A standard I²C transfer consists of the following:

- START condition
- Slave target address transmission
- Data transfer
- STOP condition

The following figure shows the interaction of these four parts with the calling address, data byte, and new calling address components of the I²C protocol. The details of the protocol are described in the following sections.

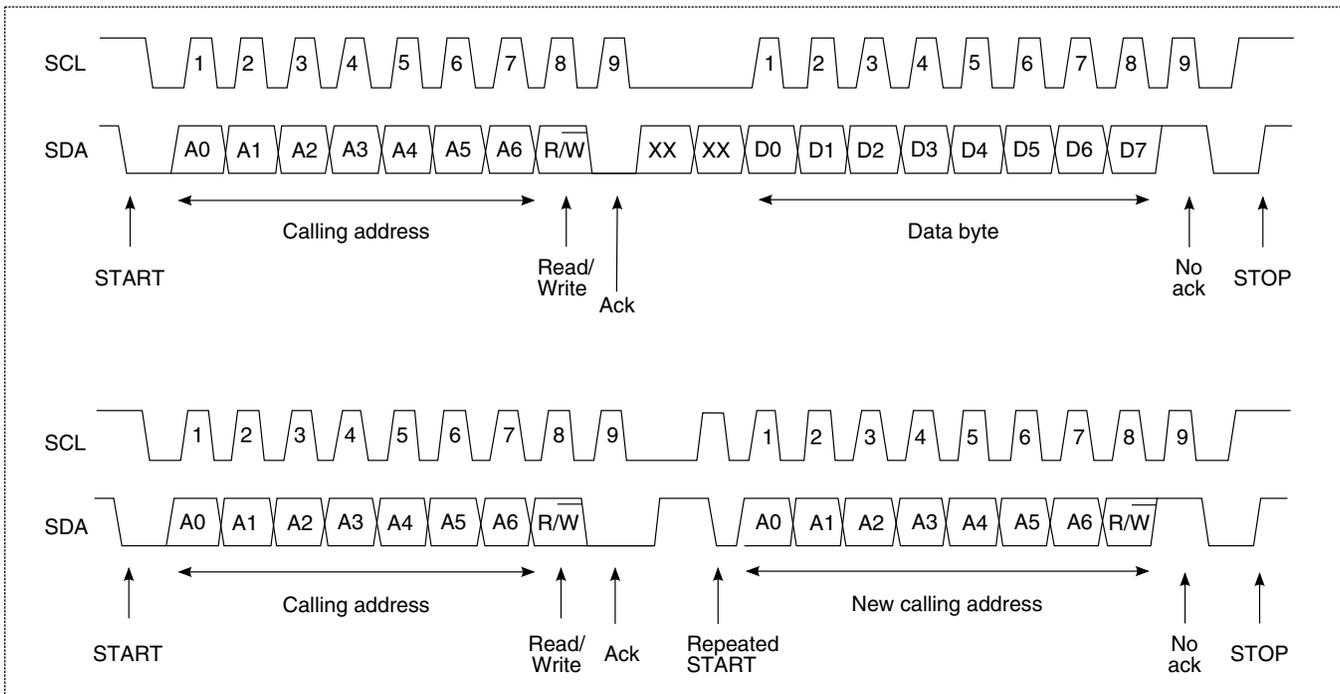


Figure 10-20. I²C interface transaction protocol

10.5.1.1 START condition

When the I²C bus is not engaged (both SDA and SCL lines are at logic high), a master can initiate a transfer by sending a START condition.

As shown in [Figure 10-20](#), a START condition is defined as a high-to-low transition of SDA while SCL is high. This condition denotes the beginning of a new data transfer. Each data transfer can contain several bytes and awakens all slaves. The START condition is initiated by a software write that sets I2CCR[MSTA].

10.5.1.2 Slave address transmission

The first byte of data is transferred by the master immediately after the START condition is the slave address.

This is a seven-bit calling address followed by a R/W_B bit, which indicates the direction of the data being transferred to the slave. Each slave in the system has a unique address. In addition, when the I²C module is operating as a master, it must not transmit an address that is the same as its slave address. An I²C device cannot be master and slave at the same time; if this is attempted, the results are boundedly undefined.

Only the slave with a calling address that matches the one transmitted by the master responds by returning an acknowledge bit (pulling the SDA signal low at the 9th clock) as shown in [Figure 10-20](#). If no slave acknowledges the address, the master should generate a STOP condition or a repeated START condition.

When slave addressing is successful (and SCL returns to zero), the data transfer can proceed on a byte-to-byte basis in the direction specified by the R/W_B bit sent by the calling master.

The I²C module responds to a general call (broadcast) command when I2CCR[BCST] is set. A broadcast address is always zero; however the I²C module does not check the R/W_B bit. The second byte of the broadcast message is the master address. Because the second byte is automatically acknowledged by hardware, the receiver device software must verify that the broadcast message is intended for itself by reading the second byte of the message. If the master address is for another receiver device and the third byte is a write command, software can ignore the third byte during the broadcast. If the master address is for another receiver device and the third byte is a read command, software must write 0xFF to I2CDR with I2CCR[TXAK] = 1, so that it does not interfere with the data written from the addressed device.

Each data byte is 8 bits long. Data bits can be changed only while SCL is low and must be held stable while SCL is high, as shown in [Figure 10-20](#). There is one clock pulse on SCL for each data bit, and the most significant bit (msb) is transmitted first. Each byte of data must be followed by an acknowledge bit, which is signaled from the receiving device by pulling the SDA line low at the 9th clock. Therefore, one complete data byte transfer takes 9 clock pulses. Several bytes can be transferred during a data transfer session.

If the slave receiver does not acknowledge the master, the SDA line must be left high by the slave. The master can then generate a stop condition to abort the data transfer or a START condition (repeated START) to begin a new calling.

If the master receiver does not acknowledge the slave transmitter after a byte of transmission, the slave interprets that the end-of-data has been reached. Then the slave releases the SDA line for the master to generate a STOP or a START condition.

10.5.1.3 Repeated START condition

Figure 10-20 shows a repeated START condition, which is generated without a STOP condition that can terminate the previous transfer. The master uses this method to communicate with another slave or with the same slave in a different mode (transmit/receive mode) without releasing the bus.

10.5.1.4 STOP condition

The master can terminate the transfer by generating a STOP condition to free the bus.

A STOP condition is defined as a low-to-high transition of the SDA signal while SCL is high. For more information, see Figure 10-20. Note that a master can generate a STOP even if the slave has transmitted an acknowledge bit, at which point the slave must release the bus. The STOP condition is initiated by a software write that clears I2CCR[MSTA].

As described in [Repeated START condition](#), the master can generate a START condition followed by a calling address without generating a STOP condition for the previous transfer. This is called a repeated START condition.

10.5.1.5 Protocol implementation details

The following sections provide details about how aspects of the protocol are implemented in this I²C module.

10.5.1.5.1 Transaction monitoring-implementation details

The different conditions of the I²C data transfers are monitored as follows:

- START conditions are detected when an SDA fall occurs while SCL is high.
- STOP conditions are detected when an SDA rise occurs while SCL is high.
- Data transfers in progress are canceled when a STOP condition is detected or if there is a slave address mismatch. Cancellation of data transactions resets the clock module.
- The bus is detected to be busy upon the detection of a START condition, and idle upon the detection of a STOP condition.

10.5.1.5.2 Control transfer-implementation details

The I²C module contains logic that controls the output to the serial data (SDA) and serial clock (SCL) lines of the I²C.

The SCL output is pulled low as determined by the internal clock generated in the clock module. The SDA output can only change at the midpoint of a low cycle of the SCL, unless it is performing a START, STOP, or restart condition. Otherwise, the SDA output is held constant.

The SDA signal is pulled low when one or more of the following conditions are true in either master or slave mode:

- Master mode
 - Data bit (transmit)
 - Ack bit (receive)
 - START condition
 - STOP condition
 - Restart condition
- Slave mode
 - Acknowledging address match
 - Data bit (transmit)
 - Ack bit (receive)

The SCL signal corresponds to the internal SCL signal when one or more of the following conditions are true in either master or slave mode:

- Master mode
 - Bus owner
 - Lost arbitration
 - START condition
 - STOP condition
 - Restart condition begin
 - Restart condition end
- Slave mode
 - Address cycle
 - Transmit cycle
 - Ack cycle

10.5.1.6 Address compare-implementation details

Address compare block determines if a slave has been properly addressed, either by its slave address or by the general broadcast address (which addresses all slaves).

The three performed address comparisons are described as follows:

- Whether a broadcast message has been received, to update the I2CSR

- Whether the module has been addressed as a slave, to update the I2CSR and to generate an interrupt
- If the address transmitted by the current master matches the general broadcast address

10.5.2 Arbitration procedure

The I²C interface is a true multiple-master bus that allows more than one master device to be connected on it.

If two or more masters simultaneously try to control the bus, each master's clock synchronization procedure (including the I²C module) determines the bus clock—the low period is equal to the longest clock low period and the high is equal to the shortest one among the masters. A bus master loses arbitration if it transmits a logic 1 on SDA while another master transmits a logic 0. The losing masters immediately switch to slave-receive mode and stop driving the SDA line. In this case, the transition from master to slave mode does not generate a STOP condition. Meanwhile, the I²C unit sets the I2CSR[MAL] status bit to indicate the loss of arbitration and, as a slave, services the transaction if it is directed to itself.

If the I²C module is enabled in the middle of an ongoing byte transfer, the interface behaves as follows:

- Slave mode—The I²C module ignores the current transfer on the bus and starts operating whenever a subsequent START condition is detected.
- Master mode—The I²C module cannot tell whether the bus is busy; therefore, if a START condition is initiated, the current bus cycle can be corrupted. This ultimately results in the current bus master of the I²C interface losing arbitration, after which bus operations return to normal.

10.5.2.1 Arbitration control

The arbitration control block controls the arbitration procedure of the master mode.

A loss of arbitration occurs whenever the master detects a 0 on the external SDA line while attempting to drive a 1, tries to generate a START or restart at an inappropriate time, or detects an unexpected STOP request on the line.

In master mode, arbitration by the master is lost (and I2CSR[MAL] is set) under the following conditions:

- SDA samples low when the master drives high during an address or data-transmit cycle (transmit).
- SDA samples low when the master drives high during a data-receive cycle of the acknowledge (Ack) bit (receive).
- A START condition is attempted when the bus is busy.
- A repeated START condition is requested in slave mode.
- A start condition is attempted when the requesting device is not the bus owner
- Unexpected STOP condition detected

Note that the I²C module does not automatically retry a failed transfer attempt.

10.5.3 Handshaking

The clock synchronization mechanism can be used as a handshake in data transfer.

Slave devices can hold SCL low after completion of a 1-byte transfer (9 bits). In such cases, it halts the bus clock and forces the master clock into wait states until the slave releases the SCL line.

10.5.4 Clock control

The clock control block handles requests from the clock signal for transferring and controlling data for multiple tasks.

A 9-cycle data transfer clock is requested for the following conditions:

- Master mode
 - Transmit slave address after START condition
 - Transmit slave address after restart condition
 - Transmit data
 - Receive data
- Slave mode
 - Transmit data
 - Receive data
 - Receive slave address after START or restart condition

10.5.4.1 Clock synchronization

Due to the wire AND logic on the SCL line, a high-to-low transition on the SCL line affects all devices connected on the bus.

The devices begin counting their low period when the master drives the SCL line low. After a device has driven SCL low, it holds the SCL line low until the clock high state is reached. However, the change of low-to-high in a device clock may not change the state of the SCL line if another device is still within its low period. Therefore, the synchronized clock signal, SCL, is held low by the device with the longest low period. Devices with shorter low periods enter a high wait state during this time. When all devices concerned have counted off their low period, the synchronized SCL line is released and pulled high. Then there is no difference between the devices' clocks and the state of the SCL line, and all the devices begin counting their high periods. The first device to complete its high period pulls the SCL line low again.

10.5.4.2 Input synchronization and digital filter

The following sections describes the synchronizing of the input signals and the filtering of the SCL and SDA lines in detail.

10.5.4.2.1 Input signal synchronization

The input synchronization block synchronizes the input SCL and SDA signals to the system clock and detects transitions of these signals.

10.5.4.2.2 Filtering of SCL and SDA lines

The SCL and SDA inputs are filtered to eliminate noise.

Three consecutive samples of the SCL and SDA lines are compared to a pre-determined sampling rate. If they are all high, the output of the filter is high. If they are all low, the output is low. If they are any combination of highs and lows, the output is whatever the value of the line was in the previous clock cycle.

The sampling rate is equal to a binary value stored in the frequency register I2CDFSRR. The duration of the sampling cycle is controlled by a down counter. This allows a software write to the frequency register to control the filtered sampling rate.

10.5.4.3 Clock stretching

Slaves can use the clock synchronization mechanism to slow down the transfer bit rate.

After the master has driven the SCL line low, the slave can drive SCL low for the required period and then release it. If the slave SCL low period is greater than the master SCL low period, then the resulting SCL bus signal low period is stretched.

10.5.5 Boot sequencer mode

If boot sequencer mode is selected on POR (by the settings on the `cfg_boot_seq[0:1]` reset configuration signals, as described in [Boot sequencer configuration](#)), the I²C1 module communicates with one or more EEPROMs through the I²C interface on IIC1_SCL and IIC1_SDA. The boot sequencer accesses the I²C1 serial ROM device at a serial bit clock frequency equal to the platform (CCB) clock frequency divided by 2560. The EEPROM(s) can be programmed to initialize one or more configuration registers of this integrated device.

If the boot sequencer is enabled for normal I²C addressing mode, the I²C interface initiates the following sequence during reset:

1. Generate RESET sequence (START then 9 SCL cycles) to the EEPROM twice. This clears any transactions that may have been in progress prior to the reset.
2. Generate START
3. Transmit 0xA0 which is the 7-bit calling address (0b101_0000) with a write command appended (0 as the least significant bit).
4. Transmit 0x00 which is the 8-bit starting address
5. Generate a repeated START
6. Transmit 0xA1 which is the 7-bit calling address (0b101_0000) with a read command appended (1 as the least significant bit).
7. Receive 256 bytes of data from the EEPROM (unless the CONT bit is cleared in the data structure).
8. Generate a repeated START
9. Transmit 0xA2 which is the 7-bit calling address of the second target (0b101_0001) with a write command appended (0 as the least significant bit).
10. Transmit 0x00 which is the 8-bit starting address for the second target.
11. Generate a repeated START
12. Transmit 0xA3 which is the 7-bit calling address (0b101_0001) with a read command appended (1 as the least significant bit).
13. Receive another 256 bytes of data from the second EEPROM (unless the CONT bit is cleared in the data structure).

The sequence repeats with successive targets until the CONT bit in the data structure is cleared and the CRC check is executed. If the last register is not detected (that is, the CONT bit is never cleared) before wrapping back to the first address, an error condition is detected, causing the device to hang and the HRESET_REQ_B signal to assert externally. The I²C module continues to read from the EEPROM(s) as long as the

continue (CONT) bit is set in the EEPROM(s). The CONT bit resides in the address/attributes field that is transferred from the EEPROM, as described in [EEPROM calling address](#). There should be no other I²C traffic when the boot sequencer is active.

The boot sequencer mode also supports an extension of the standard I²C interface that uses more address bits to allow for EEPROM devices that have more than 256 bytes, and this extended addressing mode is selectable during POR with a different encoding on the `cfg_boot_seq[0:1]` reset configuration signals. In this mode, only one EEPROM device may be used, and the maximum number of registers is limited by the size of the EEPROM. If the boot sequencer is enabled for extended I²C addressing mode, the I²C interface initiates the following sequence during reset:

1. Generate RESET sequence (START then 9 SCL cycles) to the EEPROM twice. This clears any transactions that may have been in progress prior to the reset.
2. Generate START
3. Transmit 0xA0 which is the 7-bit calling address (0b101_0000) with a write command appended (0 as the least significant bit).
4. Transmit 0x00 which is the high-order starting address
5. Transmit 0x00 which is the low-order starting address
6. Generate a repeated START
7. Transmit 0xA1 which is the 7-bit calling address (0b101_0000) with a read command appended (1 as the least significant bit).
8. Receive data continuously from the EEPROM until the CONT bit is cleared and the CRC check is executed. See [EEPROM data format](#), for more information.

Note that as described in [Boot sequencer configuration](#), the default value for the `cfg_boot_seq[0:1]` reset configuration pins is 0b11, which corresponds to the I²C boot sequencer being disabled at power-up.

10.5.5.1 EEPROM calling address

The device uses 0b101_0000 for the EEPROM calling address.

The first EEPROM to be addressed must be programmed to respond to this address, or an error is generated. If more EEPROMs are used, they are addressed in sequential order.

10.5.5.2 EEPROM data format

The I²C module expects that a particular data format be used for data in the EEPROM.

A preamble should be the first 3 bytes programmed into the EEPROM. It should have a value of 0xAA_55AA. The I²C module checks to ensure that this preamble is correctly detected before proceeding further. Following the preamble, there should be a series of configuration registers (known as register preloads) programmed into the EEPROM. Each configuration register should be programmed according to a particular format, as shown in [Table 10-21](#). The first 3 bytes hold the attributes and address offset, as follows. The attributes contained are alternate configuration space (ACS), byte enables, and continue (CONT). The boot sequencer expects the address offset to be a 32-bit (word) offset, that is, the 2 low-order bits are not included in the boot sequencer command. For example, to access LAWBAR0 (byte offset of 0x00C08), the boot sequencer ADDR[0:17] should be set to 0x00302.

After the first 3 bytes, 4 bytes of data should hold the desired value of the configuration register, regardless of the size of the transaction. Byte enables should be asserted for any byte that is written to the configuration register, and they should be asserted contiguously, creating a 1-, 2-, or 4-byte write to a register. The boot sequencer assumes that a big-endian address is stored in the EEPROM. In addition, byte enable bit 0 (bit 1 of the byte) corresponds to the most-significant byte of data (data[0:7]), and byte enable bit 3 (bit 4 of the byte) corresponds to the LSB of data (data[24:31]).

By setting ACS, an alternate configuration space address is prepended to the write request from the boot sequencer. Otherwise, CCSRBAR is prepended to the EEPROM address.

If CONT is cleared, the first 3 bytes, including ACS, the byte enables, and the address, must also be cleared. Also, the data contains the final cyclic redundancy check (CRC). A CRC-32 algorithm is used to check the integrity of the data. The polynomial used is:

$$1 + x^1 + x^2 + x^4 + x^5 + x^7 + x^8 + x^{10} + x^{11} + x^{12} + x^{16} + x^{22} + x^{23} + x^{26} + x^{32}$$

CRC values are calculated using the above polynomial with a start value of 0xFFFF_FFFF and an XOR with 0x0000_0000. The CRC should cover all bytes stored in the EEPROM prior to the CRC. This includes the preamble, all register preloads, and the first 3 bytes of the last 7-byte preload (which should be all zeros). If a preamble or CRC fail is detected, the device hangs and the external HRESET_REQ_B signal asserts. If there is a preamble fail, the boot sequencer may continue to pull I²C pins low until a hard reset occurs.

Table 10-21. EEPROM data format for one register preload command

0	1			4	5	6	7
ACS	BYTE_EN				CONT	ADDR[0-1]	
ADDR[2-9]							
ADDR[10-17]							

Table continues on the next page...

Table 10-21. EEPROM data format for one register preload command (continued)

DATA[0-7]
DATA[8-15]
DATA[16-23]
DATA[24-31]

Table 10-22 shows an example of the EEPROM contents, including the preamble, data format, and CRC.

Table 10-22. EEPROM contents

0	1	2	3	4	5	6	7	
1	0	1	0	1	0	1	0	Preamble
0	1	0	1	0	1	0	1	
1	0	1	0	1	0	1	0	
ACS	BYTE_EN			1	ADDR[0-1]			First configuration Preload command
ADDR[2-9]								
ADDR[10-17]								
DATA[0-7]								
DATA[8-15]								
DATA[16-23]								
DATA[24-31]								
ACS	BYTE_EN			1	ADDR[0-1]			Second configuration Preload command
ADDR[2-9]								
ADDR[10-17]								
DATA[0-7]								
DATA[8-15]								
DATA[16-23]								
DATA[24-31]								
.								
.								
.								
ACS	BYTE_EN			1	ADDR[0-1]			Last configuration Preload command
ADDR[2-9]								
ADDR[10-17]								
DATA[0-7]								
DATA[8-15]								
DATA[16-23]								
DATA[24-31]								
0	0	0	0	0	0	0	0	End command
0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	

Table continues on the next page...

Table 10-22. EEPROM contents (continued)

0	1	2	3	4	5	6	7	
CRC[0-7]								Cyclic Redundancy Check
CRC[8-15]								
CRC[16-23]								
CRC[24-31]								

10.6 Initialization/application information

This section describes some programming guidelines recommended for the I²C interface. See [Figure 10-21](#) for a recommended flowchart for I²C interrupt service routines.

The I²C registers in this chapter are shown in big-endian format. If the system is in little-endian mode, software must swap the bytes appropriately. This appropriate byte swapping is needed as I²C registers are byte registers. Also, an **msync** assembly instruction must be executed after each I²C register read/write access to guarantee in-order execution.

The I²C controller does not guarantee its recovery from all illegal I²C bus activity. In addition, a malfunctioning device may hold the bus captive. A good programming practice is for software to rely on a watchdog timer to help recover from I²C bus hangs. The recovery routine should also handle the case when the status bits returned after an interrupt are not consistent with what was expected due to illegal I²C bus protocol behavior.

10.6.1 Initialization sequence

A hard reset initializes all the I²C registers to their default states.

The following initialization sequence initializes the I²C unit:

1. All I²C registers must be located in a cache-inhibited page.
2. Update I2CFDR[FDR] and select the required division ratio to obtain the SCL frequency from the CCB (platform) clock. Note that the platform frequency must first be divided by two; see [I2C frequency divider register \(I2C_I2CFDR\)](#), for more details.
3. Update I2CADR to define the slave address for this device.
4. Modify I2CCR to select master/slave mode, transmit/receive mode, and interrupt-enable or disable.
5. Set the I2CCR[MEN] to enable the I²C interface.

10.6.2 Generation of START

After initialization, the following sequence can be used to generate START:

1. If the device is connected to a multimaster I²C system, test the state of I2CSR[MBB] to check whether the serial bus is free (I2CSR[MBB] = 0) before switching to master mode.
2. Select master mode (set I2CCR[MSTA]) to transmit serial data and select transmit mode (set I2CCR[MTX]) for the address cycle.
3. Write the slave address being called into I2CDR. The data written to I2CDR[0-6] comprises the slave calling address. I2CCR[MTX] indicates the direction of transfer (transmit/receive) required from the slave.

The scenario above assumes that the I²C interrupt bit (I2CSR[MIF]) is cleared. If MIF is set at any time, an I²C interrupt is generated (provided interrupt reporting is enabled with I2CCR[MIE] = 1) so that the I²C interrupt handler can handle the interrupt. Note that the interrupts for I²C1 and I²C2 are combined into one interrupt, which is sourced by the dual I²C controller.

10.6.3 Post-transfer software response

Transmission or reception of a byte automatically sets the data transferring bit (I2CSR[MCF]), which indicates that one byte has been transferred.

The I²C interrupt bit (I2CSR[MIF]) is also set and an interrupt is generated to the processor if the interrupt function is enabled during the initialization sequence (I2CCR[MIE] is set). In the interrupt handler, software must take the following steps:

1. Clear I2CSR[MIF].
2. Read the contents of the I²C data register (I2CDR) in receive mode or write to I2CDR in transmit mode. Note that this causes I2CSR[MCF] to be cleared. See [Interrupt service routine flowchart](#).

When an interrupt occurs at the end of the address cycle, the master remains in transmit mode. If master receive mode is required, I2CCR[MTX] must be toggled at this stage. See [Interrupt service routine flowchart](#).

If the interrupt function is disabled, software can service the I2CDR in the main program by monitoring I2CSR[MIF]. In this case, I2CSR[MIF] must be polled rather than I2CSR[MCF] because MCF behaves differently when arbitration is lost. Note that

interrupt or other bus conditions may be detected before the I²C signals have time to settle. Thus, when polling I2CSR[MIF] (or any other I2CSR bits), software delays may be needed in order to give the I²C signals sufficient time to settle.

During slave-mode address cycles (I2CSR[MAAS] is set), I2CSR[SRW] should be read to determine the direction of the subsequent transfer and I2CCR[MTX] should be programmed accordingly. For slave-mode data cycles (MAAS is cleared), I2CSR[SRW] is not valid and I2CCR[MTX] must be read to determine the direction of the current transfer. See [Interrupt service routine flowchart](#), for more details.

10.6.4 Generation of STOP

A data transfer ends with a STOP condition generated by the master device.

A master transmitter can generate a STOP condition after all the data has been transmitted.

If a master receiver wants to terminate a data transfer, it must inform the slave transmitter by not acknowledging the last byte of data (by setting the transmit acknowledge bit (I2CCR[TXAK])) before reading the next-to-last byte of data. At this time, the next-to-last byte of data has already been transferred on the I²C interface, so the last byte does not receive the data acknowledge (because I2CCR[TXAK] is set). Before the interrupt service routine reads the last byte of data, a STOP condition must first be generated.

I2CCR[TXAK] must be set before allowing the I²C module to receive the last data byte on the I²C bus. Eventually, I2CCR[TXAK] must be cleared again for subsequent I²C transactions. This can be accomplished when setting up the I2CCR for the next transfer.

10.6.5 Generation of repeated START

At the end of a data transfer, if the master still wants to communicate on the bus, it can generate another START condition followed by another slave address without first generating a STOP condition.

This is accomplished by setting I2CCR[RSTA].

10.6.6 Generation of SCL when SDA low

It is sometimes necessary to force the I²C module to become the I²C bus master out of reset and drive SCL (even though SDA may already be driven, which indicates that the bus is busy).

This can occur when a system reset does not cause all I²C devices to be reset. Thus, SDA can be driven low by another I²C device while this I²C module is coming out of reset and stays low indefinitely. The following procedure can be used to force this I²C module to generate SCL so that the device driving SDA can finish its transaction:

1. Disable the I²C module and set the master bit by setting I2CCR to 0x20
2. Enable the I²C module by setting I2CCR to 0xA0
3. Read the I2CDR
4. Return the I²C module to slave mode by setting I2CCR to 0x80

10.6.7 Slave mode interrupt service routine

In the slave interrupt service routine, the module addressed as a slave should be tested to check if a calling of its own address has been received.

If I2CSR[MAAS] is set, software should set the transmit/receive mode select bit (I2CCR[MTX]) according to the R/W_B command bit (I2CSR[SRW]). Writing to I2CCR clears MAAS automatically. MAAS is read as set only in the interrupt handler at the end of that address cycle where an address match occurred; interrupts resulting from subsequent data transfers clear MAAS. A data transfer can then be initiated by writing to I2CDR for slave transmits or dummy reading from I2CDR in slave-receive mode. The slave drives SCL low between byte transfers. SCL is released when the I2CDR is accessed in the required mode.

10.6.7.1 Slave transmitter and received acknowledge

In the slave transmitter routine, the received acknowledge bit (I2CSR[RXAK]) must be tested before sending the next byte of data.

The master signals an end-of-data by not acknowledging the data transfer from the slave. When no acknowledge is received (I2CSR[RXAK] is set), the slave transmitter interrupt routine must clear I2CCR[MTX] to switch the slave from transmitter to receiver mode. A dummy read of I2CDR then releases SCL so that the master can generate a STOP condition. See [Interrupt service routine flowchart](#).

10.6.7.2 Loss of arbitration and forcing of slave mode

When a master loses arbitration the following conditions all occur:

- I2CSR[MAL] is set

- I2CCR[MSTA] is cleared (changing the master to slave mode)
- An interrupt occurs (if enabled) at the falling edge of the 9th clock of this transfer

Thus, the slave interrupt service routine should first test I2CSR[MAL] and software should clear it if it is set. See [Arbitration control](#), for more information.

10.6.8 Interrupt service routine flowchart

The following figure shows an example algorithm for an I²C interrupt service routine.

Deviation from the flowchart may result in unpredictable I²C bus behavior. However, in the slave receive mode the interrupt service routine may need to set I2CCR[TXAK] when the next-to-last byte is to be accepted. It is recommended that an **msync** instruction follow each I²C register read or write to guarantee in-order instruction execution.

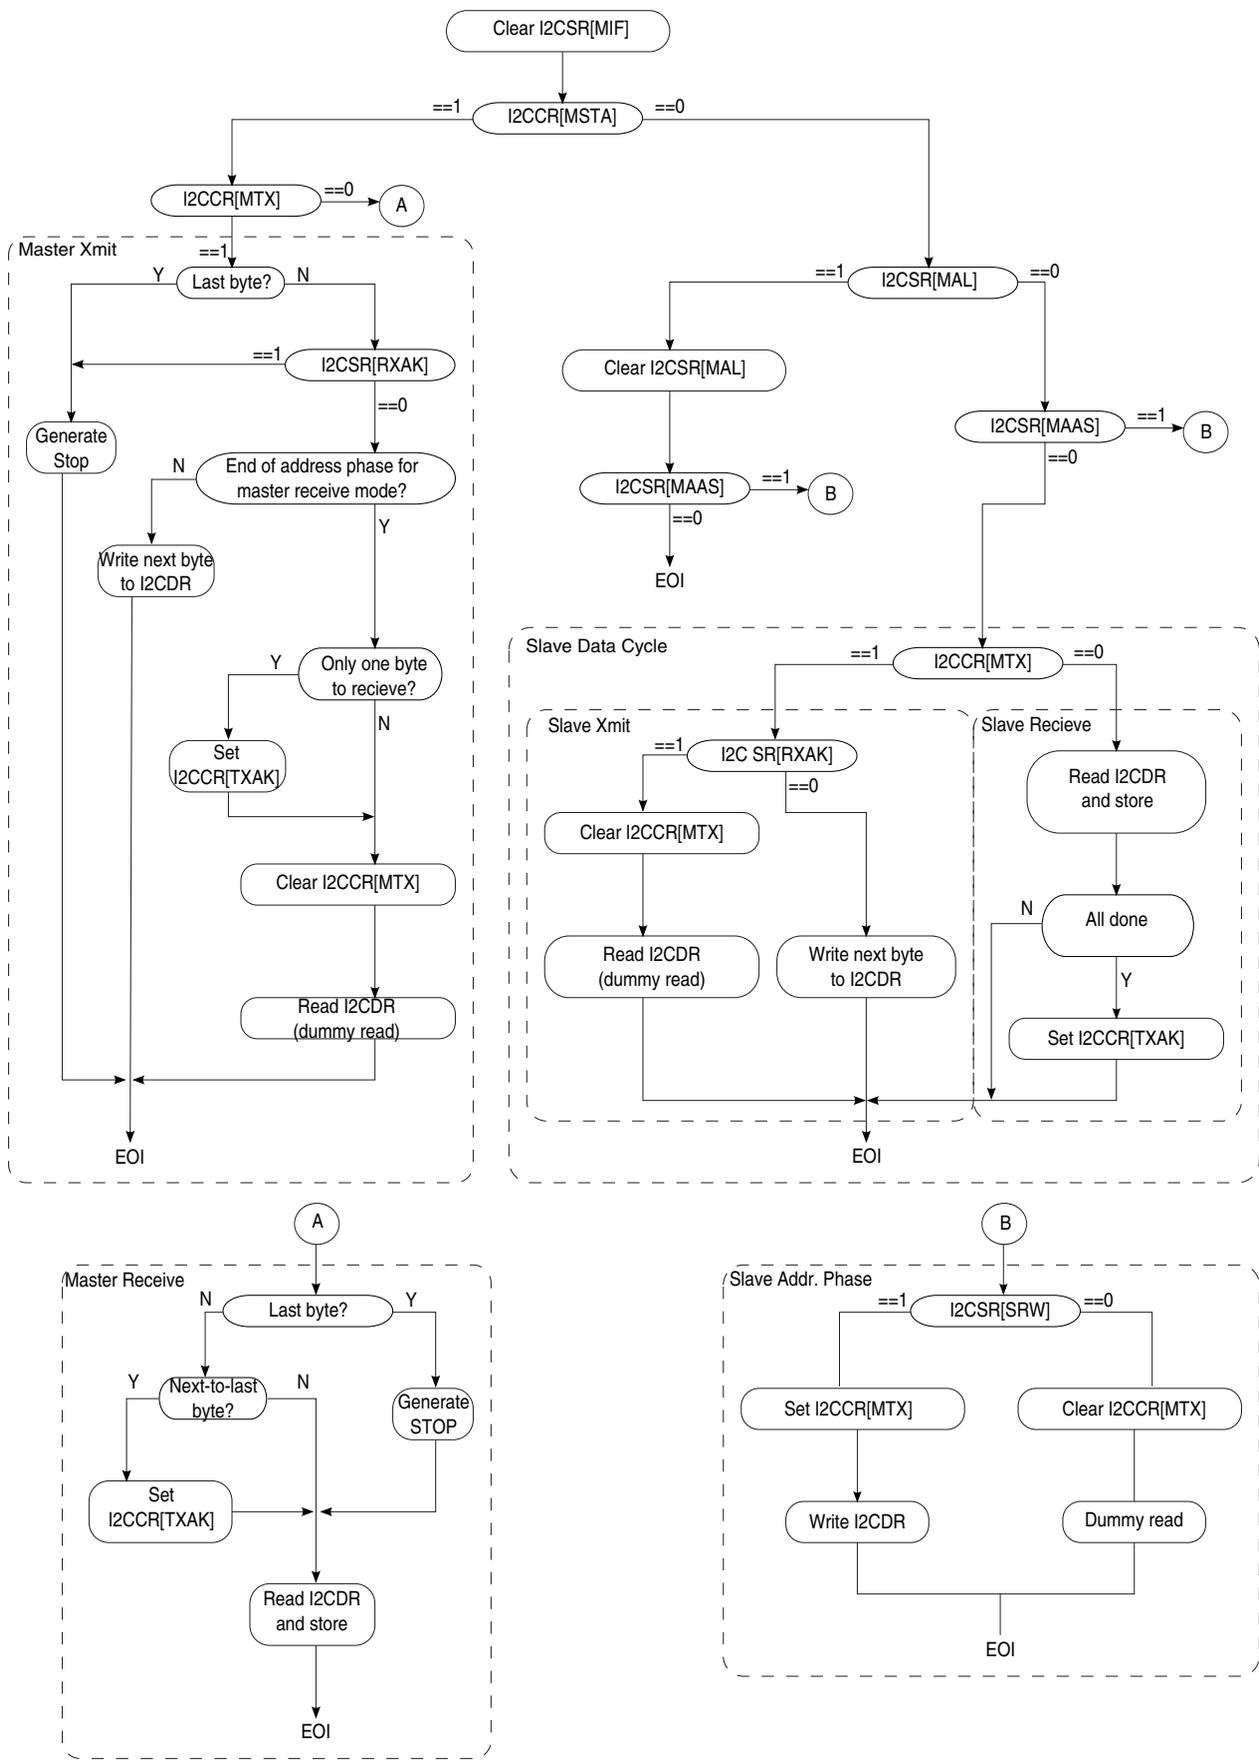


Figure 10-21 Example I2C interrupt service routine flowchart



Chapter 11

DUART

This chapter describes the dual universal asynchronous receiver/transmitters (DUART). It describes the functional operation, the initialization sequence, and the programming details for the DUART registers and features.

11.1 Introduction

This chapter describes the dual universal asynchronous receiver/transmitters (DUART). It describes the functional operation, the initialization sequence, and the programming details for the DUART registers and features.

11.1.1 Overview

The DUART consists of two universal asynchronous receiver/transmitters (UARTs).

The UARTs act independently; all references to UART refer to one of these receiver/transmitters. Each UART is clocked by the platform (CCB) clock. The DUART programming model is compatible with the PC16552D.

The UART interface is point to point, meaning that only two UART devices are attached to the connecting signals. As shown in [Figure 11-1](#), each UART module consists of the following:

- Receive and transmit buffers
- Clear to send (CTS_B) input port and request to send (RTS_B) output port for data flow control
- 16-bit counter for baud rate generation
- Interrupt control logic

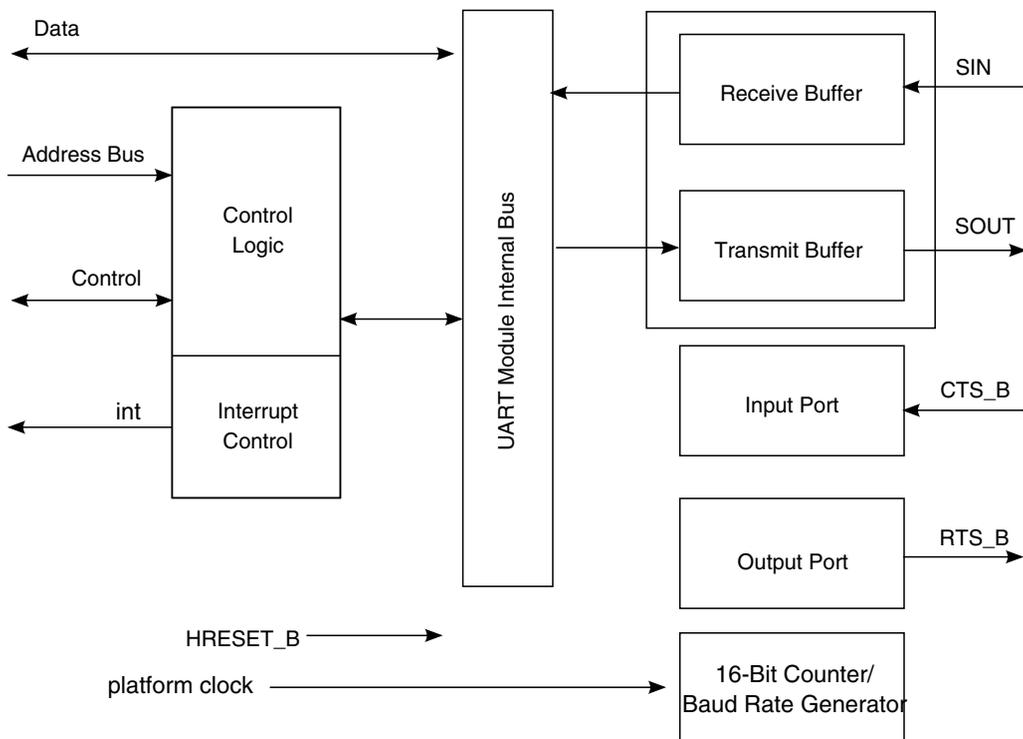


Figure 11-1. UART block diagram

11.1.1.1 Features

The DUART includes these distinctive features:

- Full-duplex operation
- Programming model compatible with original PC16450 UART and PC16550D (improved version of PC16450 that also operates in FIFO mode)
- PC16450 register reset values
- FIFO mode for both transmitter and receiver, providing 16-byte FIFOs
- Serial data encapsulation and decapsulation with standard asynchronous communication bits (START, STOP, and parity)
- Maskable transmit, receive, line status, and modem status interrupts
- Software-programmable baud generators that divide the platform clock by 1 to $(2^{16} - 1)$ and generate a 16x clock for the transmitter and receiver engines
- Clear to send (CTS_B and ready to send (RTS_B) modem control functions
- Software-selectable serial interface data format (data length, parity, 1/1.5/2 STOP bit, baud rate)
- Line and modem status registers
- Line-break detection and generation

- Internal diagnostic support, local loopback, and break functions
- Prioritized interrupt reporting
- Overrun, parity, and framing error detection

11.1.1.2 Modes of operation

The communication channel provides a full-duplex asynchronous receiver and transmitter using an operating frequency derived from the platform clock.

The transmitter accepts parallel data from a write to the transmitter holding register (UTHR). In FIFO mode, the data is placed directly into an internal transmitter shift register of the transmitter FIFO. The transmitter converts the data to a serial bit stream inserting the appropriate start, stop, and optional parity bits. Finally, it outputs a composite serial data stream on the channel transmitter serial data output signal (SOUT). The transmitter status may be polled or interrupt driven.

The receiver accepts serial data bits on the channel receiver serial data input signal (SIN), converts it to parallel format, checks for a start bit, parity (if any), stop bits, and transfers the assembled character (with start, stop, parity bits removed) from the receiver buffer (or FIFO) in response to a read of the UART's receiver buffer register (URBR). The receiver status may be polled or interrupt driven.

11.2 DUART external signal descriptions

The DUART signals are described in [Table 11-1](#). Note that although the actual device signal names are prepended with the UART_ prefix as shown in the table, the functional (abbreviated) signal names are often used throughout this chapter.

Table 11-1. DUART signals-detailed signal descriptions

Signal	I/O	Description	
UART_SIN[0:1]	I	Serial data in. Data is received on the receivers of UART0 and UART1 through the respective serial data input signal, with the least-significant bit received first.	
		State meaning	Asserted/Negated-Represents the data being received on the UART interface.
		Timing	Assertion/Negation-An internal logic sample signal, <i>rxcnt</i> , uses the frequency of the baud-rate generator to sample the data on SIN.

Table continues on the next page...

Table 11-1. DUART signals-detailed signal descriptions (continued)

Signal	I/O	Description
UART_SOUT[0:1]	O	Serial data out. The serial data output signals for the UART0 and UART1 are set ('mark' condition) when the transmitter is disabled, idle, or operating in the local loopback mode. Data is shifted out on these signals, with the least significant bit transmitted first.
		State meaning Asserted/Negated-Represents the data being transmitted on the respective UART interface.
		Timing Assertion/Negation- An internal logic sample signal, <i>rxcnt</i> , uses the frequency of the baud-rate generator to update and drive the data on SOUT.
UART_CTS_B[0:1]	I	Clear to send. These active-low inputs are the clear-to-send inputs. They are connected to the respective RTS_B outputs of the other UART devices on the bus. They can be programmed to generate an interrupt on change-of-state of the signal.
		State meaning Asserted/Negated-Represent the clear to send condition for their respective UART.
		Timing Assertion/Negation-Sampled at the rising edge of every platform clock.
UART_RTS_B[0:1]	O	Request to send. UART_RTS_Bx are active-low output signals that can be programmed to be automatically negated and asserted by either the receiver or transmitter. When connected to the clear-to-send (CTS_B) input of a transmitter, this signal can be used to control serial data flow.
		State meaning Asserted/Negated-Represents the data being transmitted on the respective UART interface.
		Timing Assertion/Negation-Updated and driven at the rising edge of every platform clock.

11.3 DUART memory map/register definition

The table below lists the DUART registers and their offsets. It lists the address, name, and a cross-reference to the complete description of each register. Note that the full register address is comprised of CCSRBAR together with the block base address and offset listed in the table below.

There are two complete sets of DUART registers (one for each UART). The two UARTs on the device are identical, except that the registers for each UART are located at different offsets. Throughout this chapter, the registers are described by a singular acronym: for example, LCR represents the line control register for either UART0 or UART1 .

The registers in each UART interface are used for configuration, control, and status. The divisor latch access bit, ULCR[DLAB], is used to access the divisor latch least- and most-significant bit registers and the alternate function register. Refer to [Line Control Registers \(DUART_ULCR_n\)](#) , for more information on ULCR[DLAB].

All the DUART registers are one-byte wide. Reads and writes to these registers must be byte-wide operations. The table below provides a register summary with references to the section and page that contains detailed information about each register. Undefined byte address spaces within offset 0x000-0xFFF are reserved.

DUART memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
4500	Receiver Buffer Registers (DUART_URBR1)	8	R	00h	11.3.1/506
4500	Transmitter Holding Registers (DUART_UTHR1)	8	W	00h	11.3.2/506
4500	Divisor Least Significant Byte Registers (DUART_UDLB1)	8	R/W	00h	11.3.3/507
4501	Divisor Most Significant Byte Registers (DUART_UDMB1)	8	R/W	00h	11.3.4/508
4501	Interrupt Enable Register (DUART_UIER1)	8	R/W	00h	11.3.5/509
4502	Interrupt ID Registers (DUART_UIIR1)	8	R	01h	11.3.6/510
4502	FIFO Control Registers (DUART_UFCR1)	8	W	00h	11.3.7/511
4502	Alternate Function Registers (DUART_UAFR1)	8	R/W	00h	11.3.8/513
4503	Line Control Registers (DUART_ULCR1)	8	R/W	00h	11.3.9/513
4504	Modem Control Registers (DUART_UMCR1)	8	R/W	00h	11.3.10/515
4505	Line Status Registers (DUART_ULSR1)	8	R	60h	11.3.11/516
4506	Modem Status Registers (DUART_UMSR1)	8	R	00h	11.3.12/517
4507	Scratch Registers (DUART_USCR1)	8	R/W	00h	11.3.13/518
4510	DMA Status Registers (DUART_UDSR1)	8	R	01h	11.3.14/518
4600	Receiver Buffer Registers (DUART_URBR2)	8	R	00h	11.3.1/506
4600	Transmitter Holding Registers (DUART_UTHR2)	8	W	00h	11.3.2/506
4600	Divisor Least Significant Byte Registers (DUART_UDLB2)	8	R/W	00h	11.3.3/507
4601	Divisor Most Significant Byte Registers (DUART_UDMB2)	8	R/W	00h	11.3.4/508
4601	Interrupt Enable Register (DUART_UIER2)	8	R/W	00h	11.3.5/509
4602	Interrupt ID Registers (DUART_UIIR2)	8	R	01h	11.3.6/510
4602	FIFO Control Registers (DUART_UFCR2)	8	W	00h	11.3.7/511
4602	Alternate Function Registers (DUART_UAFR2)	8	R/W	00h	11.3.8/513
4603	Line Control Registers (DUART_ULCR2)	8	R/W	00h	11.3.9/513
4604	Modem Control Registers (DUART_UMCR2)	8	R/W	00h	11.3.10/515
4605	Line Status Registers (DUART_ULSR2)	8	R	60h	11.3.11/516
4606	Modem Status Registers (DUART_UMSR2)	8	R	00h	11.3.12/517
4607	Scratch Registers (DUART_USCR2)	8	R/W	00h	11.3.13/518

Table continues on the next page...

DUART memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
4610	DMA Status Registers (DUART_UDSR2)	8	R	01h	11.3.14/518

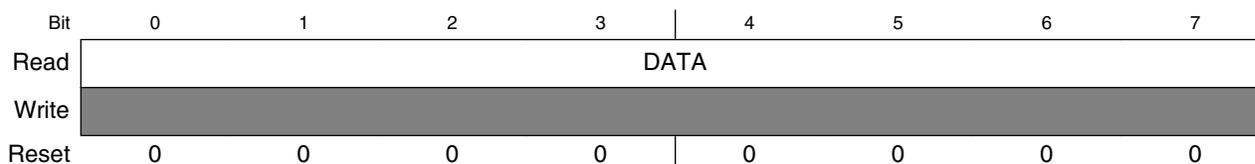
11.3.1 Receiver Buffer Registers (DUART_URBR_n)

This register is accessible when ULCR[DLAB] = 0.

These registers contain the data received from the transmitter on the UART buses. In FIFO mode, when read, they return the first byte received. For FIFO status information, refer to the UDSR[RXRDY] description.

Except for the case when there is an overrun, URBR returns the data in the order it was received from the transmitter. Refer to the ULSR[OE] description, [Line Status Registers \(DUART_ULSR_n\)](#). Note that these registers have same offset as the UTHR_s.

Address: 4000h base + 500h offset + (256d × i), where i=0d to 1d



DUART_URBR_n field descriptions

Field	Description
0-7 DATA	Data received from the transmitter on the UART bus (read only)

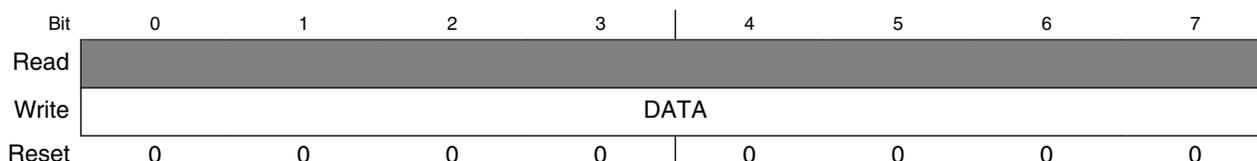
11.3.2 Transmitter Holding Registers (DUART_UTHR_n)

This register is accessible when ULCR[DLAB] = 0.

A write to these 8-bit registers causes the UART devices to transfer 5-8 data bits on the UART bus in the format set up in the ULCR (line control register). In FIFO mode, data written to UTHR is placed into the FIFO. The data written to UTHR is the data sent onto the UART bus, and the first byte written to UTHR is the first byte onto the bus.

UDSR[TXRDY_B] indicates when the FIFO is full. Refer to the tables in [DMA Status Registers \(DUART_UDSR_n\)](#) for more details.

Address: 4000h base + 500h offset + (256d × i), where i=0d to 1d



DUART_UTHR_n field descriptions

Field	Description
0-7 DATA	Data that is written to UTHR (write only)

11.3.3 Divisor Least Significant Byte Registers (DUART_UDLB_n)

This register is accessible when ULCR[DLAB] = 1.

The divisor least significant byte register (UDLB) is concatenated with the divisor most significant byte register (UDMB) to create the divisor used to divide the input clock into the DUART. The output frequency of the baud generator is 16 times the baud rate; therefore the desired baud rate = platform clock frequency ÷ (16 × [UDMB||UDLB]). Equivalently, [UDMB||UDLB:0b0000] = platform clock frequency ÷ desired baud rate. Baud rates that can be generated by specific input clock frequencies are shown in the table below.

The following table shows examples of baud rate generation based on common input clock frequencies. Many other target baud rates are also possible. Note that because only integer values can be used as divisors, the actual baud rate differs slightly from the desired (target) baud rate; for this reason, both target and actual baud rates are given, along with the percentage of error.

Table 11-7. Baud Rate Examples

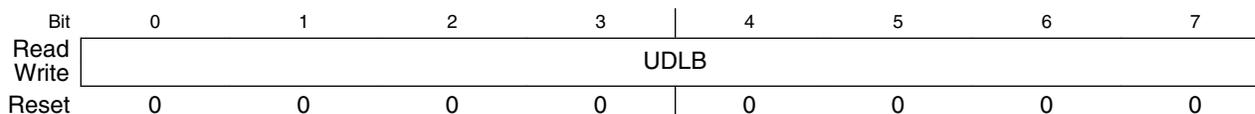
Target Baud Rate (Decimal)	Divisor		Platform Clock (CCB) Frequency (MHz)	Actual Baud Rate (Decimal)	Percent Error (Decimal)
	Decimal	Hex			
9,600	1736	6C8	266	9600.61444	0.0064
19,200	868	364	266	19,201.22888	0.0064
38,400	434	1B2	266	38,402.45776	0.0064
57,600	289	121	266	57,670.12673	0.1217
115,200	145	91	266	114,942.52844	0.2235
230,400	72	48	266	231,481.48090	0.4694
9,600	2170	87A	333	9600.61444	0.0064
19,200	1085	43D	333	19,201.22888	0.0064

Table continues on the next page...

Table 11-7. Baud Rate Examples (continued)

Target Baud Rate (Decimal)	Divisor		Platform Clock (CCB) Frequency (MHz)	Actual Baud Rate (Decimal)	Percent Error (Decimal)
	Decimal	Hex			
38,400	543	21F	333	38,367.09638	0.0858
57,600	362	16A	333	57,550.64457	0.0857
115,200	181	B5	333	115,101.28913	0.0857
230,400	90	5A	333	231,481.48148	0.4694
9,600	3472	D90	533	9600.61444	0.0064
19,200	1736	6C8	533	19,201.22888	0.0064
38,400	868	364	533	38,402.45776	0.0064
57,600	579	243	533	57,570.52389	0.0512
115,200	289	121	533	115,340.25375	0.1217
230,400	145	91	533	229,885.05747	0.2235
9,600	3906	F42	600	9600.6144	0.0064
19,200	1953	7A1	600	19,201.2288	0.0064
38,400	977	3D1	600	38382.8045	0.0448
57,600	651	28B	600	57603.6866	0.0064
115,200	326	146	600	115030.6748	0.1470
230,400	163	A3	600	230061.3497	0.1470

Address: 4000h base + 500h offset + (256d × i), where i=0d to 1d



DUART_UDLBn field descriptions

Field	Description
0-7 UDLB	Divisor least significant byte. This is concatenated with UDMB.

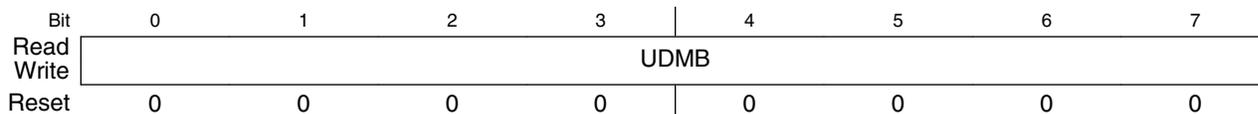
11.3.4 Divisor Most Significant Byte Registers (DUART_UDMBn)

This register is accessible when ULCR[DLAB] = 1.

The divisor least significant byte register (UDLB) is concatenated with the divisor most significant byte register (UDMB) to create the divisor used to divide the input clock into the DUART. The output frequency of the baud generator is 16 times the baud rate; therefore the desired baud rate = platform clock frequency ÷ (16 × [UDMB||UDLB]).

Equivalently, [UDMB||UDLB:0b0000] = platform clock frequency ÷ desired baud rate. Baud rates that can be generated by specific input clock frequencies are shown in the table in [Divisor Least Significant Byte Registers \(DUART_UDLBn\)](#) .

Address: 4000h base + 501h offset + (256d × i), where i=0d to 1d



DUART_UDMBn field descriptions

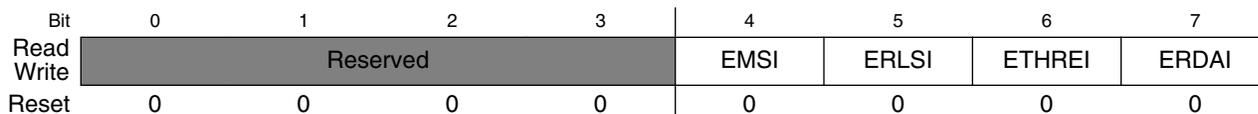
Field	Description
0-7 UDMB	Divisor most significant byte

11.3.5 Interrupt Enable Register (DUART_UIERn)

This register is accessible when ULCR[DLAB] = 0.

The UIER gives the user the ability to mask specific UART interrupts to the programmable interrupt controller (PIC).

Address: 4000h base + 501h offset + (256d × i), where i=0d to 1d



DUART_UIERn field descriptions

Field	Description
0-3 -	This field is reserved. Reserved.
4 EMSI	Enable modem status interrupt. 0 Mask interrupts caused by UMSR[DCTS] being set 1 Enable and assert interrupts when the clear-to-send bit in the UART modem status register (UMSR) changes state
5 ERLSI	Enable receiver line status interrupt. 0 Mask interrupts when ULSR's overrun, parity error, framing error or break interrupt bits are set 1 Enable and assert interrupts when ULSR's overrun, parity error, framing error or break interrupt bits are set
6 ETHREI	Enable transmitter holding register empty interrupt. 0 Mask interrupt when ULSR[THRE] is set 1 Enable and assert interrupts when ULSR[THRE] is set

Table continues on the next page...

DUART_UIERn field descriptions (continued)

Field	Description
7 ERDAI	Enable received data available interrupt. 0 Mask interrupt when new receive data is available or receive data time out has occurred 1 Enable and assert interrupts when a new data character is received from the external device and/or a time-out interrupt occurs in the FIFO mode

11.3.6 Interrupt ID Registers (DUART_UIIRn)

This register is accessible when ULCR[DLAB] = 0.

The UIIRs indicate when an interrupt is pending from the corresponding UART and what type of interrupt is active. They also indicate if the FIFOs are enabled.

The DUART prioritizes interrupts into four levels and records these in the corresponding UIIR. The four levels of interrupt conditions in order of priority are:

1. Receiver line status
2. Received data ready/character time-out
3. Transmitter holding register empty
4. Modem status

When the UIIR is read, the associated DUART serial channel freezes all interrupts and indicates the highest priority pending interrupt. While this read transaction is occurring, the associated DUART serial channel records new interrupts, but does not change the contents of UIIR until the read access is complete.

Table 11-15. UIIR IID Bits Summary

IID Bits IID[3-0]	Priority Level	Interrupt Type	Interrupt Description	How To Reset Interrupt
0b0001	-	-	-	-
0b0110	Highest	Receiver line status	Overrun error, parity error, framing error, or break interrupt	Read the line status register.
0b0100	Second	Received data available	Receiver data available or trigger level reached in FIFO mode	Read the receiver buffer register or interrupt is automatically reset if the number of bytes in the receiver FIFO drops below the trigger level.

Table continues on the next page...

Table 11-15. UIIR IID Bits Summary (continued)

IID Bits IID[3-0]	Priority Level	Interrupt Type	Interrupt Description	How To Reset Interrupt
0b1100	Second	Character time-out	No characters have been removed from or input to the receiver FIFO during the last 4 character times and there is at least one character in the receiver FIFO during this time.	Read the receiver buffer register.
0b0010	Third	UTHR empty	Transmitter holding register is empty	Read the UIIR or write to the UTHR.
0b0000	Fourth	Modem status	CTS_B input value changed since last read of UMSR	Read the UMSR.

Address: 4000h base + 502h offset + (256d × i), where i=0d to 1d



DUART_UIIRn field descriptions

Field	Description
0–1 FE	FIFOs enabled. Reflects the setting of UFCR[FEN]
2–3 -	This field is reserved. Reserved
4 IID3	Interrupt ID bits identify the highest priority interrupt that is pending as indicated in Table 11-15 . IID3 is set along with IID2 only when a timeout interrupt is pending for FIFO mode.
5 IID2	Interrupt ID bits identify the highest priority interrupt that is pending as indicated in Table 11-15 .
6 IID1	Interrupt ID bits identify the highest priority interrupt that is pending as indicated in Table 11-15 .
7 IID0	IID0 indicates when an interrupt is pending. 0 The UART has an active interrupt ready to be serviced. 1 No interrupt is pending.

11.3.7 FIFO Control Registers (DUART_UFCRn)

This register is accessible when ULCR[DLAB] = 0.

DUART memory map/register definition

The UFCR, a write-only register, is used to enable and clear the receiver and transmitter FIFOs, set a receiver FIFO trigger level to control the received data available interrupt, and select the type of DMA signaling.

When the UFCR bits are written, the FIFO enable bit must also be set or else the UFCR bits are not programmed. When changing from FIFO mode to 16450 mode (non-FIFO mode) and vice versa, data is automatically cleared from the FIFOs.

After all the bytes in the receiver FIFO are cleared, the receiver internal shift register is not cleared. Similarly, the bytes are cleared in the transmitter FIFO, but the transmitter internal shift register is not cleared. Both TFR and RFR are self-clearing bits.

Address: 4000h base + 502h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7
Read	Reserved			Reserved				
Write	RTL		Reserved		DMS	TFR	RFR	FEN
Reset	0	0	0	0	0	0	0	0

DUART_UFCR_n field descriptions

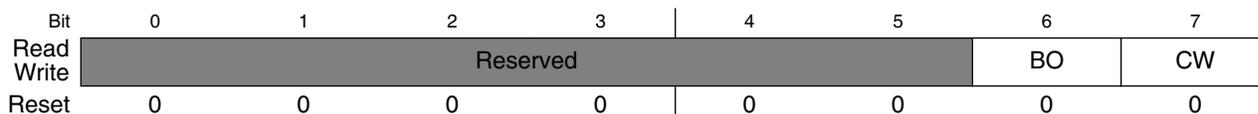
Field	Description
0–1 RTL	Receiver trigger level. A received data available interrupt occurs when UIER[ERDAI] is set and the number of bytes in the receiver FIFO equals the designated interrupt trigger level as follows: 00 1 byte 01 4 bytes 10 8 bytes 11 14 bytes
2–3 -	This field is reserved. Reserved
4 DMS	DMA mode select. See DMA mode select for more information. 0 UDSR[RXRDY] and UDSR[TXRDY] bits are in mode 0. 1 UDSR[RXRDY] and UDSR[TXRDY] bits are in mode 1 if UFCR[FEN] = 1.
5 TFR	Transmitter FIFO reset 0 No action 1 Clears all bytes in the transmitter FIFO and resets the FIFO counter/pointer to 0
6 RFR	Receiver FIFO reset 0 No action 1 Clears all bytes in the receiver FIFO and resets the FIFO counter/pointer to 0
7 FEN	FIFO enable 0 FIFOs are disabled and cleared 1 Enables the transmitter and receiver FIFOs

11.3.8 Alternate Function Registers (DUART_UAFR_n)

This register is accessible when ULCR[DLAB] = 1.

The UAFRs give software the ability to gate off the baud clock and write to both UART0/UART1 registers simultaneously with the same write operation.

Address: 4000h base + 502h offset + (256d × i), where i=0d to 1d



DUART_UAFR_n field descriptions

Field	Description
0–5 -	This field is reserved. Reserved.
6 BO	Baud clock select. 0 The baud clock is not gated off. 1 The baud clock is gated off.
7 CW	Concurrent write enable. 0 Disables writing to both UART0 and UART1 1 Enables concurrent writes to corresponding UART registers. A write to a register in UART0 is also a write to the corresponding register in UART1 and vice versa . The user needs to ensure that the LCR[DLAB] of both UARTs are in the same state before executing a concurrent write to register addresses 0x n 00, 0x n 01 and 0x n 02, where n is the offset of the corresponding UART.

11.3.9 Line Control Registers (DUART_ULCR_n)

This register is accessible when ULCR[DLAB] = x.

The ULCRs specify the data format for the UART bus and set the divisor latch access bit ULCR[DLAB], which controls the ability to access the divisor latch least and most significant bit registers and the alternate function register.

After initializing the ULCR, the software should not re-write the ULCR when valid transfers on the UART bus are active. The software should not re-write the ULCR until the last STOP bit has been received and there are no new characters being transferred on the bus.

The stick parity bit, ULCR[SP], assigns a set parity value for the parity bit time slot sent on the UART bus. The set value is defined as mark parity (logic 1) or space parity (logic 0). ULCR[PEN] and ULCR[EPS] help determine the set parity value. See [Table 11-23](#)

DUART memory map/register definition

for more information. ULCR[NSTB], defines the number of STOP bits to be sent at the end of the data transfer. The receiver only checks the first STOP bit, regardless of the number of STOP bits selected. The word length select bits (1 and 0) define the number of data bits that are transmitted or received as a serial character. The word length does not include START, parity, and STOP bits.

Table 11-23. Parity Selection Using ULCR[PEN], ULCR[SP], and ULCR[EPS]

PEN	SP	EPS	Parity Selected
0	0	0	No parity
0	0	1	No parity
0	1	0	No parity
0	1	1	No parity
1	0	0	Odd parity
1	0	1	Even parity
1	1	0	Mark parity
1	1	1	Space parity

Address: 4000h base + 503h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7
Read	DLAB	SB	SP	EPS	PEN	NSTB	WLS	
Write								
Reset	0	0	0	0	0	0	0	0

DUART_ULCRn field descriptions

Field	Description
0 DLAB	Divisor latch access bit. 0 Access to all registers except UDLB, UAFR, and UDMB 1 Ability to access divisor latch least and most significant byte registers and alternate function register (UAFR)
1 SB	Set break. 0 Send normal UTHR data onto the serial output (SOUT) signal 1 Force logic 0 to be on the SOUT signal. Data in the UTHR is not affected
2 SP	Stick parity. 0 Stick parity is disabled. 1 If PEN = 1 and EPS = 1, space parity is selected. And if PEN = 1 and EPS = 0, mark parity is selected.
3 EPS	Even parity select. See Table 11-23 for more information. 0 If PEN = 1 and SP = 0, odd parity is selected. 1 If PEN = 1 and SP = 0, even parity is selected.
4 PEN	Parity enable.

Table continues on the next page...

DUART_ULCRn field descriptions (continued)

Field	Description
	0 No parity generation and checking 1 Generate parity bit as a transmitter, and check parity as a receiver
5 NSTB	Number of STOP bits. 0 One STOP bit is generated in the transmitted data. 1 When a 5-bit data length is selected, 1½ STOP bits are generated. When either a 6-, 7-, or 8-bit word length is selected, two STOP bits are generated.
6–7 WLS	Word length select. Number of bits that comprise the character length. The word length select values are as follows: 00 5 bits 01 6 bits 10 7 bits 11 8 bits

11.3.10 Modem Control Registers (DUART_UMCRn)

This register is accessible when ULCR[DLAB] = x.

The UMCRs control the interface with the external peripheral device on the UART bus.

Address: 4000h base + 504h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7
Read				LOOP			RTS	
Write								
Reset	0	0	0	0	0	0	0	0

DUART_UMCRn field descriptions

Field	Description
0–2 -	This field is reserved. Reserved.
3 LOOP	Local loopback mode. 0 Normal operation 1 Functionally, the data written to UTHR can be read from URBR of the same UART , and UMCR[RTS] is tied to UMSR[CTS] .
4–5 -	This field is reserved. Reserved.
6 RTS	Ready to send. 0 Negates corresponding RTS_B output 1 Assert corresponding RTS_B output. Informs external modem or peripheral that the UART is ready for sending/receiving data
7 -	This field is reserved. Reserved.

11.3.11 Line Status Registers (DUART_ULSRn)

This register is accessible when ULCR[DLAB] = x.

The ULSRs are read-only registers that monitor the status of the data transfer on the UART buses. To isolate the status bits from the proper character received through the UART bus, software should read the ULSR and then the URBR.

Address: 4000h base + 505h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7
Read	RFE	TEMT	THRE	BI	FE	PE	OE	DR
Write								
Reset	0	1	1	0	0	0	0	0

DUART_ULSRn field descriptions

Field	Description
0 RFE	Receiver FIFO error. 0 This bit is cleared when there are no errors in the receiver FIFO or on a read of the ULSR with no remaining receiver FIFO errors. 1 Set to one when one of the characters in the receiver FIFO encounters an error (framing, parity, or break interrupt)
1 TEMT	Transmitter empty. 0 Either or both the UTHR or the internal transmitter shift register has a data character. In FIFO mode, a data character is in the transmitter FIFO or the internal transmitter shift register. 1 Both the UTHR and the internal transmitter shift register are empty. In FIFO mode, both the transmitter FIFO and the internal transmitter shift register are empty.
2 THRE	Transmitter holding register empty. 0 The UTHR is not empty. 1 A data character has transferred from the UTHR into the internal transmitter shift register. In FIFO mode, the transmitter FIFO contains no data character.
3 BI	Break interrupt. 0 This bit is cleared when the ULSR is read or when a valid data transfer is detected (that is, STOP bit is received). 1 Received data of logic 0 for more than START bit + Data bits + Parity bit + one STOP bits length of time. A new character is not loaded until SIN returns to the mark state (logic 1) and a valid START is detected. In FIFO mode, a zero character is encountered in the FIFO (the zero character is at the top of the FIFO). In FIFO mode, only one zero character is stored.
4 FE	Framing error. 0 This bit is cleared when ULSR is read or when a new character is loaded into the URBR from the receiver shift register. 1 Invalid STOP bit for receive data (only the first STOP bit is checked). In FIFO mode, this bit is set when the character that detected a framing error is encountered in the FIFO (that is the character at the top of the FIFO). An attempt to resynchronize occurs after a framing error. The UART assumes that the framing error (due to a logic 0 being read when a logic 1 (STOP) was expected) was due to a

Table continues on the next page...

DUART_ULSR_n field descriptions (continued)

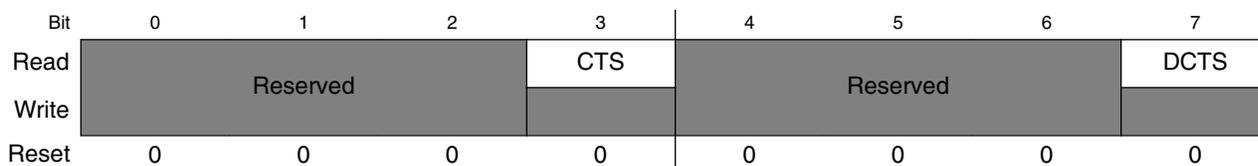
Field	Description
	STOP bit overlapping with the next START bit, so it assumes this logic 0 sample is a true START bit and then receives the following new data.
5 PE	Parity error. 0 This bit is cleared when ULSR is read or when a new character is loaded into the URBR. 1 Unexpected parity value encountered when receiving data. In FIFO mode, the character with the error is at the top of the FIFO .
6 OE	Overrun error. 0 This bit is cleared when ULSR is read. 1 Before the URBR is read, the URBR was overwritten with a new character. The old character is loss. In FIFO mode, the receiver FIFO is full (regardless of the receiver FIFO trigger level setting) and a new character has been received into the internal receiver shift register. The old character was overwritten by the new character. Data in the receiver FIFO was not overwritten.
7 DR	Data ready. 0 This bit is cleared when URBR is read or when all of the data in the receiver FIFO is read. 1 A character has been received in the URBR or the receiver FIFO.

11.3.12 Modem Status Registers (DUART_UMSR_n)

This register is accessible when ULCR[DLAB] = x.

The UMSRs track the status of the modem (or external peripheral device) clear to send (CTS_B) signal for the corresponding UART .

Address: 4000h base + 506h offset + (256d × i), where i=0d to 1d



DUART_UMSR_n field descriptions

Field	Description
0–2 -	This field is reserved. Reserved.
3 CTS	Clear to send. Represents the inverted value of the CTS_B input pin from the external peripheral device 0 Corresponding CTS_B <i>n</i> is negated 1 Corresponding CTS_B <i>n</i> is asserted. The modem or peripheral device is ready for data transfers.
4–6 -	This field is reserved. Reserved.
7 DCTS	Clear to send.

Table continues on the next page...

DUART_UMSR n field descriptions (continued)

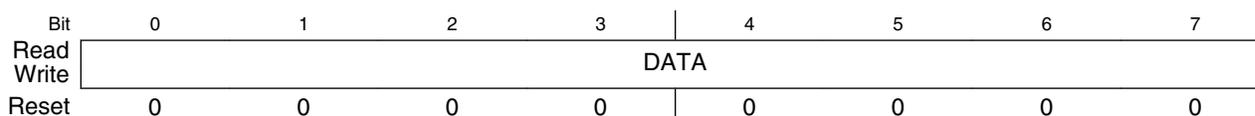
Field	Description
0	No change on the corresponding CTS_B n signal since the last read of UMSR[CTS]
1	The CTS_B n value has changed, since the last read of UMSR[CTS]. Causes an interrupt if UIER[EMSI] is set to detect this condition

11.3.13 Scratch Registers (DUART_USCR n)

This register is accessible when ULCR[DLAB] = x.

The USCR registers are for debugging software or the DUART hardware. The USCRs do not affect the operation of the DUART.

Address: 4000h base + 507h offset + (256d × i), where i=0d to 1d



DUART_USCR n field descriptions

Field	Description
0–7 DATA	Data

11.3.14 DMA Status Registers (DUART_UDSR n)

This register is accessible when ULCR[DLAB] = x.

The DMA status registers (UDSRs) are read-only registers that return transmitter and receiver FIFO status. UDSRs also provide the ability to assist DMA data operations to and from the FIFOs.

Table 11-35. UDSR[TXRDY] Set Conditions

DMS	FEN	DMA Mode	Meaning
0	0	0	TXRDY is set after the first character is loaded into the transmitter FIFO or UTHR.
0	1	0	
1	0	0	
1	1	1	TXRDY is set when the transmitter FIFO is full.

Table 11-36. UDSR[TXRDY] Cleared Conditions

DMS	FEN	DMA Mode	Meaning
0	0	0	TXRDY is cleared when there are no characters in the transmitter FIFO or UTHR.
0	1	0	
1	0	0	
1	1	1	TXRDY is cleared when there are no characters in the transmitter FIFO or UTHR. TXRDY remains clear when the transmitter FIFO is not yet full.

Table 11-37. UDSR[RXRDY] Set Conditions

DMS	FEN	DMA Mode	Meaning
0	0	0	RXRDY is set when there are no characters in the receiver FIFO or URBR.
0	1	0	
1	0	0	
1	1	1	RXRDY is set when the trigger level has not been reached and there has been no time out.

Table 11-38. UDSR[RXRDY] Cleared Conditions

DMS	FEN	DMA Mode	Meaning
0	0	0	RXRDY is cleared when there is at least one character in the receiver FIFO or URBR.
0	1	0	
1	0	0	
1	1	1	RXRDY is cleared when the trigger level or a time-out has been reached. RXRDY remains cleared until the receiver FIFO is empty.

Address: 4000h base + 510h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7
Read	Reserved						TXRDY	RXRDY
Write								
Reset	0	0	0	0	0	0	0	1

DUART_UDSRn field descriptions

Field	Description
0–5 -	This field is reserved. Reserved
6 TXRDY	Transmitter ready. This read-only bit reflects the status of the transmitter FIFO or the UTHR. The status depends on the DMA mode selected, which is determined by the DMS and FEN bits in the UFCR. 0 This bit is cleared, as shown in Table 11-36 . 1 This bit is set, as shown in Table 11-35 .
7 RXRDY	Receiver ready. This read-only bit reflects the status of the receiver FIFO or URBR. The status depends on the DMA mode selected, which is determined by the DMS and FEN bits in the UFCR. 0 This bit is cleared, as shown in Table 11-38 . 1 This bit is set, as shown in Table 11-37 .

11.4 Functional description

The communication channel provides a full-duplex asynchronous receiver and transmitter using an operating frequency derived from the platform clock signal.

The transmitter accepts parallel data with a write access to the transmitter holding register (UTHR). In FIFO mode, the data is placed directly into an internal transmitter shift register, or into the transmitter FIFO-see [FIFO mode](#). The transmitting registers convert the data to a serial bit stream, by inserting the appropriate START, STOP, and optional parity bits. Finally, the registers output a composite serial data stream on the channel transmitter serial data output (SOUT). The transmitter status may be polled or interrupt-driven.

The receiver accepts serial data on the channel receiver serial data input (SIN), converts the data into parallel format, and checks for START, STOP, and parity bits. In FIFO mode, the receiver removes the START, STOP, and parity bits and then transfers the assembled character from the receiver buffer, or receiver FIFO. This transfer occurs in response to a read of the UART receiver buffer register (URBR). The receiver status may be polled or interrupt driven.

11.4.1 Serial interface

The UART bus is a serial, full-duplex, point-to-point bus as shown in [Figure 11-44](#). Therefore, only two devices are attached to the same signals and there is no need for address or arbitration bus cycles.

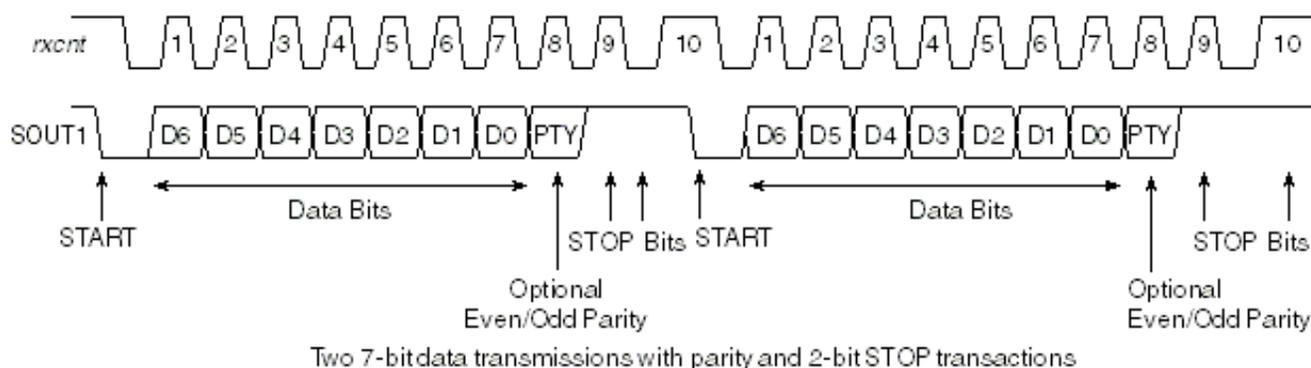


Figure 11-44. UART bus interface transaction protocol example

A standard UART bus transfer is composed of either three or four parts:

- START bit
- Data transfer bits (least-significant bit is first data bit on the bus)
- Parity bit (optional)
- STOP bits

An internal logic sample signal, *rxcnt*, uses the frequency of the baud-rate generator to drive the bits on SOUT.

The following sections describe the four components of the serial interface, the baud-rate generator, local loopback mode, different errors, and FIFO mode.

11.4.1.1 START bit

A write to the transmitter holding register (UTHR) generates a START bit on the SOUT signal.

Figure 11-44 shows that the START bit is defined as a logic 0. The START bit denotes the beginning of a new data transfer which is limited to the bit length programmed in the UART line control register (ULCR). When the bus is idle, SOUT is high.

11.4.1.2 Data transfer

Each data transfer contains 5-8 bits of data.

The ULCR data bit length for the transmitter and receiver UART devices must agree before a transfer begins; otherwise, a parity or framing error may occur. A transfer begins when UTHR is written. At that time a START bit is generated followed by 5-8 of the data

bits previously written to the UTHR. The data bits are driven from the least significant to the most significant bits. After the parity and STOP bits, a new data transfer can begin if new data is written to the UTHR.

11.4.1.3 Parity bit

The user has the option of using even, odd, no parity, or stick parity.

See [Line Control Registers \(DUART_ULCR \$n\$ \)](#). Both the receiver and transmitter parity definition must agree before attempting to transfer data. When receiving data a parity error can occur if an unexpected parity value is detected. See [Line Status Registers \(DUART_ULSR \$n\$ \)](#).

11.4.1.4 STOP bit

The transmitter device ends the write transfer by generating a STOP bit.

The STOP bit is always high. The user can program the length of the STOP bit(s) in the ULCR. Both the receiver and transmitter STOP bit length must agree before attempting to transfer data. A framing error can occur if an invalid STOP bit is detected.

11.4.2 Baud-rate generator logic

Each UART contains an independent programmable baud-rate generator, that is capable of taking the platform clock input and dividing the input by any divisor from 1 to $2^{16} - 1$.

The baud rate is defined as the number of bits per second that can be sent over the UART bus. The formula for calculating baud rate is as follows:

$$\text{Baud rate} = (1/16) \times (\text{platform clock frequency} \div \text{divisor value})$$

Therefore, the output frequency of the baud-rate generator is 16 times the baud rate.

The divisor value is determined by the following two 8-bit registers to form a 16-bit binary number:

- UART divisor most significant byte register (UDMB)
- UART divisor least significant byte register (UDLB)

Upon loading either of the divisor latches, a 16-bit baud-rate counter is loaded.

The divisor latches must be loaded during initialization to ensure proper operation of the baud-rate generator. Both UART devices on the same bus must be programmed for the same baud-rate before starting a transfer.

The baud clock can be passed to the performance monitor by enabling the UAFR[BO] bit. This can be used to determine baud rate errors.

11.4.3 Local loopback mode

Local loopback mode is provided for diagnostic testing.

The data written to UTHR can be read from the receiver buffer register (URBR) of the same UART. In this mode, the modem control register UMCR[RTS] is internally tied to the modem status register UMSR[CTS]. The transmitter SOUT is set to a logic 1 and the receiver SIN is disconnected. The output of the transmitter shift register is looped back into the receiver shift register input. The CTS_B (input signal) is disconnected, RTS_B is internally connected to CTS_B, and the RTS_B (output signal) becomes inactive. In this diagnostic mode, data that is transmitted is immediately received. In local loopback mode the transmit and receive data paths of the DUART can be verified. Note that in local loopback mode, the transmit/receive interrupts are fully operational and can be controlled by the interrupt enable register (UIER).

11.4.4 Errors

The following sections describe framing, parity, and overrun errors which may occur while data is transferred on the UART bus.

Each of the error bits are usually cleared, as described below, when the line status register (ULSR) is read.

11.4.4.1 Framing error

When an invalid STOP bit is detected, a framing error occurs and ULSR[FE] is set.

Note that only the first STOP bit is checked. In FIFO mode, ULSR[FE] is set when the character at the top of the FIFO detects a framing error. An attempt to re-synchronize occurs after a framing error. The UART assumes that the framing error (due to a logic 0 being read when a logic 1 (STOP) was expected) was due to a STOP bit overlapping with the next START bit. ULSR[FE] is cleared when ULSR is read or when a new character is loaded into the URBR from the receiver shift register.

11.4.4.2 Parity error

A parity error occurs, and ULSR[PE] is set, when unexpected parity values are encountered while receiving data.

In FIFO mode, ULSR[PE] is set when the character with the error is at the top of the FIFO. ULSR[PE] is cleared when ULSR is read or when a new character is loaded into the URBR.

11.4.4.3 Overrun error

When a new (overwriting character) STOP bit is detected and the old character is lost, an overrun error occurs and ULSR[OE] is set.

In FIFO mode, ULSR[OE] is set after the receiver FIFO is full (despite the receiver FIFO trigger level setting) and a new character has been received into the internal receiver shift register. Data in the FIFO is not overwritten; only the shift register data is overwritten. Therefore, the interrupt occurs immediately. ULSR[OE] is cleared when ULSR is read.

11.4.5 FIFO mode

The UARTs use an alternate mode (FIFO mode) to relieve the processor core from excessive software overhead.

The FIFO control register (UFCR) is used to enable and clear the receiver and transmitter FIFOs and set the FIFO receiver trigger level UFCR[RTL] to control the received data available interrupt UIER[ERDAI].

The UFCR also selects the type of DMA signaling. The UDSR[RXRDY] indicates the status of the receiver FIFO. The DMA status registers (UDSR[TXRDY]) indicate when the transmitter FIFO is full. When in FIFO mode, data written to UTHR is placed into the transmitter FIFO. The first byte written to UTHR is the first byte onto the UART bus.

11.4.5.1 FIFO interrupts

In FIFO mode, the UIER[ERDAI] is set when a time-out interrupt occurs. When a receive data time-out occurs there is a maskable interrupt condition (through UIER[ERDAI]). See [Interrupt Enable Register \(DUART_UIER \$n\$ \)](#), for more details on interrupt enables.

The interrupt ID register (UIIR) indicates if the FIFOs are enabled. Interrupt ID3 UIIR[IID3] bit is only set for FIFO mode interrupts. The character time-out interrupt occurs when no characters have been removed from or input to the receiver FIFO during the last four character times and there is at least one character in the receiver FIFO during this time. The character time-out interrupt (controlled by UIIR[IID n]) is cleared when the URBR is read. See [Interrupt ID Registers \(DUART_UIIR \$n\$ \)](#), for more information.

The UIIR[FE] bits indicate if FIFO mode is enabled.

11.4.5.2 DMA mode select

The UDSR[RXRDY] bit reflects the status of the receiver FIFO or URBR.

In mode 0 (UFCR[DMS] is cleared), UDSR[RXRDY] is cleared when there is at least one character in the receiver FIFO or URBR and it is set when there are no more characters in the receiver FIFO or URBR. This occurs regardless of the setting of the UFCR[FEN] bit. In mode 1 (UFCR[DMS] and UFCR[FEN] are set), UDSR[RXRDY] is cleared when the trigger level or a time-out has been reached and it is set when there are no more characters in the receiver FIFO.

The UDSR[TXRDY] bit reflects the status of the transmitter FIFO or UTHR. In mode 0 (UFCR[DMS] is cleared), UDSR[TXRDY] is cleared when there are no characters in the transmitter FIFO or UTHR and it is set after the first character is loaded into the transmitter FIFO or UTHR. This occurs regardless of the setting of the UFCR[FEN] bit. In mode 1 (UFCR[DMS] and UFCR[FEN] are set), UDSR[TXRDY] is cleared when there are no characters in the transmitter FIFO or UTHR and it is set when the transmitter FIFO is full.

See [DMA Status Registers \(DUART_UDSR \$n\$ \)](#), for a complete description of the UDSR[RXRDY] and UDSR[TXRDY] bits.

11.4.5.3 Interrupt control logic

An interrupt is active when DUART interrupt ID register bit 7 (UIIR[IID0]), is cleared.

The interrupt enable register (UIER) is used to mask specific interrupt types. For more details refer to the description of UIER in [Interrupt Enable Register \(DUART_UIER \$n\$ \)](#).

When the interrupts are disabled in UIER, polling software cannot use UIIR[IID0] to determine whether the UART is ready for service. The software must monitor the appropriate bits in the line status (ULSR) and/or the modem status (UMSR) registers. UIIR[IID0] can be used for polling if the interrupts are enabled in UIER.

11.5 DUART initialization/application information

The following requirements must be met for DUART accesses:

- All DUART registers must be mapped to a cache-inhibited and guarded area. (That is, the WIMG setting in the MMU needs to be 0b01X1.)
- All DUART registers are 1 byte wide. Reads and writes to these registers must be byte-wide operations.

A system reset puts the DUART registers to a default state. Before the interface can transfer serial data, the following initialization steps are recommended:

1. Update the programmable interrupt controller (PIC) DUART channel interrupt vector source registers.
2. Set data attributes and control bits in the ULCR, UFCR, UAFR, UMCR, UDLB, and UDMB.
3. Set the data attributes and control bits of the external modem or peripheral device.
4. Set the interrupt enable register (UIER).
5. To start a write transfer, write to the UTHR.
6. Poll UIIR if the interrupts generated by the DUART are masked.

Chapter 12

Enhanced local bus controller (eLBC)

This chapter describes the external signals and the memory-mapped registers as well as a functional description of the general-purpose chip-select machine (GPCM), NAND Flash control machine (FCM), and user-programmable machines (UPMs) of the eLBC. Finally, it includes an initialization and applications information section with many specific examples of its use.

12.1 eLBC introduction

This chapter describes the enhanced local bus controller (eLBC) block.

It describes the external signals and the memory-mapped registers as well as a functional description of the general-purpose chip-select machine (GPCM), NAND Flash control machine (FCM), and user-programmable machines (UPMs) of the eLBC. Finally, it includes an initialization and applications information section with many specific examples of its use.

The figure below is a functional block diagram of the eLBC, which supports three interfaces: GPCM, FCM, and UPM controllers. Note that PLL signals, LSYNC_OUT and LSYNC_IN are also part of this block diagram.

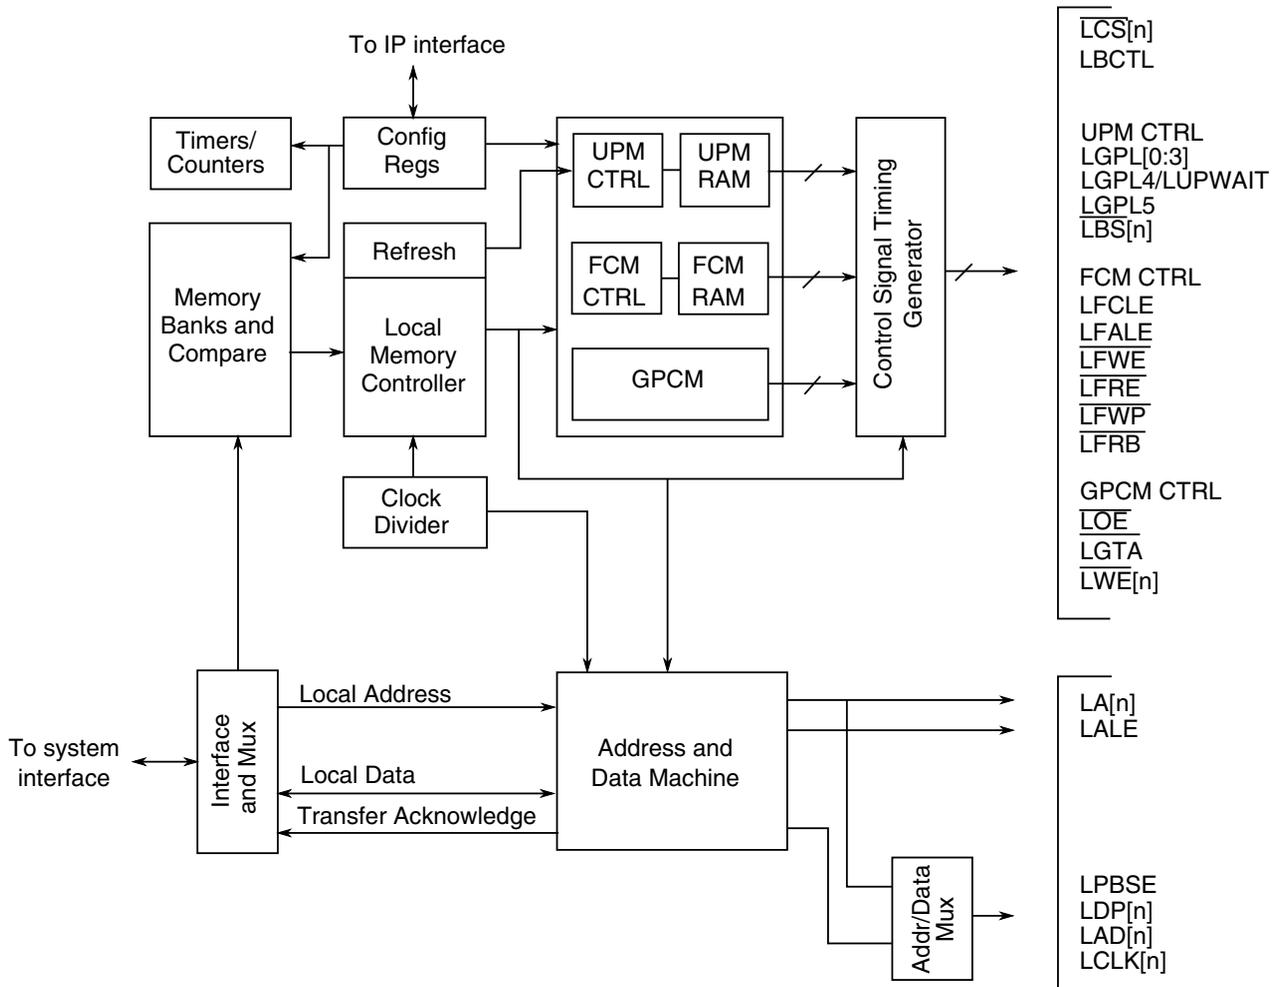


Figure 12-1. Enhanced local bus controller block diagram

12.1.1 Overview

The main component of the eLBC is its memory controller, which provides a seamless 16-bit interface to many types of memory devices and peripherals.

The memory controller is responsible for controlling eight memory banks shared by a GPCM, an FCM, and up to three UPMs. As such, it supports a minimal glue logic interface to SRAM, EPROM, NOR Flash EEPROM, NAND Flash EEPROM, burstable RAM, regular DRAM devices, extended data output DRAM devices, and other peripherals. The external address latch signal (LALE) allows multiplexing of addresses with data signals to reduce the device pin count.

The eLBC also includes a number of data checking and protection features such as data parity generation and checking, write protection, and a bus monitor to ensure that each bus cycle is terminated within a user-specified period.

12.1.2 Features

The eLBC main features are as follows:

- Memory controller with eight memory banks
 - 32-bit address decoding with mask
 - Variable memory block sizes (32 Kbytes to 4 Gbytes)
 - Selection of control signal generation on a per-bank basis
 - Data buffer controls activated on a per-bank basis
 - Automatic segmentation of large transactions into memory accesses optimized for bus width and addressing capability
 - Odd/even parity checking for single access
 - Write-protection capability
 - Parity byte-select
- General-purpose chip-select machine (GPCM)
 - Compatible with SRAM, EPROM, NOR Flash EEPROM, and peripherals
 - Global (boot) chip-select available at system reset
 - Boot chip-select support for 8- and 16-bit devices
 - Minimum three-clock access to external devices
 - Two byte-write-enable signals (LWE_B[0:1])
 - Output enable signal (LOE_B)
 - External access termination signal (LGTA_B)
- NAND Flash control machine (FCM)
 - Compatible with small (512 + 16 bytes) and large (2048 + 64 bytes) page parallel NAND Flash EEPROM
 - Global (boot) chip-select available at system reset, with 4-Kbyte boot block buffer for execute-in-place boot loading
 - ECC checking enable/disable feature supported during boot
 - Read-only ECC registers to verify after write operation
 - Boot chip-select support for 8-bit devices
 - Dual 2-Kbyte/eight 512-byte buffers allow simultaneous data transfer during Flash reads and programming
 - Interrupt-driven block transfer for reads and writes
 - Programmable command and data transfer sequences of up to eight steps supported
 - Generic command and address registers support proprietary Flash interfaces
 - Block write locking to ensure system security and integrity

- Three user-programmable machines (UPMs)
 - Programmable-array-based machine controls external signal timing with a granularity of up to one quarter of an external bus clock period
 - User-specified control-signal patterns run when an internal master requests a single-beat or burst read or write access.
 - UPM refresh timer runs a user-specified control signal pattern to support refresh
 - User-specified control-signal patterns can be initiated by software
 - Each UPM can be defined to support DRAM devices with depths of 64, 128, 256, and 512 Kbytes, and 1, 2, 4, 8, 16, 32, 64, 128, and 256 Mbytes
 - Support for 8- and 16-bit devices
 - Page mode support for successive transfers within a burst
 - Internal address multiplexing supporting 64-, 128-, 256-, and 512-Kbyte, and 1-, 2-, 4-, 8-, 16-, 32-, 64-, 128-, and 256-Mbyte page banks
- Optional monitoring of transfers between local bus internal masters and local bus slaves (local bus error reporting on interrupt and status registers)
- Different machines (FCM/GPCM/UPM) share the address, data, and control signals. While the eLBC is servicing a transaction, subsequent transactions are queued until the current transaction has completed.
- Support for phase-locked loop (PLL) with software-configurable bypass for low frequency bus clocks
- 16-bit interface to peripherals

12.1.3 Modes of operation

The eLBC provides one GPCM, one FCM, and three UPMs for the local bus, with no restriction on how many of the eight banks (chip selects) can be programmed to operate with any given machine.

The internal transaction address is limited to 32 bits, so all chip selects must fall within the 4-Gbyte window addressed by the internal transaction address. When a memory transaction is dispatched to the eLBC, the internal transaction address is compared with the address information of each bank (chip select). The corresponding machine assigned to that bank (GPCM, FCM, or UPM) then takes ownership of the external signals that control the access and maintains control until the transaction ends. Thus, with the eLBC in GPCM or FCM, or UPM mode, only one of the eight chip selects is active at any time for the duration of the transaction except in the case of UPM refresh where all UPM machines that are enabled for refresh have concurrent chip select assertion.

12.1.3.1 eLBC bus clock and clock ratios

The eLBC supports ratios of 4, 8, and 16 between the faster internal (platform) clock and slower external bus clock (LCLK[0:1]).

This ratio is software programmable through the clock ratio register (LCRR[CLKDIV]).

This ratio affects the resolution of signal timing shifts in GPCM and FCM modes and the interpretation of UPM array words in UPM mode. The bus clock is driven identically onto pins, LCLK[0:1], to allow the clock load to be shared equally across a set of signal nets, thereby enhancing the edge rates of the bus clock.

12.1.3.2 Source ID debug mode

The eLBC provides the ID of a transaction source on external device pins. When those pins are selected, the 5-bit internal ID of the current transaction source appears on MSRCID[0:4] whenever valid address or data is available on the eLBC external pins. The reserved value of 0x1F, which indicates invalid address or data, appears on the source ID pins at all other times. The combination of a valid source ID (any value except 0x1F) and the value of external address latch enable (LALE) and data valid (MDVAL) facilitate capturing useful debug data as follows:

- If a valid source ID is detected on MSRCID[0:4] and LALE is asserted, a valid full 32-bit address may be latched from LAD[0:15] and combined with LA[16:31].
- If a valid source ID is detected on MSRCID[0:4] and MDVAL is asserted, valid data may be latched from LAD.

The MSRCID[0:4] and MDVAL signals are multiplexed with other functions sharing the same external pins. See , and [Signal Descriptions](#), and [Reset Clocking, and Initialization](#) , and to learn how to enable the MSRCID/MDVAL pins.

12.2 eLBC external signal descriptions

[Table 12-1](#) contains a list of external signals related to the eLBC and summarizes their function. Note that during assertion of HRESET_B, the PLL is initially unlocked, so the LCLK and LSYNC_OUT values are likely to be unstable/jittery for several microseconds; after the PLL locks, stable clock signals are driven on these signals.

Table 12-1. Signal properties-summary

Name	Alternate Function(s)	Mode	Descriptions	No. of Signals	I/O
LALE	-	-	External address latch enable	1	O
LCS_B[0]	-	-	Chip select 0	1	O
LCS_B[1:7]	-	-	Chip selects [1-7]	7	O
LWE0_B/ LFW_E_B LBS_B0	LWE_B LFW_E_B LBS_B	GPCM FCM UPM	Write enable 0 Write enable Byte (lane) select 0	1 1 1	O
LWE1_B /LBS1_B	LWE_B LBS_B	GPCM UPM	Write enable 1 Byte (lane) select 1	1 1	O
LGPL0/ LFCLE	LGPL0 LFCLE	UPM FCM	General purpose line 0 Flash command latch enable	1 1	O
LGPL1/ LFALE	LGPL1 LFALE	UPM FCM	General purpose line 1 Flash address latch enable	1 1	O
LOE_B/ LGPL2/ LFRE_B	LOE_B LFRE_B LGPL2	GPCM FCM UPM	Output enable Flash read enable General purpose line 2	1 1 1	O
LGPL3/ LFWP_B	LGPL3 LFWP_B	UPM FCM	General purpose line 3 Flash write protect	1 1	O
LGPL4/ LFRB_B/LGTA_B/ LUPWAIT/ LPBSE	LGPL4 LFRB_B LGTA_B LUPWAIT LPBSE	UPM FCM GPCM UPM -	General purpose line 4 Flash ready/busy , open-drain shared pin Transaction termination External device wait Local bus parity byte select	1 1 1 1 1	O I I I O
LGPL5	-	UPM	General purpose line 5	1	O
LBCTL	-	-	Data buffer control	1	O
LA[16:31]	-	-	Non-multiplexed address bus	16	O
LAD[0:15]	-	-	Multiplexed address/data bus	16	I/O
LDP[0:1]	-	-	Local bus data parity	2	I/O
LCLK[0:1]	-	-	Local bus clocks	2	O
LSYNC_IN	-	-	PLL synchronize input	1	I
LSYNC_OUT	-	-	PLL synchronize output	1	O

Table 12-2. Enhanced local bus controller detailed signal descriptions

Signal	I/O	Description		
LALE	O	External address latch enable. The local bus memory controller provides control for an external address latch, which allows address and data to be multiplexed on the device pins. <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 15%;">State Meaning</td> <td>Asserted/Negated-LALE is asserted with the address at the beginning of each memory controller transaction. The number of cycles for which it is asserted is governed by the ORn[EAD] and LCRR[EADC] fields. Note that no other control signals are asserted during the assertion of LALE.</td> </tr> </table>	State Meaning	Asserted/Negated-LALE is asserted with the address at the beginning of each memory controller transaction. The number of cycles for which it is asserted is governed by the ORn[EAD] and LCRR[EADC] fields. Note that no other control signals are asserted during the assertion of LALE.
State Meaning	Asserted/Negated-LALE is asserted with the address at the beginning of each memory controller transaction. The number of cycles for which it is asserted is governed by the ORn[EAD] and LCRR[EADC] fields. Note that no other control signals are asserted during the assertion of LALE.			

Table continues on the next page...

Table 12-2. Enhanced local bus controller detailed signal descriptions (continued)

Signal	I/O	Description
LCS_B[0:7]	O	Chip selects. Eight chip selects are provided that are mutually exclusive.
		State Meaning Asserted/Negated-Used to enable specific memory devices or peripherals connected to the eLBC. LCS_B[0:7] are provided on a per-bank basis with LCS0_B corresponding to the chip select for memory bank 0, which has the memory type and attributes defined by BR0 and OR0.
LWE0_B/ LWE1_B/ LBS0_B, LBS1_B	O	GPCM write enable 0/FCM write enable/UPM byte select 0. These signals select or validate each byte lane of the data bus. For an 8-bit port size, bit 0 is the only defined signal. The least-significant address bits of each access also determine which byte lanes are considered valid for a given data transfer.
		State Meaning Asserted/Negated-For GPCM operation, LWE_B[0:1] assert for each byte lane enabled for writing. LFWE_B enables command, address, and data writes to NAND Flash EEPROMs controlled by FCM. LBS_B[0:1] are programmable byte-select signals in UPM mode. See RAM array , for programming details about LBS_B[0:1].
		Timing Assertion/Negation-See General-purpose chip-select machine (GPCM) , for details regarding the timing of LWE_B[0:1].
LGPL0/ LFCLE	O	General purpose line 0/FCM command latch enable.
		State Meaning Asserted/Negated-In UPM mode, LGPL0 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode, LFCLE enables command cycles to NAND Flash EEPROMs.
LGPL1/ LFALE	O	General-purpose line 1/FCM address latch enable.
		State Meaning Asserted/Negated-In UPM mode, LGPL1 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode, LFALE enables address cycles to NAND Flash EPROMs.
LOE_B/ LGPL2/ LFRE_B	O	GPCM output enable/General-purpose line 2/FCM read enable.
		State Meaning Asserted/Negated-Controls the output buffer of memory when accessing memory/devices in GPCM mode. In UPM mode, LGPL2 is one of six general purpose signals; it is driven with a value programmed into the UPM array. LFRE_B enables data read cycles from NAND Flash EEPROMs controlled by FCM.
LGPL3/ LFWP_B	O	General-purpose line 3/FCM write protect.
		State Meaning Asserted/Negated-In UPM mode, LGPL3 is one of six general purpose signals; it is driven with a value programmed into the UPM array. In FCM mode LFWP_B protects NAND Flash EEPROMs from accidental erasure and programming when LFWP_B is asserted low-see Flash mode register (eLBC_FMR) , for programming of FCM operations to control LFWP_B .
LGTA_B/ LGPL4/ LFRB_B LUPWAIT/ LPBSE	I/O	GPCM transfer acknowledge/General-purpose line 4/FCM Flash ready-busy/UPM wait/parity byte select.
		State Meaning Asserted/Negated-Input in GPCM or FCM modes used for transaction termination. It may also be configured as one of six general-purpose output signals when in UPM mode or as an input to force the UPM controller to wait for the memory/device. FCM uses LFRB_B to stall during long-latency read and programming operations, continuing once LFRB_B returns high. When configured as LPBSE, it disables any use in GPCM, FCM, or UPM modes. Because systems that use read-modify-write parity require an additional memory device, they must generate a byte-select like a normal data device. ANDing LBSn_B through external logic to achieve the logical function of this byte-select can affect memory access timing. The LBC provides this optional byte-select signal connection to RMW-parity devices.
LGPL5	O	General-purpose line 5
		State Meaning Asserted/Negated-One of six general purpose signals when in UPM mode, and drives a value programmed in the UPM array.

Table continues on the next page...

Table 12-2. Enhanced local bus controller detailed signal descriptions (continued)

Signal	I/O	Description		
LBCTL	O	Data buffer control. The memory controller activates LBCTL for the local bus when a GPCM-, UPM-, or FCM-controlled bank is accessed. Buffer control is disabled by setting ORn[BCTLD].		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-The LBCTL pin normally functions as a write/read_B control for a bus transceiver connected to the LAD lines. Note that an external data buffer must not drive the LAD lines in conflict with the eLBC when LBCTL is high, because LBCTL remains high after reset and during address phases.</td> </tr> </table>	State Meaning	Asserted/Negated-The LBCTL pin normally functions as a write/read_B control for a bus transceiver connected to the LAD lines. Note that an external data buffer must not drive the LAD lines in conflict with the eLBC when LBCTL is high, because LBCTL remains high after reset and during address phases.
State Meaning	Asserted/Negated-The LBCTL pin normally functions as a write/read_B control for a bus transceiver connected to the LAD lines. Note that an external data buffer must not drive the LAD lines in conflict with the eLBC when LBCTL is high, because LBCTL remains high after reset and during address phases.			
LA[16:31]	O	Nonmultiplexed address bus. All bits driven are defined for 8-bit port sizes. For 16-bit port sizes LA[31] is a don't care.		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-LA is the address bus used to transmit addresses to external RAM devices. Refer to eLBC initialization/application information, for address signal multiplexing.</td> </tr> </table>	State Meaning	Asserted/Negated-LA is the address bus used to transmit addresses to external RAM devices. Refer to eLBC initialization/application information , for address signal multiplexing.
State Meaning	Asserted/Negated-LA is the address bus used to transmit addresses to external RAM devices. Refer to eLBC initialization/application information , for address signal multiplexing.			
LAD[0:15]	I/O	Multiplexed address/data bus. For a port size of 16 bits, LAD[0:7] connect to the most-significant byte lane (at address offset 0), while LAD[8:15] connect to the least-significant byte lane (at address offset 1). For a port size of 8 bits, only LAD[0:7] are connected to the external RAM.		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-LAD is the shared 16-bit address/data bus through which external RAM devices transfer data and receive addresses.</td> </tr> </table>	State Meaning	Asserted/Negated-LAD is the shared 16-bit address/data bus through which external RAM devices transfer data and receive addresses.
		State Meaning	Asserted/Negated-LAD is the shared 16-bit address/data bus through which external RAM devices transfer data and receive addresses.	
<table border="1"> <tr> <td>Timing</td> <td>Assertion/Negation-During assertion of LALE, LAD are driven with the RAM address for the access to follow. External logic should propagate the address on LAD while LALE is asserted, and latch the address upon negation of LALE. After LALE is negated, LAD are either driven by write data or are made high-impedance by the eLBC in order to sample read data driven by an external device. Following the last data transfer of a write access, LAD are again taken into a high-impedance state.</td> </tr> </table>	Timing	Assertion/Negation-During assertion of LALE, LAD are driven with the RAM address for the access to follow. External logic should propagate the address on LAD while LALE is asserted, and latch the address upon negation of LALE. After LALE is negated, LAD are either driven by write data or are made high-impedance by the eLBC in order to sample read data driven by an external device. Following the last data transfer of a write access, LAD are again taken into a high-impedance state.		
Timing	Assertion/Negation-During assertion of LALE, LAD are driven with the RAM address for the access to follow. External logic should propagate the address on LAD while LALE is asserted, and latch the address upon negation of LALE. After LALE is negated, LAD are either driven by write data or are made high-impedance by the eLBC in order to sample read data driven by an external device. Following the last data transfer of a write access, LAD are again taken into a high-impedance state.			
LDP[0: 1]	I/O	Local bus data parity. Drives and receives the data parity corresponding with the data phases on LAD for GPCM and UPM controlled banks.		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-During write accesses, a parity bit is generated for each 8 bits of LAD[0:15], such that LDP0 is even/odd parity for LAD[0:7], while LDP[1] is even/odd parity for LAD[8:15]. Unused byte lanes for port sizes less than 16 bits have undefined parity.</td> </tr> </table>	State Meaning	Asserted/Negated-During write accesses, a parity bit is generated for each 8 bits of LAD[0:15], such that LDP0 is even/odd parity for LAD[0:7], while LDP[1] is even/odd parity for LAD[8:15]. Unused byte lanes for port sizes less than 16 bits have undefined parity.
		State Meaning	Asserted/Negated-During write accesses, a parity bit is generated for each 8 bits of LAD[0:15], such that LDP0 is even/odd parity for LAD[0:7], while LDP[1] is even/odd parity for LAD[8:15]. Unused byte lanes for port sizes less than 16 bits have undefined parity.	
<table border="1"> <tr> <td>Timing</td> <td>Assertion/Negation-Drive and receive the data parity corresponding with the data phases on LAD. For read accesses, the parity bits for each byte lane are sampled on LDP[0: 1] with the same timing that read data is sampled on LAD.</td> </tr> </table>	Timing	Assertion/Negation-Drive and receive the data parity corresponding with the data phases on LAD. For read accesses, the parity bits for each byte lane are sampled on LDP[0: 1] with the same timing that read data is sampled on LAD.		
Timing	Assertion/Negation-Drive and receive the data parity corresponding with the data phases on LAD. For read accesses, the parity bits for each byte lane are sampled on LDP[0: 1] with the same timing that read data is sampled on LAD.			
LCLK[0:1]	O	Local bus clocks		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-LCLK[0:1] drive an identical bus clock signal for distributed loads.If the eLBC PLL is enabled (see LCRR[PBYYP], Clock ratio register (eLBC_LCRR)), the bus clock phase is shifted earlier than transitions on other eLBC signals (such as LADn and LCSn_B) by a time delay matching the delay of the PLL timing loop set up between LSYNC_OUT and LSYNC_IN.</td> </tr> </table>	State Meaning	Asserted/Negated-LCLK[0:1] drive an identical bus clock signal for distributed loads.If the eLBC PLL is enabled (see LCRR[PBYYP], Clock ratio register (eLBC_LCRR)), the bus clock phase is shifted earlier than transitions on other eLBC signals (such as LADn and LCSn_B) by a time delay matching the delay of the PLL timing loop set up between LSYNC_OUT and LSYNC_IN.
State Meaning	Asserted/Negated-LCLK[0:1] drive an identical bus clock signal for distributed loads.If the eLBC PLL is enabled (see LCRR[PBYYP], Clock ratio register (eLBC_LCRR)), the bus clock phase is shifted earlier than transitions on other eLBC signals (such as LADn and LCSn_B) by a time delay matching the delay of the PLL timing loop set up between LSYNC_OUT and LSYNC_IN.			
LSYNC_OUT_B	O	PLL synchronization out.		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-A replica of the bus clock, appearing on LSYNC_OUT, should be propagated through a passive timing loop and returned to LSYNC_IN for achieving correct PLL lock.</td> </tr> </table>	State Meaning	Asserted/Negated-A replica of the bus clock, appearing on LSYNC_OUT, should be propagated through a passive timing loop and returned to LSYNC_IN for achieving correct PLL lock.
		State Meaning	Asserted/Negated-A replica of the bus clock, appearing on LSYNC_OUT, should be propagated through a passive timing loop and returned to LSYNC_IN for achieving correct PLL lock.	
<table border="1"> <tr> <td>Timing</td> <td>Assertion/Negation-The time delay of the timing loop should be such that it compensates for the round-trip flight time of LCLK[0:1] and clocked drivers in the system. No load other than a timing loop should be placed on LSYNC_OUT.</td> </tr> </table>	Timing	Assertion/Negation-The time delay of the timing loop should be such that it compensates for the round-trip flight time of LCLK[0:1] and clocked drivers in the system. No load other than a timing loop should be placed on LSYNC_OUT.		
Timing	Assertion/Negation-The time delay of the timing loop should be such that it compensates for the round-trip flight time of LCLK[0:1] and clocked drivers in the system. No load other than a timing loop should be placed on LSYNC_OUT.			
LSYNC_IN_B	I	PLL synchronization in.		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-See description of LSYNC_OUT.</td> </tr> </table>	State Meaning	Asserted/Negated-See description of LSYNC_OUT.
State Meaning	Asserted/Negated-See description of LSYNC_OUT.			

12.3 Enhanced Local Bus Controller (eLBC) Memory Map

The table below shows the memory mapped registers of the eLBC. Undefined 4-byte address spaces within offset 0x0000-0xFFFF are reserved and should not be accessed for reading or writing. Similarly, only zero should be written to reserved bits of defined registers, as writing ones can have unpredictable results in some cases.

Bits designated as write-one-to-clear are cleared only by writing ones to them. Writing zeros to them has no effect.

eLBC memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5000	Base register 0 (eLBC_BR0)	32	R/W	See section	12.3.1/538
5004	Options register 0 layout for GPCM Mode (eLBC_ORg0)	32	R/W	0000_0FF7h	12.3.2/540
5004	Options register 0 layout for FCM Mode (eLBC_ORf0)	32	R/W	See section	12.3.3/543
5004	Options register 0 layout for UPM Mode (eLBC_ORu0)	32	R/W	See section	12.3.4/547
5008	Base register n (eLBC_BR1)	32	R/W	0000_0000h	12.3.5/551
500C	Options register n layout for GPCM Mode (eLBC_ORg1)	32	R/W	0000_0000h	12.3.6/552
500C	Options register n layout for FCM Mode (eLBC_ORf1)	32	R/W	0000_0000h	12.3.7/556
500C	Options register n layout for UPM Mode (eLBC_ORu1)	32	R/W	0000_0000h	12.3.8/560
5010	Base register n (eLBC_BR2)	32	R/W	0000_0000h	12.3.5/551
5014	Options register n layout for GPCM Mode (eLBC_ORg2)	32	R/W	0000_0000h	12.3.6/552
5014	Options register n layout for FCM Mode (eLBC_ORf2)	32	R/W	0000_0000h	12.3.7/556
5014	Options register n layout for UPM Mode (eLBC_ORu2)	32	R/W	0000_0000h	12.3.8/560
5018	Base register n (eLBC_BR3)	32	R/W	0000_0000h	12.3.5/551
501C	Options register n layout for GPCM Mode (eLBC_ORg3)	32	R/W	0000_0000h	12.3.6/552
501C	Options register n layout for FCM Mode (eLBC_ORf3)	32	R/W	0000_0000h	12.3.7/556
501C	Options register n layout for UPM Mode (eLBC_ORu3)	32	R/W	0000_0000h	12.3.8/560
5020	Base register n (eLBC_BR4)	32	R/W	0000_0000h	12.3.5/551
5024	Options register n layout for GPCM Mode (eLBC_ORg4)	32	R/W	0000_0000h	12.3.6/552
5024	Options register n layout for FCM Mode (eLBC_ORf4)	32	R/W	0000_0000h	12.3.7/556
5024	Options register n layout for UPM Mode (eLBC_ORu4)	32	R/W	0000_0000h	12.3.8/560
5028	Base register n (eLBC_BR5)	32	R/W	0000_0000h	12.3.5/551
502C	Options register n layout for GPCM Mode (eLBC_ORg5)	32	R/W	0000_0000h	12.3.6/552
502C	Options register n layout for FCM Mode (eLBC_ORf5)	32	R/W	0000_0000h	12.3.7/556
502C	Options register n layout for UPM Mode (eLBC_ORu5)	32	R/W	0000_0000h	12.3.8/560
5030	Base register n (eLBC_BR6)	32	R/W	0000_0000h	12.3.5/551
5034	Options register n layout for GPCM Mode (eLBC_ORg6)	32	R/W	0000_0000h	12.3.6/552
5034	Options register n layout for FCM Mode (eLBC_ORf6)	32	R/W	0000_0000h	12.3.7/556

Table continues on the next page...

eLBC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
5034	Options register n layout for UPM Mode (eLBC_ORu6)	32	R/W	0000_0000h	12.3.8/560
5038	Base register n (eLBC_BR7)	32	R/W	0000_0000h	12.3.5/551
503C	Options register n layout for GPCM Mode (eLBC_ORg7)	32	R/W	0000_0000h	12.3.6/552
503C	Options register n layout for FCM Mode (eLBC_ORf7)	32	R/W	0000_0000h	12.3.7/556
503C	Options register n layout for UPM Mode (eLBC_ORu7)	32	R/W	0000_0000h	12.3.8/560
5068	UPM address register (eLBC_MAR)	32	R/W	0000_0000h	12.3.9/562
5070	UPMn mode register (eLBC_MAMR)	32	R/W	0000_0000h	12.3.10/563
5074	UPMn mode register (eLBC_MBMR)	32	R/W	0000_0000h	12.3.10/563
5078	UPMn mode register (eLBC_MCMR)	32	R/W	0000_0000h	12.3.10/563
5084	Memory refresh timer prescaler register (eLBC_MRTPR)	32	R/W	0000_0000h	12.3.11/566
5088	UPM data register (eLBC_MDRu)	32	R/W	0000_0000h	12.3.12/566
5088	FCM data register (eLBC_MDRf)	32	R/W	0000_0000h	12.3.13/567
5090	Special operation initiation register (eLBC_LSOR)	32	R/W	0000_0000h	12.3.14/567
50A0	UPM refresh timer (eLBC_LURT)	32	R/W	0000_0000h	12.3.15/568
50B0	Transfer error status register (eLBC_LTESR)	32	w1c	0000_0000h	12.3.16/569
50B4	Transfer error disable register (eLBC_LTEDR)	32	R/W	0000_0000h	12.3.17/571
50B8	Transfer error interrupt register (eLBC_LTEIR)	32	R/W	0000_0000h	12.3.18/572
50BC	Transfer error attributes register (eLBC_LTEATR)	32	R/W	0000_0000h	12.3.19/573
50C0	Transfer error address register (eLBC_LTEAR)	32	R/W	0000_0000h	12.3.20/574
50C4	Transfer error ECC register (eLBC_LTECCR)	32	w1c	0000_0000h	12.3.21/575
50D0	Configuration register (eLBC_LBCR)	32	R/W	0000_0000h	12.3.22/576
50D4	Clock ratio register (eLBC_LCRR)	32	R/W	8000_0000h	12.3.23/577
50E0	Flash mode register (eLBC_FMR)	32	R/W	See section	12.3.24/579
50E4	Flash instruction register (eLBC_FIR)	32	R/W	0000_0000h	12.3.25/581

Table continues on the next page...

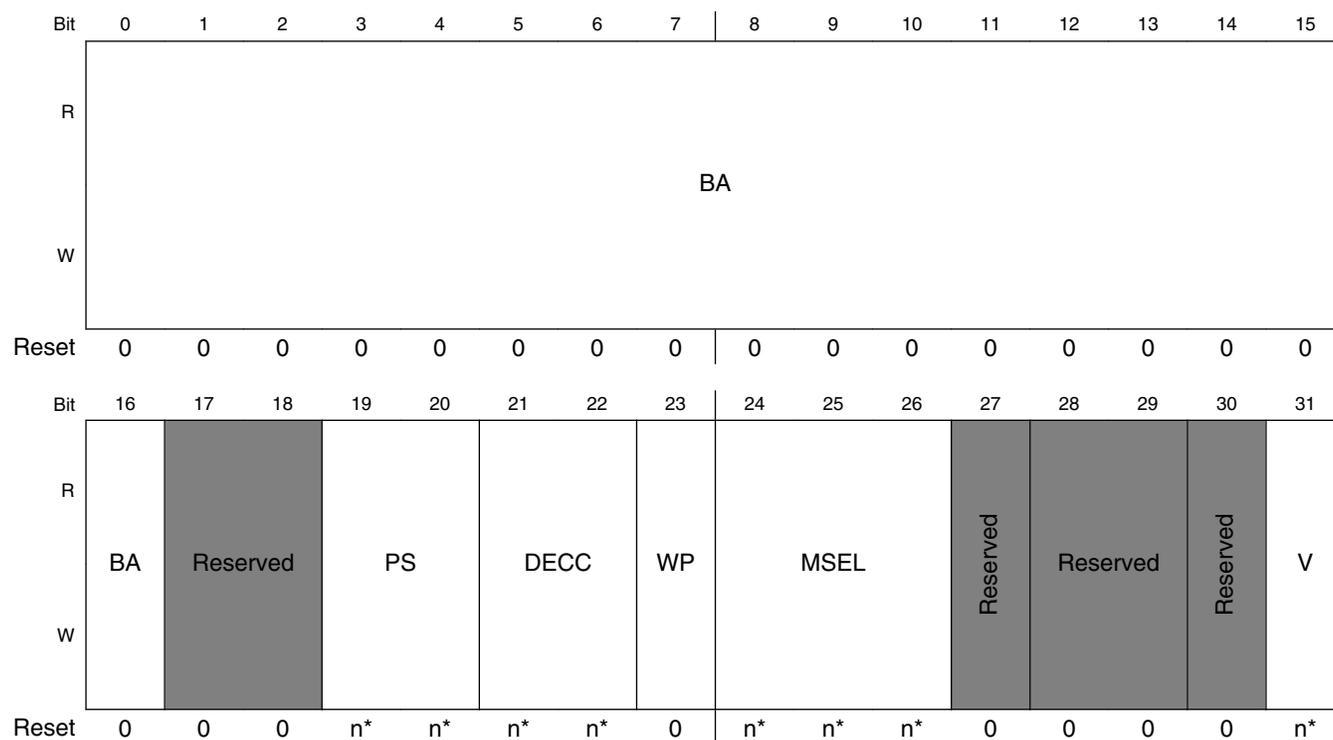
eLBC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
50E8	Flash command register (eLBC_FCR)	32	R/W	0000_0000h	12.3.26/582
50EC	Flash block address register (eLBC_FBAR)	32	R/W	0000_0000h	12.3.27/583
50F0	Flash page address register [Large Page Device (ORx[PGS] = 1)] (eLBC_FPARI)	32	R/W	0000_0000h	12.3.28/583
50F0	Flash page address register [Small Page Device (ORx[PGS] = 0)] (eLBC_FPARs)	32	R/W	0000_0000h	12.3.29/584
50F4	Flash byte count register (eLBC_FBCR)	32	R/W	0000_0000h	12.3.30/585
5100	Flash ECC block n registers (eLBC_FECC0)	32	R	0000_0000h	12.3.31/586
5104	Flash ECC block n registers (eLBC_FECC1)	32	R	0000_0000h	12.3.31/586
5108	Flash ECC block n registers (eLBC_FECC2)	32	R	0000_0000h	12.3.31/586
510C	Flash ECC block n registers (eLBC_FECC3)	32	R	0000_0000h	12.3.31/586

12.3.1 Base register 0 (eLBC_BR0)

The base registers (BR *n*) contain the base address and address types for each memory bank. The memory controller uses this information to compare the address bus value with the current address accessed. Each register (bank) includes a memory attribute and selects the machine for memory operation handling. Note that after system reset, BR0[V] is set, BR1[V]-BR7[V] are cleared, and the value of BR0[PS] reflects the initial port size configured by the boot ROM location power-on configuration settings .

Address: 5000h base + 0h offset = 5000h



* Notes:

- V field: BR0 has its valid bit (V) set at reset .
- MSEL field: M = 0 for MSEL of GPCM, 1 for MSEL of FCM at boot.
- DECC field: The reset value for DECC is determined by cfg_elbc_ecc .
- PS field: Thus bank 0 is valid with the port size (PS) configured from cfg_rom_loc[0:3] at power-on reset .

eLBC_BR0 field descriptions

Field	Description
0–16 BA	Base address. The upper 17 bits of each base register are compared to the address on the address bus to determine if the bus master is accessing a memory bank controlled by the memory controller. Used with the address mask bits OR <i>n</i> [AM].

Table continues on the next page...

eLBC_BR0 field descriptions (continued)

Field	Description
17–18 -	This field is reserved. Reserved
19–20 PS	Port size. Specifies the port size of this memory region. For BR0, PS is configured from the boot ROM location power-on reset configuration setting . For all other banks the value is reset to 00 (port size not defined). 00 Reserved 01 8-bit (supported for GPCM, UPM, FCM) 10 16-bit (supported for GPCM, UPM) 11 Reserved
21–22 DECC	Specifies the method for data error checking. 00 Data error checking disabled , but normal parity generation for GPCM and UPM . No ECC generation for FCM. 01 ECC checking is enabled, but ECC generation is disabled, for FCM on full-page transfers. Normal parity generation and checking for GPCM and UPM. 10 ECC checking and generation are enabled for FCM on full-page transfers. 11 Reserved
23 WP	Write protect. 0 Read and write accesses are allowed. 1 Only read accesses are allowed. The memory controller does not assert \overline{LCSn} on write cycles to this memory bank. LTESR[WP] is set (if WP is set) if a write to this memory bank is attempted, and a local bus error interrupt is generated (if enabled), terminating the cycle.
24–26 MSEL	Machine select. Specifies the machine to use for handling memory operations. 000 GPCM (possible reset value) 001 FCM (possible reset value) 010 Reserved 011 Reserved 100 UPMA 101 UPMB 110 UPMC 111 Reserved
27 -	This field is reserved. Reserved
28–29 -	This field is reserved. Reserved
30 -	This field is reserved. Reserved
31 V	Valid bit. Indicates that the contents of the BR <i>n</i> and OR <i>n</i> pair are valid. \overline{LCSn} does not assert unless V is set (an access to a region that has no valid bit set may cause a bus time-out). After a system reset, only BR0[V] is set. 0 This bank is invalid. 1 This bank is valid.

12.3.2 Options register 0 layout for GPCM Mode (eLBC_ORg0)

The OR *n* registers define the sizes of memory banks and access attributes. The OR *n* attribute bits support the following three modes of operation as defined by BR *n* [MSEL]:

- GPCM mode
- FCM mode
- UPM mode

The OR *n* registers are interpreted differently depending on which of the three machine types is selected for that bank; this section describes the interpretation when the GPCM controls bank *n*.

Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options, specifically the `cfg_rom_loc` power-on-reset pins.

Table 12-5. Reset value of OR0 Register

Boot Source	OR0 Reset Value
FCM (small page NAND Flash)	0000_03AE
FCM (large page NAND Flash)	0000_07AE
GPCM	0000_0FF7
eLBC not used as a boot source	0000_0F07

The address mask field of the option registers (OR *n* [AM]) masks up to 17 corresponding BR *n* [BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. The table below shows memory bank sizes from 32 Kbytes to 4 Gbytes.

Table 12-6. Memory Bank Sizes in Relation to Address Mask

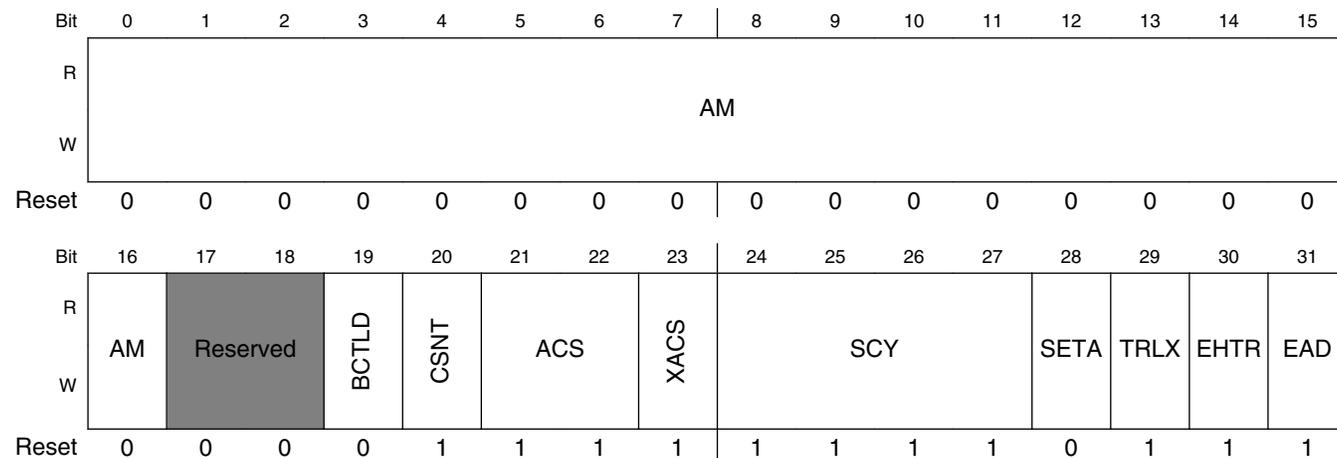
AM	Memory Bank Size
0000_0000_0000_0000_0	4 Gbytes
1000_0000_0000_0000_0	2 Gbytes
1100_0000_0000_0000_0	1 Gbyte
1110_0000_0000_0000_0	512 Mbytes
1111_0000_0000_0000_0	256 Mbytes
1111_1000_0000_0000_0	128 Mbytes
1111_1100_0000_0000_0	64 Mbytes

Table continues on the next page...

Table 12-6. Memory Bank Sizes in Relation to Address Mask (continued)

AM	Memory Bank Size
1111_1110_0000_0000_0	32 Mbytes
1111_1111_0000_0000_0	16 Mbytes
1111_1111_1000_0000_0	8 Mbytes
1111_1111_1100_0000_0	4 Mbytes
1111_1111_1110_0000_0	2 Mbytes
1111_1111_1111_0000_0	1 Mbyte
1111_1111_1111_1000_0	512 Kbytes
1111_1111_1111_1100_0	256 Kbytes
1111_1111_1111_1110_0	128 Kbytes
1111_1111_1111_1111_0	64 Kbytes
1111_1111_1111_1111_1	32 Kbytes

Address: 5000h base + 4h offset = 5004h



eLBC_ORg0 field descriptions

Field	Description
0–16 AM	GPCM address mask. Masks corresponding BRn bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked and therefore don't care for address checking. 1 Corresponding address bits are used in the comparison between base and transaction addresses.
17–18 -	This field is reserved. Reserved
19 BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.

Table continues on the next page...

eLBC_ORg0 field descriptions (continued)

Field	Description
20 CSNT	<p>Chip select negation time. Determines when $\overline{\text{LCSn}}$ and $\overline{\text{LWE}}$ are negated during an external memory write access handled by the GPCM, provided that $\text{ACS} \neq 00$ (when $\text{ACS} = 00$, only $\overline{\text{LWE}}$ is affected by the setting of CSNT). This helps meet address/data hold times for slow memories and peripherals.</p> <p>0 $\overline{\text{LCSn}}$ and $\overline{\text{LWE}}$ are negated normally. 1 $\overline{\text{LCSn}}$ and $\overline{\text{LWE}}$ are negated one quarter of a bus clock cycle earlier.</p>
21–22 ACS	<p>Address to chip-select setup. Determines the delay of the $\overline{\text{LCSn}}$ assertion relative to the address change when the external memory access is handled by the GPCM. At system reset, $\text{OR0}[\text{ACS}] = 11$.</p> <p>00 $\overline{\text{LCSn}}$ is output at the same time as the address lines. Note that this overrides the value of CSNT such that $\text{CSNT} = 0$. 01 Reserved. 10 $\overline{\text{LCSn}}$ is output one quarter bus clock cycle after the address lines. 11 $\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.</p>
23 XACS	<p>Extra address to chip-select setup. Setting this bit increases the delay of the $\overline{\text{LCSn}}$ assertion relative to the address change when the external memory access is handled by the GPCM. After a system reset, $\text{OR0}[\text{XACS}] = 1$.</p> <p>0 Address to chip-select setup is determined by $\text{ORx}[\text{ACS}]$. 1 Address to chip-select setup is extended (see Table 12-190 and Table 12-191).</p>
24–27 SCY	<p>Cycle length in bus clocks. Determines the number of wait states inserted in the bus cycle, when the GPCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings. After a system reset, $\text{OR0}[\text{SCY}] = 1111$.</p> <p>0000 No wait states 0001 1 bus clock cycle wait state ... 1111 15 bus clock cycle wait states</p>
28 SETA	<p>External address termination.</p> <p>0 Access is terminated internally by the memory controller unless the external device asserts $\overline{\text{LGTA}}$ earlier to terminate the access. 1 Access is terminated externally by asserting the $\overline{\text{LGTA}}$ external pin. (Only $\overline{\text{LGTA}}$ can terminate the access).</p>
29 TRLX	<p>Timing relaxed. Modifies the settings of timing parameters for slow memories or peripherals.</p> <p>0 Normal timing is generated by the GPCM. 1 Relaxed timing on the following parameters: <ul style="list-style-type: none"> • Adds an additional cycle between the address and control signals (only if ACS is not equal to 00). • Doubles the number of wait states specified by SCY, providing up to 30 wait states. • Works in conjunction with EHTR to extend hold time on read accesses. • $\overline{\text{LCSn}}$ (only if ACS is not equal to 00) and $\overline{\text{LWE}}$ signals are negated one cycle earlier during writes. </p>
30 EHTR	<p>Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access.</p>

Table continues on the next page...

eLBC_ORg0 field descriptions (continued)

Field	Description		
	Table 12-7. TRLX and EHTR		
	TRLX	EHTR	Meaning
	0	0	The memory controller generates normal timing. No additional cycles are inserted.
	0	1	1 idle clock cycle is inserted.
	1	0	4 idle clock cycles are inserted.
	1	1	8 idle clock cycles are inserted.
31 EAD	External address latch delay. Allow extra bus clock cycles when using external address latch (LALE).		
	0	No additional bus clock cycles (LALE asserted for one bus clock cycle only)	
	1	Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by LCRR[EADC]).	

12.3.3 Options register 0 layout for FCM Mode (eLBC_ORf0)

The OR n registers define the sizes of memory banks and access attributes. The OR n attribute bits support the following three modes of operation as defined by BR n [MSEL]:

- GPCM mode
- FCM mode
- UPM mode

The OR n registers are interpreted differently depending on which of the three machine types is selected for that bank; this section describes the interpretation when the FCM controls bank n .

Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options, specifically the `cfg_rom_loc` power-on-reset pins.

Table 12-9. Reset value of OR0 Register

Boot Source	OR0 Reset Value
FCM (small page NAND Flash)	0000_03AE
FCM (large page NAND Flash)	0000_07AE
GPCM	0000_0FF7
eLBC not used as a boot source	0000_0F07

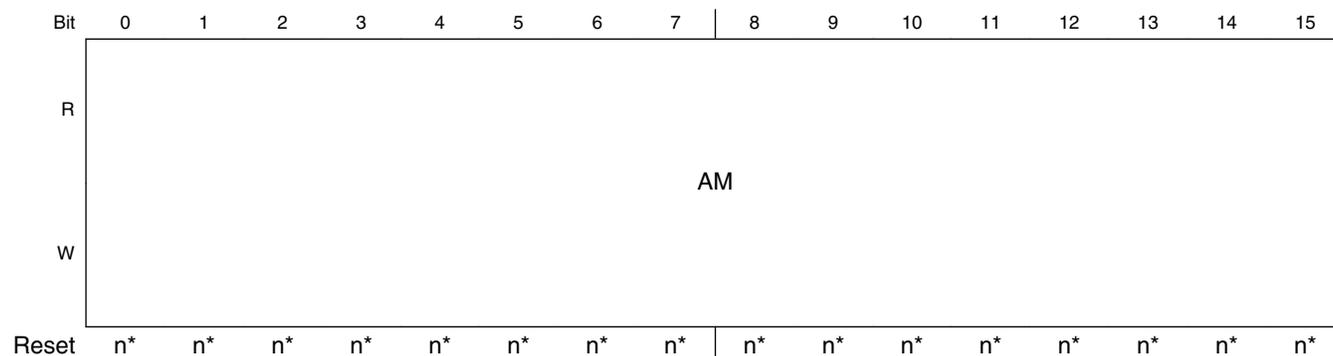
The address mask field of the option registers (OR n [AM]) masks up to 17 corresponding BR n [BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking

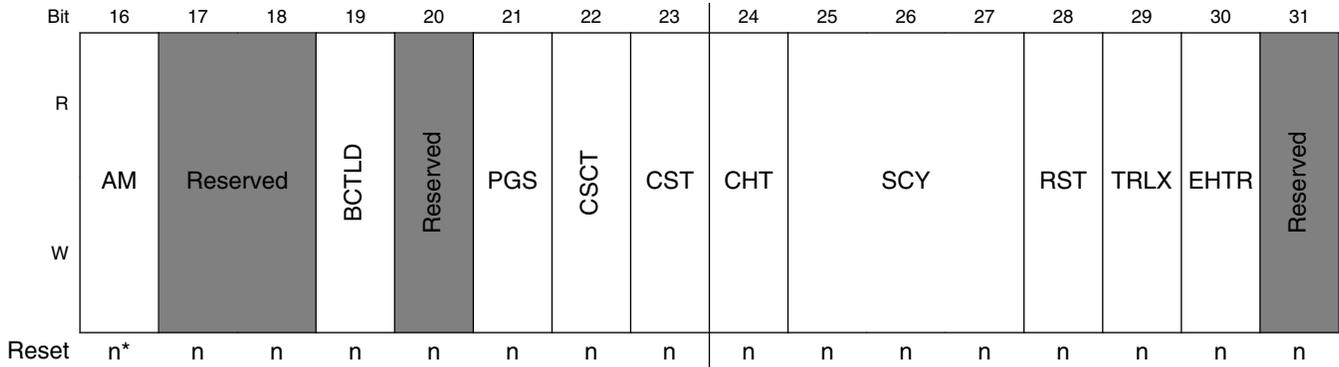
address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. The table below shows memory bank sizes from 32 Kbytes to 4 Gbytes.

Table 12-10. Memory Bank Sizes in Relation to Address Mask

AM	Memory Bank Size
0000_0000_0000_0000_0	4 Gbytes
1000_0000_0000_0000_0	2 Gbytes
1100_0000_0000_0000_0	1 Gbyte
1110_0000_0000_0000_0	512 Mbytes
1111_0000_0000_0000_0	256 Mbytes
1111_1000_0000_0000_0	128 Mbytes
1111_1100_0000_0000_0	64 Mbytes
1111_1110_0000_0000_0	32 Mbytes
1111_1111_0000_0000_0	16 Mbytes
1111_1111_1000_0000_0	8 Mbytes
1111_1111_1100_0000_0	4 Mbytes
1111_1111_1110_0000_0	2 Mbytes
1111_1111_1111_0000_0	1 Mbyte
1111_1111_1111_1000_0	512 Kbytes
1111_1111_1111_1100_0	256 Kbytes
1111_1111_1111_1110_0	128 Kbytes
1111_1111_1111_1111_0	64 Kbytes
1111_1111_1111_1111_1	32 Kbytes

Address: 5000h base + 4h offset = 5004h





* Notes:

- AM field: See Table "Reset value of OR0 Register" above for the OR0 reset value. All other option registers have all bits cleared.

eLBC_ORf0 field descriptions

Field	Description
0–16 AM	FCM address mask. Masks corresponding BRn bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked. 1 Corresponding address bits are used in the comparison between base and transaction addresses.
17–18 -	This field is reserved. Reserved
19 BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.
20 -	This field is reserved. Reserved
21 PGS	NAND Flash EEPROM page size, buffer size, and block size. 0 Page size of 512 main area bytes plus 16 spare area bytes (small page devices); FCM RAM buffers are 1 Kbyte each; Flash block size of 16 Kbytes. 1 Page size of 2048 main area bytes plus 64 spare area bytes (large page devices); FCM RAM buffers are 4 Kbytes each; Flash block size of 128 Kbytes.
22 CSCT	Chip select to command time. Determines how far in advance \overline{LCSn} is asserted prior to any bus activity during a NAND Flash access handled by the FCM. This helps meet chip-select setup times for slow memories.

Table 12-15. TRLX and CSCT

TRLX	CSCT	Meaning
0	0	The chip-select is asserted 1 clock cycle before any command.
0	1	The chip-select is asserted 4 clock cycles before any command.
1	0	The chip-select is asserted 2 clock cycles before any command.
1	1	The chip-select is asserted 8 clock cycles before any command.

Table continues on the next page...

eLBC_ORf0 field descriptions (continued)

Field	Description															
23 CST	<p>Command setup time. Determines the delay of $\overline{\text{LFW}}\overline{\text{E}}$ assertion relative to the command, address, or data change when the external memory access is handled by the FCM.</p> <p style="text-align: center;">Table 12-14. TRLX and CST</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>CST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is asserted coincident with any command.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is asserted 0.25 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is asserted 0.5 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The write-enable is asserted 1 clock cycle after any command, address, or data.</td> </tr> </tbody> </table>	TRLX	CST	Meaning	0	0	The write-enable is asserted coincident with any command.	0	1	The write-enable is asserted 0.25 clock cycles after any command, address, or data.	1	0	The write-enable is asserted 0.5 clock cycles after any command, address, or data.	1	1	The write-enable is asserted 1 clock cycle after any command, address, or data.
TRLX	CST	Meaning														
0	0	The write-enable is asserted coincident with any command.														
0	1	The write-enable is asserted 0.25 clock cycles after any command, address, or data.														
1	0	The write-enable is asserted 0.5 clock cycles after any command, address, or data.														
1	1	The write-enable is asserted 1 clock cycle after any command, address, or data.														
24 CHT	<p>Command hold time. Determines the $\overline{\text{LFW}}\overline{\text{E}}$ negation prior to the command, address, or data change when the external memory access is handled by the FCM.</p> <p style="text-align: center;">Table 12-13. TRLX and CHT</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>CHT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is negated 0.5 clock cycles before any command, address, or data change.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is negated 1 clock cycle before any command, address, or data change.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is negated 1.5 clock cycles before any command, address, or data change.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The write-enable is negated 2 clock cycles before any command, address, or data change.</td> </tr> </tbody> </table>	TRLX	CHT	Meaning	0	0	The write-enable is negated 0.5 clock cycles before any command, address, or data change.	0	1	The write-enable is negated 1 clock cycle before any command, address, or data change.	1	0	The write-enable is negated 1.5 clock cycles before any command, address, or data change.	1	1	The write-enable is negated 2 clock cycles before any command, address, or data change.
TRLX	CHT	Meaning														
0	0	The write-enable is negated 0.5 clock cycles before any command, address, or data change.														
0	1	The write-enable is negated 1 clock cycle before any command, address, or data change.														
1	0	The write-enable is negated 1.5 clock cycles before any command, address, or data change.														
1	1	The write-enable is negated 2 clock cycles before any command, address, or data change.														
25–27 SCY	<p>Cycle length in bus clocks. Determines the following:</p> <ul style="list-style-type: none"> the number of wait states inserted in command, address, or data transfer bus cycles, when the FCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings. the delay between command/address writes and data write cycles, or the delay between write cycles and read cycles from NAND Flash EEPROM. A delay of $4x(2+SCY)$ clock cycles (TRLX = 0) or $8x(2+SCY)$ clock cycles (TRLX = 1) is inserted between the last write and the first data transfer to/from NAND Flash devices. the delay between a command write and the first sample point of the RDY/$\overline{\text{BSY}}$ pin (connected to $\overline{\text{LFRB}}$). $\overline{\text{LFRB}}$ is not sampled until $8x(2+SCY)$ clock cycles (TRLX = 0) or $16x(2+SCY)$ clock cycles (TRLX = 1) have elapsed following the command. <p>000 No extra wait states 001 1 bus clock cycle wait state ... 111 7 bus clock cycle wait states</p>															
28 RST	<p>Read setup time. Determines the delay of $\overline{\text{LFRE}}$ assertion relative to sampling of read data when the external memory access is handled by the FCM.</p> <p style="text-align: center;">Table 12-12. TRLX and RST</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>RST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The read-enable is asserted 0.75 clock cycles prior to any wait states.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The read-enable is asserted 1 clock cycle prior to any wait states.</td> </tr> </tbody> </table>	TRLX	RST	Meaning	0	0	The read-enable is asserted 0.75 clock cycles prior to any wait states.	0	1	The read-enable is asserted 1 clock cycle prior to any wait states.						
TRLX	RST	Meaning														
0	0	The read-enable is asserted 0.75 clock cycles prior to any wait states.														
0	1	The read-enable is asserted 1 clock cycle prior to any wait states.														

Table continues on the next page...

eLBC_ORf0 field descriptions (continued)

Field	Description		
	Table 12-12. TRLX and RST (continued)		
	1	0	The read-enable is asserted 0.5 clock cycles prior to any wait states.
	1	1	The read-enable is asserted 1 clock cycle prior to any wait states.
29 TRLX	Timing relaxed. Modifies the settings of timing parameters for slow memories. 0 Normal timing is generated by the FCM. 1 Relaxed timing on the following parameters: <ul style="list-style-type: none"> • Doubles the number of clock cycles between $\overline{\text{LCSn}}$ assertion and commands. • Doubles the number of wait states specified by SCY, providing up to 14 wait states. • Works in conjunction with CST and RST to extend command/address/data setup times. • Adds one clock cycle to the command/address/data hold times. • Works in conjunction with CBT to extend the wait time for read/busy status sampling by 16 clock cycles. • Works in conjunction with EHTR to double hold time on read accesses. 		
30 EHTR	Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. Table 12-11. TRLX and EHTR		
	TRLX	EHTR	Meaning
	0	0	1 idle clock cycle is inserted.
	0	1	2 idle clock cycles are inserted.
	1	0	4 idle clock cycles are inserted.
	1	1	8 idle clock cycles are inserted.
31 -	This field is reserved. Reserved		

12.3.4 Options register 0 layout for UPM Mode (eLBC_ORu0)

The OR n registers define the sizes of memory banks and access attributes. The OR n attribute bits support the following three modes of operation as defined by BR n [MSEL]:

- GPCM mode
- FCM mode
- UPM mode

The OR n registers are interpreted differently depending on which of the three machine types is selected for that bank; this section describes the interpretation when the UPM controls bank n .

Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options , specifically the `cfg_rom_loc` power-on-reset pins .

Table 12-17. Reset value of OR0 Register

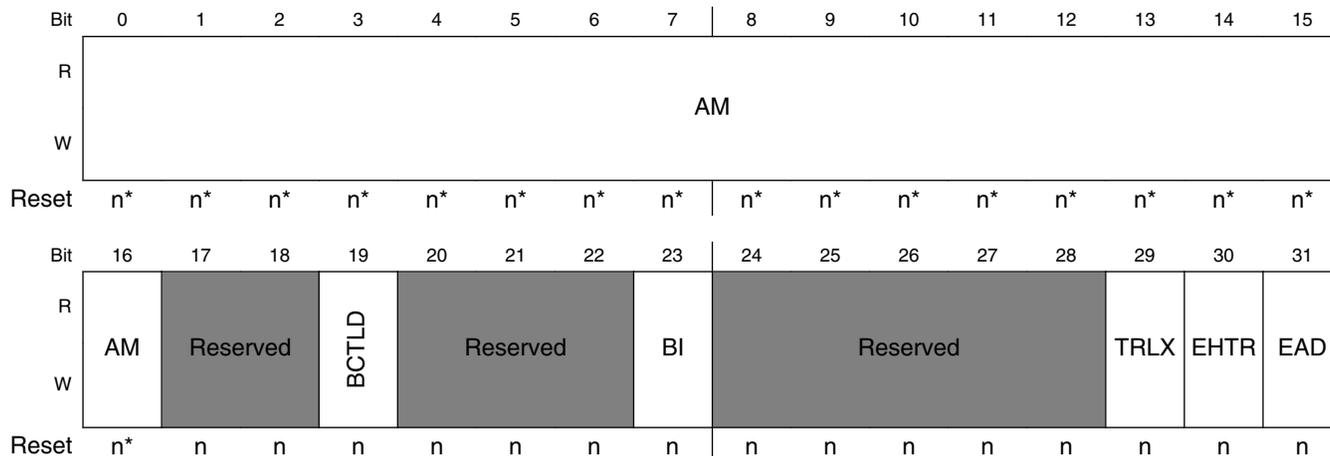
Boot Source	OR0 Reset Value
FCM (small page NAND Flash)	0000_03AE
FCM (large page NAND Flash)	0000_07AE
GPCM	0000_0FF7
eLBC not used as a boot source	0000_0F07

The address mask field of the option registers (OR *n* [AM]) masks up to 17 corresponding BR *n* [BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. The table below shows memory bank sizes from 32 Kbytes to 4 Gbytes.

Table 12-18. Memory Bank Sizes in Relation to Address Mask

AM	Memory Bank Size
0000_0000_0000_0000_0	4 Gbytes
1000_0000_0000_0000_0	2 Gbytes
1100_0000_0000_0000_0	1 Gbyte
1110_0000_0000_0000_0	512 Mbytes
1111_0000_0000_0000_0	256 Mbytes
1111_1000_0000_0000_0	128 Mbytes
1111_1100_0000_0000_0	64 Mbytes
1111_1110_0000_0000_0	32 Mbytes
1111_1111_0000_0000_0	16 Mbytes
1111_1111_1000_0000_0	8 Mbytes
1111_1111_1100_0000_0	4 Mbytes
1111_1111_1110_0000_0	2 Mbytes
1111_1111_1111_0000_0	1 Mbyte
1111_1111_1111_1000_0	512 Kbytes
1111_1111_1111_1100_0	256 Kbytes
1111_1111_1111_1110_0	128 Kbytes
1111_1111_1111_1111_0	64 Kbytes
1111_1111_1111_1111_1	32 Kbytes

Address: 5000h base + 4h offset = 5004h



* Notes:

- AM field: See Table "Reset value of OR0 Register" above for the OR0 reset value. All other option registers have all bits cleared.

eLBC_ORu0 field descriptions

Field	Description
0–16 AM	UPM address mask. Masks corresponding BR <i>n</i> bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked. 1 The corresponding address bits are used in the comparison with address pins.
17–18 -	This field is reserved. Reserved
19 BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.
20–22 -	This field is reserved. Reserved
23 BI	Burst inhibit. Indicates if this memory bank supports burst accesses. 0 The bank supports burst accesses. 1 The bank does not support burst accesses. The selected UPM executes burst accesses as a series of single accesses.
24–28 -	This field is reserved. Reserved
29 TRLX	Timing relaxed. Works in conjunction with EHTR to extend hold time on read accesses.
30 EHTR	Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access.

Table continues on the next page...

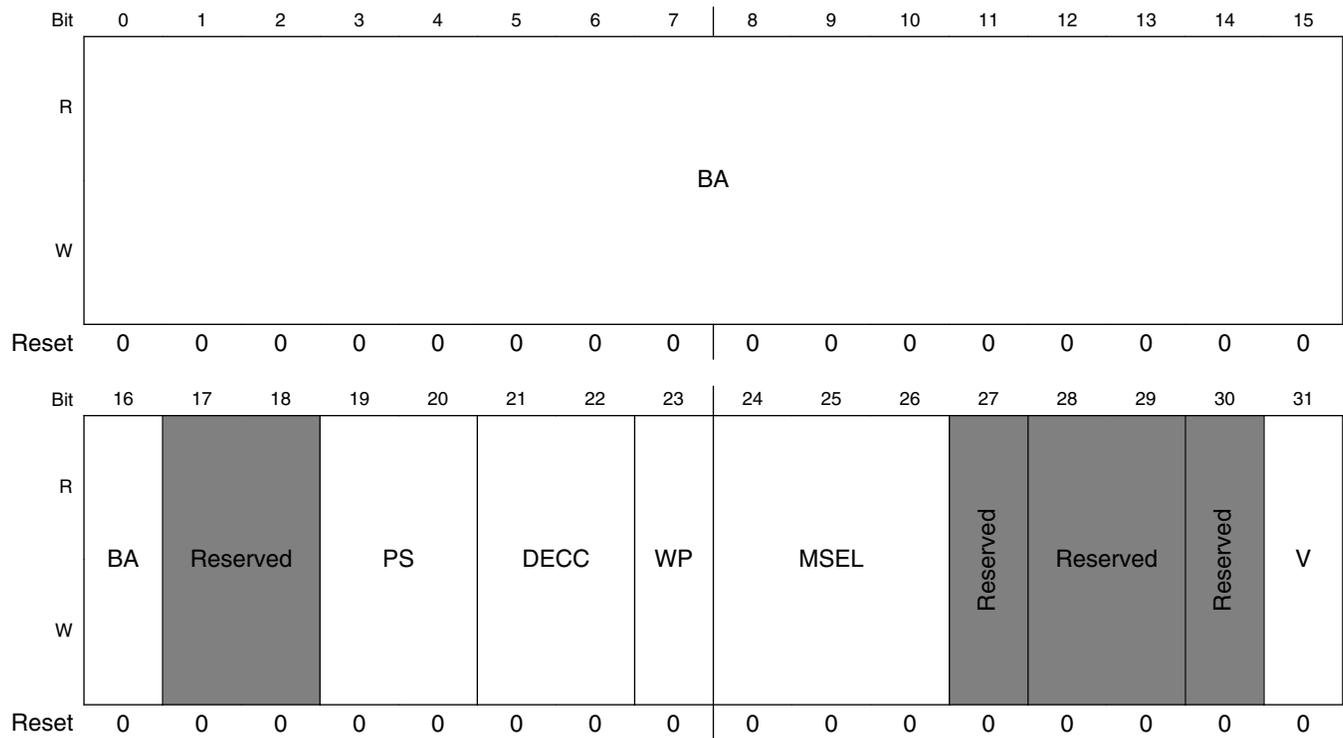
eLBC_ORu0 field descriptions (continued)

Field	Description		
	Table 12-19. TRLX and EHTR		
	TRLX	EHTR	Meaning
	0	0	The memory controller generates normal timing. No additional cycles are inserted.
	0	1	1 idle clock cycle is inserted.
	1	0	4 idle clock cycles are inserted.
	1	1	8 idle clock cycles are inserted.
31 EAD	External address latch delay. Allow extra bus clock cycles when using external address latch (LALE).		
	0	No additional bus clock cycles (LALE asserted for one bus clock cycle only)	
	1	Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by LCRR[EADC]).	

12.3.5 Base register n (eLBC_BRn)

The base registers (BR *n*) contain the base address and address types for each memory bank. The memory controller uses this information to compare the address bus value with the current address accessed. Each register (bank) includes a memory attribute and selects the machine for memory operation handling. Note that after system reset, BR0[V] is set, BR1[V]-BR7[V] are cleared, and the value of BR0[PS] reflects the initial port size configured by the boot ROM location power-on configuration settings .

Address: 5000h base + 8h offset + (8d × i), where i=0d to 6d



eLBC_BRn field descriptions

Field	Description
0–16 BA	Base address. The upper 17 bits of each base register are compared to the address on the address bus to determine if the bus master is accessing a memory bank controlled by the memory controller. Used with the address mask bits OR <i>n</i> [AM].
17–18 -	This field is reserved. Reserved
19–20 PS	Port size. Specifies the port size of this memory region. For BR0, PS is configured from the boot ROM location power-on reset configuration setting . For all other banks the value is reset to 00 (port size not defined). 00 Reserved

Table continues on the next page...

eLBC_BR n field descriptions (continued)

Field	Description
	01 8-bit (supported for GPCM, UPM, FCM) 10 16-bit (supported for GPCM, UPM) 11 Reserved
21–22 DECC	Specifies the method for data error checking. 00 Data error checking disabled , but normal parity generation for GPCM and UPM . No ECC generation for FCM. 01 ECC checking is enabled, but ECC generation is disabled, for FCM on full-page transfers. Normal parity generation and checking for GPCM and UPM. 10 ECC checking and generation are enabled for FCM on full-page transfers. 11 Reserved
23 WP	Write protect. 0 Read and write accesses are allowed. 1 Only read accesses are allowed. The memory controller does not assert \overline{LCSn} on write cycles to this memory bank. LTESR[WP] is set (if WP is set) if a write to this memory bank is attempted, and a local bus error interrupt is generated (if enabled), terminating the cycle.
24–26 MSEL	Machine select. Specifies the machine to use for handling memory operations. 000 GPCM (possible reset value) 001 FCM (possible reset value) 010 Reserved 011 Reserved 100 UPMA 101 UPMB 110 UPMC 111 Reserved
27 -	This field is reserved. Reserved
28–29 -	This field is reserved. Reserved
30 -	This field is reserved. Reserved
31 V	Valid bit. Indicates that the contents of the BR n and OR n pair are valid. \overline{LCSn} does not assert unless V is set (an access to a region that has no valid bit set may cause a bus time-out). After a system reset, only BR0[V] is set. 0 This bank is invalid. 1 This bank is valid.

12.3.6 Options register n layout for GPCM Mode (eLBC_OR gn)

The OR n registers define the sizes of memory banks and access attributes. The OR n attribute bits support the following three modes of operation as defined by BR n [MSEL]:

- GPCM mode

- FCM mode
- UPM mode

The OR *n* registers are interpreted differently depending on which of the three machine types is selected for that bank; this section describes the interpretation when the GPCM controls bank *n*.

Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options , specifically the `cfg_rom_loc` power-on-reset pins .

Table 12-23. Reset value of OR0 Register

Boot Source	OR0 Reset Value
FCM (small page NAND Flash)	0000_03AE
FCM (large page NAND Flash)	0000_07AE
GPCM	0000_0FF7
eLBC not used as a boot source	0000_0F07

The address mask field of the option registers (OR *n* [AM]) masks up to 17 corresponding BR *n* [BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. The table below shows memory bank sizes from 32 Kbytes to 4 Gbytes.

Table 12-24. Memory Bank Sizes in Relation to Address Mask

AM	Memory Bank Size
0000_0000_0000_0000_0	4 Gbytes
1000_0000_0000_0000_0	2 Gbytes
1100_0000_0000_0000_0	1 Gbyte
1110_0000_0000_0000_0	512 Mbytes
1111_0000_0000_0000_0	256 Mbytes
1111_1000_0000_0000_0	128 Mbytes
1111_1100_0000_0000_0	64 Mbytes
1111_1110_0000_0000_0	32 Mbytes
1111_1111_0000_0000_0	16 Mbytes
1111_1111_1000_0000_0	8 Mbytes
1111_1111_1100_0000_0	4 Mbytes
1111_1111_1110_0000_0	2 Mbytes
1111_1111_1111_0000_0	1 Mbyte
1111_1111_1111_1000_0	512 Kbytes
1111_1111_1111_1100_0	256 Kbytes

Table continues on the next page...

Table 12-24. Memory Bank Sizes in Relation to Address Mask (continued)

AM	Memory Bank Size
1111_1111_1111_1110_0	128 Kbytes
1111_1111_1111_1111_0	64 Kbytes
1111_1111_1111_1111_1	32 Kbytes

Address: 5000h base + Ch offset + (8d × i), where i=0d to 6d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	AM															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	AM	Reserved	BCTLD	CSNT	ACS	XACS	SCY			SETA	TRLX	EHTR	EAD			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_ORgn field descriptions

Field	Description
0–16 AM	GPCM address mask. Masks corresponding BRn bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked and therefore don't care for address checking. 1 Corresponding address bits are used in the comparison between base and transaction addresses.
17–18 -	This field is reserved. Reserved
19 BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.
20 CSNT	Chip select negation time. Determines when \overline{LCSn} and \overline{LWE} are negated during an external memory write access handled by the GPCM, provided that $ACS \neq 00$ (when $ACS = 00$, only \overline{LWE} is affected by the setting of CSNT). This helps meet address/data hold times for slow memories and peripherals. 0 \overline{LCSn} and \overline{LWE} are negated normally. 1 \overline{LCSn} and \overline{LWE} are negated one quarter of a bus clock cycle earlier.
21–22 ACS	Address to chip-select setup. Determines the delay of the \overline{LCSn} assertion relative to the address change when the external memory access is handled by the GPCM. At system reset, $ORO[ACS] = 11$. 00 \overline{LCSn} is output at the same time as the address lines. Note that this overrides the value of CSNT such that $CSNT = 0$. 01 Reserved.

Table continues on the next page...

eLBC_ORgn field descriptions (continued)

Field	Description															
	10 $\overline{\text{LCSn}}$ is output one quarter bus clock cycle after the address lines. 11 $\overline{\text{LCSn}}$ is output one half bus clock cycle after the address lines.															
23 XACS	Extra address to chip-select setup. Setting this bit increases the delay of the $\overline{\text{LCSn}}$ assertion relative to the address change when the external memory access is handled by the GPCM. After a system reset, $\text{OR0}[\text{XACS}] = 1$. 0 Address to chip-select setup is determined by $\text{ORx}[\text{ACS}]$. 1 Address to chip-select setup is extended (see Table 12-190 and Table 12-191).															
24–27 SCY	Cycle length in bus clocks. Determines the number of wait states inserted in the bus cycle, when the GPCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings. After a system reset, $\text{OR0}[\text{SCY}] = 1111$. 0000 No wait states 0001 1 bus clock cycle wait state ... 1111 15 bus clock cycle wait states															
28 SETA	External address termination. 0 Access is terminated internally by the memory controller unless the external device asserts $\overline{\text{LGTA}}$ earlier to terminate the access. 1 Access is terminated externally by asserting the $\overline{\text{LGTA}}$ external pin. (Only $\overline{\text{LGTA}}$ can terminate the access).															
29 TRLX	Timing relaxed. Modifies the settings of timing parameters for slow memories or peripherals. 0 Normal timing is generated by the GPCM. 1 Relaxed timing on the following parameters: <ul style="list-style-type: none"> • Adds an additional cycle between the address and control signals (only if ACS is not equal to 00). • Doubles the number of wait states specified by SCY, providing up to 30 wait states. • Works in conjunction with EHTR to extend hold time on read accesses. • $\overline{\text{LCSn}}$ (only if ACS is not equal to 00) and $\overline{\text{LWE}}$ signals are negated one cycle earlier during writes. 															
30 EHTR	Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. <p style="text-align: center;">Table 12-25. TRLX and EHTR</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 15%;">TRLX</th> <th style="width: 15%;">EHTR</th> <th style="width: 70%;">Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The memory controller generates normal timing. No additional cycles are inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table>	TRLX	EHTR	Meaning	0	0	The memory controller generates normal timing. No additional cycles are inserted.	0	1	1 idle clock cycle is inserted.	1	0	4 idle clock cycles are inserted.	1	1	8 idle clock cycles are inserted.
TRLX	EHTR	Meaning														
0	0	The memory controller generates normal timing. No additional cycles are inserted.														
0	1	1 idle clock cycle is inserted.														
1	0	4 idle clock cycles are inserted.														
1	1	8 idle clock cycles are inserted.														
31 EAD	External address latch delay. Allow extra bus clock cycles when using external address latch (LALE). 0 No additional bus clock cycles (LALE asserted for one bus clock cycle only) 1 Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by $\text{LCRR}[\text{EADC}]$).															

12.3.7 Options register n layout for FCM Mode (eLBC_ORfn)

The OR *n* registers define the sizes of memory banks and access attributes. The OR *n* attribute bits support the following three modes of operation as defined by BR *n* [MSEL]:

- GPCM mode
- FCM mode
- UPM mode

The OR *n* registers are interpreted differently depending on which of the three machine types is selected for that bank; this section describes the interpretation when the FCM controls bank *n*.

Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options, specifically the `cfg_rom_loc` power-on-reset pins.

Table 12-31. Reset value of OR0 Register

Boot Source	OR0 Reset Value
FCM (small page NAND Flash)	0000_03AE
FCM (large page NAND Flash)	0000_07AE
GPCM	0000_0FF7
eLBC not used as a boot source	0000_0F07

The address mask field of the option registers (OR *n* [AM]) masks up to 17 corresponding BR *n* [BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. The table below shows memory bank sizes from 32 Kbytes to 4 Gbytes.

Table 12-32. Memory Bank Sizes in Relation to Address Mask

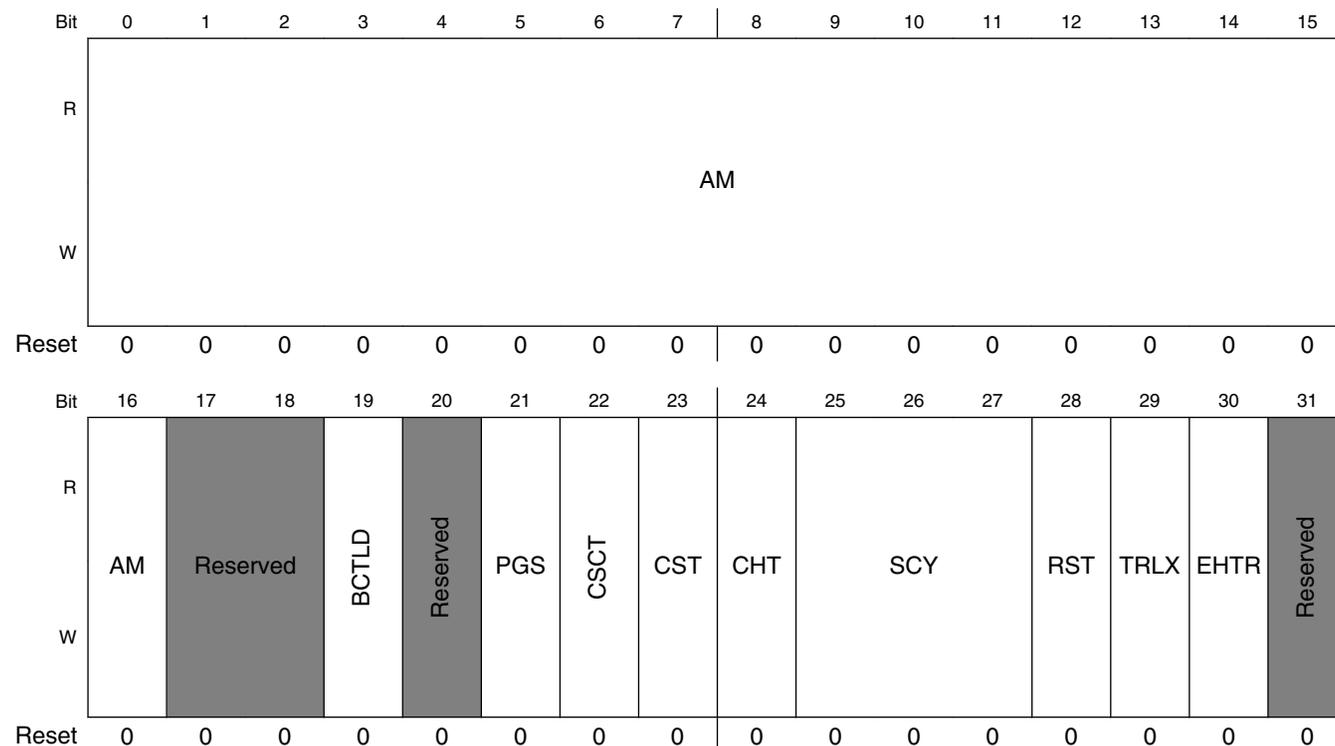
AM	Memory Bank Size
0000_0000_0000_0000_0	4 Gbytes
1000_0000_0000_0000_0	2 Gbytes
1100_0000_0000_0000_0	1 Gbyte
1110_0000_0000_0000_0	512 Mbytes
1111_0000_0000_0000_0	256 Mbytes
1111_1000_0000_0000_0	128 Mbytes
1111_1100_0000_0000_0	64 Mbytes

Table continues on the next page...

Table 12-32. Memory Bank Sizes in Relation to Address Mask (continued)

AM	Memory Bank Size
1111_1110_0000_0000_0	32 Mbytes
1111_1111_0000_0000_0	16 Mbytes
1111_1111_1000_0000_0	8 Mbytes
1111_1111_1100_0000_0	4 Mbytes
1111_1111_1110_0000_0	2 Mbytes
1111_1111_1111_0000_0	1 Mbyte
1111_1111_1111_1000_0	512 Kbytes
1111_1111_1111_1100_0	256 Kbytes
1111_1111_1111_1110_0	128 Kbytes
1111_1111_1111_1111_0	64 Kbytes
1111_1111_1111_1111_1	32 Kbytes

Address: 5000h base + Ch offset + (8d × i), where i=0d to 6d



eLBC_ORfn field descriptions

Field	Description
0–16 AM	FCM address mask. Masks corresponding BRn bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. 0 Corresponding address bits are masked. 1 Corresponding address bits are used in the comparison between base and transaction addresses.

Table continues on the next page...

eLBC_ORfn field descriptions (continued)

Field	Description															
17-18 -	This field is reserved. Reserved															
19 BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.															
20 -	This field is reserved. Reserved															
21 PGS	NAND Flash EEPROM page size, buffer size, and block size. 0 Page size of 512 main area bytes plus 16 spare area bytes (small page devices); FCM RAM buffers are 1 Kbyte each; Flash block size of 16 Kbytes. 1 Page size of 2048 main area bytes plus 64 spare area bytes (large page devices); FCM RAM buffers are 4 Kbytes each; Flash block size of 128 Kbytes.															
22 CSCT	Chip select to command time. Determines how far in advance \overline{LCSn} is asserted prior to any bus activity during a NAND Flash access handled by the FCM. This helps meet chip-select setup times for slow memories. Table 12-37. TRLX and CSCT															
<table border="1"> <thead> <tr> <th>TRLX</th> <th>CSCT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The chip-select is asserted 1 clock cycle before any command.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The chip-select is asserted 4 clock cycles before any command.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The chip-select is asserted 2 clock cycles before any command.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The chip-select is asserted 8 clock cycles before any command.</td> </tr> </tbody> </table>		TRLX	CSCT	Meaning	0	0	The chip-select is asserted 1 clock cycle before any command.	0	1	The chip-select is asserted 4 clock cycles before any command.	1	0	The chip-select is asserted 2 clock cycles before any command.	1	1	The chip-select is asserted 8 clock cycles before any command.
TRLX	CSCT	Meaning														
0	0	The chip-select is asserted 1 clock cycle before any command.														
0	1	The chip-select is asserted 4 clock cycles before any command.														
1	0	The chip-select is asserted 2 clock cycles before any command.														
1	1	The chip-select is asserted 8 clock cycles before any command.														
23 CST	Command setup time. Determines the delay of \overline{LFWE} assertion relative to the command, address, or data change when the external memory access is handled by the FCM. Table 12-36. TRLX and CST															
<table border="1"> <thead> <tr> <th>TRLX</th> <th>CST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is asserted coincident with any command.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is asserted 0.25 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is asserted 0.5 clock cycles after any command, address, or data.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The write-enable is asserted 1 clock cycle after any command, address, or data.</td> </tr> </tbody> </table>		TRLX	CST	Meaning	0	0	The write-enable is asserted coincident with any command.	0	1	The write-enable is asserted 0.25 clock cycles after any command, address, or data.	1	0	The write-enable is asserted 0.5 clock cycles after any command, address, or data.	1	1	The write-enable is asserted 1 clock cycle after any command, address, or data.
TRLX	CST	Meaning														
0	0	The write-enable is asserted coincident with any command.														
0	1	The write-enable is asserted 0.25 clock cycles after any command, address, or data.														
1	0	The write-enable is asserted 0.5 clock cycles after any command, address, or data.														
1	1	The write-enable is asserted 1 clock cycle after any command, address, or data.														
24 CHT	Command hold time. Determines the \overline{LFWE} negation prior to the command, address, or data change when the external memory access is handled by the FCM. Table 12-35. TRLX and CHT															
<table border="1"> <thead> <tr> <th>TRLX</th> <th>CHT</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The write-enable is negated 0.5 clock cycles before any command, address, or data change.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The write-enable is negated 1 clock cycle before any command, address, or data change.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The write-enable is negated 1.5 clock cycles before any command, address, or data change.</td> </tr> </tbody> </table>		TRLX	CHT	Meaning	0	0	The write-enable is negated 0.5 clock cycles before any command, address, or data change.	0	1	The write-enable is negated 1 clock cycle before any command, address, or data change.	1	0	The write-enable is negated 1.5 clock cycles before any command, address, or data change.			
TRLX	CHT	Meaning														
0	0	The write-enable is negated 0.5 clock cycles before any command, address, or data change.														
0	1	The write-enable is negated 1 clock cycle before any command, address, or data change.														
1	0	The write-enable is negated 1.5 clock cycles before any command, address, or data change.														

Table continues on the next page...

eLBC_ORfn field descriptions (continued)

Field	Description																
	Table 12-35. TRLX and CHT (continued)																
	1	1	The write-enable is negated 2 clock cycles before any command, address, or data change.														
25–27 SCY	<p>Cycle length in bus clocks. Determines the following:</p> <ul style="list-style-type: none"> the number of wait states inserted in command, address, or data transfer bus cycles, when the FCM handles the external memory access. Thus it is the main parameter for determining cycle length. The total cycle length depends on other timing attribute settings. the delay between command/address writes and data write cycles, or the delay between write cycles and read cycles from NAND Flash EEPROM. A delay of 4x(2+SCY) clock cycles (TRLX = 0) or 8x(2+SCY) clock cycles (TRLX = 1) is inserted between the last write and the first data transfer to/from NAND Flash devices. the delay between a command write and the first sample point of the RDY/\overline{BSY} pin (connected to LFRB). LFRB is not sampled until 8x(2+SCY) clock cycles (TRLX = 0) or 16x(2+SCY) clock cycles (TRLX = 1) have elapsed following the command. <p>000 No extra wait states 001 1 bus clock cycle wait state ... 111 7 bus clock cycle wait states</p>																
28 RST	<p>Read setup time. Determines the delay of \overline{LFRE} assertion relative to sampling of read data when the external memory access is handled by the FCM.</p> <p style="text-align: center;">Table 12-34. TRLX and RST</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>RST</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The read-enable is asserted 0.75 clock cycles prior to any wait states.</td> </tr> <tr> <td>0</td> <td>1</td> <td>The read-enable is asserted 1 clock cycle prior to any wait states.</td> </tr> <tr> <td>1</td> <td>0</td> <td>The read-enable is asserted 0.5 clock cycles prior to any wait states.</td> </tr> <tr> <td>1</td> <td>1</td> <td>The read-enable is asserted 1 clock cycle prior to any wait states.</td> </tr> </tbody> </table>		TRLX	RST	Meaning	0	0	The read-enable is asserted 0.75 clock cycles prior to any wait states.	0	1	The read-enable is asserted 1 clock cycle prior to any wait states.	1	0	The read-enable is asserted 0.5 clock cycles prior to any wait states.	1	1	The read-enable is asserted 1 clock cycle prior to any wait states.
TRLX	RST	Meaning															
0	0	The read-enable is asserted 0.75 clock cycles prior to any wait states.															
0	1	The read-enable is asserted 1 clock cycle prior to any wait states.															
1	0	The read-enable is asserted 0.5 clock cycles prior to any wait states.															
1	1	The read-enable is asserted 1 clock cycle prior to any wait states.															
29 TRLX	<p>Timing relaxed. Modifies the settings of timing parameters for slow memories.</p> <p>0 Normal timing is generated by the FCM. 1 Relaxed timing on the following parameters:</p> <ul style="list-style-type: none"> Doubles the number of clock cycles between \overline{LCSn} assertion and commands. Doubles the number of wait states specified by SCY, providing up to 14 wait states. Works in conjunction with CST and RST to extend command/address/data setup times. Adds one clock cycle to the command/address/data hold times. Works in conjunction with CBT to extend the wait time for read/busy status sampling by 16 clock cycles. Works in conjunction with EHTR to double hold time on read accesses. 																
30 EHTR	<p>Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access.</p> <p style="text-align: center;">Table 12-33. TRLX and EHTR</p> <table border="1"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>1 idle clock cycle is inserted.</td> </tr> </tbody> </table>		TRLX	EHTR	Meaning	0	0	1 idle clock cycle is inserted.									
TRLX	EHTR	Meaning															
0	0	1 idle clock cycle is inserted.															

Table continues on the next page...

eLBC_OR n field descriptions (continued)

Field	Description		
Table 12-33. TRLX and EHTR (continued)			
	TRLX	EHTR	Meaning
	0	1	2 idle clock cycles are inserted.
	1	0	4 idle clock cycles are inserted.
	1	1	8 idle clock cycles are inserted.
31 -	This field is reserved. Reserved		

12.3.8 Options register n layout for UPM Mode (eLBC_OR n)

The OR n registers define the sizes of memory banks and access attributes. The OR n attribute bits support the following three modes of operation as defined by BR n [MSEL]:

- GPCM mode
- FCM mode
- UPM mode

The OR n registers are interpreted differently depending on which of the three machine types is selected for that bank; this section describes the interpretation when the UPM controls bank n .

Because bank 0 can be used to boot, the reset value of OR0 may be different depending on power-on configuration options, specifically the `cfg_rom_loc` power-on-reset pins.

Table 12-47. Reset value of OR0 Register

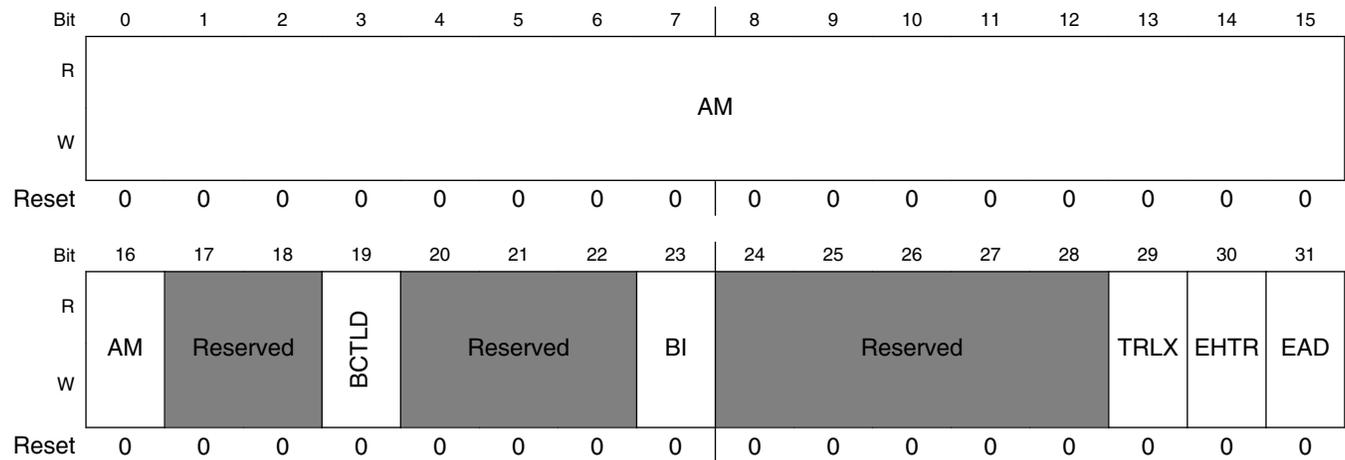
Boot Source	OR0 Reset Value
FCM (small page NAND Flash)	0000_03AE
FCM (large page NAND Flash)	0000_07AE
GPCM	0000_0FF7
eLBC not used as a boot source	0000_0F07

The address mask field of the option registers (OR n [AM]) masks up to 17 corresponding BR n [BA] fields. The 15 LSBs of the 32-bit internal transaction address do not participate in bank address matching in selecting a bank for access. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map. The table below shows memory bank sizes from 32 Kbytes to 4 Gbytes.

Table 12-48. Memory Bank Sizes in Relation to Address Mask

AM	Memory Bank Size
0000_0000_0000_0000_0	4 Gbytes
1000_0000_0000_0000_0	2 Gbytes
1100_0000_0000_0000_0	1 Gbyte
1110_0000_0000_0000_0	512 Mbytes
1111_0000_0000_0000_0	256 Mbytes
1111_1000_0000_0000_0	128 Mbytes
1111_1100_0000_0000_0	64 Mbytes
1111_1110_0000_0000_0	32 Mbytes
1111_1111_0000_0000_0	16 Mbytes
1111_1111_1000_0000_0	8 Mbytes
1111_1111_1100_0000_0	4 Mbytes
1111_1111_1110_0000_0	2 Mbytes
1111_1111_1111_0000_0	1 Mbyte
1111_1111_1111_1000_0	512 Kbytes
1111_1111_1111_1100_0	256 Kbytes
1111_1111_1111_1110_0	128 Kbytes
1111_1111_1111_1111_0	64 Kbytes
1111_1111_1111_1111_1	32 Kbytes

Address: 5000h base + Ch offset + (8d × i), where i=0d to 6d


eLBC_ORun field descriptions

Field	Description
0–16 AM	UPM address mask. Masks corresponding BR <i>n</i> bits. Masking address bits independently allows external devices of different size address ranges to be used. Address mask bits can be set or cleared in any order in the field, allowing a resource to reside in more than one area of the address map.
	0 Corresponding address bits are masked. 1 The corresponding address bits are used in the comparison with address pins.

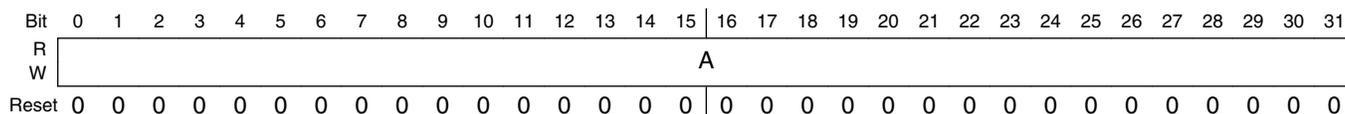
Table continues on the next page...

eLBC_ORun field descriptions (continued)

Field	Description															
17–18 -	This field is reserved. Reserved															
19 BCTLD	Buffer control disable. Disables assertion of LBCTL during access to the current memory bank. 0 LBCTL is asserted upon access to the current memory bank. 1 LBCTL is not asserted upon access to the current memory bank.															
20–22 -	This field is reserved. Reserved															
23 BI	Burst inhibit. Indicates if this memory bank supports burst accesses. 0 The bank supports burst accesses. 1 The bank does not support burst accesses. The selected UPM executes burst accesses as a series of single accesses.															
24–28 -	This field is reserved. Reserved															
29 TRLX	Timing relaxed. Works in conjunction with EHTR to extend hold time on read accesses.															
30 EHTR	Extended hold time on read accesses. Indicates with TRLX how many cycles are inserted between a read access from the current bank and the next access. Table 12-49. TRLX and EHTR															
<table border="1"> <thead> <tr> <th>TRLX</th> <th>EHTR</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>The memory controller generates normal timing. No additional cycles are inserted.</td> </tr> <tr> <td>0</td> <td>1</td> <td>1 idle clock cycle is inserted.</td> </tr> <tr> <td>1</td> <td>0</td> <td>4 idle clock cycles are inserted.</td> </tr> <tr> <td>1</td> <td>1</td> <td>8 idle clock cycles are inserted.</td> </tr> </tbody> </table>		TRLX	EHTR	Meaning	0	0	The memory controller generates normal timing. No additional cycles are inserted.	0	1	1 idle clock cycle is inserted.	1	0	4 idle clock cycles are inserted.	1	1	8 idle clock cycles are inserted.
TRLX	EHTR	Meaning														
0	0	The memory controller generates normal timing. No additional cycles are inserted.														
0	1	1 idle clock cycle is inserted.														
1	0	4 idle clock cycles are inserted.														
1	1	8 idle clock cycles are inserted.														
31 EAD	External address latch delay. Allow extra bus clock cycles when using external address latch (LALE). 0 No additional bus clock cycles (LALE asserted for one bus clock cycle only) 1 Extra bus clock cycles are added (LALE is asserted for the number of bus clock cycles specified by LCRR[EADC]).															

12.3.9 UPM address register (eLBC_MAR)

Address: 5000h base + 68h offset = 5068h



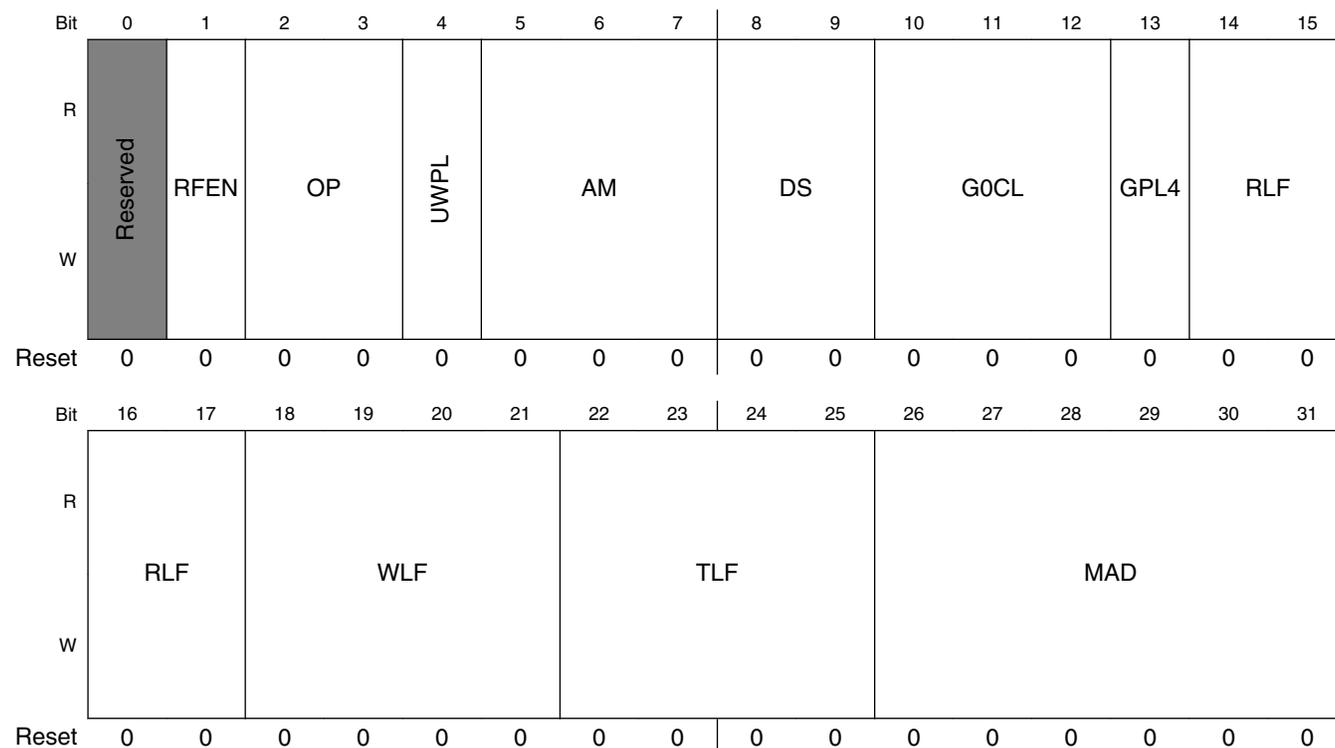
eLBC_MAR field descriptions

Field	Description
0–31 A	Address that can be output to the address signals under control of the AMX bits in the UPM RAM word.

12.3.10 UPMn mode register (eLBC_MnMR)

The UPM machine mode registers (MAMR, MBMR, and MCMR) contain the configuration for the three UPMs.

Address: 5000h base + 70h offset + (4d × i), where i=0d to 2d



eLBC_MnMR field descriptions

Field	Description
0 -	This field is reserved. Reserved
1 RFEN	Refresh enable. Indicates that the UPM needs refresh services. This bit must be set for UPMA (refresh executor) if refresh services are required on any UPM assigned chip selects. If MAMR[RFEN] = 0, no refresh services can be provided, even if UPMB and/or UPMC have their RFEN bit set. 0 Refresh services are not required 1 Refresh services are required

Table continues on the next page...

eLBC_MnMR field descriptions (continued)

Field	Description
2–3 OP	<p>Command opcode. Determines the command executed by the UPM <i>n</i> when a memory access hits a UPM assigned bank.</p> <p>00 Normal operation</p> <p>01 Write to UPM array. On the next memory access that hits a UPM assigned bank, write the contents of the MDR into the RAM location pointed to by MAD. After the access, MAD is automatically incremented.</p> <p>10 Read from UPM array. On the next memory access that hits a UPM assigned bank, read the contents of the RAM location pointed to by MAD into the MDR. After the access, MAD is automatically incremented.</p> <p>11 Run pattern. On the next memory access that hits a UPM assigned bank, run the pattern written in the RAM array. The pattern run starts at the location pointed to by MAD and continues until the LAST bit is set in the RAM word.</p>
4 UWPL	<p>LUPWAIT polarity active low. Sets the polarity of the LUPWAIT pin when in UPM mode.</p> <p>0 LUPWAIT is active high.</p> <p>1 LUPWAIT is active low.</p>
5–7 AM	<p>Address multiplex size. Determines how the address of the current memory cycle can be output on the address pins. This field is needed when interfacing with devices requiring row and column addresses multiplexed on the same pins. See Address multiplexing (AMX) for more information.</p> <p>000 Internal transaction address a[8:23] driven on LA [16:31]; LAD[0:15] driven low.</p> <p>001 Internal transaction address a[7:22] driven on LA [16:31]; LAD[0:15] driven low.</p> <p>010 Internal transaction address a[6:21] driven on LA [16:31]; LAD[0:15] driven low.</p> <p>011 Internal transaction address a[5:20] driven on LA [16:31]; LAD[0:15] driven low.</p> <p>100 Internal transaction address a[4:19] driven on LA [16:31]; LAD[0:15] driven low.</p> <p>101 Internal transaction address a[3:18] driven on LA [16:31]; LAD[0:15] driven low.</p> <p>110 Reserved</p> <p>111 Reserved</p>
8–9 DS	<p>Disable timer period. Guarantees a minimum time between accesses to the same memory bank controlled by UPM <i>n</i>. The disable timer is turned on by the TODT bit in the RAM array word, and when expired, the UPM <i>n</i> allows the machine access to handle a memory pattern to the same bank. Accesses to a different bank by the same UPM <i>n</i> is also allowed. To avoid conflicts between successive accesses to different banks, the minimum pattern in the RAM array for a request serviced, should not be shorter than the period established by DS.</p> <p>00 1-bus clock cycle disable period</p> <p>01 2-bus clock cycle disable period</p> <p>10 3-bus clock cycle disable period</p> <p>11 4-bus clock cycle disable period</p>
10–12 GOCL	<p>General line 0 control. Determines which logical address line can be output to the LGPL0 pin when the UPM <i>n</i> is selected to control the memory access.</p> <p>000 A12</p> <p>001 A11</p> <p>010 A10</p> <p>011 A9</p> <p>100 A8</p> <p>101 A7</p>

Table continues on the next page...

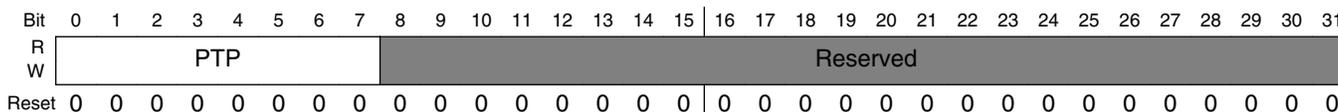
eLBC_MnMR field descriptions (continued)

Field	Description														
	110 A6 111 A5														
13 GPL4	LGPL4 output line disable. Determines how the LGPL4/LUPWAIT pin is controlled by the corresponding bits in the UPM <i>n</i> array. See Table 12-201 . Table 12-158. Effect of GPL4 Settings <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th rowspan="2">Value</th> <th rowspan="2">LGPL4/LUPWAIT Pin Function</th> <th colspan="2">Interpretation of UPM Word Bits</th> </tr> <tr> <th>G4T1/DLT3</th> <th>G4T3/WAEN</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>LGPL4 (output)</td> <td>G4T1</td> <td>G4T3</td> </tr> <tr> <td>1</td> <td>LUPWAIT (input)</td> <td>DLT3</td> <td>WAEN</td> </tr> </tbody> </table>	Value	LGPL4/LUPWAIT Pin Function	Interpretation of UPM Word Bits		G4T1/DLT3	G4T3/WAEN	0	LGPL4 (output)	G4T1	G4T3	1	LUPWAIT (input)	DLT3	WAEN
Value	LGPL4/LUPWAIT Pin Function			Interpretation of UPM Word Bits											
		G4T1/DLT3	G4T3/WAEN												
0	LGPL4 (output)	G4T1	G4T3												
1	LUPWAIT (input)	DLT3	WAEN												
14–17 RLF	Read loop field. Determines the number of times a loop defined in the UPM <i>n</i> will be executed for a burst- or single-beat read pattern or when $M \times MR[OP] = 11$ (run command) 0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15														
18–21 WLF	Write loop field. Determines the number of times a loop defined in the UPM <i>n</i> will be executed for a burst- or single-beat write pattern. 0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15														
22–25 TLF	Refresh loop field. Determines the number of times a loop defined in the UPM <i>n</i> will be executed for a refresh service pattern. 0000 16 0001 1 0010 2 0011 3 ... 1110 14 1111 15														
26–31 MAD	Machine address. RAM address pointer for the command executed. This field is incremented by 1, each time the UPM is accessed and the OP field is set to WRITE or READ. Address range is 64 words per UPM <i>n</i> .														

12.3.11 Memory refresh timer prescaler register (eLBC_MRTPR)

The refresh timer prescaler register (MRTPR) is used to divide the platform clock to provide the UPM refresh timers clock.

Address: 5000h base + 84h offset = 5084h



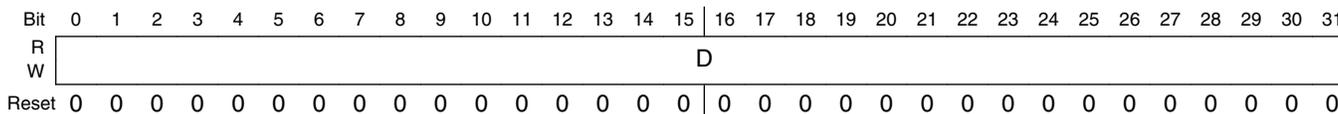
eLBC_MRTPR field descriptions

Field	Description
0–7 PTP	Refresh timers prescaler. Determines the period of the refresh timers input clock. The platform clock is divided by PTP except when the value is 00000_0000, which represents the maximum divider of 256.
8–31 -	This field is reserved. Reserved

12.3.12 UPM data register (eLBC_MDRu)

The memory data register (MDR) contains data written to or read from the RAM array for UPM read or write commands. MDR also contains data written to or read from an external NAND Flash EEPROM for FCM write address, write data, and read status commands. MDR must be set up before issuing a write command to the UPM, or before issuing a FCM operation sequence that uses MDR to source address or data bytes.

Address: 5000h base + 88h offset = 5088h



eLBC_MDRu field descriptions

Field	Description
0–31 D	D is the data to be read or written into the RAM array when a write or read command is supplied to the UPM (M x MR[OP] = 01 or M x MR[OP] = 10).

12.3.13 FCM data register (eLBC_MDRf)

The memory data register (MDR) contains data written to or read from the RAM array for UPM read or write commands. MDR also contains data written to or read from an external NAND Flash EEPROM for FCM write address, write data, and read status commands. MDR must be set up before issuing a write command to the UPM, or before issuing a FCM operation sequence that uses MDR to source address or data bytes.

Address: 5000h base + 88h offset = 5088h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	AS3								AS2								AS1								AS0							
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_MDRf field descriptions

Field	Description
0–7 AS3	AS3 is the fourth byte of address sent by a custom address write operation, or the fourth byte of data read from a read status operation.
8–15 AS2	AS2 is the third byte of address sent by a custom address write operation, or the third byte of data read from a read status operation.
16–23 AS1	AS1 is the second byte of address sent by a custom address write operation, or the second byte of data read from a read status operation.
24–31 AS0	AS0 is the first byte of address sent by a custom address write operation, or the first byte of data read from a read status operation.

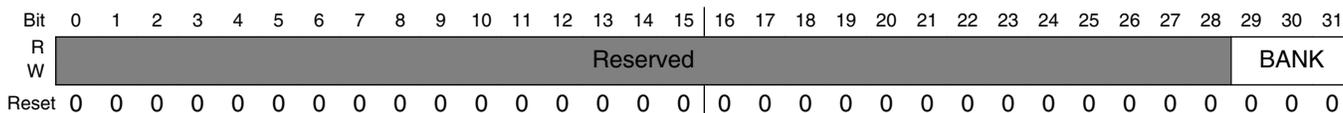
12.3.14 Special operation initiation register (eLBC_LSOR)

The special operation initiation register (LSOR) is used by software to trigger a special operation on the indicated bank. Writing to LSOR activates a special operation on bank LSOR[BANK] provided that the bank is valid and controlled by a memory controller whose mode OP field is set to a value other than "normal operation." If eLBC is currently busy with a memory transaction, writing LSOR completes immediately, but the special operation request is queued until eLBC can service it. To avoid race conditions between software and a busy eLBC, registers that affect currently running special operation and LSOR must not be rewritten before a pending special operation has been completed. The UPM and FCM have different indications of when such special operations are completed. The behavior of eLBC is unpredictable if special operation modes are altered between LSOR being written and the relevant memory controller completing that access.

Enhanced Local Bus Controller (eLBC) Memory Map

UPM special operation modes are set in registers $M \times MR[OP]$, see [UPMn mode register \(eLBC_MnMR\)](#). FCM special operation modes are set in $FMR[OP]$, see [Flash mode register \(eLBC_FMR\)](#). Writing LSOR has the same effect as setting a special controller mode and performing a dummy access to a bank associated with the controller in question, but use of LSOR avoids changing settings for the address space occupied by the bank. More details of special operation sequences appear in [UPM programming example \(two sequential writes to the RAM array\)](#).

Address: 5000h base + 90h offset = 5090h



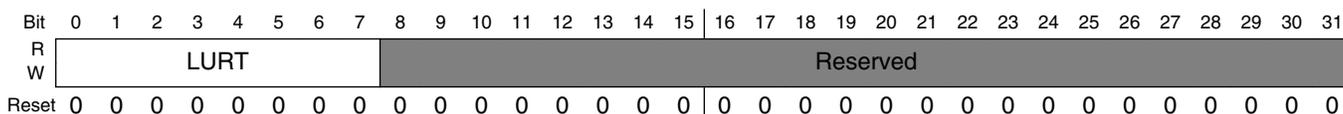
eLBC_LSOR field descriptions

Field	Description
0–28 -	This field is reserved. Reserved
29–31 BANK	Bank on which a special operation is initiated. If the bank identified by BANK is marked valid (BRn[V] set) and the bank is controlled by a memory controller whose current mode OP is non-zero-or a special operation-eLBC will request the special operation to be activated on the selected bank when this field is written. Otherwise, writing this field has no effect. 000 Bank 0 is triggered for special operation ... 111 Bank 7 is triggered for special operation

12.3.15 UPM refresh timer (eLBC_LURT)

The UPM refresh timer (LURT) generates a refresh request for all valid banks that selected a UPM machine and are refresh-enabled ($M \times MR[RFEN] = 1$). Each time the timer expires, a qualified bank generates a refresh request using the selected UPM. The qualified banks rotate their requests.

Address: 5000h base + A0h offset = 50A0h



eLBC_LURT field descriptions

Field	Description
0–7 LURT	<p>UPM refresh timer period. Determines, along with the timer prescaler (MRTPR), the timer period according to the following equation:</p> $\text{TimerPeriod} = \frac{\text{LURT}}{\left(\frac{\text{InputClock}}{\text{MRTPR}[\text{PTP}]}\right)}$ <p>Example: For a 266-MHz system clock and a required service rate of 15.6 μs, given MRTPR[PTP] = 32, the LURT value should be 128 decimal. 128/(266 MHz/32) = 15.4 μs, which is less than the required service period of 15.6 μs.</p> <p>Note that the reset value (0x00) sets the maximum period to 256 x MRTPR[PTP] system clock cycles.</p>
8–31 -	This field is reserved. Reserved

12.3.16 Transfer error status register (eLBC_LTESR)

The transfer error status register (LTESR) indicates the cause of an error or event. LTESR is a write-1-to-clear register. Reading LTESR occurs normally; however, write operations can clear but not set bits. A bit is cleared whenever the register is written, and the data in the corresponding bit location is a 1. For example, to clear only the write protect error bit (LTESR[WP]) without affecting other LTESR bits, 0x0400_0000 should be written to the register. After any error/event reported by LTESR, LTEATR[V] must be cleared for LTESR to updated again. Note that error statuses are only reflected in LTESR if they have been enabled in LTEDR.

Address: 5000h base + B0h offset = 50B0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	BM	FCT	PAR	Reserved		WP	Reserved					CS	Reserved			
W	w1c	w1c	w1c			w1c						w1c				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved												UCC	CC		
W													w1c	w1c		
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_LTESR field descriptions

Field	Description
0 BM	Bus monitor time-out

Table continues on the next page...

eLBC_LTESR field descriptions (continued)

Field	Description
	<p>0 No local bus monitor time-out occurred.</p> <p>1 Local bus monitor time-out occurred. No data beat was acknowledged on the bus within LBCR[BMT] x LBCR[BMTPS] bus clock cycles from the start of a transaction.</p>
1 FCT	<p>FCM command time-out</p> <p>0 No FCM command time-out occurred.</p> <p>1 A CW0, CW1, CW2, or CW3 command issued to FCM timed-out with respect to the timer configured by FMR[CWTO].</p>
2 PAR	<p>Parity or ECC error</p> <p>0 No local bus parity error</p> <p>1 Local bus parity error (GPCM or UPM), or non-correctable ECC error (FCM). LTEATR[PB] indicates the byte lane that caused the error and LTEATR[BNK] indicates which memory controller bank was accessed.</p>
3–4 -	<p>This field is reserved. Reserved</p>
5 WP	<p>Write protect error</p> <p>0 No write protect error occurred.</p> <p>1 A write was attempted to a local bus memory region that was defined as read-only in the memory controller. Usually, in this case, a bus monitor time-out will occur (as the cycle is not automatically terminated).</p>
6–11 -	<p>This field is reserved. Reserved</p>
12 CS	<p>Chip select error</p> <p>0 No chip select error occurred.</p> <p>1 A transaction was sent to the eLBC that did not hit any memory bank.</p>
13–29 -	<p>This field is reserved. Reserved</p>
30 UCC	<p>UPM Run pattern (MxMR[OP]=11) command completion event</p> <p>0 No UPM Run pattern operation in progress, or operation pending.</p> <p>1 UPM Run pattern operation has completed, allowing software to continue processing of results.</p>
31 CC	<p>FCM command completion event</p> <p>0 No FCM operation in progress, or operation pending.</p> <p>1 FCM operation has completed, allowing software to continue processing of results.</p>

12.3.17 Transfer error disable register (eLBC_LTEDR)

The transfer error check disable register (LTEDR) is used to disable error/event checking, which are reported in LTESR. Note that control of error/event checking is independent of control of reporting of errors/events (LTEIR) through the interrupt mechanism.

Address: 5000h base + B4h offset = 50B4h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_LTEDR field descriptions

Field	Description
0 BMD	Bus monitor disable 0 Bus monitor is enabled. 1 Bus monitor is disabled, but internal bus time-outs can still occur.
1 FCTD	FCM command time-out disable 0 FCM command timer is enabled. 1 FCM command time-out is disabled, but internal FCM command timer can terminate command waits.
2 PARD	Parity and ECC error checking disabled. Note that non-correctable read errors may cause the assertion of <i>core_fault_in</i> , which causes the core to generate a machine check interrupt, unless it is disabled. 0 Parity and ECC error checking is enabled. 1 Parity and ECC error checking is disabled.
3–4 -	This field is reserved. Reserved
5 WPD	Write protect error checking disable. 0 Write protect error checking is enabled. 1 Write protect error checking is disabled.
6–11 -	This field is reserved. Reserved

Table continues on the next page...

eLBC_LTEDR field descriptions (continued)

Field	Description
12 CSD	Chip select error checking disable. 0 Chip select error checking is enabled. 1 Chip select error checking is disabled.
13–29 -	This field is reserved. Reserved
30 UCCD	UPM Run pattern command completion checking disable. 0 UPM Run pattern command completion checking is enabled. 1 UPM Run pattern command completion checking is disabled.
31 CCD	FCM command completion checking disable. 0 Command completion checking is enabled. 1 Command completion checking is disabled.

12.3.18 Transfer error interrupt register (eLBC_LTEIR)

The transfer error interrupt enable register (LTEIR) is used to send or block error/event reporting through the eLBC internal interrupt mechanism. Software should clear pending errors/events in LTESR before enabling interrupts. After an interrupt has occurred, clearing relevant LTESR error/event bits negates the interrupt.

Address: 5000h base + B8h offset = 50B8h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	BMI	FCTI	PARI	Reserved	WPI	Reserved						CSI	Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved														UCCI	CCI
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_LTEIR field descriptions

Field	Description
0 BMI	Bus monitor error interrupt enable. 0 Bus monitor error reporting is disabled. 1 Bus monitor error reporting is enabled.
1 FCTI	FCM command time-out interrupt enable. 0 FCM command time-out error reporting is disabled. 1 FCM command time-out error reporting is enabled.

Table continues on the next page...

eLBC_LTEIR field descriptions (continued)

Field	Description
2 PARI	Parity and ECC error interrupt enable. 0 Parity and ECC error reporting is disabled. 1 Parity and ECC error reporting is enabled.
3–4 -	This field is reserved. Reserved
5 WPI	Write protect error interrupt enable. 0 Write protect error reporting is disabled. 1 Write protect error reporting is enabled.
6–11 -	This field is reserved. Reserved
12 CSI	Chip select error interrupt enable. 0 Chip select error reporting is disabled. 1 Chip select error reporting is enabled.
13–29 -	This field is reserved. Reserved
30 UCCI	UPM Run pattern command completion Event interrupt enable. 0 UPM Run pattern command completion reporting is disabled. 1 UPM Run pattern command completion reporting is enabled.
31 CCI	FCM command completion Event interrupt enable. 0 Command completion reporting is disabled. 1 Command completion reporting is enabled.

12.3.19 Transfer error attributes register (eLBC_LTEATR)

The transfer error attributes register (LTEATR) captures source attributes of an error/event. After LTEATR[V] has been set, software must clear this bit to allow LTESR, LTEATR, and LTEAR to update following any subsequent events/errors.

NOTE

LTEATR may not capture accurate information for errors that occur when an FCM special operation is in progress.

Address: 5000h base + BCh offset = 50BCh

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W																
	Reserved			RWB	Reserved						SRCID					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W																
	PB				BNK								Reserved			V
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_LTEATR field descriptions

Field	Description
0–2 -	This field is reserved. Reserved
3 RWB	Transaction type for the error: 0 The transaction for the error was a write transaction. 1 The transaction for the error was a read transaction.
4–10 -	This field is reserved. Reserved
11–15 SRCID	Captures the source of the transaction when this information is provided on the internal interface to the eLBC.
16–19 PB	Parity error on byte or block . For GPCM and UPM, there are four parity error status bits, one per byte lane. A bit is set for the byte that had a parity error (bit 16 represents byte 0, the most significant byte lane). For FCM, there are at most four 512-byte page blocks (for a large page device) checked by ECC. A bit is set for the 512-byte block that had an non-correctable ECC error on read (bit 16 represents block 0, the first 512 bytes of a page; if ORx[PGS] = 0, bits 17-19 are always 0).
20–27 BNK	Memory controller bank. There is one error status bit per memory controller bank (bit 20 represents bank 0). A bit is set for the local bus memory controller bank that had an error.
28–30 -	This field is reserved. Reserved
31 V	Error attribute capture is valid. Indicates that the captured error information is valid. 0 Captured error attributes and address are not valid. 1 Captured error attributes and address are valid.

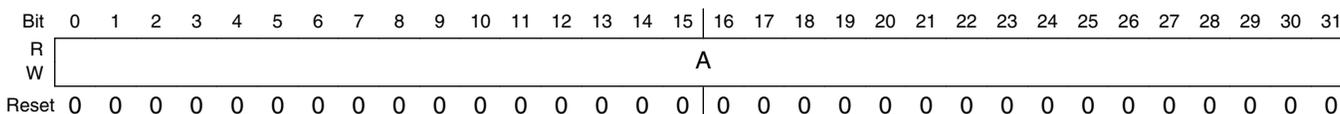
12.3.20 Transfer error address register (eLBC_LTEAR)

The transfer error address register (LTEAR) captures the address of a transaction that caused an error/event.

NOTE

LTEAR may not capture accurate information for errors that occur when an FCM special operation is in progress.

Address: 5000h base + C0h offset = 50C0h



eLBC_LTEAR field descriptions

Field	Description
0–31 A	Transaction address for the error. For GPCM and UPM, holds the 32-bit address of the transaction resulting in an error. For FCM, this register is undefined.

12.3.21 Transfer error ECC register (eLBC_LTECCR)

The transfer error ECC register (LTECCR) captures single bit and multibit errors per 512-byte sector in FCM mode. LTECCR is a write-1-to-clear register. Write operations can clear but not set bits. It captures the errors during full page read transfers on FCM command completion event, provided ECC check is enabled in BRx[DECC].

Address: 5000h base + C4h offset = 50C4h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved												SBCE			Reserved												MBUE				
W	Reserved												w1c			Reserved												w1c				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_LTECCR field descriptions

Field	Description
0–11 -	This field is reserved. Reserved
12–15 SBCE	Single bit correctable error. There are at most four 512-byte page blocks (for a large page device) checked by ECC. A bit is set for the 512-byte block that had a single bit correctable ECC error on read (bit 12 represents block 0, the first 512 bytes of a page; if ORx[PGS] = 0, bits 13-15 are always 0).
16–27 -	This field is reserved. Reserved
28–31 MBUE	Multi bit uncorrectable error There are at most four 512-byte page blocks (for a large page device) checked by ECC. A bit is set for the 512-byte block that had an uncorrectable ECC error on read (bit 28 represents block 0, the first 512 bytes of a page; if ORx[PGS] = 0, bits 29-31 are always 0).

12.3.22 Configuration register (eLBC_LBCR)

Address: 5000h base + D0h offset = 50D0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R		Reserved										Reserved				
W	LDIS	Reserved							BCTLC	AHD	Reserved			LPBSE	EPAR	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	BMT								Reserved				BMTPS			
W	BMT								Reserved				BMTPS			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_LBCR field descriptions

Field	Description
0 LDIS	Local bus disable 0 Local bus is enabled. 1 Local bus is disabled. No internal transactions will be acknowledged.
1–7 -	This field is reserved. Reserved
8–9 BCTLC	Defines the use of LBCTL 00 LBCTL is used as $\overline{W/\overline{R}}$ control for GPCM or UPM accesses (buffer control). 01 LBCTL is used as \overline{LOE} for GPCM accesses only. 10 LBCTL is used as \overline{LWE} for GPCM accesses only. 11 Reserved
10 AHD	Address hold disable. Removes part of the hold time for LAD with respect to LALE in order to lengthen the LALE pulse. 0 During address phases on the local bus, the LALE signal negates one platform clock period prior to the address being invalidated. For instance, at 333 MHz, this provides 3 ns of additional address hold time at the external address latch. Running the platform at a lower frequency improves the hold time. 1 During address phases on the local bus, the LALE signal negates 0.5 platform clock period prior to the address being invalidated. This halves the address hold time, but extends the latch enable duration. This may be necessary for very high frequency designs.
11–13 -	This field is reserved. Reserved.
14 LPBSE	Enables parity byte select on $\overline{LGTA/LFRB/LGPL4/LUPWAIT/LPBSE}$ signal.

Table continues on the next page...

eLBC_LBCR field descriptions (continued)

Field	Description
	0 Parity byte select is disabled. $\overline{\text{LGTA}}$ /LGPL4/LPBSE signal is available for memory control as LGPL4 (output) or $\overline{\text{LGTA}}$ /LFRB /LUPWAIT (input). 1 Parity byte select is enabled. LPBSE signal is dedicated as the parity byte select output, and $\overline{\text{LGTA}}$ /LFRB /LUPWAIT is disabled.
15 EPAR	Determines odd or even parity. Writing GPCM or UPM controlled memory with EPAR = 1 and reading the memory with EPAR = 0 generates parity errors for testing. 0 Odd parity; normal, odd-parity ECC 1 Even parity; inverted, even-parity ECC
16–23 BMT	Bus monitor timing. Defines the bus monitor time-out period. The number of LCLK clock cycles to count down before a time-out error is generated is given by: if BMT = 0, then bus cycles = 256 x PS if BMT \neq 0, then bus cycles = BMT x PS where PS is set according to LBCR[BMTPS]. The value of bus cycles must not be less than 40 bus cycles for reliable operation. When BMT = 0, bus cycles = 256 x PS.
24–27 -	This field is reserved. Reserved
28–31 BMTPS	Bus monitor timer prescale. Defines the multiplier, PS, to scale LBCR[BMT] for determining bus time-outs. 0000 PS = 8 0001 PS = 16 0010 PS = 32 0011 PS = 64 0100 PS = 128 0101 PS = 256 0110 PS = 512 0111 PS = 1024 1000 PS = 2048 1001 PS = 4096 1010 PS = 8192 1011 PS = 16,384 1100 PS = 32,768 1101 PS = 65,536 1110 PS = 131,072 1111 PS = 262,144

12.3.23 Clock ratio register (eLBC_LCRR)

The clock ratio register sets the platform clock to eLBC bus frequency ratio. It also provides configuration bits for extra delay cycles for address and control signals.

NOTE

For proper operation of the system, it is required that this register setting will not be altered while local bus memories or

Enhanced Local Bus Controller (eLBC) Memory Map

devices are being accessed. Special care needs to be taken when running instructions from an eLBC memory.

Address: 5000h base + D4h offset = 50D4h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	PBYP	Reserved													EADC	
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved											CLKDIV				
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

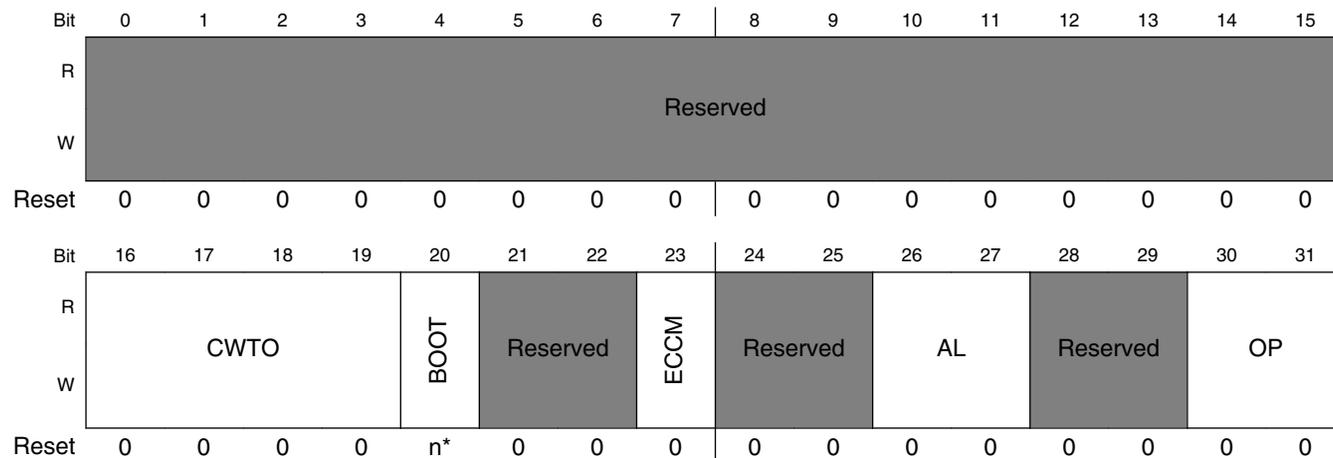
eLBC_LCRR field descriptions

Field	Description
0 PBYP	<p>PLL bypass. This bit should be set when using low bus clock frequencies (See device hardware specifications for applicable frequencies.) . When in PLL bypass mode, incoming data is captured in the middle of the bus clock cycle.</p> <p>0 The PLL is enabled. 1 The PLL is bypassed.</p>
1–13 -	<p>This field is reserved. Reserved.</p>
14–15 EADC	<p>External address delay cycles of LCLK. Defines the number of cycles for the assertion of LALE.</p> <p>00 4 01 1 10 2 11 3</p>
16–26 -	<p>This field is reserved. Reserved</p>
27–31 CLKDIV	<p>Clock divider. Sets the frequency ratio between the platform clock and the local bus clock. Only the values shown below are allowed.</p> <p>NOTE: It is critical that no transactions are being executed via the local bus while CLKDIV is being modified. As such, prior to modification, the user must ensure that code is not executing out of the local bus. Once LCRR[CLKDIV] is written, the register should be read, and then an isync should be executed.</p> <p>00000-00001 Reserved 00010 4 00011 Reserved 00100 8(default for boot ROM set to FCM) 00101-00111 Reserved 01000 16(default for boot ROM set to GPCM) 01001-11111 Reserved</p>

12.3.24 Flash mode register (eLBC_FMR)

The local bus Flash mode register (FMR) controls global operation of the FCM.

Address: 5000h base + E0h offset = 50E0h



* Notes:

- BOOT field: Bit is set if power-on-reset configuration selects FCM as the boot ROM target.

eLBC_FMR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–19 CWTO	Command wait time-out. For FCM commands that wait on $\overline{\text{LFRB}}$ being sampled high (CW0, CW1, RBW, and RSW), FCM pauses execution of the instruction sequence until either $\overline{\text{LFRB}}$ is sampled high, or a timer controlled by CTO expires, whichever occurs first. The time-out in the latter case is as follows: 0000 256 cycles of LCLK 0001 512 cycles of LCLK 0010 1024 cycles of LCLK 0011 2048 cycles of LCLK 0100 4096 cycles of LCLK 0101 8192 cycles of LCLK 0110 16,384 cycles of LCLK 0111 32,768 cycles of LCLK 1000 65,536 cycles of LCLK 1001 131,072 cycles of LCLK 1010 262,144 cycles of LCLK 1011 524,288 cycles of LCLK 1100 1,048,576 cycles of LCLK 1101 2,097,152 cycles of LCLK

Table continues on the next page...

eLBC_FMR field descriptions (continued)

Field	Description
	1110 4,194,304 cycles of LCLK 1111 8,388,608 cycles of LCLK
20 BOOT	Flash auto-boot load mode. During system boot from NAND Flash EEPROM, this bit remains set to alter the use of the FCM buffer RAM. Software should clear BOOT once FCM is to be restored to normal operation. Setting BOOT without auto-boot in progress only alters the mapping of the buffer RAM. 0 FCM is operating in normal functional mode, with an 8 Kbyte FCM buffer RAM. 1 eLBC has been configured-either from reset or by a special operation OP = 01-to autoload a 4-Kbyte boot block into the FCM buffer RAM, which maps only the 4 Kbytes of NAND Flash main data region comprising the boot block. Any access to the buffer RAM is delayed until the entire boot block has been loaded.
21-22 -	This field is reserved. Reserved
23 ECCM	ECC mode. When hardware checking and/or generation of error correcting codes (ECC) is enabled (that is, when BRn[DECC] is 01 or 10, and full page transfers are specified with FBCR[BC] = 0), ECCM sets the ECC block size and position of the ECC code word(s) in the NAND Flash spare region for both checking and generation functions. The format of the ECC code word conforms with the Samsung/Toshiba spare region assignment specifications. 0 ECC is checked/calculated over 512-Byte blocks. A 24-bit ECC is assigned to spare region bytes at offsets (N x 16) + 6 through (N x 16) + 8 for spare region N, N = 0-3. 1 ECC is checked/calculated over 512-Byte blocks. A 24-bit ECC is assigned to spare region bytes at offsets (N x 16) + 8 through (N x 16) + 10 for spare region N, N = 0-3.
24-25 -	This field is reserved. Reserved
26-27 AL	Address length. AL sets the number of address bytes issued during page address (PA) operations. However, the number of address bytes issued for column address (CA) operations is determined by the device page size (for ORn[PGS] = 0, 1 CA byte is issued; for ORn[PGS] = 1, 2 CA bytes are issued). 00 2 bytes are issued for page addresses, thus a total of 3 (ORn[PGS] = 0) or 4 (ORn[PGS] = 1) address bytes are issued for a {CA,PA} sequence 01 3 bytes are issued for page addresses, thus a total of 4 (ORn[PGS] = 0) or 5 (ORn[PGS] = 1) address bytes are issued for a {CA,PA} sequence 10 4 bytes are issued for page addresses, thus a total of 5 (ORn[PGS] = 0) or 6 (ORn[PGS] = 1) address bytes are issued for a {CA,PA} sequence 11 Reserved
28-29 -	This field is reserved. Reserved
30-31 OP	Flash operation. For OP not equal to 00, a special operation is triggered on the next write to LSOR or dummy access to a bank controlled by FCM. Once a special operation has commenced, OP is automatically reset to 00 by FCM. Individual blocks may be temporarily unlocked for erase and reprogramming operations. 00 Normal operation. All read and write accesses to banks controlled by FCM access the shared FCM buffer RAM. No bus activity is caused by this operation. 01 Simulate auto-boot block loading, and set FMR[BOOT]. Boot block loading occurs from the bank triggered on the special operation, therefore the appropriate bank configuration must be initialized prior to issuing this operation.

Table continues on the next page...

eLBC_FMR field descriptions (continued)

Field	Description
10	Execute the command sequence contained in FIR, but with write protection enabled (pin $\overline{\text{FWP}}$ asserted low) so that all Flash blocks are protected from accidental erasure and reprogramming.
11	Execute the command sequence contained in FIR, but permit the single block identified by FBAR[BLK] to be erased or reprogrammed, with pin $\overline{\text{FWP}}$ remaining high during the access.

12.3.25 Flash instruction register (eLBC_FIR)

The local bus Flash instruction register (FIR) holds a sequence of up to eight instructions for issue by the FCM. Setting FMR[OP] non-zero and writing LSOR or accessing a bank controlled by FCM causes FCM to read FIR 4 bits at a time, starting at bit 0 and continuing with adjacent 4-bit opcodes, until only NOP opcodes remain. The programmed instruction sequence of OP0, OP1,..., OP7 is performed on the activated bank, using the data buffer addressed by FPAR. If LTEIR[CCI] = 1 and LTEDR[CCD] = 0, eLBC will generate an interrupt once the entire sequence has completed, and software should examine LTEATR and clear its V bit.

Software must not alter the contents of the addressed FCM buffer, FIR, MDR, FCR, FBAR, FPAR, or FBCR while an operation is in progress-or eLBC will behave unpredictably-but software can freely modify the contents of any currently unused FCM RAM buffer in preparation for the next operation.

Address: 5000h base + E4h offset = 50E4h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_FIR field descriptions

Field	Description
0-3 OP0	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. 0000 NOP-No-operation and end of operation sequence 0001 CA-Issue current column address as set in FPAR, with length set by ORx[PGS] 0010 PA-Issue current block+page address as set in FBAR and FPAR, with length set by FMR[AL] 0011 UA-Issue user-defined address byte from next AS field in MDR 0100 CM0-Issue command from FCR[CMD0] 0101 CM1-Issue command from FCR[CMD1] 0110 CM2-Issue command from FCR[CMD2] 0111 CM3-Issue command from FCR[CMD3] 1000 WB-Write FBCR bytes of data from current FCM buffer to Flash device 1001 WS-Write one byte (8b port) of data from next AS field of MDR to Flash device 1010 RB-Read FBCR bytes of data from Flash device into current FCM RAM buffer 1011 RS-Read one byte (8b port) of data from Flash device into next AS field of MDR

Table continues on the next page...

eLBC_FIR field descriptions (continued)

Field	Description
1100	CW0-Wait for $\overline{\text{LFRB}}$ to return high or time-out, then issue command from FCR[CMD0]
1101	CW1-Wait for $\overline{\text{LFRB}}$ to return high or time-out, then issue command from FCR[CMD1]
1110	RBW-Wait for $\overline{\text{LFRB}}$ to return high or time-out, then read FBCR bytes of data from Flash device into current FCM RAM buffer
1111	RSW-Wait for $\overline{\text{LFRB}}$ to return high or time-out, then read one byte (8b port) of data from Flash device into next AS field of MDR
4–7 OP1	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. See description for OP0.
8–11 OP2	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. See description for OP0.
12–15 OP3	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. See description for OP0.
16–19 OP4	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. See description for OP0.
20–23 OP5	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. See description for OP0.
24–27 OP6	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. See description for OP0.
28–31 OP7	FCM operation codes. OP0 is executed first, followed by OP1, through to OP7. See description for OP0.

12.3.26 Flash command register (eLBC_FCR)

The local bus Flash command register (FCR) holds up to four NAND Flash EEPROM command bytes that may be referenced by opcodes in FIR during FCM operation. The values of the commands should follow the manufacturer's datasheet for the relevant NAND Flash device.

Address: 5000h base + E8h offset = 50E8h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_FCR field descriptions

Field	Description
0–7 CMD0	General purpose FCM Flash command byte 0. Opcodes in FIR that issue command index 0 write CMD0 to the NAND Flash command/data bus.
8–15 CMD1	General purpose FCM Flash command byte 1. Opcodes in FIR that issue command index 1 write CMD1 to the NAND Flash command/data bus.

Table continues on the next page...

eLBC_FCR field descriptions (continued)

Field	Description
16–23 CMD2	General purpose FCM Flash command byte 2. Opcodes in FIR that issue command index 2 write CMD2 to the NAND Flash command/data bus.
24–31 CMD3	General purpose FCM Flash command byte 3. Opcodes in FIR that issue command index 3 write CMD3 to the NAND Flash command/data bus.

12.3.27 Flash block address register (eLBC_FBAR)

The local bus Flash block address register (FBAR) locates the NAND Flash block index for the page currently accessed.

Address: 5000h base + ECh offset = 50ECh

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								BLK																							
W	Reserved								BLK																							
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_FBAR field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–31 BLK	Flash block address. The size of the NAND Flash, as configured in ORn[PGS] and FMR[AL], determines the number of bits of BLK that are issued to the EEPROM during block address phases.

12.3.28 Flash page address register [Large Page Device (ORx[PGS] = 1)] (eLBC_FPARI)

The local bus Flash page address register (FPAR) locates the current NAND Flash page in both the external NAND Flash device and FCM buffer RAM.

Address: 5000h base + F0h offset = 50F0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved														PI	
W	Reserved														PI	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	PI				MS	CI										
W	PI				MS	CI										
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_FPARI field descriptions

Field	Description
0–13 -	This field is reserved. Reserved
14–19 PI	Page index. PA indexes the page in NAND Flash EEPROM at the current block defined by FBAR, and locates the corresponding transfer buffer in the FCM buffer RAM. The LSB of PI indexes one of the two 4 Kbyte buffers in the FCM buffer RAM as follows: 0 The page is transferred to/from FCM buffer 0, address offsets 0x0000-0x0FFF 1 The page is transferred to/from FCM buffer 1, address offsets 0x1000-0x1FFF
20 MS	Main/spare region locator. In the case that FBCR[BC] = 0, MS is treated as 0. 0 Data is transferred to/from the main region of the FCM buffer; that is, the first 2048 bytes of the buffer are used as the starting address. 1 Data is transferred to/from the spare region of the FCM buffer; that is, the second 2048 bytes of the buffer are used as the starting address, but only an initial 64 bytes of spare region are defined.
21–31 CI	Column index. CI indexes the first byte to transfer to/from the main or spare region of the NAND Flash EEPROM and corresponding transfer buffer. In the case that FBCR[BC] = 0, CI is treated as 0. For MS = 0, CI can range 0x000-0x7FF; for MS = 1, CI can range 0x000-0x03F.

12.3.29 Flash page address register [Small Page Device (ORx[PGS] = 0)] (eLBC_FPARRs)

The local bus Flash page address register (FPAR) locates the current NAND Flash page in both the external NAND Flash device and FCM buffer RAM.

Address: 5000h base + F0h offset = 50F0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved	PI					MS	CI								
W	Reserved	PI					MS	CI								
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

eLBC_FPARRs field descriptions

Field	Description
0–16 -	This field is reserved. Reserved
17–21 PI	Page index. PI indexes the page in NAND Flash EEPROM at the current block defined by FBAR, and locates the corresponding transfer buffer in the FCM buffer RAM. The 3 LSBs of PI index one of the eight 1 Kbyte buffers in the FCM buffer RAM as follows:

Table continues on the next page...

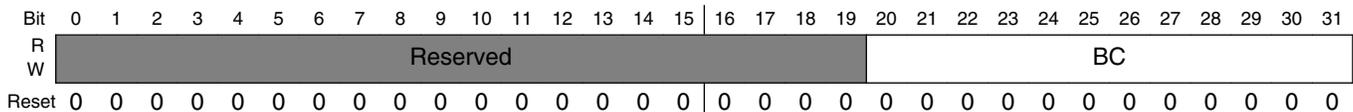
eLBC_FPARs field descriptions (continued)

Field	Description
	000 The page is transferred to/from FCM buffer 0, address offsets 0x0000-0x03FF 001 The page is transferred to/from FCM buffer 1, address offsets 0x0400-0x07FF 010 The page is transferred to/from FCM buffer 2, address offsets 0x0800-0x0BFF 011 The page is transferred to/from FCM buffer 3, address offsets 0x0C00-0x0FFF 100 The page is transferred to/from FCM buffer 4, address offsets 0x1000-0x13FF 101 The page is transferred to/from FCM buffer 5, address offsets 0x1400-0x17FF 110 The page is transferred to/from FCM buffer 6, address offsets 0x1800-0x1BFF 111 The page is transferred to/from FCM buffer 7, address offsets 0x1C00-0x1FFF
22 MS	Main/spare region locator. In the case that FBCR[BC] = 0, MS is treated as 0. 0 Data is transferred to/from the main region of the FCM buffer; that is, the first 512 bytes of the buffer are used as the starting address. 1 Data is transferred to/from the spare region of the FCM buffer; that is, the second 512 bytes of the buffer are used as the starting address, but only an initial 16 bytes of spare region are defined.
23–31 CI	Column index. CI indexes the first byte to transfer to/from the main or spare region of the NAND Flash EEPROM and corresponding transfer buffer. In the case that FBCR[BC] = 0, CI is treated as 0. For MS = 0, CI can range 0x000-0x1FF; for MS = 1, CI can range 0x000-0x00F.

12.3.30 Flash byte count register (eLBC_FBCR)

The local bus Flash byte count register (FBCR) defines the size of FCM block transfers for reads and writes to the NAND Flash EEPROM.

Address: 5000h base + F4h offset = 50F4h



eLBC_FBCR field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 BC	Byte count determines how many bytes are transferred by the FCM during data read (RB) or data write (WB) opcodes. The first byte accessed in the NAND Flash EEPROM is located by the FPAR register, and successive bytes are transferred until either BC bytes have been counted, or the end of the spare region of the currently addressed Flash page has been reached. If BC = 0, an entire Flash page and its spare region will be transferred by FCM, in which case FPAR[MS] and FPAR[CI] are treated as zero regardless of their values. BC = 0 is the only setting that permits FCM to generate and check ECC.

12.3.31 Flash ECC block n registers (eLBC_FECCn)

The local bus Flash ECC block *n* register (FECC *n*) specifies the ECC value calculated during writes or reads by eLBC. It can be used for verify after write feature in software. Note that the valid bit sets before the command completion event and hence the correct ECC could be read before actual completion of writes/reads.

Address: 5000h base + 100h offset + (4d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R	V	Reserved							ECC								
W		Reserved							Reserved								
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	ECC																
W	Reserved																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

eLBC_FECCn field descriptions

Field	Description
0 V	Valid bit. This bit denotes that the ECC stored in this register is valid. It is set for full page write/read transfers if ECC generation/checking is enabled in BRn[DECC].
1–7 -	This field is reserved. Reserved
8–31 ECC	24 bit ECC; For <i>n</i> th 512 bytes of a page in case of large page or for (4 <i>k</i> + <i>n</i>) th 512 byte page for small page where <i>k</i> = 0,1,2,...). It stores calculated ECC value during writes/reads.

12.4 eLBC functional description

The eLBC allows the implementation of memory systems with very specific timing requirements.

- The GPCM provides interfacing for simpler, lower-performance memories and memory-mapped devices. It has inherently lower performance because it does not support bursting. For this reason, GPCM-controlled banks are used primarily for boot-loading from NVRAM or NOR Flash, and access to low-performance memory-mapped peripherals.
- The FCM interfaces the eLBC to NAND Flash EEPROMs with 8-bit data bus. The FCM has an automatic boot-loading feature that allows the CPU to boot from high

density EEPROM, loading the boot block into 4 Kbytes of RAM for execution of the first level boot code. Following boot, FCM provides a flexible instruction sequencer that allows a user-defined command, address, and data transfer sequence of up to 8 steps to be executed against a memory-mapped buffer RAM. Programmable set-up time, hold time, and wait states permit the FCM to maximize the performance of NAND Flash block transfers, which can proceed in parallel with software processing of the multiple RAM buffers. A single-pass ECC engine in the FCM permits zero-overhead error checking, reporting, and correction in both boot blocks and page data transfers if enabled.

- The UPM supports refresh timers, address multiplexing of the external bus, and generation of programmable control signals for row address and column address strobes, to allow for a minimal glue logic interface to DRAMs, burstable SRAMs, and almost any other kind of peripheral with asynchronous timing or single data rate clocking. The UPM can be used to generate flexible, user-defined timing patterns for control signals that govern a memory device. These patterns define how the external control signals behave during a read, write, burst-read, or burst-write access. Refresh timers are also available to periodically initiate user-defined refresh patterns.

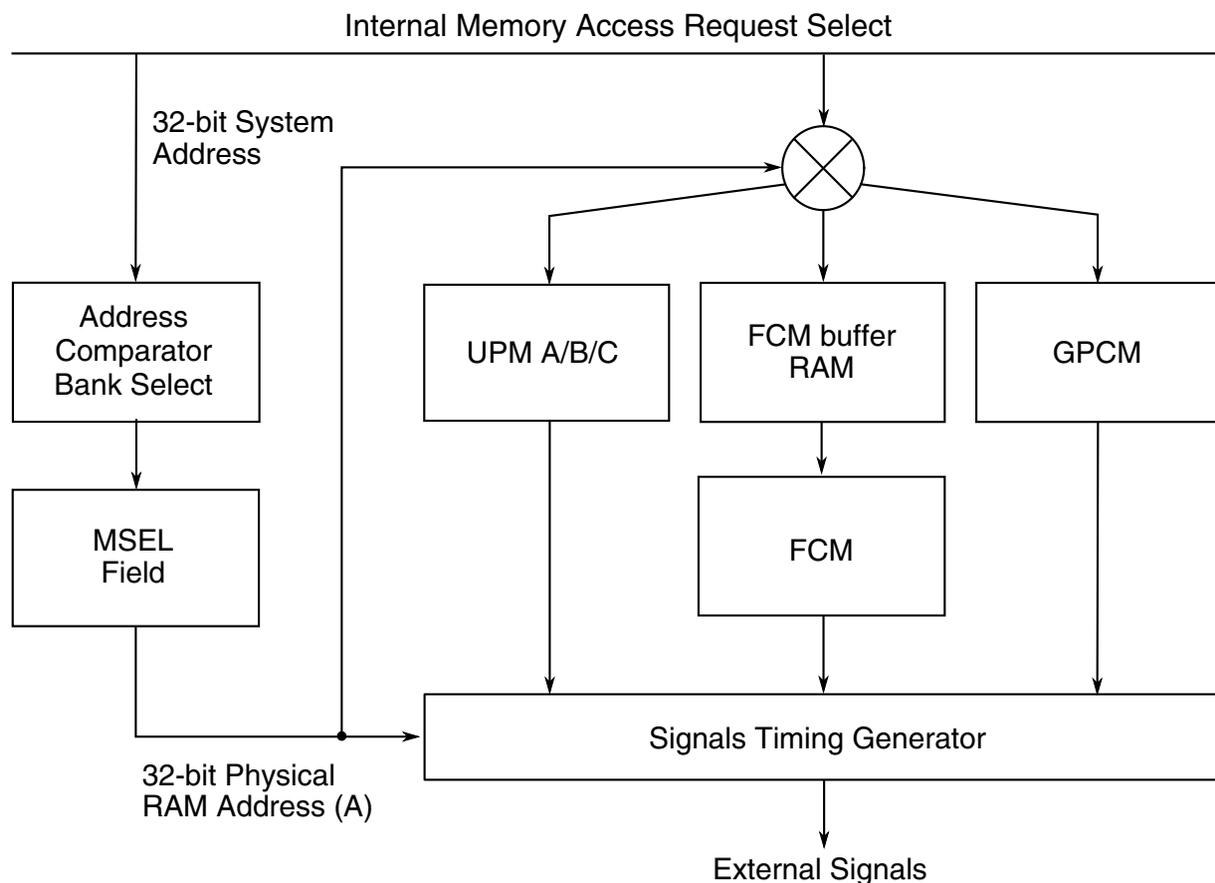


Figure 12-68. Basic operation of memory controllers in the eLBC

Each memory bank (chip select) can be assigned to any of these three types of machines through the machine select bits of the base register for that bank ($BR_n[MSEL]$), as illustrated in the above figure. If a bank match occurs, the corresponding machine (GPCM, FCM, or UPM) then takes ownership of the external signals that control the access and maintains control until the transaction ends.

NOTE

Different machines (FCM/GPCM/UPM) share the address, data, and control signals. While the eLBC is servicing a transaction, subsequent transactions are queued until the current transaction has completed.

12.4.1 Basic architecture

The following subsections describe the basic architecture of the eLBC.

12.4.1.1 Address and address space checking

The defined base addresses are written to the BR_n registers, while the corresponding address masks are written to the OR_n registers.

Each time a local bus access is requested, the internal transaction address is compared with each bank. Addresses are decoded by comparing the 17 MSBs of the address, masked by $OR_n[AM]$, with the base address for each bank ($BR_n[BA]$). If a match is found on a memory controller bank, the attributes defined in the BR_n and OR_n for that bank are used to control the memory access. If a match is found in more than one bank, the lowest-numbered bank handles the memory access (that is, bank 0 has priority over bank 1).

12.4.1.2 External address latch enable signal (LALE)

The local bus uses a multiplexed address/data bus.

Therefore the eLBC must distinguish between address and data phases, which take place on the same bus (LAD pins). The LALE signal, when asserted, signifies an address phase during which the eLBC drives the memory address on the LAD pins. An external address latch uses this signal to capture the address and provide it to the address pins of the memory or peripheral device. When LALE is negated, LAD then serves as the (bi-directional) data bus for the access. Any address phase initiates the assertion of LALE, which has a programmable duration of between 1 and 4 bus clock cycles.

To ensure adequate hold time on the external address latch, LALE negates earlier than the address changes on LAD during address phases. By default, LALE negates earlier by 1 platform clock period. For example, if the platform clock is operating at 533 MHz, then 1.8 ns of address hold time is introduced. However, at higher frequencies, the duration of the shortened LALE pulse may not meet the minimum latch enable pulse width specifications of some latches. In such cases, setting LBCR[AHD] = 1 increases the LALE pulse width by $\hat{A}1/2$ platform clock cycle, but decreases the address hold time by the same amount. If both longer hold time and longer LALE pulse duration are needed, then the address phase can be extended using the ORn[EAD] and LCRR[EADC] fields, and the LBCR[AHD] bit can be left at 0. However, this will add latency to all address tenures.

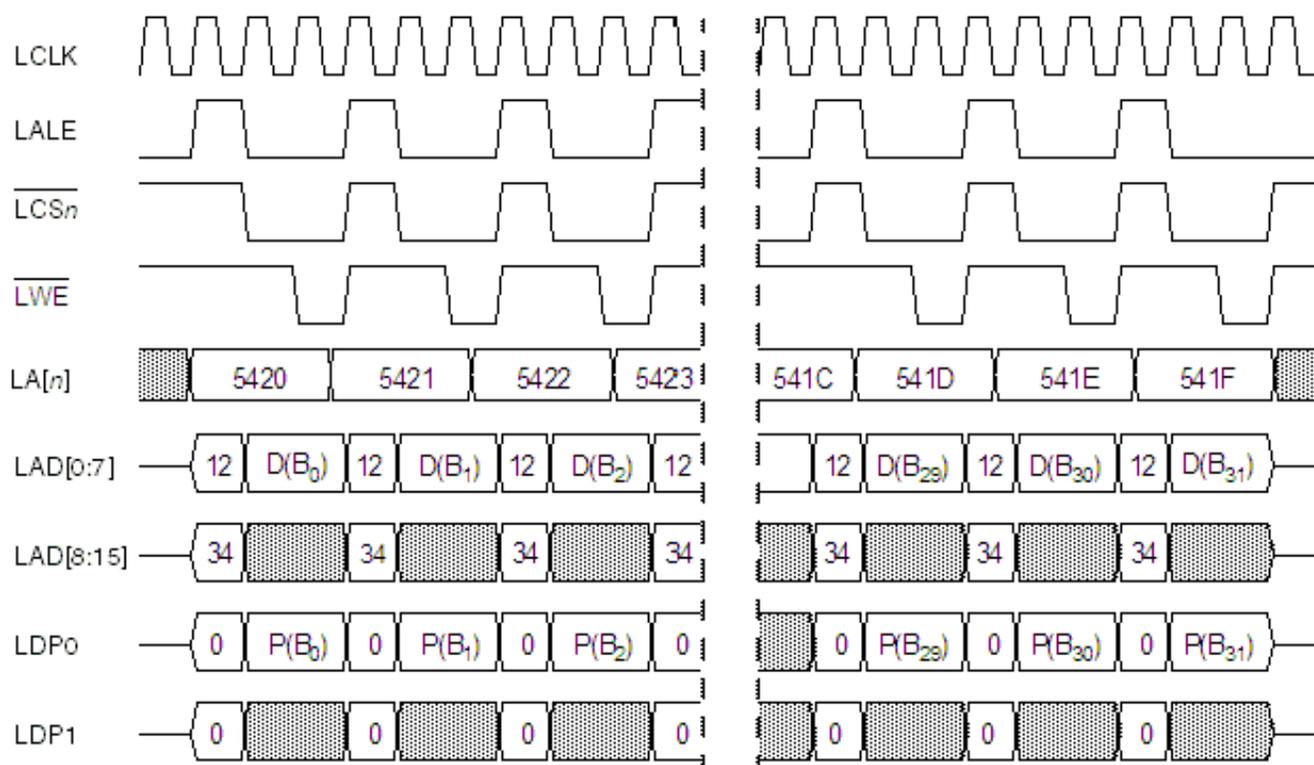
The frequency of LALE assertion varies across the three memory controllers:

- For GPCM, every assertion of LCSn_B is considered an independent access, and accordingly, LALE asserts prior to each such access. For example, GPCM driving an 8-bit port would assert LALE and LCSn_B 32 times in order to satisfy a 32-byte cache line transfer.
- For FCM, LALE asserts prior to each multi-command operation sequence, but LALE can be ignored on NAND Flash EEPROM accesses as the signal does *not* enable address latching in such devices. The value on the LAD and LA pins during LALE assertion is driven low-impedance, but otherwise not defined for FCM banks.
- In the case of UPM, the frequency of LALE assertion depends on how the UPM RAM is programmed. UPM single accesses typically assert LALE once, upon commencement, but it is possible to program UPM to assert LALE several times, and to change the values of LAN with and without LALE being involved.

In general, when using the GPCM controller it is typically required to use both LA and the latched version of LAD to capture the entire address during LALE phases. The UPMs may require LA if the eLBC is generating its own burst address sequence.

To illustrate how a large transaction is handled by the eLBC, the following figure shows eLBC signals for the GPCM performing a 32-byte write starting at address 0x1234_5420. Note that during each of the 32 assertions of LALE, LA[16:31] and LAD[0:15] together drive the address, but during data phases, only LAD[0:7] and LDP[0] are driven with valid data and parity, respectively.

eLBC functional description



Note: All address and signal values are shown in hexadecimal.
 $D(B_k)$ = k^{th} of 32 data bytes, $P(B_k)$ = parity bit of k^{th} data byte.

Figure 12-69. Example of 8-bit GPCM writing 32 bytes to address 0x1234_5420 (LCRR[PBYP] = 0)

12.4.1.3 Data transfer acknowledge (TA)

The three memory controllers in the eLBC generate an internal transfer acknowledge signal, TA, to allow data on LAD to be either sampled (for reads) or changed (on writes).

The data sampling/data change always occurs at the end of the bus cycle in which the eLBC asserts TA internally. In eLBC debug mode, TA is also visible externally on the MDVAL pin. The GPCM controller automatically generates TA according to the timing parameters programmed for them in the option and mode registers; FCM generates TA whenever data read and write instructions are executed out of register FIR; a UPM generates TA only when a UPM pattern has the UTA RAM word bit set. [Figure 12-70](#) and [Figure 12-71](#) shows LALE, TA (internal), and LCSn_B . Note that TA and LALE are never asserted together, and that for the duration of LALE, LCSn_B (or any other control signal) remains negated or frozen.

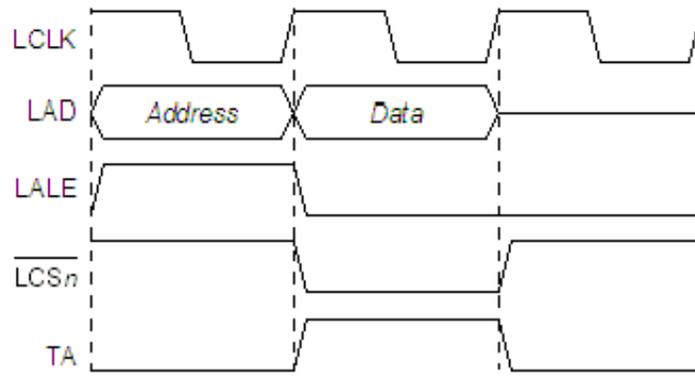


Figure 12-70. Basic eLBC bus cycle with LALE, TA, and LCSn_B (PLL enable)

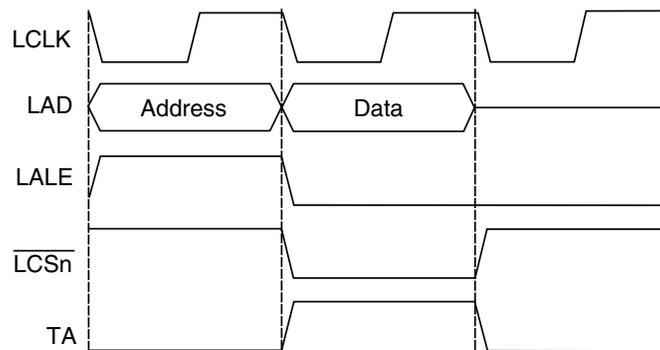


Figure 12-71. Basic eLBC bus cycle with LALE, TA, and LCSn_B (PLL bypass)

12.4.1.4 Data buffer control (LBCTL)

The memory controller provides a data buffer control signal for the local bus (LBCTL).

This signal is activated when a GPCM-, FCM-, or UPM-controlled bank is accessed. LBCTL can be disabled by setting $OR_n[BCTL D]$. LBCTL can be further configured by $LBCR[BCTL C]$ to act as an extra LWE_B or an extra LOE_B signal when in GPCM mode.

If LBCTL is configured as a data buffer control ($LBCR[BCTL C] = 00$), the signal is asserted (high) on the rising edge of the bus clock on the first cycle of the memory controller operation, coincident with LALE. If the access is a write, LBCTL remains high for the whole duration. However, if the access is a read, LBCTL is negated (low) with the negation of LALE so that the memory device is able to drive the bus. If back-to-back read accesses are pending, LBCTL is asserted (high) one bus clock cycle before the next transaction starts (that is, one bus clock cycle before LALE) to allow a whole bus cycle for the bus to turn around before the next address is driven.

12.4.1.5 Parity generation and checking (LDP)

Parity can be configured for any GPCM or UPM bank by programming $BR_n[DECC]$. Parity is generated and checked on a per-byte basis using LDP[0:1] for the bank if $BR_n[DECC] = 01$ (normal parity). Byte lane parity on LDP[0:1] is generated regardless of the $BR_n[DECC]$ setting.

FCM calculates an ECC over 512-byte blocks, and hence does not use the LDP[0:1] pins. The setting of $BR_n[DECC] = 01$ enables ECC checking only, while $BR_n[DECC] = 10$ enables ECC generation and checking; in either case, LBCR[EPAR] determines the global type of block parity for ECC (odd or even).

12.4.1.6 Bus monitor

A bus monitor is provided to ensure that each transaction is terminated within a reasonable (user defined) period.

When a transaction starts, the bus monitor starts counting down from the time-out value ($LBCR[BMT] \times LBCR[BMTPS]$) until a data beat is acknowledged on the bus. It then reloads the time-out value and resumes the countdown until the data tenure completes and then idles if there is no pending transaction. Setting LTEDR[BMD] disables bus monitor error checking (that is, the LTESR[BM] bit is not set by a bus monitor time-out); however, the bus monitor is still active and can generate a UPM exception (as noted in [Exception requests](#)), or terminate a GPCM access.

It is very important to ensure that the value of LBCR[BMT] is not set too low; otherwise spurious bus time-outs may occur during normal operation-resulting in incomplete data transfers. Accordingly, the time-out value represented by the LBCR[BMT], LBCR[BMTPS] pair must not be set below 40 bus cycles for time-out under any circumstances.

NOTE

When the FCM is in the middle of a long transaction (such as NAND erase, write, and so on), another transaction on the GPCM or UPM will trigger the bus monitor to start, even though the GPCM or UPM is waiting for the FCM to finish. If the bus monitor times out, it could corrupt the current NAND Flash operation as well as terminate the GPCM or UPM operation. To avoid such cases, it is recommended that while using FCM the bus monitor timeout be programmed to its

maximum setting of LBCR[BMT] = 0 and LBCR[BMTPS] = 0xF.

12.4.1.7 PLL Bypass mode

At LCLK frequencies in excess of 83MHz the local bus PLL is used to provide improved hold times at external receivers, and ease set-up margins for read data captured by eLBC. A wire loop between pins LSYNC_OUT and LSYNC_IN establishes the amount of LCLK skewing achieved by the PLL, which locks so as to produce edges on LCLK before the transition of other eLBC control and data signals.

At lower frequencies, the PLL may be unable to lock or provide sufficient hold time improvement for particularly slow devices. Accordingly, LCRR[PBYP] should be set to 1 to bypass the PLL at low frequencies, with the eLBC generating LCLK directly, while skewing it by half a bus clock cycle. An illustration of GPCM or UPM timing both with and without the PLL activated are shown in [Figure 12-72](#) and [Figure 12-73](#). When LCRR[PBYP] = 0, the skew, t_{LSKEW} , matches the round-trip propagation delay of the timing loop between LSYNC_OUT and LSYNC_IN, and data is generated or sampled on the next rising edge of LCLK. The timing diagrams shown normally in this chapter assume that LCRR[PBYP] = 0. When LCRR[PBYP] = 1, the skew equals half the period of LCLK to maximize hold time at the external receiver; in this bypass mode, eLBC drives new address, data, and control signals effectively on falling edges of LCLK, but continues to sample synchronous read data on rising edges of LCLK to maximize the set-up margin for reads.

In PLL bypass mode, eLBC drives new address, data, and control signals effectively on falling edges of LCLK, but continues to sample synchronous read data on rising edges of LCLK to maximize the set-up margin for reads.

NOTE

Because LCLK is not used for NAND Flash EEPROMs controlled by FCM, the eLBC drives and samples data on the same edge (rising edge when LCRR[PBYP] = 0 and falling edge when LCRR[PBYP] = 1) on FCM controlled banks.

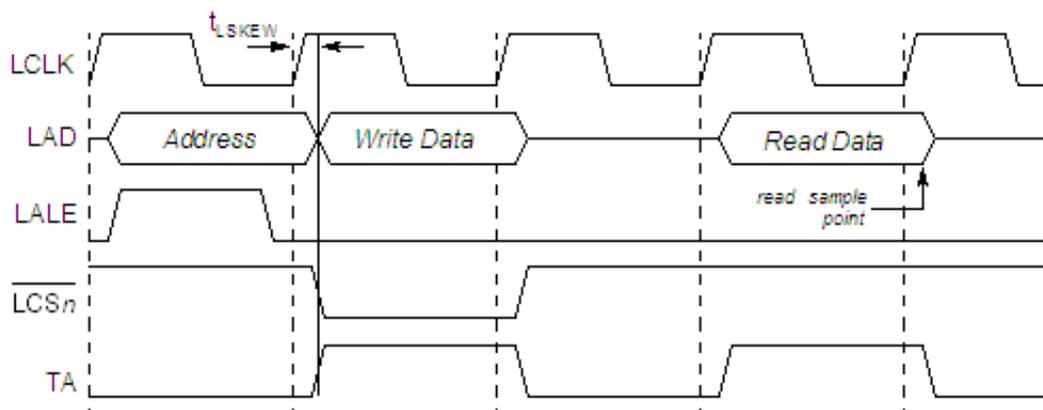


Figure 12-72. eLBC bus cycles in PLL mode (GPCM and UPM only)

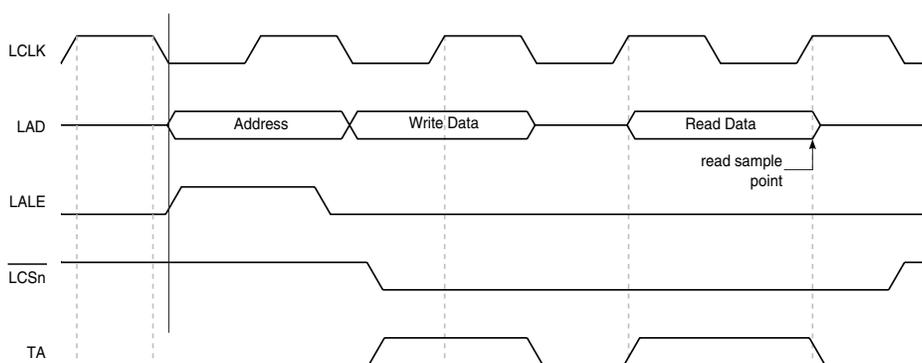


Figure 12-73. eLBC bus cycles in PLL-bypassed mode (GPCM and UPM only)

12.4.2 General-purpose chip-select machine (GPCM)

The GPCM allows a minimal glue logic and flexible interface to SRAM, EPROM, FEPRM, ROM devices, and external peripherals.

The GPCM contains two basic configuration register groups- BR_n and OR_n .

Byte-write enable signals (LWE_B) are available for each byte written to memory. Also, the output enable signal (LOE_B) is provided to minimize external glue logic. On system reset, a global (boot) chip-select is available that provides a boot ROM chip-select (LCS0_B) prior to the system being fully configured.

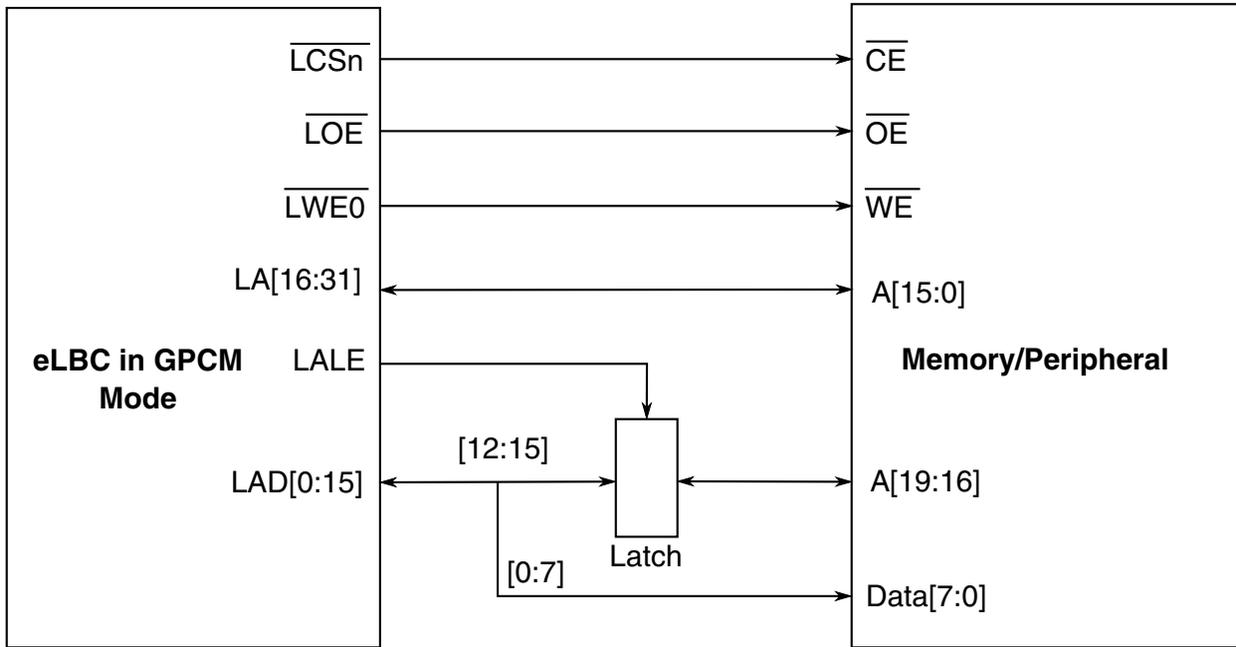


Figure 12-74. Enhanced local bus to GPCM device interface

The figure below shows LCS_B as defined by the setup time required between the address lines and CE_B . The user can configure $ORn[ACS]$ to specify LCS_B to meet this requirement. Generally, the attributes for the memory cycle are taken from ORn . These attributes include the $CSNT$, ACS , $XACS$, SCY , $TRLX$, $EHTR$, and $SETA$ fields.

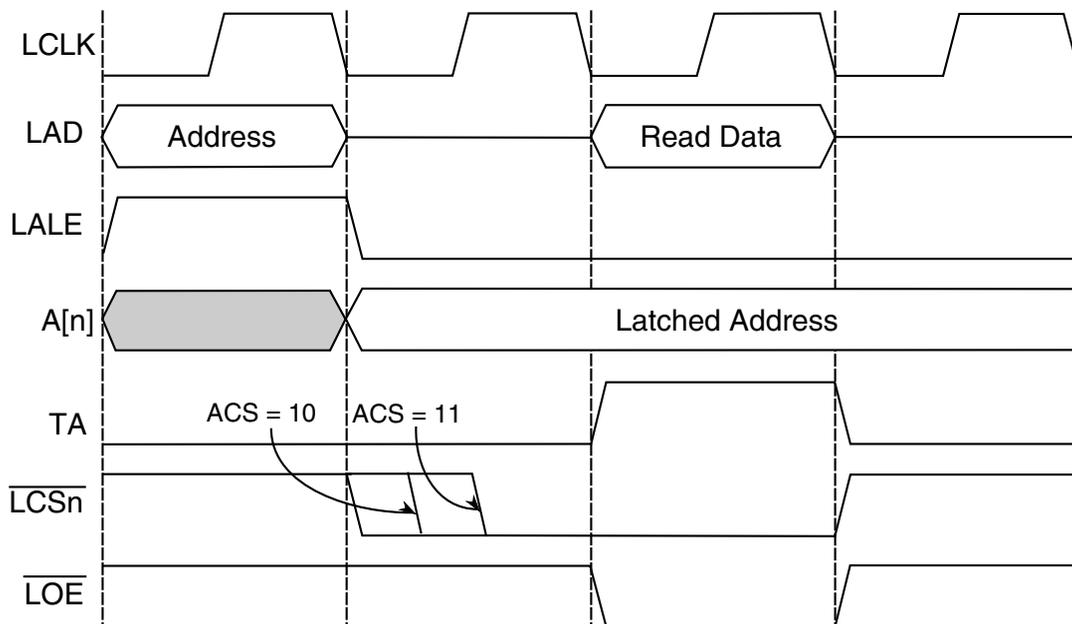


Figure 12-75. GPCM basic read timing ($XACS = 0$, $ACS = 1x$, $TRLX = 0$)

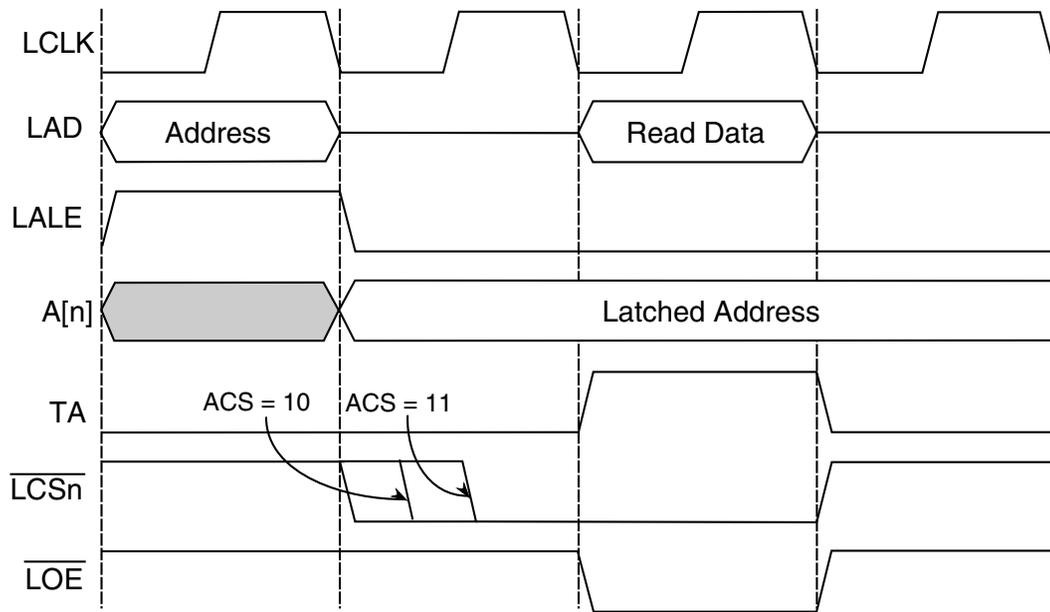
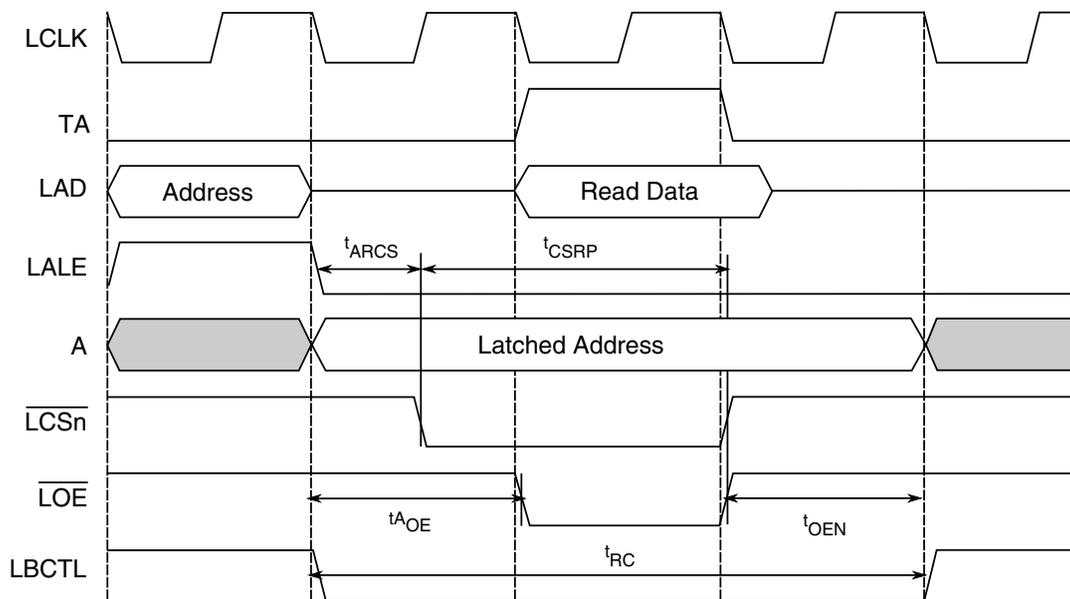


Figure 12-76. GPCM basic read timing (XACS = 0, ACS = 1x, TRLX = 0)

12.4.2.1 GPCM read signal timing

The basic GPCM read timing parameters that may be set by the OR_n attributes are shown below.

The read access cycle commences upon latching of the memory address (LALE negated), and concludes when LBCTL returns high to turn the local bus around for a subsequent address phase. Read data is captured on the rising edge of LCLK when TA is asserted. LOE_B and LCSn_B negate high simultaneously, in some cases before the end of the read access to provide additional hold time for the external memory.

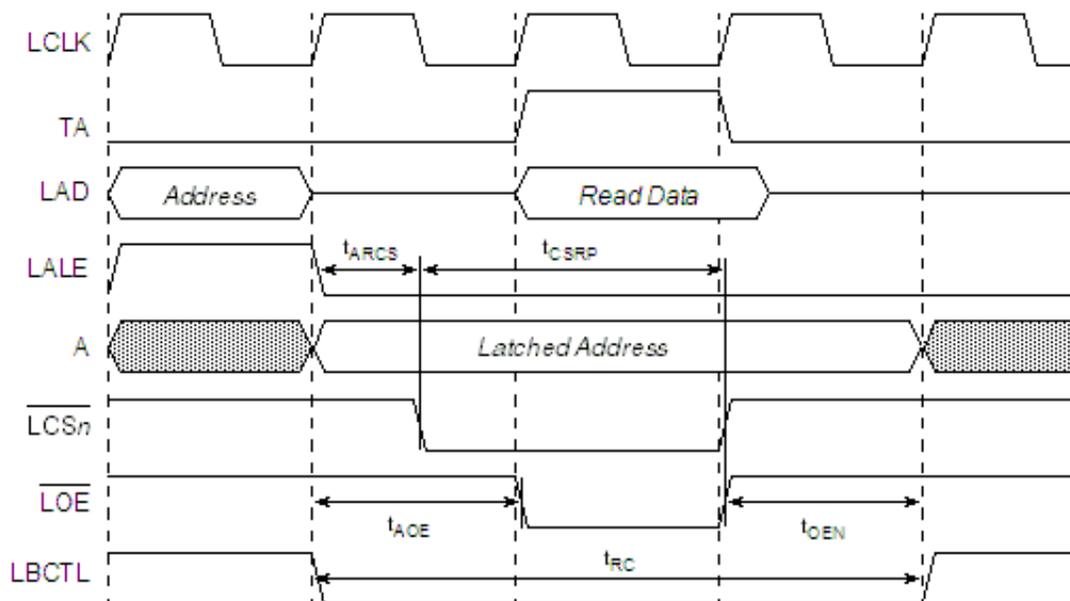

Notes:
 t_{RC} = Read cycle time.

 t_{CSRP} = Read chip-select assertion period.

 t_{ARCS} = Address valid to read chip-select time.

 t_{OEN} = Output enable negated time.

 t_{AOE} = Address valid to output enable time.

Figure 12-77. GPCM general read timing parameters with PLL bypass

Notes:
 t_{RC} = Read cycle time.

 t_{CSRP} = Read chip-select assertion period.

 t_{ARCS} = Address valid to read chip-select time.

 t_{OEN} = Output enable negated time.

 t_{AOE} = Address valid to output enable time.

Figure 12-78. GPCM general read timing parameters

The table below lists the signal timing parameters for a GPCM read access as the option register attributes are varied.

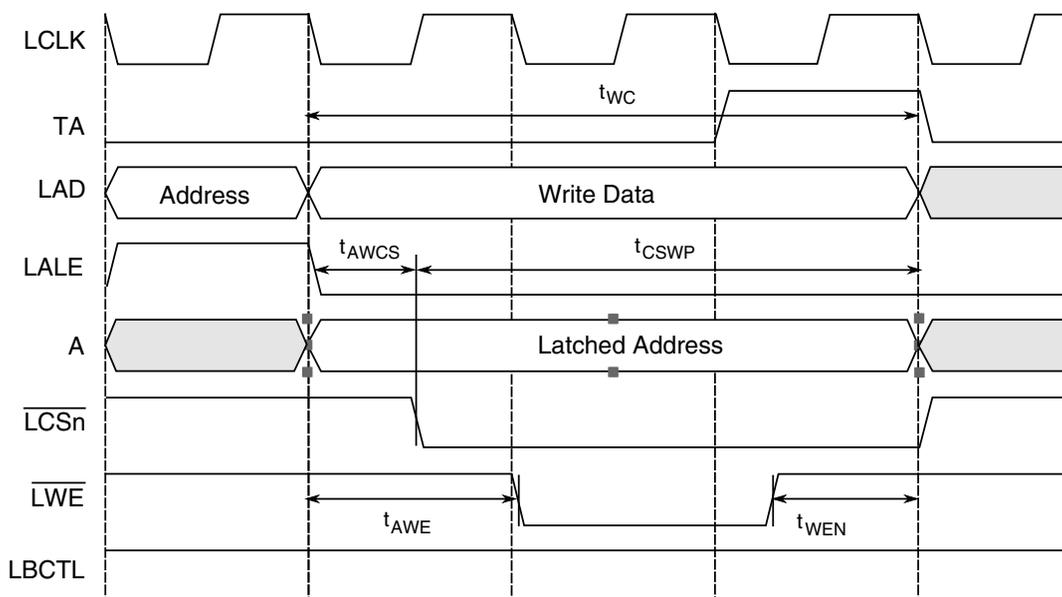
Table 12-190. GPCM read control signal timing

Option Register Attributes				Signal Timing (LCLK clock cycles)				
TRLX	EHTR	XACS	ACS	t _{ARCS}	t _{CSRP}	t _{AOE}	t _{OEN}	t _{RC}
0	0	0	0X	0	2 + SCY	1	0	2 + SCY
0	0	0	10	¼	1¾ + SCY	1	0	2 + SCY
0	0	0	11	½	1½ + SCY	1	0	2 + SCY
0	0	1	0X	0	2 + SCY	1	0	2 + SCY
0	0	1	10	1	1 + SCY	1	0	2 + SCY
0	0	1	11	2	1 + SCY	2	0	3 + SCY
0	1	0	0X	0	2 + SCY	1	1	3 + SCY
0	1	0	10	¼	1¾ + SCY	1	1	3 + SCY
0	1	0	11	½	1½ + SCY	1	1	3 + SCY
0	1	1	0X	0	2 + SCY	1	1	3 + SCY
0	1	1	10	1	1 + SCY	1	1	3 + SCY
0	1	1	11	2	1 + SCY	2	1	4 + SCY
1	0	0	0X	0	2 + 2 x SCY	1	4	6 + 2 x SCY
1	0	0	10	¼	1¾ + 2 x SCY	2	4	7 + 2 x SCY
1	0	0	11	½	1½ + 2 x SCY	2	4	7 + 2 x SCY
1	0	1	0X	0	2 + 2 x SCY	1	4	6 + 2 x SCY
1	0	1	10	2	1 + 2 x SCY	2	4	7 + 2 x SCY
1	0	1	11	3	1 + 2 x SCY	3	4	8 + 2 x SCY
1	1	0	0X	0	2 + 2 x SCY	1	8	10 + 2 x SCY
1	1	0	10	¼	1¾ + 2 x SCY	2	8	11 + 2 x SCY
1	1	0	11	½	1½ + 2 x SCY	2	8	11 + 2 x SCY
1	1	1	0X	0	2 + 2 x SCY	1	8	10 + 2 x SCY
1	1	1	10	2	1 + 2 x SCY	2	8	11 + 2 x SCY
1	1	1	11	3	1 + 2 x SCY	3	8	12 + 2 x SCY

12.4.2.2 GPCM write signal timing

The basic GPCM write timing parameters that may be set by the OR_n attributes are shown in the figure below.

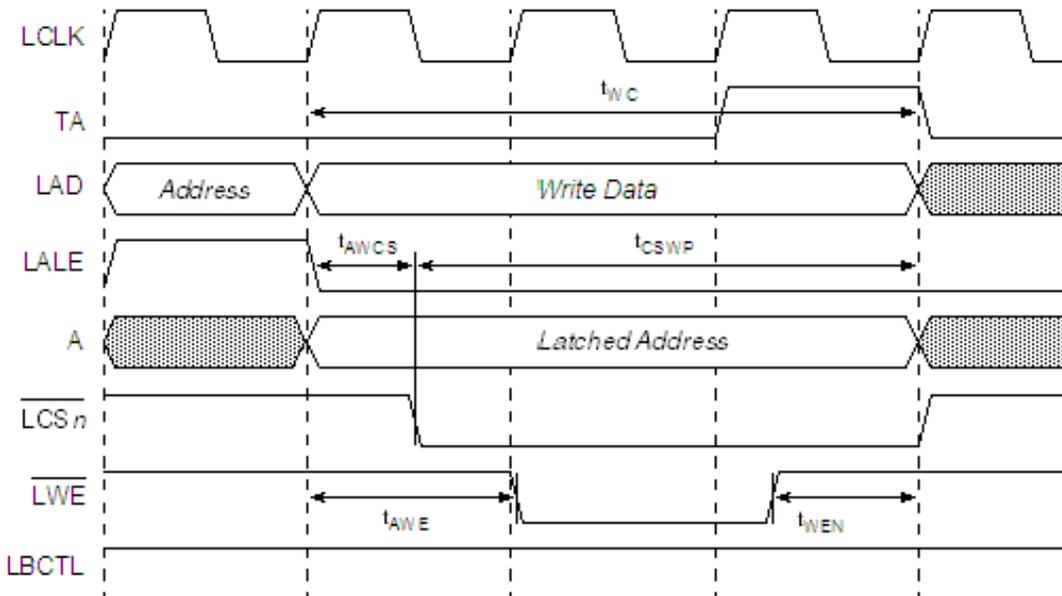
The write access cycle commences upon latching of the memory address (LALE negated), and concludes when LCS_n_B returns high. LBCTL remains stable for the entire cycle to drive data onto any secondary data bus. Write data becomes invalid following the falling edge of TA. LWE_B may, in some cases, negate high before the end of the write access to provide additional hold time for the external memory.



Notes:

t_{WC} = Write cycle time. t_{CSWP} = Write chip-select assertion period
 t_{AWCS} = Address valid to write chip-select time. t_{WEN} = Write enable negated time wrt chip-select negation time
 t_{AWE} = Address valid to write enable time.

Figure 12-79. GPCM general write timing parameters



Notes:

t_{WC} = Write cycle time. t_{CSWP} = Write chip-select assertion period.
 t_{AWCS} = Address valid to write chip-select time. t_{WEN} = Write enable negated time wrt chip-select negation time.
 t_{AWE} = Address valid to write enable time.

Figure 12-80. GPCM general write timing parameters

Table 12-191. GPCM write control signal timing

Option register attributes				Signal Timing (LCLK clock cycles)				
TRLX	XACS	ACS	CSNT	t _{AWCS}	t _{CSWP}	t _{AWE}	t _{WEN}	t _{wc}
0	0	00	0	0	2 + SCY	1	0	2 + SCY
0	0	10	0	¼	1¾ + SCY	1	0	2 + SCY
0	0	11	0	½	1½ + SCY	1	0	2 + SCY
0	1	00	0	0	2 + SCY	1	0	2 + SCY
0	1	10	0	1	1 + SCY	1	0	2 + SCY
0	1	11	0	2	1 + SCY	2	0	3 + SCY
0	0	00	1	0	2 + SCY	1	¼	2 + SCY
0	0	10	1	¼	1½ + SCY	1	0	1¾ + SCY
0	0	11	1	½	1¼ + SCY	1	0	1¾ + SCY
0	1	00	1	0	2 + SCY	1	¼	2 + SCY
0	1	10	1	1	¾ + SCY	1	0	1¾ + SCY
0	1	11	1	2	¾ + SCY	2	0	2¾ + SCY
1	0	00	0	0	2 + 2 x SCY	1	0	2 + 2 x SCY
1	0	10	0	1¼	1¾ + 2 x SCY	2	0	3 + 2 x SCY
1	0	11	0	1½	1½ + 2 x SCY	2	0	3 + 2 x SCY
1	1	00	0	0	2 + 2 x SCY	1	0	2 + 2 x SCY
1	1	10	0	2	1 + 2 x SCY	2	0	3 + 2 x SCY
1	1	11	0	3	1 + 2 x SCY	3	0	4 + 2 x SCY
1	0	00	1	0	3 + 2 x SCY	1	1¼	3 + 2 x SCY
1	0	10	1	1¼	1½ + 2 x SCY	2	0	2¾ + 2 x SCY
1	0	11	1	1½	1¼ + 2 x SCY	2	0	2¾ + 2 x SCY
1	1	00	1	0	3 + 2 x SCY	1	1¼	3 + 2 x SCY
1	1	10	1	2	¾ + 2 x SCY	2	0	2¾ + 2 x SCY
1	1	11	1	3	¾ + 2 x SCY	3	0	3¾ + 2 x SCY

12.4.2.3 Chip-select assertion timing

The banks selected to work with the GPCM support an option to drive the LCS_B *n* signal with different timings (with respect to the external address/data bus).

LCS_B *n* can be driven in any of the following ways:

- Simultaneous with the latched memory address. (This refers to the externally latched address and not the address timing on LAD. That is, the chip select does not assert during LALE.)
- One quarter of a clock cycle later.
- One half of a clock cycle later.

- One clock cycle later (for $LCRR[CLKDIV] = 2$ (clock ratio of 4)), when $OR_n[XACS] = 1$.
- Two clock cycles later, when $OR_n[XACS] = 1$.
- Three clock cycles later, when $OR_n[XACS] = 1$ and $OR_n[TRLX] = 1$.

The timing diagram in [Figure 12-76](#) shows two chip-select assertion timings.

12.4.2.3.1 Programmable wait state configuration

The GPCM supports internal generation of transfer acknowledge.

It allows between zero and 30 wait states to be added to an access by programming $OR_n[SCY]$ and $OR_n[TRLX]$. Internal generation of transfer acknowledge is enabled if $OR_n[SETA] = 0$. If $LGTA_B$ is asserted externally two bus clock cycles or more before the wait state counter has expired (to allow for synchronization latency), the current memory cycle is terminated by $LGTA_B$; otherwise it is terminated by the expiration of the wait state counter. Regardless of the setting of $OR_n[SETA]$, wait states prolong the assertion duration of both LOE_B and LWE_B n in the same manner. When $TRLX = 1$, the number of wait states inserted by the memory controller is doubled from $OR_n[SCY]$ cycles to $2 \times OR_n[SCY]$ cycles, allowing a maximum of 30 wait states.

12.4.2.3.2 Chip-select and write enable negation timing

The figure below shows a basic connection between the local bus and a static memory device.

In this case, LCS_B n is connected directly to CE_B of the memory device. The $LWE_B[0:1]$ signals are connected to the respective $WE_B[1:0]$ signals on the memory device where each $LWE_B[0:1]$ signal corresponds to a different data byte.

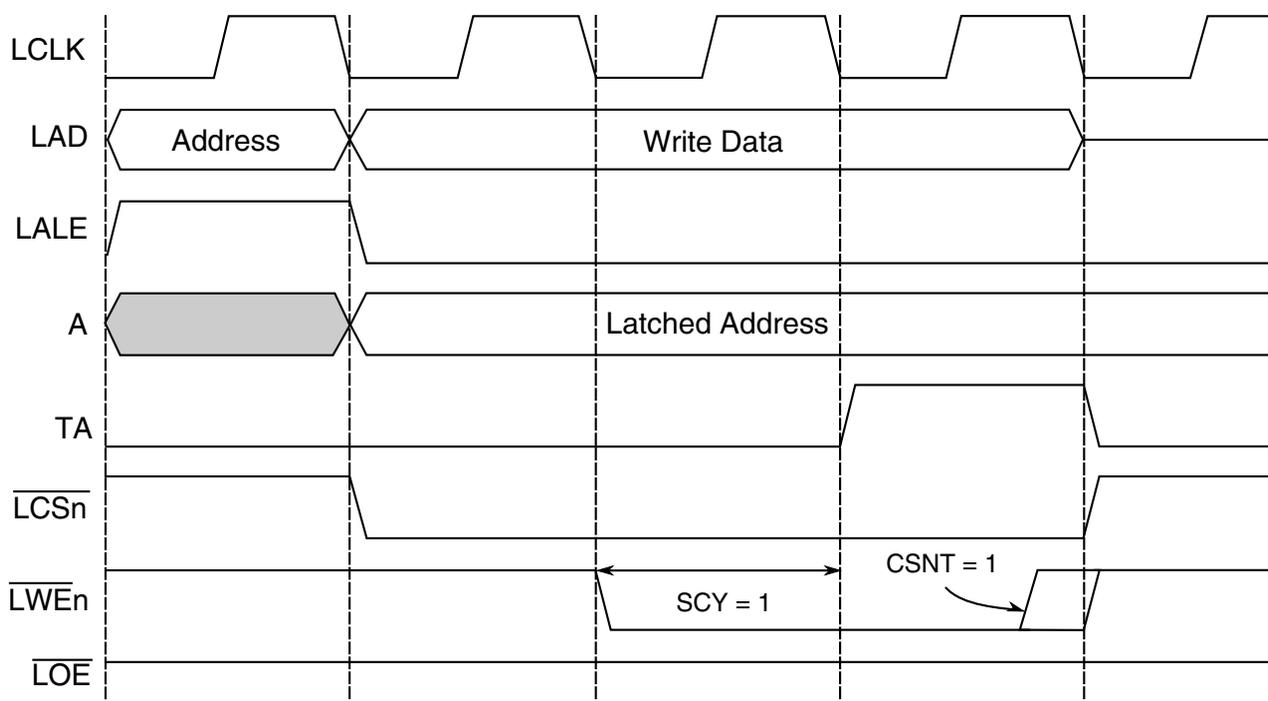


Figure 12-81. GPCM basic write timing (XACS = 0, ACS = 00, CSNT = 1, SCY = 1, TRLX = 0)

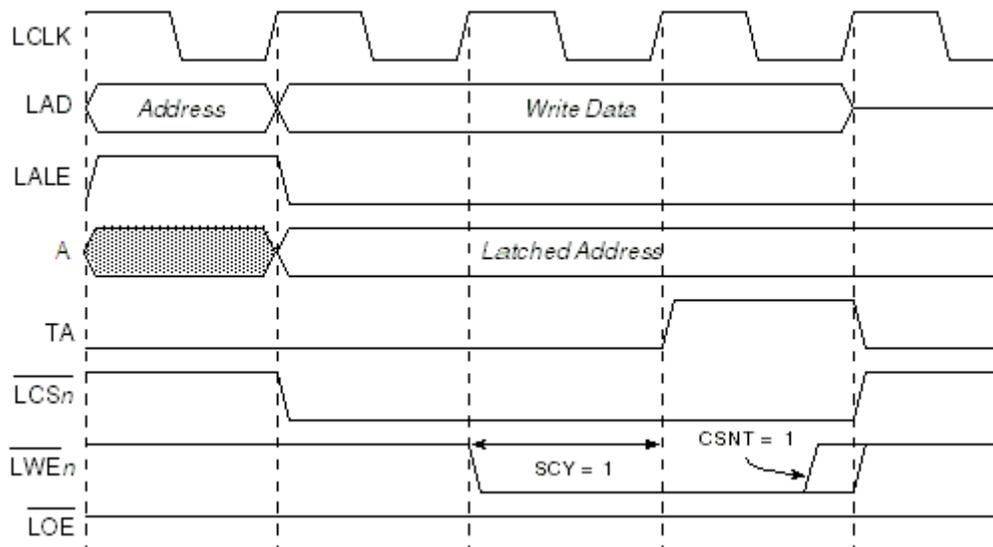


Figure 12-82. GPCM basic write timing (XACS = 0, ACS = 00, CSNT = 1, SCY = 1, TRLX = 0, CLKDIV = 4, 8)

The timing for LCS_{B n} is the same as for the latched address. The strobes for the transaction are supplied by LOE_B or LWE_{B n}, depending on the transaction direction—read or write (write case shown in the figure). OR_n[CSNT], along with OR_n[TRLX], control the timing for the appropriate strobe negation in write cycles. When this attribute

is asserted, the strobe is negated one quarter of a clock before the normal case. For example, when $ACS = 00$ and $CSNT = 1$, LWE_B_n is negated one quarter of a clock earlier, as shown in [Figure 12-81](#).

1. LCS_B_n is affected by $CSNT$ and $TRLX$ only if $ACS[0]$ is non zero. However, LWE_B_n is affected independent of ACS .
2. When $CSNT$ attribute is asserted, the strobe is negated one quarter of a clock before the normal case.
3. $TRLX = 1$ in conjunction with $CSNT = 1$, negates the LCS_B_n and $LWE_B_{n+1} + 1/4$ cycle earlier.

For example, when $ACS = 00$, $CSNT = 1$ and $TRLX = 0$, LWE_B_n is negated one quarter of a clock earlier and LCS_B_n is negated normally as shown in [Figure 12-81](#).

12.4.2.3.3 Relaxed timing

$ORx[TRLX]$ is provided for memory systems that require more relaxed timing between signals.

Setting $TRLX = 1$ has the following effect on timing:

- An additional bus cycle is added between the address and control signals (but only if ACS is not equal to 00).
- The number of wait states specified by SCY is doubled, providing up to 30 wait states.
- The extended hold time on read accesses ($EHTR$) is extended further.
- LCS_n_B signals are negated one cycle earlier during writes (but only if ACS is not equal to 00).
- $LWE_B[0:1]$ signals are negated one cycle earlier during writes.

eLBC functional description

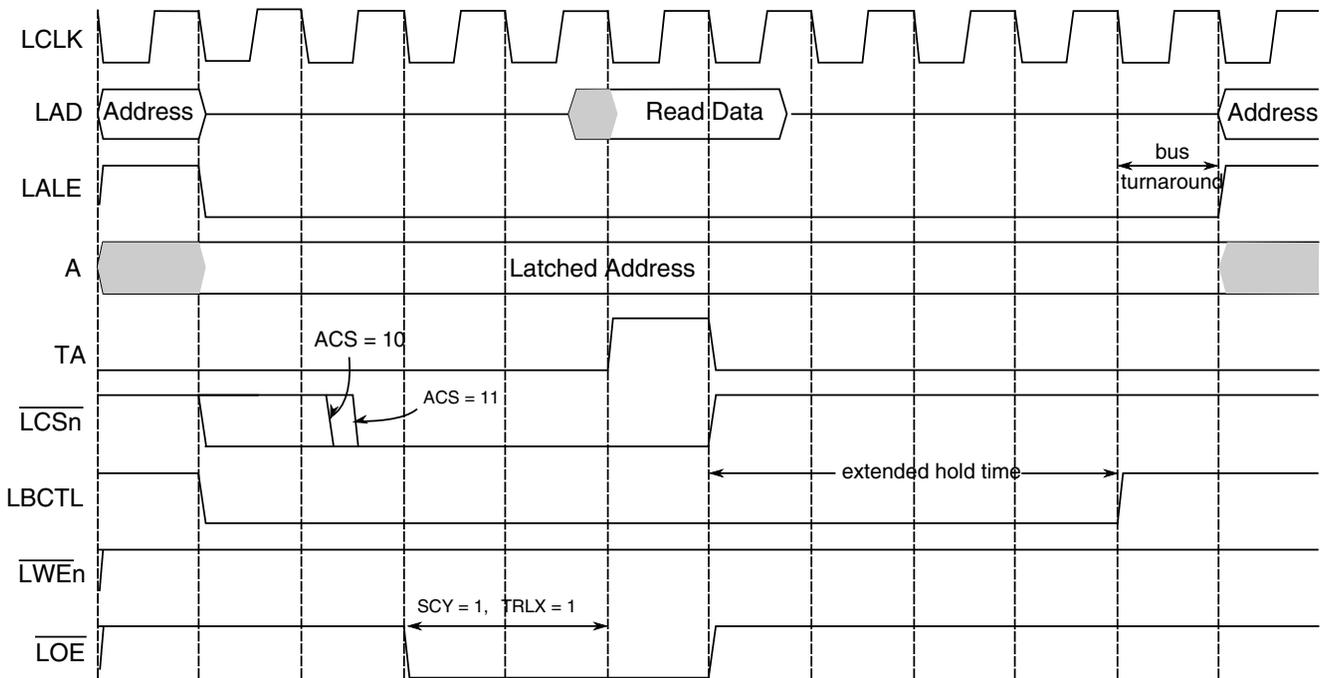


Figure 12-83. GPCM relaxed timing back-to-back reads with PLL bypass (XACS = 0, ACS = 1x, SCY = 1, CSNT = 0, TRLX = 1, EHTR = 0)

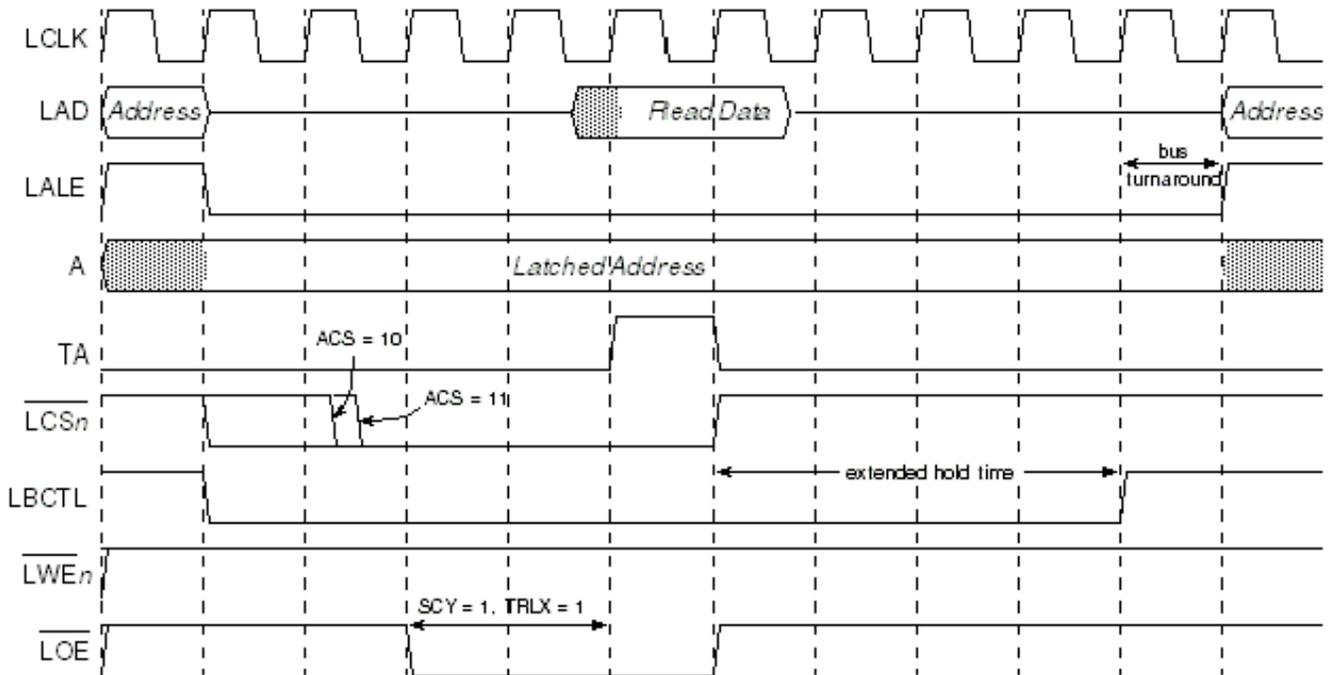


Figure 12-84. GPCM relaxed timing back-to-back reads with PLL enable (XACS = 0, ACS = 1x, SCY = 1, CSNT = 0, TRLX = 1, EHTR = 0, CLKDIV = 4, 8)

The example in [Figure 12-85](#) and [Figure 12-86](#) shows address and data multiplexing on LAD for a pair of writes issued consecutively.

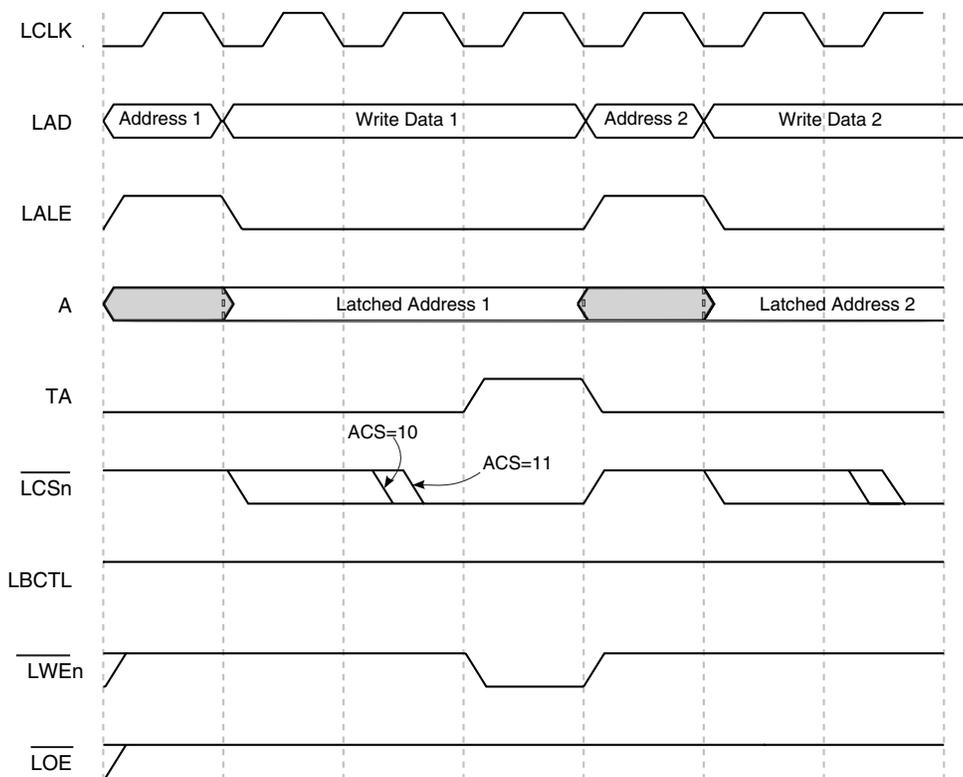


Figure 12-85. GPCM relaxed timing back-to-back writes with PLL bypass (XACS = 0, ACS = 1x, SCY = 0, CSNT = 0, TRLX = 1)

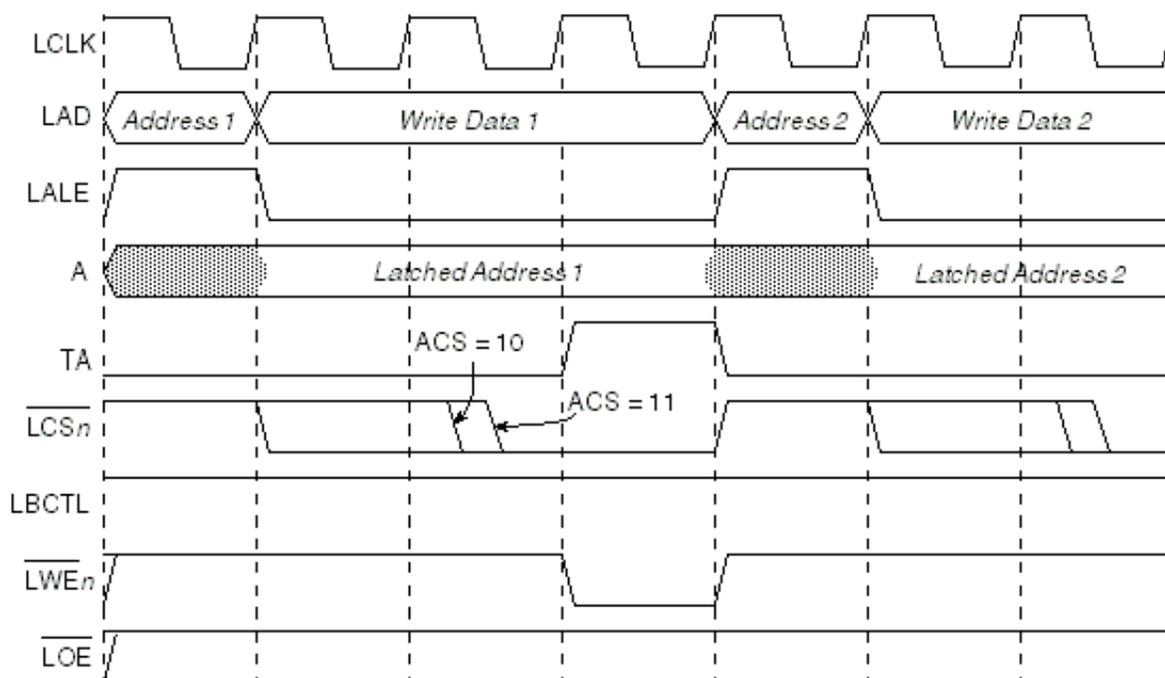


Figure 12-86. GPCM relaxed timing back-to-back writes with PLL enable (XACS = 0, ACS = 1x, SCY = 0, CSNT = 0, TRLX = 1, CLKDIV = 4, 8)

When TRLX and CSNT are set in a write access, the LWE_B[0:1] strobe signals are negated one clock earlier than in the normal case. If ACS ≠ 00, LCSn_B is also negated one clock earlier.

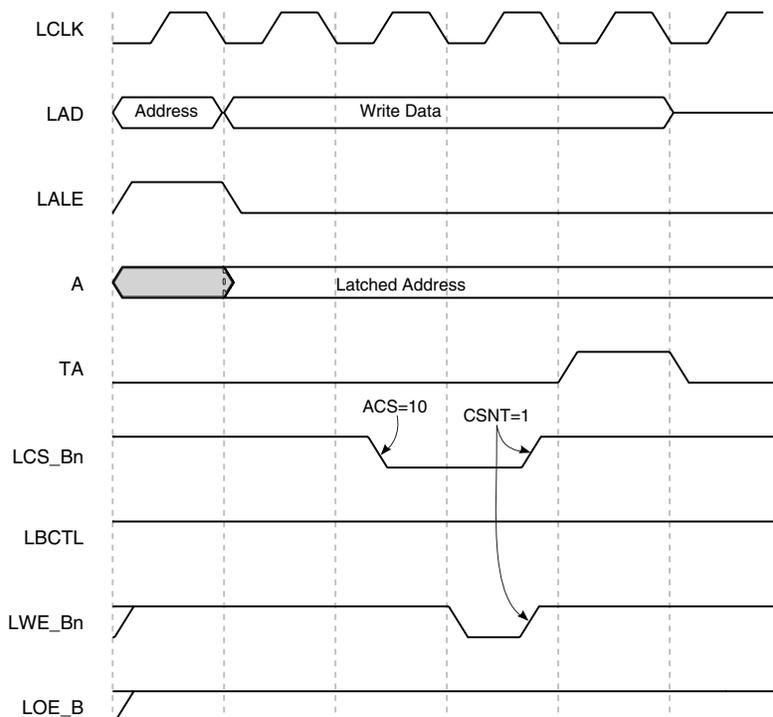


Figure 12-87. GPCM relaxed timing write (XACS = 0, ACS = 10, SCY = 0, CSNT = 1, TRLX = 1)

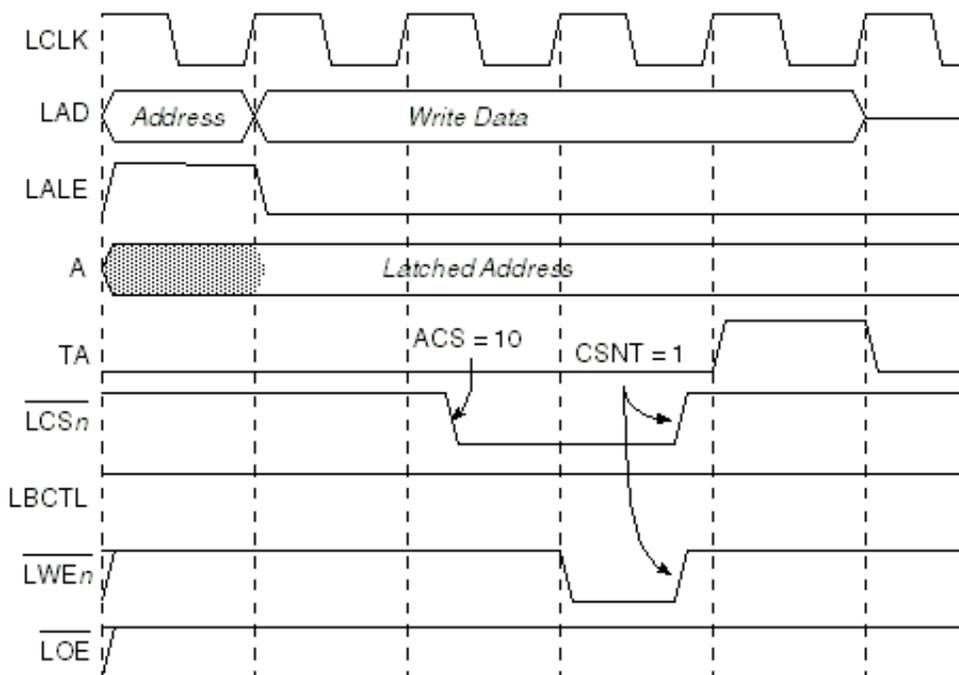


Figure 12-88. GPCM relaxed timing write (XACS = 0, ACS = 10, SCY = 0, CSNT = 1, TRLX = 1)

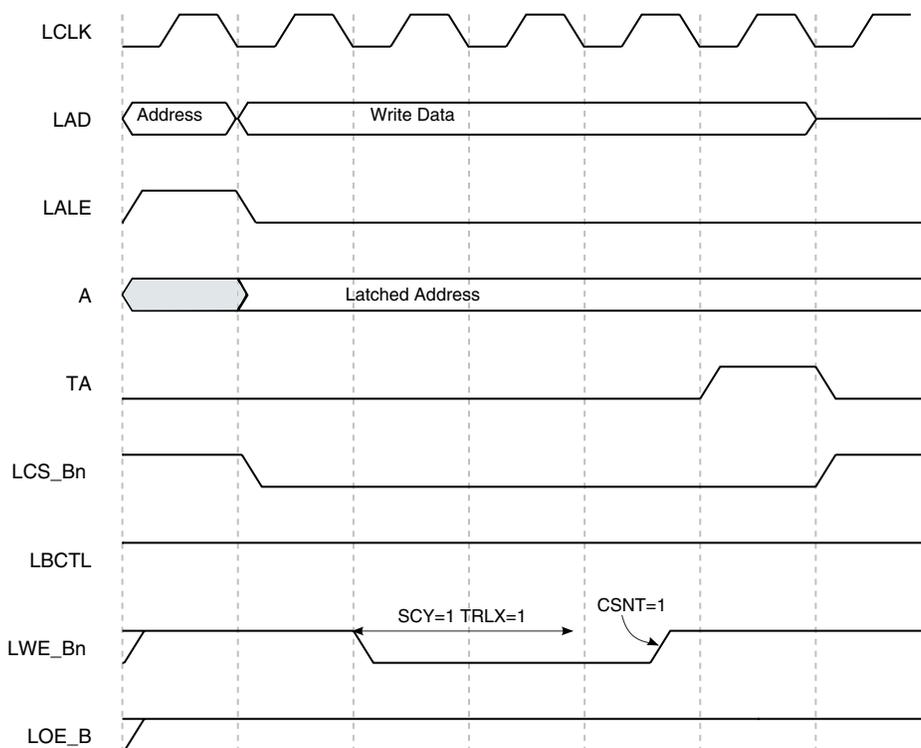


Figure 12-89. GPCM relaxed timing write ($XACS = 0$, $ACS = 00$, $SCY = 1$, $CSNT = 1$, $TRLX = 1$)

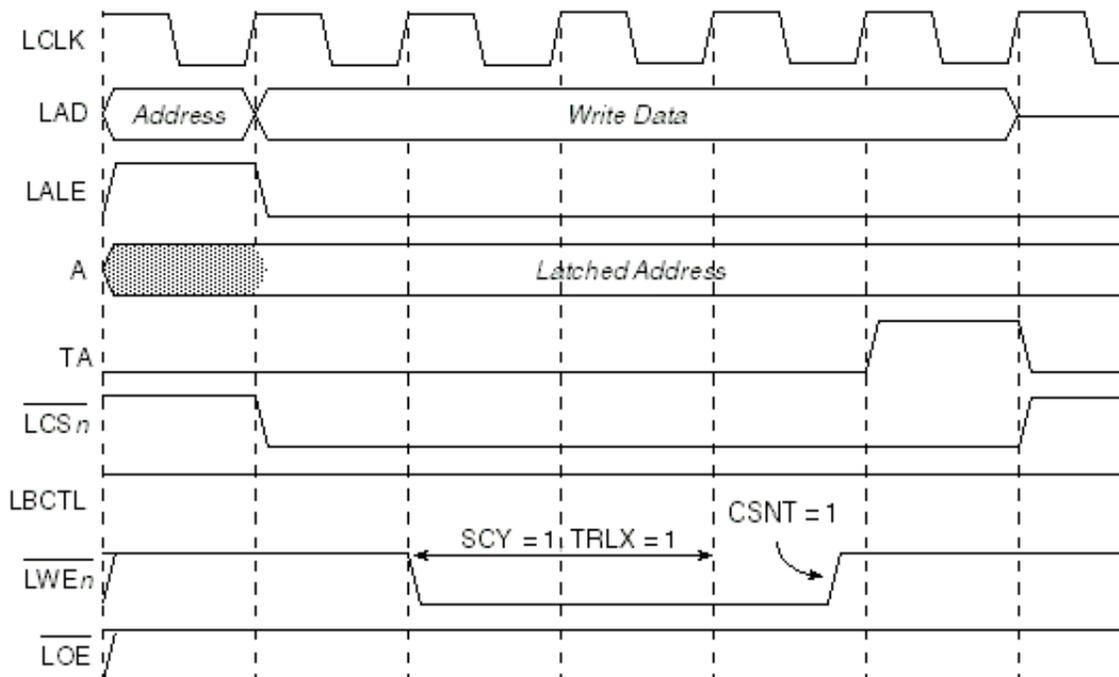


Figure 12-90. GPCM relaxed timing write ($XACS = 0$, $ACS = 00$, $SCY = 1$, $CSNT = 1$, $TRLX = 1$)

12.4.2.3.4 Output enable (LOE_B) timing

The timing of the LOE_B is affected only by TRLX.

It always asserts and negates on the rising edge of the bus clock. LOE_B asserts either on the rising edge of the bus clock after $LCSn_B$ is asserted or coinciding with $LCSn_B$ (if $XACS = 1$ and $ACS = 10$ or $ACS = 11$). Accordingly, assertion of LOE_B can be delayed (along with the assertion of $LCSn_B$ by programming $TRLX = 1$. LOE_B negates on the rising clock edge coinciding with $LCSn_B$ negation

12.4.2.3.5 Extended hold time on read accesses-GPCM

Slow memory devices that take a long time to disable their data bus drivers on read accesses should choose some combination of $ORn[TRLX, EHTR]$.

Any access following a read access to the slower memory bank is delayed by the number of clock cycles specified in [Options register 0 layout for GPCM Mode \(eLBC_ORg0\)](#) in addition to any existing bus turnaround cycle. The final bus turnaround cycle is automatically inserted by the eLBC for reads, regardless of the setting of $ORn[EHTR]$.

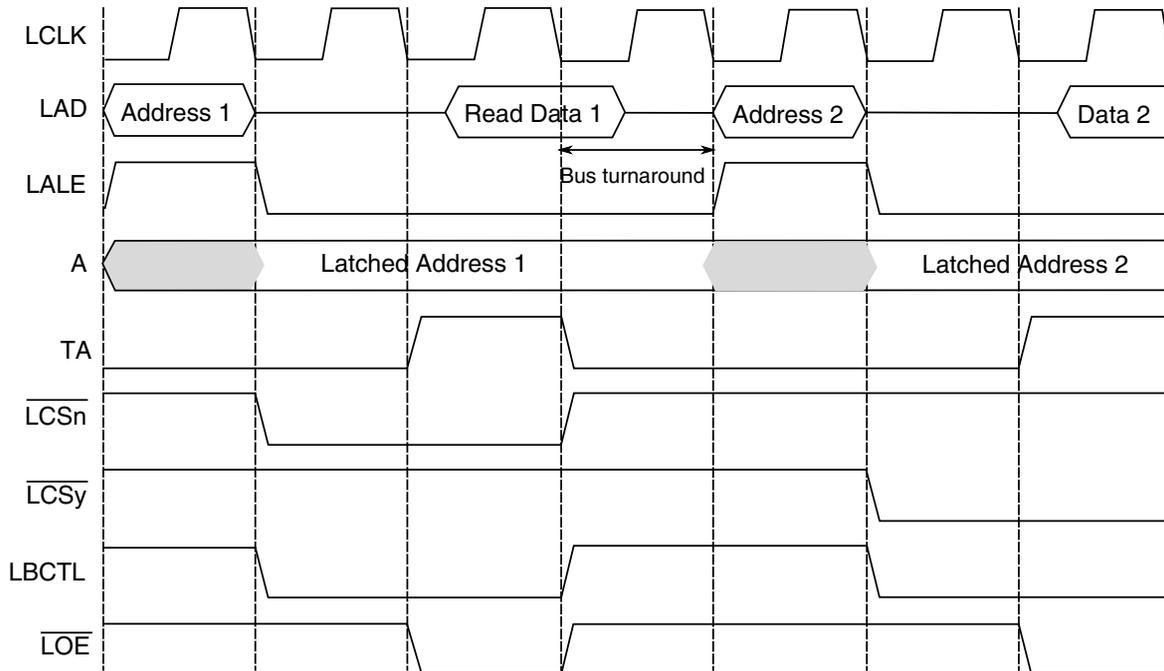


Figure 12-91. GPCM read followed by read (TRLX = 0, EHTR = 0, fastest timing)

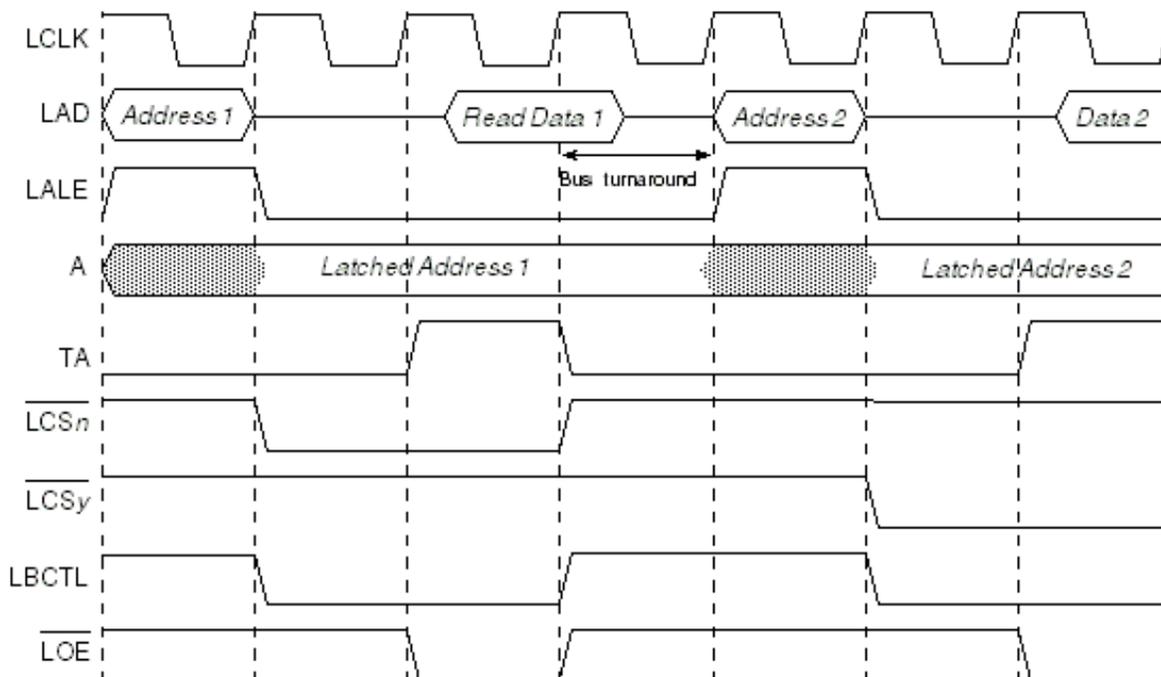


Figure 12-92. GPCM read followed by read (TRLX = 0, EHTR = 0, fastest timing)

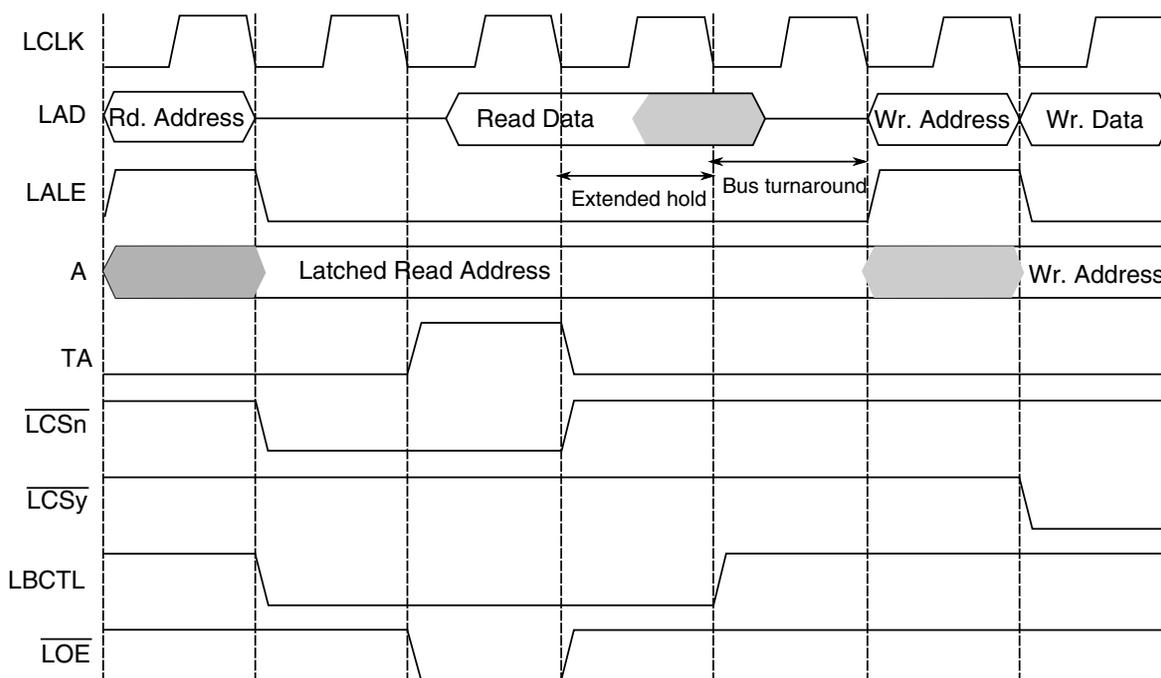


Figure 12-93. GPCM read followed by write (TRLX = 0, EHTR = 1, one-cycle extended hold time on reads)

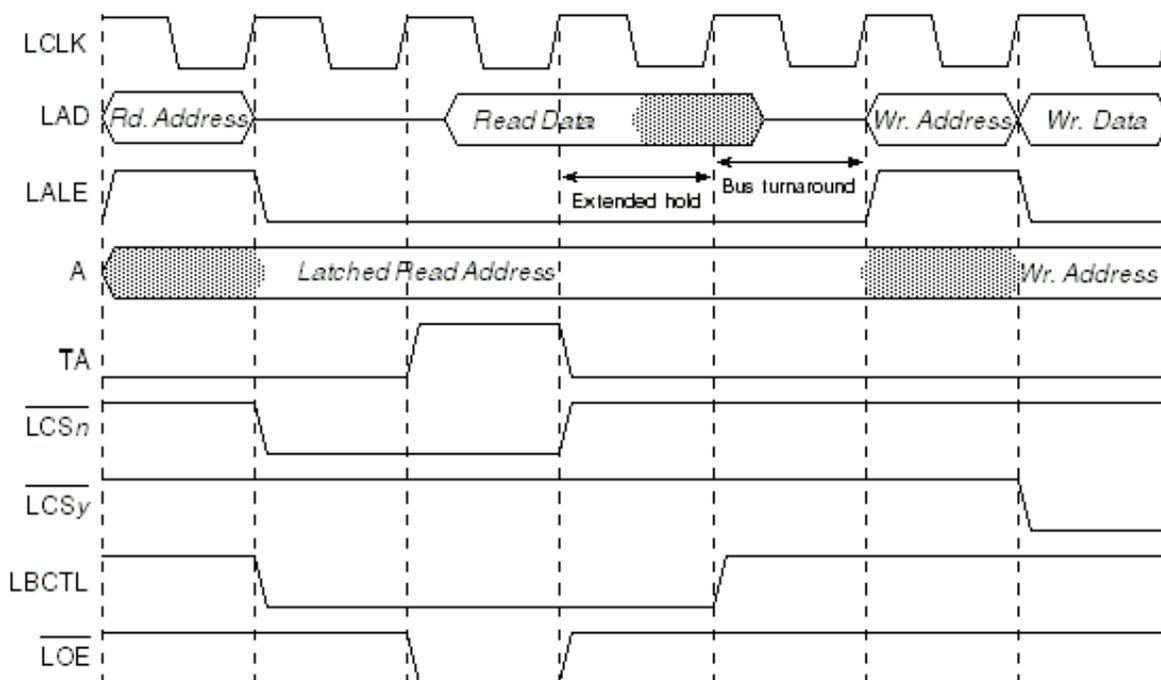


Figure 12-94. GPCM read followed by write (TRLX = 0, EHTR = 1, one-cycle extended hold time on reads)

12.4.2.4 External access termination (LGTA_B)

External access termination is supported by the GPCM using the asynchronous LGTA_B input signal, which is synchronized and sampled internally by the local bus.

If, during assertion of LCS_{n_B}, the sampled LGTA_B signal is asserted, it is converted to an internal generation of transfer acknowledge, which terminates the current GPCM access (regardless of the setting of OR_n[SETA]). LGTA_B must be asserted for at least two bus cycles to be effective. Note that because LGTA_B is synchronized, bus termination occurs two cycles after LGTA_B assertion, so in case of read cycle, the device still must drive data as long as LOE_B is asserted.

The user selects whether transfer acknowledge is generated internally or externally (LGTA_B) by programming OR_n[SETA]. Asserting LGTA_B always terminates an access, even if OR_n[SETA] = 0 (internal transfer acknowledge generation), but it is the only means by which an access can be terminated if OR_n[SETA] = 1. In PLL-enabled mode, the timing of LGTA_B is illustrated by the example in the figure below.

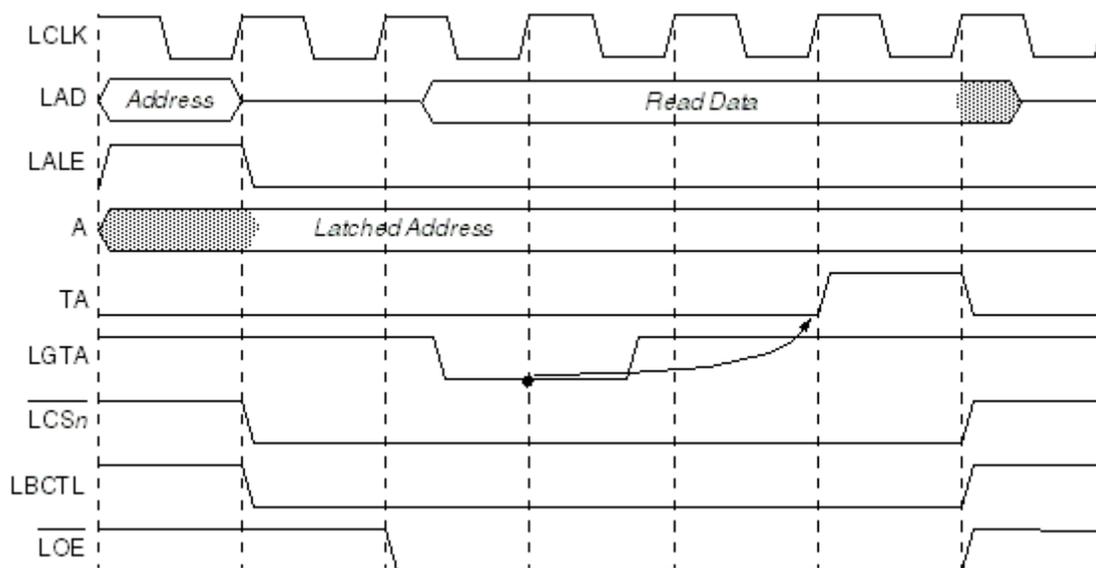


Figure 12-95. External termination of GPCM access (PLL enabled mode)

In PLL bypass mode, the timing of LGTA_B is illustrated by the example in the figure below.

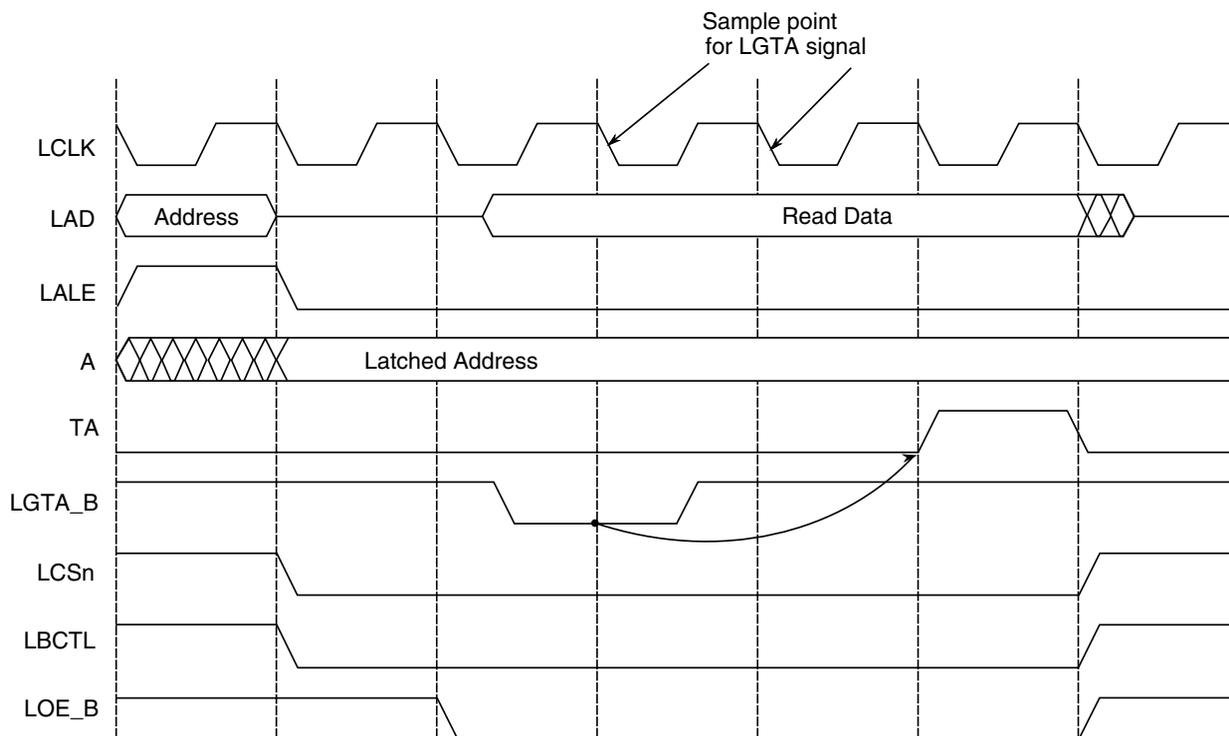


Figure 12-96. External termination of GPCM access (PLL bypass mode)

12.4.2.5 GPCM boot chip-select operation

Boot chip-select operation allows address decoding for a boot ROM before system initialization.

LCS0_B is the boot chip-select output; its operation differs from other external chip-select outputs after a system reset. When the core begins accessing memory after system reset, LCS0_B is asserted for every local bus access until BR0 or OR0 is reconfigured.

The boot chip-select also provides a programmable port size, which is configured during reset. The boot chip-select does not provide write protection. LCS0_B operates this way until the first write to OR0 and it can be used as any other chip-select register after the preferred address range is loaded into BR0. After the first write to OR0, the boot chip-select can be restarted only with a hardware reset.

Table 12-192. Boot bank field values after reset for GPCM as boot controller

Register	Field	Setting
BR0	BA	0000_0000_0000_0000_0
	PS	From cfg_rom_loc
	DECC	00
	WP	0
	MSEL	000
	-	-
	V	1
OR0	AM	0000_0000_0000_0000_0
	BCTLD	0
	CSNT	1
	ACS	11
	XACS	1
	SCY	1111
	SETA	0
	TRLX	1
	EHTR	1
	EAD	1

12.4.3 Flash control machine (FCM)

The FCM provides a glueless interface to parallel-bus NAND Flash EEPROM devices.

The FCM contains three basic configuration register groups-BR n , OR n , and FMR.

The figure below shows a simple connection between an 8-bit port size NAND Flash EEPROM and the eLBC in FCM mode. Commands, address bytes, and data are all transferred on LAD[0:7]⁸, with LFWE_B asserted for transfers written to the device, or LFRE_B asserted for transfers read from the device. eLBC signals LFCLE and LFALE determine whether writes are of type command (only LFCLE asserted), address (only LFALE asserted), or write data (neither LFCLE nor LFALE asserted). The NAND Flash RDY/BSY_B pin is normally open-drain, and should be pulled high by a 4.7-K Ω resistor. On system reset, a global (boot) chip-select is available that provides a boot ROM chip-select (LCS0_B) prior to the system being fully configured.

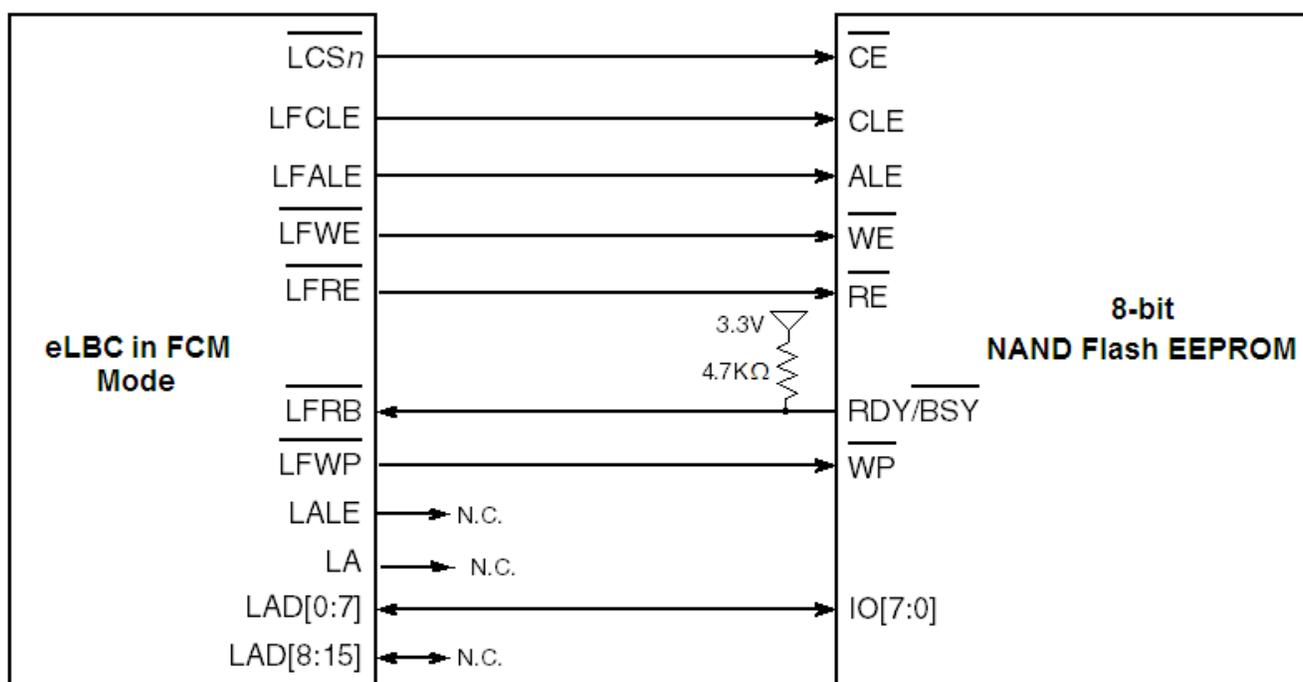


Figure 12-97. Local bus to 8-bit FCM device interface

Although LCLK is shown for reference, NAND Flash EEPROMs do not make use of the clock.

8. Note bit numbering reversal: LAD[0] (msb) connects to Flash IO[7], while LAD[7] (lsb) connects to IO[0].

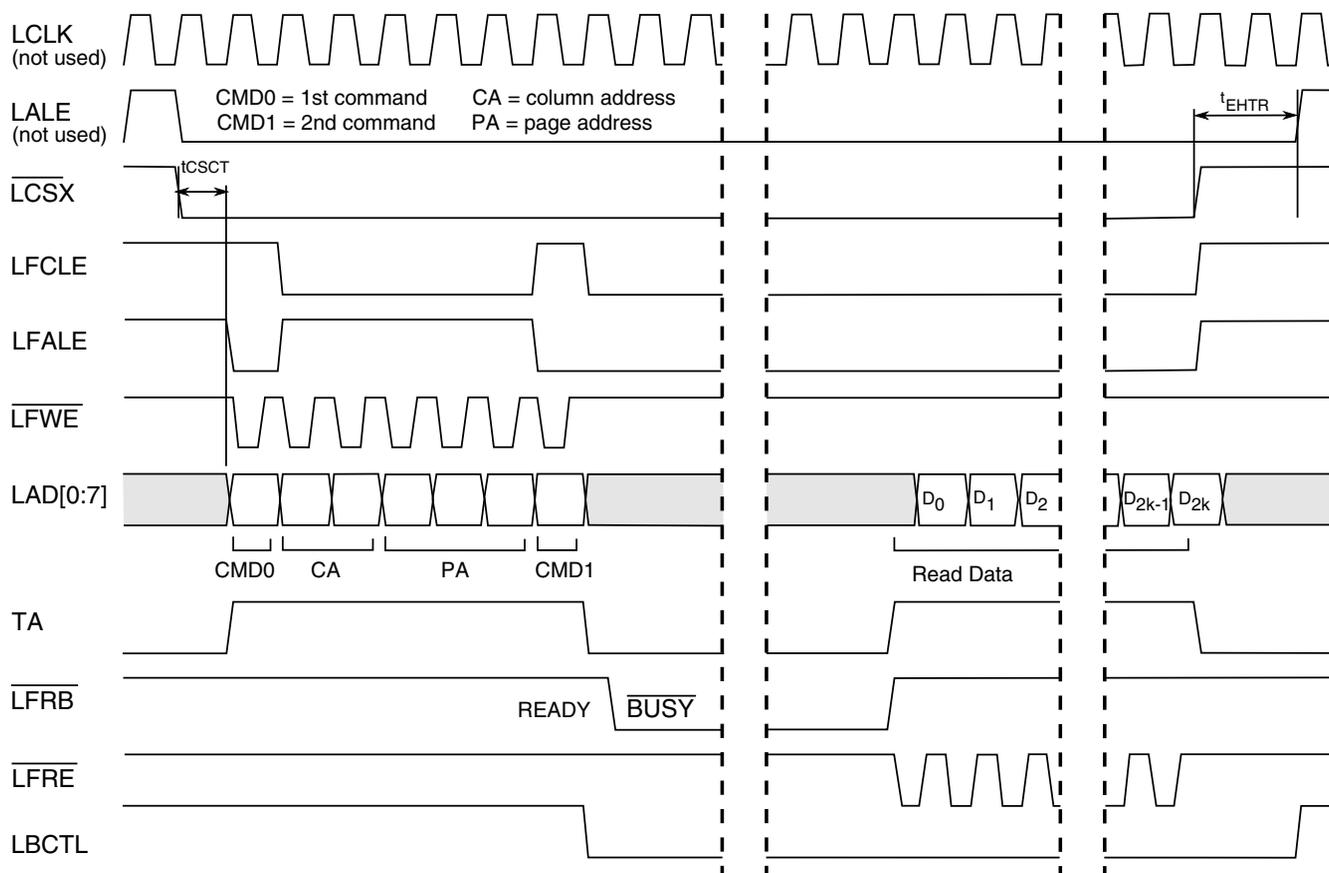


Figure 12-98. FCM basic page read timing (PGS = 1, CSCT = 0, CST = 0, CHT = 1, RST = 1, SCY = 0, TRLX = 0, EHTR = 1)

eLBC functional description

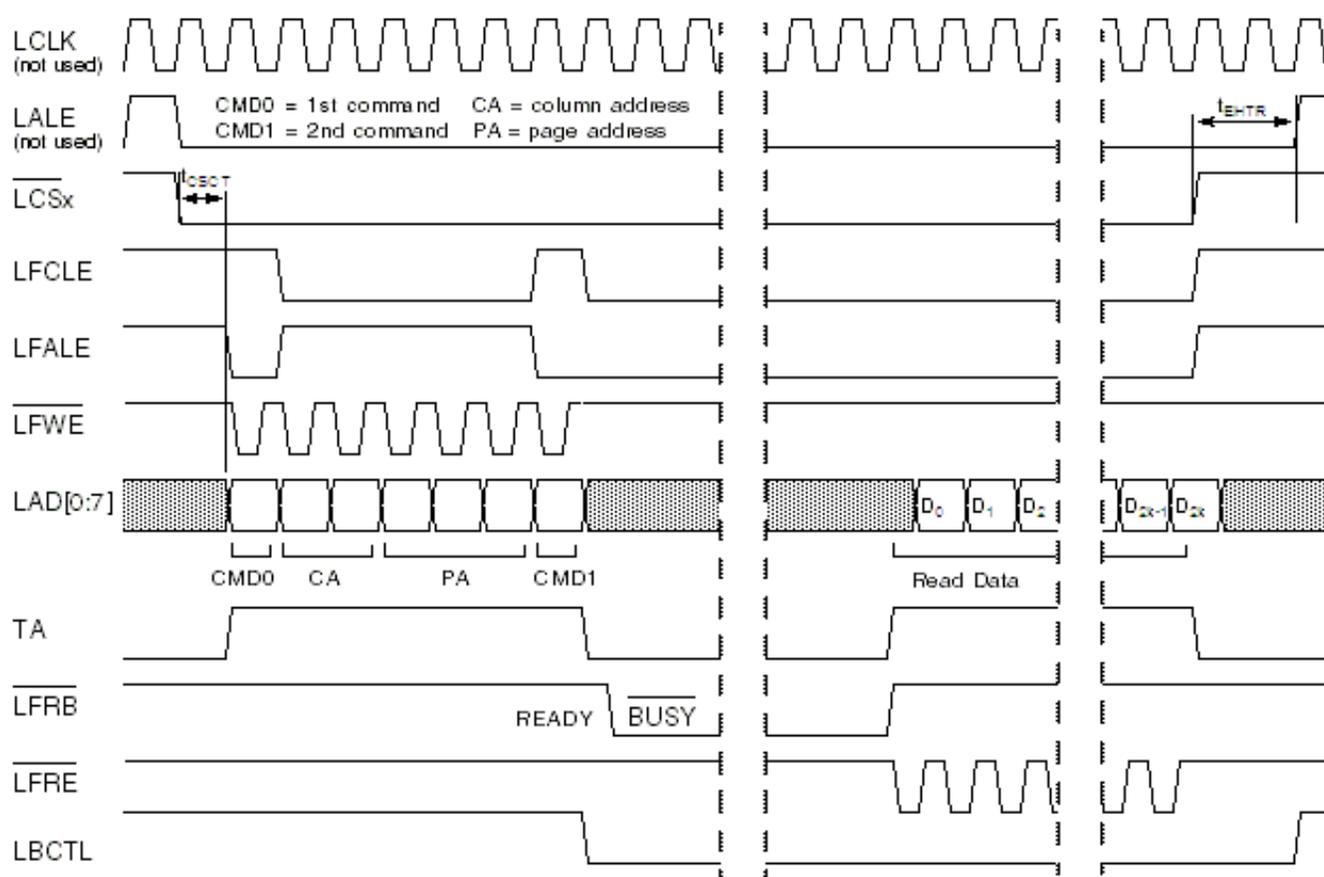


Figure 12-99. FCM basic page read timing (PGS = 1, CSCT = 0, CST = 0, CHT = 1, RST = 1, SCY = 0, TRLX = 0, EHTR = 1)

Following the assertion of LALE, FCM asserts LCSn_B to commence a command sequence to the Flash device. After a delay of t_{CSCT} , the first command can be written to the device on assertion of LFWE_B, followed by any parameters (typically address bytes and data), and concluded with a secondary command. In many cases, the second command initiates a long-running operation inside the Flash device, which pulls the wired-OR pin LFRB_B low to indicate that the device is busy. Since in [Figure 12-98](#) FCM is now expecting a read response, it takes LBCTL low to turnaround any bus transceivers that are present. Upon LFRB_B indicating ready status, FCM asserts LFRE_B repeatedly to recover bytes of read data, and the bytes are stored in eLBC's FCM buffer RAM while an ECC is optionally computed on the bytes transferred. Finally, FCM negates LCSn_B and delays eLBC by t_{EHTR} before any subsequent memory access occurs.

12.4.3.1 FCM buffer RAM

Read and write accesses to eLBC banks controlled by FCM do not access attached NAND Flash EEPROMs directly.

Rather, these accesses read and write the FCM buffer RAM—a single, shared 8-Kbyte space internal to eLBC and mapped by the base address of every FCM bank. Even though each FCM-controlled bank will have a different base address to differentiate it, all accesses to such banks will access the same buffer space. External eLBC signals, such as LALE and LCS_n_B, will not assert upon accesses to the buffer RAM. The FCM buffer RAM is logically divided into two or more buffers, depending on the setting of OR_n[PGS], with different buffers being accessible concurrently by software and FCM.

To perform a page read operation from a NAND Flash device, software initializes the FCM command, mode, and address registers, before issuing a special operation (FMR[OP] set non-zero) to a particular FCM-controlled bank. FCM will execute the sequence of op-codes held in FIR, reading data from the Flash device into the shared buffer RAM. While this read is taking place, software is free to access any data stored in other, currently inactive buffers of the FCM buffer RAM through reads or writes to any bank controlled by FCM. If command completion interrupts are enabled, an interrupt will be generated once FCM has completed the read. When FCM has completed its last command, software can switch to the newly read buffer and issue further commands.

To perform a page write operation, software first prepares data to be written in a fresh buffer. Then, the FCM command, mode, and address registers are initialized, and a special operation (FMR[OP] set non-zero) is issued to a particular FCM-controlled bank. FCM will execute the sequence of op-codes held in FIR, writing data from shared buffer RAM to the Flash device. To ensure that the device is enabled for programming, software must initialize FMR[OP] = 11, which prevents assertion of LFWP_B during the write. While this write is taking place, software is free to access any data stored in other, currently inactive buffers of the FCM buffer RAM through reads or writes to any bank controlled by FCM. When FCM has completed its last command, software can re-use the previously written buffer and issue further commands.

See [Boot block loading into the FCM buffer RAM](#) for a description of the shared buffer RAM layout during boot.

12.4.3.1.1 Buffer layout and page mapping for small-page NAND flash devices

The FCM buffer space is divided into eight 1-Kbyte buffers for small-page devices (OR_n[PGS] = 0), mapped as shown in [Figure 12-100](#).

Each page in a small-page NAND Flash comprises 528 bytes, where 512 bytes appear as main region data, and 16 bytes appear as spare region data. The EEPROM's page numbered P is associated with buffer number $(P \bmod 8)$, where $P = \text{FPAR}[PI]$. Since the bank size set by OR_n[AM] will be greater than 8 Kbytes, an identical image of the FCM

buffer RAM appears replicated every 8 Kbytes throughout the bank address space. It is recommended that the bank size be set to 32 Kbytes, which covers a single NAND Flash block for small-page devices.

For FCM commands, register FPAR sets the page address and, therefore, also the buffer number. In the case that FBCR[BC] = 0, FCM transfers an entire page, comprising the 512-byte main region followed by the 16-byte spare region; the 496-byte reserved region is not accessible, and remains undefined for software. However, for commands given a specific byte count in FBCR[BC], FPAR[MS] locates the starting address in either the main region (MS = 0) or the spare region (MS = 1). Where different eLBC banks control both small and large-page devices, a large-page 4-Kbyte buffer must be assigned to either the first 4 or last 4 small-page buffers.

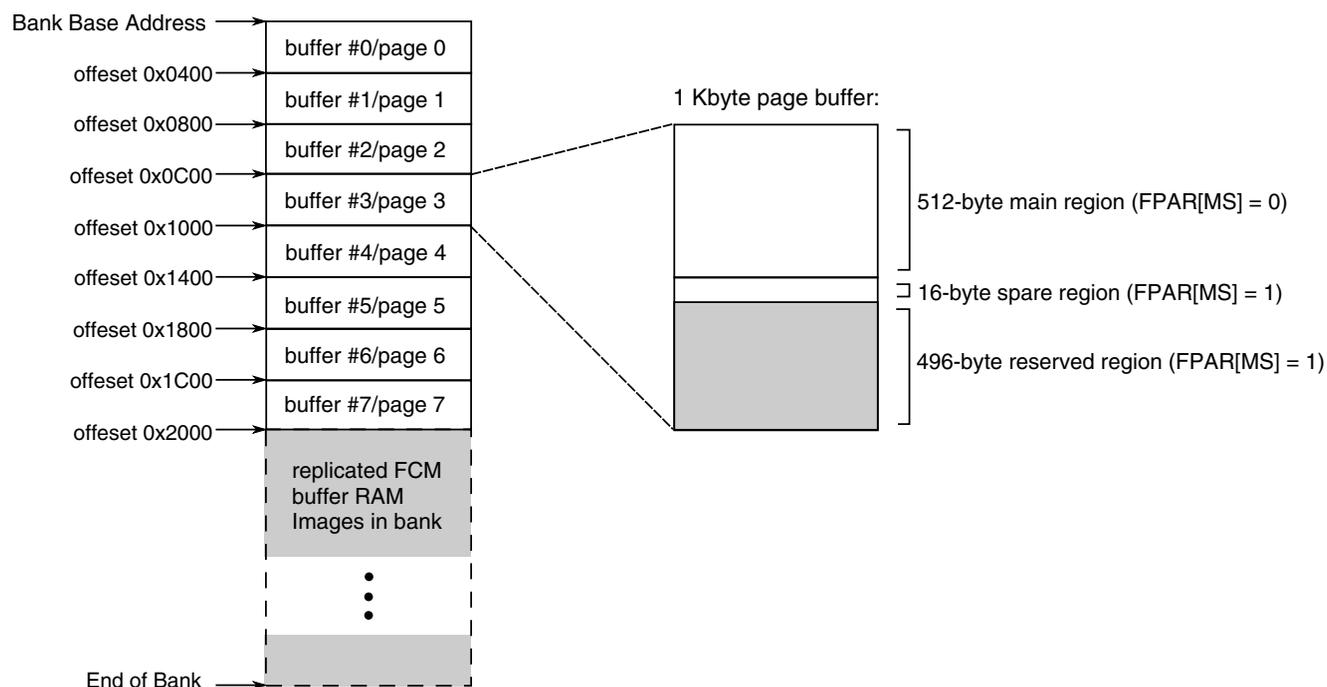


Figure 12-100. FCM buffer RAM memory map for small-page (512-byte page) NAND flash devices

12.4.3.1.2 Buffer layout and page mapping for large-page NAND flash devices

The FCM buffer space is divided into two 4 Kbyte buffers for large-page devices (ORn[PGS] = 1), mapped as shown in [Figure 12-101](#).

Each page in a large-page NAND Flash comprises 2112 bytes, where 2048 bytes appear as main region data, and 64 bytes appear as spare region data. The EEPROM's page numbered P is associated with buffer number $(P \text{ mod } 2)$, where $P = \text{FPAR}[\text{PI}]$. Since the bank size set by ORn[AM] will be greater than 8 Kbytes, an identical image of the FCM

buffer RAM appears replicated every 8 Kbytes throughout the bank address space. It is recommended that the bank size be set to 256 Kbytes, which covers a single NAND Flash block for large-page devices.

For FCM commands, register FPAR sets the page address and, therefore, also the buffer number. In the case that FBCR[BC] = 0, FCM transfers an entire page, comprising the 2048-byte main region followed by the 64-byte spare region; the 1984-byte reserved region is not accessed, and remains undefined for software. However, for commands given a specific byte count in FBCR[BC], FPAR[MS] locates the starting address in either the main region (MS = 0) or the spare region (MS = 1). Where different eLBC banks control both small and large-page devices, a large-page 4 Kbyte buffer must be assigned to either the first 4 or last 4 small-page buffers.

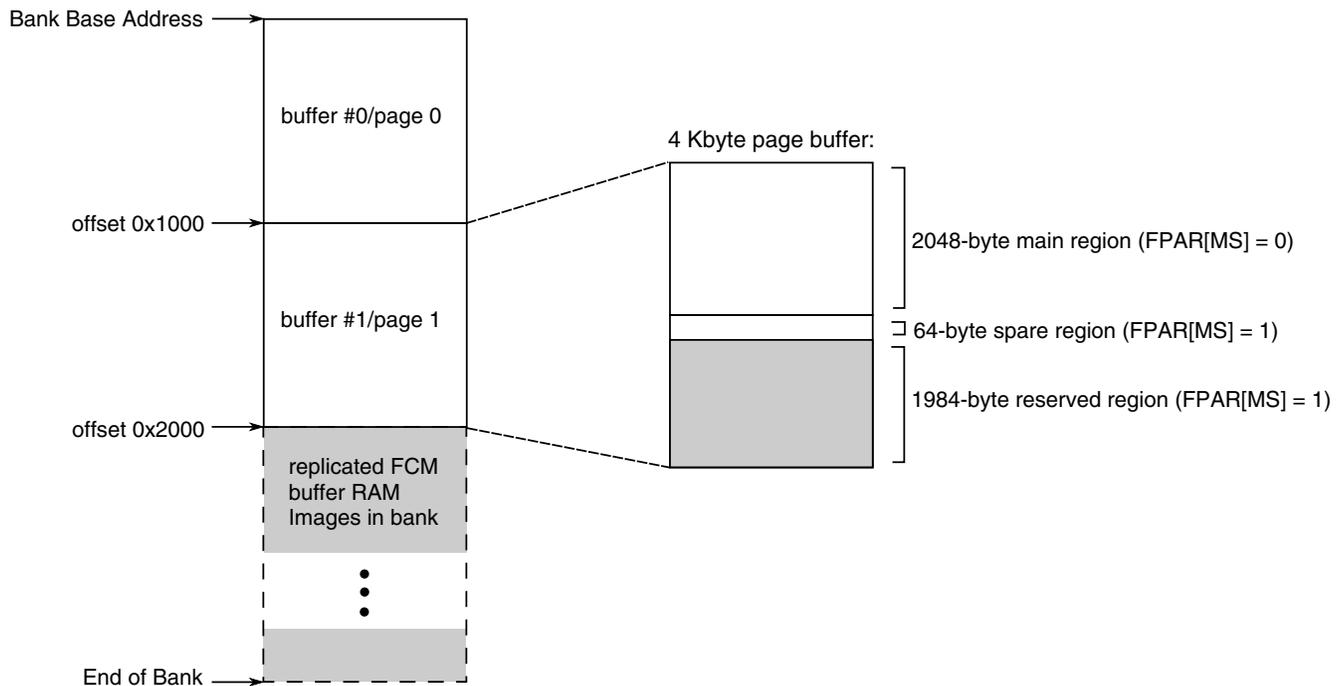


Figure 12-101. FCM buffer RAM memory map for large-page (2-Kbyte page) NAND flash devices

12.4.3.1.3 Error correcting codes and the spare region

The FCM's ECC engine makes use of data in the NAND Flash spare region to store pre-computed ECC code words.

ECC is calculated in a single pass over blocks of 512 bytes of data in the main region. The setting of FMR[ECCM] determines the location of the 24-bit ECC in the spare region.

The basic ECC algorithm is depicted in [Figure 12-102](#). The stream of data bytes is considered to form a matrix having 8 columns (corresponding with the device bus IO[7:0] or IO[15:8]) and 512 rows (corresponding with each byte in the ECC block). Six bits of parity, {P₄, P₄', P₂, P₂', P₁, P₁'}, are calculated across the columns, and at most 18 bits of parity {P₂₀₄₈, P₂₀₄₈', ..., P₁₆, P₁₆', P₈, P₈'} are calculated across the rows to create a 24-bit Hamming code for the data block. In this calculation, parity bit P_N' is the exclusive-OR of every alternate N-bit group of bits positioned at even intervals (starting at N-bit group 0, then continuing to group 2, 4, and so on), while parity bit P_N is the exclusive-OR of every alternate N-bit group of bits positioned at odd intervals (starting at N-bit group 1, then continuing to group 3, 5, and so on).

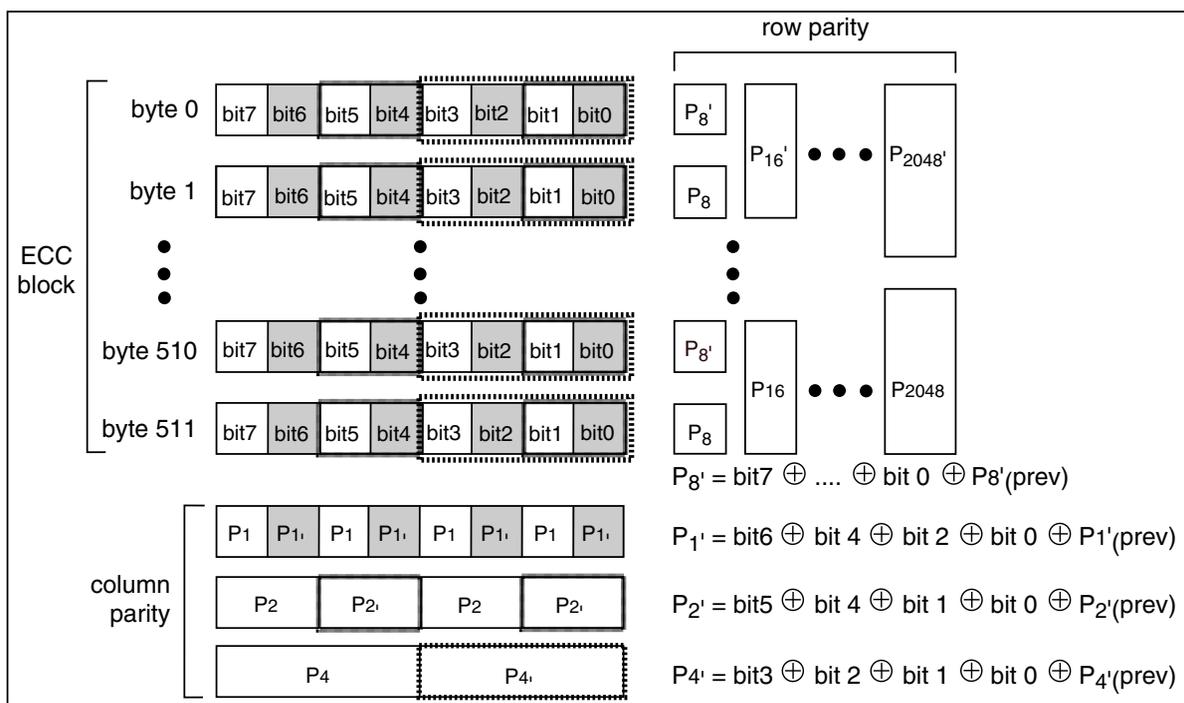


Figure 12-102. FCM ECC calculation

The 24-bit ECC code word format is shown in [Table 12-193](#) for normal ECC polarity. Setting LBCR[EPAR] = 1 changes ECC polarity, and thus omits negation of each P_N and P_N' bit.

Table 12-193. ECC layout for LBCR[EPAR] = 0 (~ represents logical negation)

	0 (MSB)	1	2	3	4	5	6	7 (LSB)
EC0	~P ₆₄	~P ₆₄ '	~P ₃₂	~P ₃₂ '	~P ₁₆	~P ₁₆ '	~P ₈	~P ₈ '
EC1	~P ₁₀₂₄	~P ₁₀₂₄ '	~P ₅₁₂	~P ₅₁₂ '	~P ₂₅₆	~P ₂₅₆ '	~P ₁₂₈	~P ₁₂₈ '
EC2	~P ₄	~P ₄ '	~P ₂	~P ₂ '	~P ₁	~P ₁ '	~P ₂₀₄₈	~P ₂₀₄₈ '

The placement of ECC code words in relation to FMR[ECCM] is shown in [Table 12-194](#). For small-page devices, only a single 512-byte main region is ECC-protected. For large-page devices, there are four adjacent main regions, and each has a 16-byte spare region-of which only one is shown in the figure. If eLBC is configured to generate ECC ($BRn[DECC] = 10$), FCM will substitute on full-page write transfers the three code word bytes in place of the spare region data originally provided at the locations shown in [Table 12-194](#) and write the same 24-bit ECC code in the appropriate FECC n register for software reference. Transfers shorter than a full page, however, require software to prepare the appropriate ECC in the spare region. Similarly, FCM can check and correct bit errors on full-page reads if $BRn[DECC] = 01$ or 10 . A correctable error is a single bit error in any 512-byte block of main region data, as judged by comparison of a regenerated ECC with the ECC retrieved from the spare region, or a single bit error in the retrieved ECC only. Bit errors in the main region are corrected before FCM completes its final read transfer and signals an event in LTESR[CC]. The bit vector in LTECCR[SBCE] can be checked on FCM CC event to find out if any 512-byte block or the corresponding ECC have single bit correctable errors. Errors that appear more complex (two or more bits in error per 512-byte block) are not corrected, but are flagged as parity errors by FCM. The bit vector in LTEATR[PB] or LTECCR[MBUE] can be checked to determine which 512-byte blocks in a large-page NAND Flash main region were found to be non-correctable.

Table 12-194. ECC placement in NAND flash spare regions in relation to FMR[ECCM]

ECCM	Byte 0	Byte 511	Other Mains	Spare 0	5	6	7	8	9	10	11	12	13	14	15
0	Main Region			-		EC0	EC1	EC2	-						
1	Main Region			-				EC0	EC1	EC2	-				

12.4.3.2 Programming FCM

FCM has a fully general command and data transfer sequencer that caters for both common and specific/proprietary NAND Flash command sequences.

The command sequencer reads a program out of the FIR register, which can hold up to 8 instructions, each represented by a 4-bit op-code, as illustrated in [Figure 12-103](#). The first instruction executed is read from FIR[OP0], the next is read from FIR[OP1], and likewise to subsequent instructions, ending at FIR[OP7] or until the only instructions remaining are NOPs. If FIR contains nothing but NOP instructions, FCM will not assert $LCSn_B$, otherwise, $LCSn_B$ is asserted prior to the first instruction and remains asserted until the last instruction has completed. If LTESR[CC] is enabled, completion of the last instruction will trigger a command completion event interrupt from eLBC.

Prior to executing a sequence, necessary operands for the instructions will need to be set in the FMR, FCR, MDR, FBCR, FBAR, and FPAR registers. The AS0-AS3 address and data pointers associated with FCM's use of MDR all reset to select AS0 at the start of the instruction sequence. A complete list of op-codes can be found in [Flash instruction register \(eLBC_FIR\)](#).

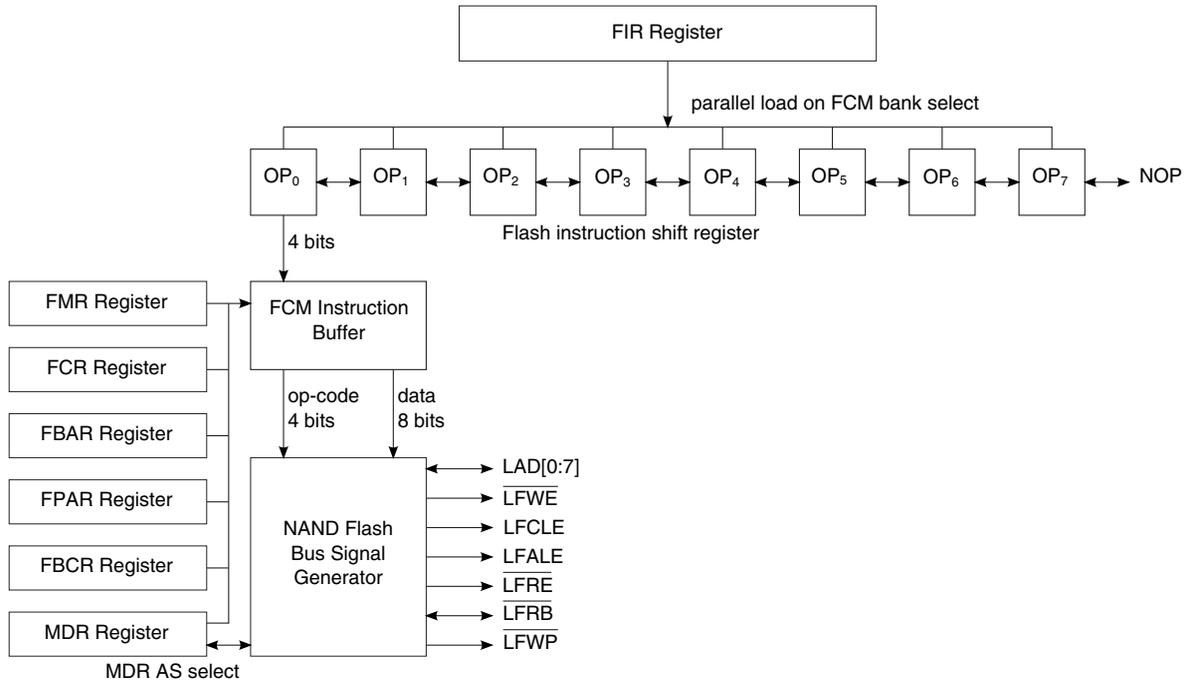


Figure 12-103. FCM instruction sequencer mechanism

12.4.3.2.1 FCM command instructions

There are two kinds of command instruction:

- Commands that issue immediately-CM0, CM1, CM2, and CM3. These commands write a single command byte by asserting LFCLE and LFWE_B while driving an 8-bit command onto LAD[0:7]. Op-code CM n sources its command byte from field FCR[CMD n], therefore up to four different commands can be issued in any FCM instruction sequence.
- Commands that wait for LFRB_B to be sampled high (EEPROM in ready state) before issuing-CW0, and CW1. These commands first poll the LFRB_B pin, waiting for it to go high, before writing a single command byte onto LAD[0:7], sourced from FCR[CMD n] for op-code CW n . It is necessary to use CW n op-codes whenever the EEPROM is expected to be in a busy state (such as following a page read, block erase, or program operation) and therefore initially unresponsive to commands. To avoid deadlock in cases where the device is already available, FCM does not expect a transition on LFRB_B. Rather, FCM waits for 8 x (2 + OR n [SCY]) clock cycles

(when $OR_n[TRLX] = 0$) or $16 \times (2 + OR_n[SCY])$ clock cycles (when $OR_n[TRLX] = 1$) before sampling the level of LFRB_B. If the level of LFRB_B does not return high before a time-out set by FMR[CWTO] occurs, FCM proceeds to issue the command normally, and a FCT event is issued to LTESR.

The manufacturer's datasheet should be consulted to determine values for programming into the FCR register, and whether a given command in the sequence is expected to initiate busy device behavior.

12.4.3.2.2 FCM no-operation instruction

A NOP instruction that appears in FIR ahead of the last instruction is executed with the timing of a regular command instruction, but neither LFCLE nor LFWE_B are asserted.

Thus a NOP instruction may be used to insert a pause matching the time taken for a regular command write.

12.4.3.2.3 FCM address instructions

Address instructions are used to issue addresses to the NAND Flash EEPROM.

A complete device address is formed from a sequence of one or more bytes, each written onto LAD[0:7] with LFALE and LFWE_B asserted together. There are three kinds of address generation provided:

- Column address-CA. A column address comprises one byte ($OR_n[PGS] = 0$) or two bytes ($OR_n[PGS] = 1$) locating the starting byte or word to be transferred in the next page read or write sequence. FPAR[CI] sets the value of the column index provided that FBCR[BC] is non-zero. In the case that FBCR[BC] = 0, a column index of zero is issued to the device, regardless of the value in FPAR[CI].
- Page address-PA. A page address comprises 2, 3, or 4 bytes, depending on the setting of FMR[AL], and locates the data page in the NAND Flash address space. The complete page address is the concatenation of the block index, read from FPAR[BLK], with the page-in-block index, read from FPAR[PI]. The page address length set in FMR[AL] should correspond with the size of EEPROM being accessed. Similarly, the block index in FPAR[BLK] must not exceed the maximum block index for the device, as most devices require reserved address bits to be written as zero.
- User-defined address-UA. This instruction allows the FCM to write a user-defined address byte, which is read from the next AS field in MDR, starting at MDR[AS0]. Each subsequent UA instruction reads an adjacent AS field in MDR, until all four AS bytes (MDR[AS0], MDR[AS1], MDR[AS2], MDR[AS3]) have been sent; a fifth and any following UA instructions send zero as the address byte. Note that each UA

instruction advances the MDR pointer for writes by one byte, and therefore a mix of UA and WS instructions can consume adjacent bytes from MDR.

12.4.3.2.4 FCM data read instructions

Data read instructions assert LFRE_B repeatedly to transfer one or more bytes of read data from the NAND Flash EEPROM.

Data read instructions are distinguished by their data destination:

- Read data to buffer RAM immediately- RB. This instruction reads FBCR[BC] bytes of data into the current FCM RAM buffer addressed by FPAR. If FBCR[BC] = 0, an entire page (including spare region) is transferred in a burst, starting at the page boundary, and the ECC calculation is checked against the ECC stored in the spare region. Correctable ECC errors are corrected and reported in LTECCR[SBCE]; other errors may cause an interrupt if enabled. If the value of FBCR[BC] takes the read pointer beyond the end of the spare region in the buffer, FCM discards any excess bytes read.
- Read data/status to MDR immediately- RS. This instruction asserts LFRE_B exactly once to read one byte (8-bit port size) of data into the next AS field of MDR. Reads beyond the fourth byte of MDR are discarded. The MDR read pointer is independent of the MDR write pointer used by UA and WS instructions.
- Read data to buffer RAM once waited on ready- RBW. This instruction first polls the LFRB_B pin, waiting for it to go high, before proceeding with a read to buffer as described for the RB instruction. Sampling and time-outs for polling the LFRB_B pin follow the behavior of CW_n instructions.
- Read data/status to MDR once waited on ready- RSW. This instruction first polls the LFRB_B pin, waiting for it to go high, before proceeding with a status read to MDR as described for the RS instruction. Sampling and time-outs for polling the LFRB_B pin follow the behavior of CW_n instructions.

12.4.3.2.5 FCM data write instructions

Data write instructions assert LFWE_B repeatedly (with LFCLE and LFALE both negated) to transfer one or more bytes of write data to the NAND Flash EEPROM.

Data write instructions are distinguished by their data source:

- Write data from FCM buffer RAM-WB. This instruction writes FBCR[BC] bytes of data from the current FCM RAM buffer addressed by FPAR. If FBCR[BC] = 0, an entire page (including spare region) is transferred in a burst, starting at the page boundary, and the ECC calculation is stored in the appropriate FECC_n registers and spare region in accordance with the setting of FMR[ECCM]. If the value of

FBCR[BC] takes the write pointer beyond the end of the spare region in the buffer, the value of data written by FCM is undefined.

- Write data/status from MDR-WS. This instruction asserts LFWE_B exactly once to write one byte (8-bit port size) of data taken from the next AS field of MDR. Attempts to write beyond four bytes of MDR has the effect of writing zeros. The MDR write pointer is independent of the MDR read pointer used by RS and RSW instructions.

12.4.3.3 FCM signal timing

If BR_n [MSEL] selects the FCM, the attributes for the memory cycle are taken from OR_n . These attributes include the CSCT, CST, CHT, RST, SCY, TRLX, and EHTR fields.

12.4.3.3.1 FCM chip-select timing

The timing of LCS_n_B assertion in FCM mode is illustrated by the timing diagram in [Figure 12-98](#).

n_B is asserted immediately following LALE negation, and remains asserted until the last instruction in FIR has completed.

The delay, t_{CSCT} , between n_B assertion and commencement of the first NAND Flash instruction is controlled by OR_n [CSCT] and OR_n [TRLX], as shown in [Table 12-195](#). OR_n [CSCT] should be set in accordance with the NAND Flash EEPROM chip-select to WE_B set-up time specification.

Table 12-195. FCM chip-select to first command timing

OR_n [TRLX]	OR_n [CSCT]	LCS_n_B to First Command Delay
0	0	1 LCLK clock cycle
0	1	4 LCLK clock cycles
1	0	2 LCLK clock cycles
1	1	8 LCLK clock cycles

12.4.3.3.2 FCM command, address, and write data timing

The FCM command (CM0-CM3, CW0, CW1), address (CA, PA, UA), and data write (WB, WS) instructions all share the same basic timing attributes.

Assertion of LFWE_B initiates transfer via LAD[0:7], and the options in OR_n for FCM mode establish the set-up, hold, and wait state timings with respect to LFWE_B, as shown in [Figure 12-105](#)[Figure 12-104](#).

eLBC functional description

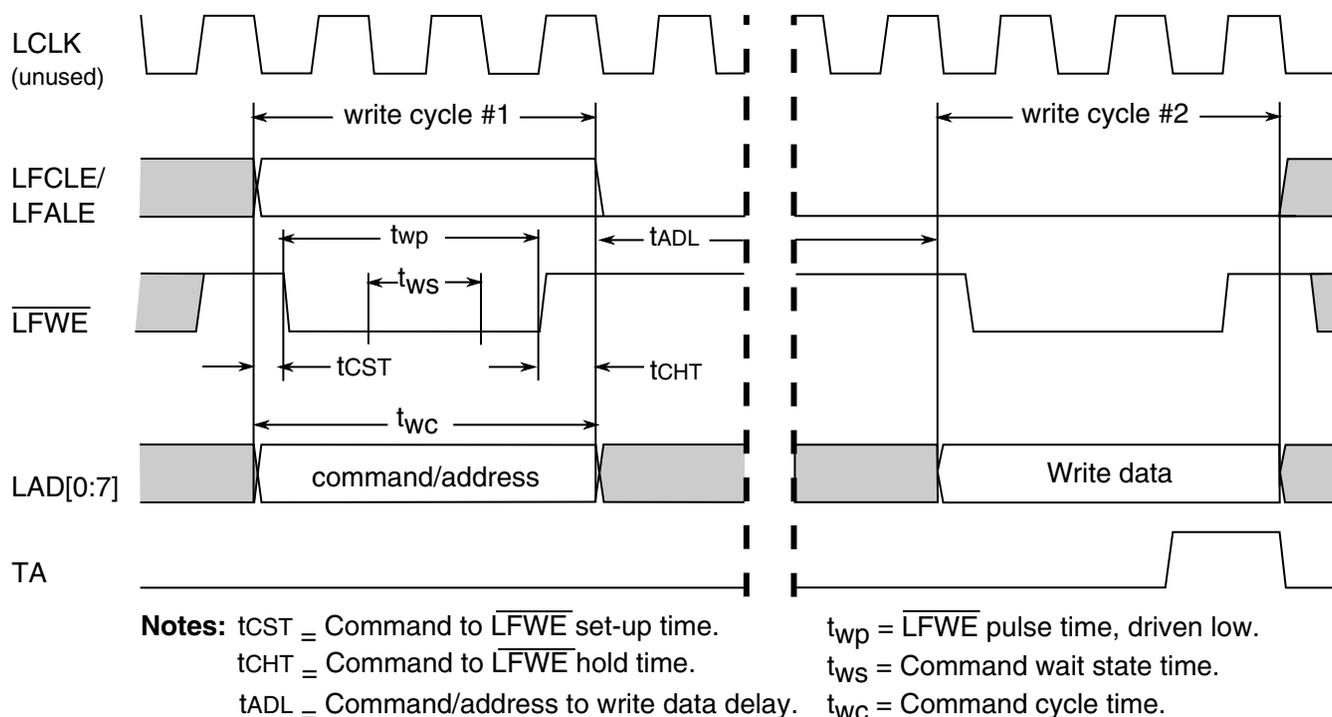


Figure 12-104. Timing of FCM command/address and write data cycles (for TRLX = 0, CHT = 0, CST = 1, SCY = 1)

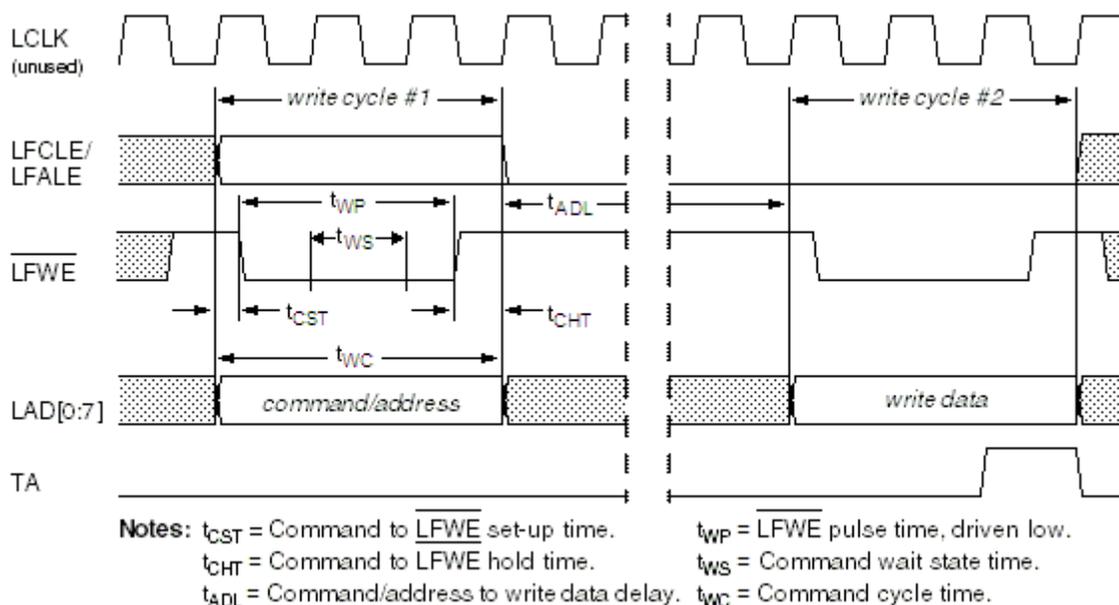


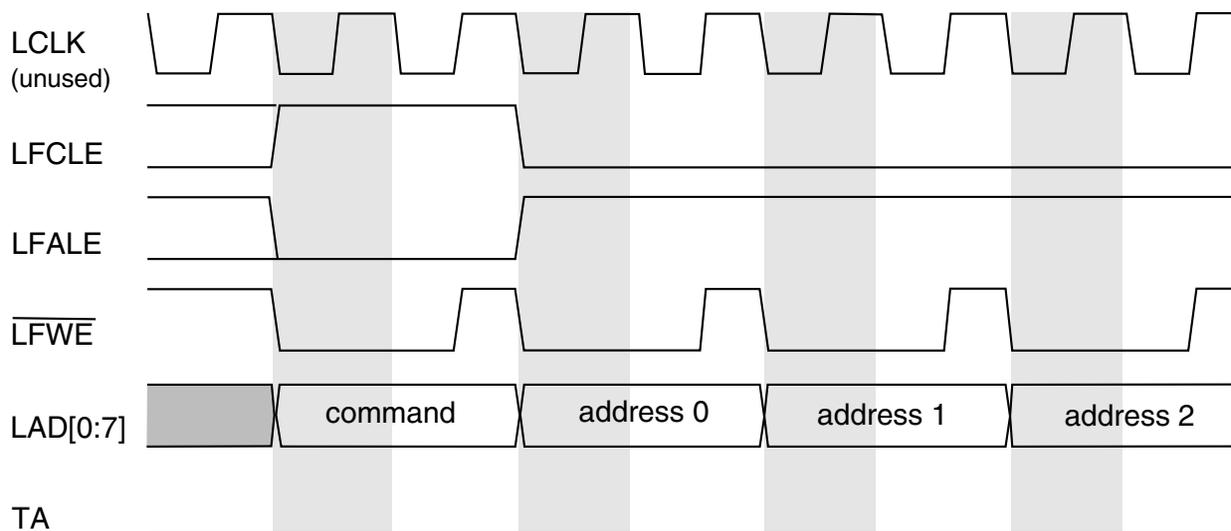
Figure 12-105. Timing of FCM command/address and write data cycles (for TRLX = 0, CHT = 0, CST = 1, SCY = 1)

Table 12-196. FCM command, address, and write data timing parameters

Option Register Attributes			Timing Parameter (LCLK Clock Cycles) ¹					
TRLX	CHT	CST	t_{CST}	t_{CHT}	t_{WS}	t_{WP}	t_{WC}	t_{ADL}
0	0	0	0	$\hat{A}\frac{1}{2}$	SCY	$1\hat{A}\frac{1}{2} + SCY$	$2 + SCY$	$4 \times (2 + SCY)$
0	0	1	$\hat{A}\frac{1}{4}$	$\hat{A}\frac{1}{2}$	SCY	$1\hat{A}\frac{1}{4} + SCY$	$2 + SCY$	$4 \times (2 + SCY)$
0	1	0	0	1	SCY	$1 + SCY$	$2 + SCY$	$4 \times (2 + SCY)$
0	1	1	$\hat{A}\frac{1}{4}$	1	SCY	$\hat{A}\frac{3}{4} + SCY$	$2 + SCY$	$4 \times (2 + SCY)$
1	0	0	$\hat{A}\frac{1}{2}$	$1\hat{A}\frac{1}{2}$	$2 \times SCY$	$1 + 2 \times SCY$	$3 + 2 \times SCY$	$8 \times (2 + SCY)$
1	0	1	1	$1\hat{A}\frac{1}{2}$	$2 \times SCY$	$\hat{A}\frac{1}{2} + 2 \times SCY$	$3 + 2 \times SCY$	$8 \times (2 + SCY)$
1	1	0	$\hat{A}\frac{1}{2}$	2	$2 \times SCY$	$\hat{A}\frac{1}{2} + 2 \times SCY$	$3 + 2 \times SCY$	$8 \times (2 + SCY)$
1	1	1	1	2	$2 \times SCY$	$2 \times SCY$	$3 + 2 \times SCY$	$8 \times (2 + SCY)$

1. In the parameters, SCY refers to a delay of ORn[SCY] clock cycles.

An example of minimum delay command timing appears in the figure below. Note that the set-up, wait-state, and hold timing of command, address, and write data cycles with respect to LFWE_B assertion are all identical, and that the minimum cycle extends for two LCLK clock cycles.


Figure 12-106. Example of FCM command and address timing with minimum delay parameters (for TRLX = 0, CHT = 0, CST = 0, SCY = 0)

eLBC functional description

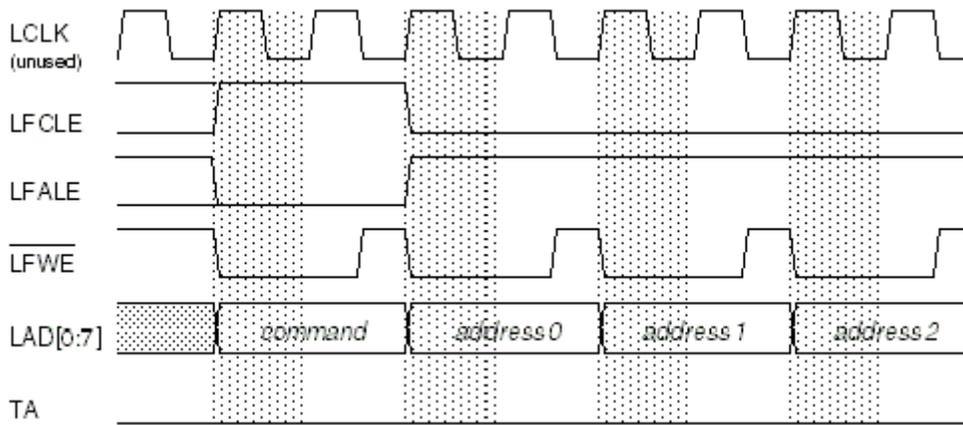


Figure 12-107. Example of FCM command and address timing with minimum delay parameters (for TRLX = 0, CHT = 0, CST = 0, SCY = 0)

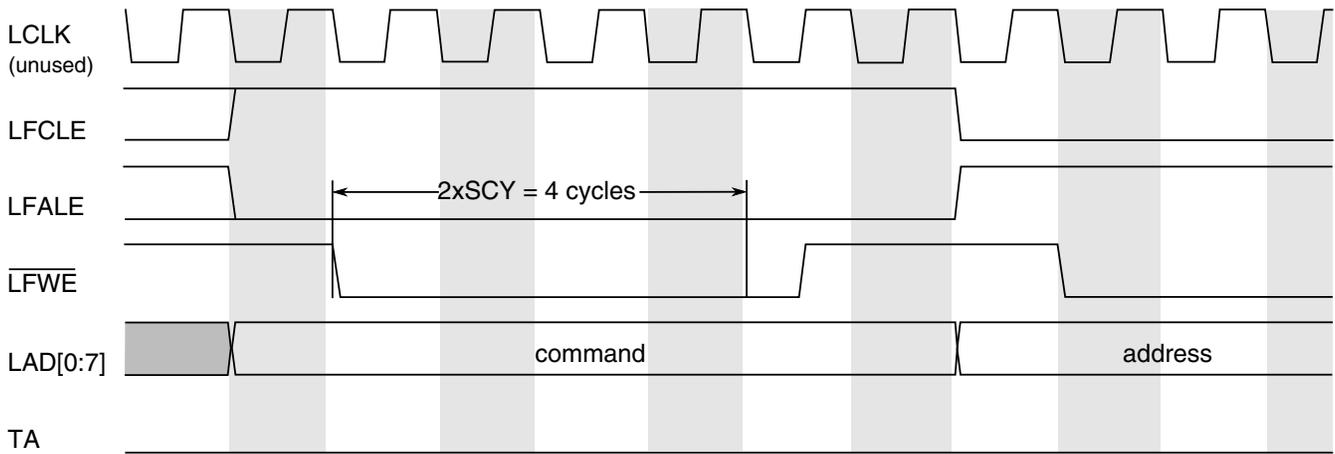


Figure 12-108. Example of FCM command and address timing with relaxed parameters (for TRLX = 1, CHT = 0, CST = 1, SCY = 2)

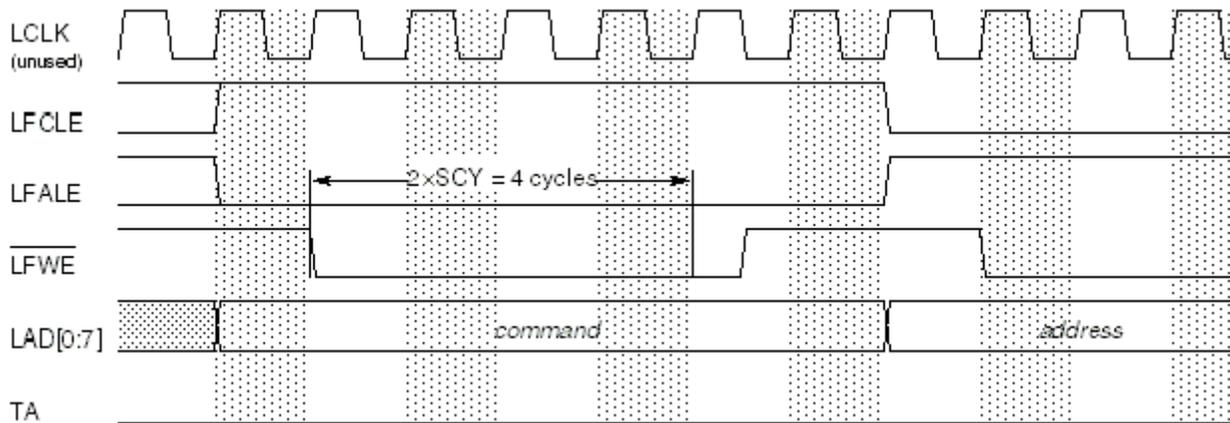


Figure 12-109. Example of FCM command and address timing with relaxed parameters (for TRLX = 1, CHT = 0, CST = 1, SCY = 2)

12.4.3.3.3 FCM ready/busy timing

Instructions CW0, CW1, RBW, and RSW force FCM to observe the state of the LFRB_B pin, which may be driven low by a long-latency NAND Flash operation, such as a page read.

Following the issue of such commands, FCM waits as shown in [Figure 12-110](#) before sampling the state of LFRB_B. This guards against observing LFRB_B before it has been properly driven low by the device, but does not preclude LFRB_B from remaining high after a command. In addition, FCM samples and compares the state of LFRB_B on two consecutive cycles of LCLK to filter out noise on this signal as it rises to the ready state (LFRB_B = 1).

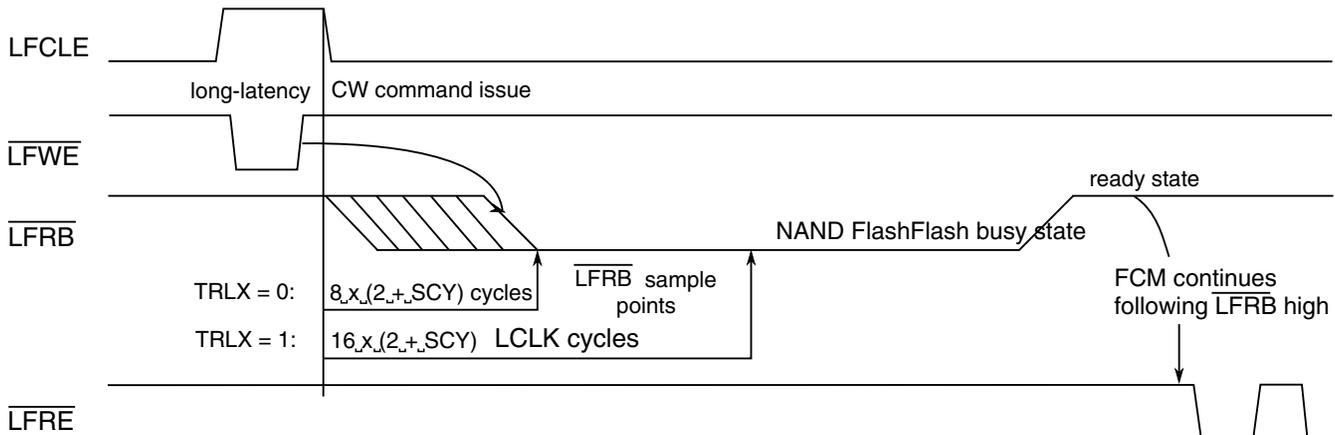


Figure 12-110. FCM delay prior to sampling LFRB_B state

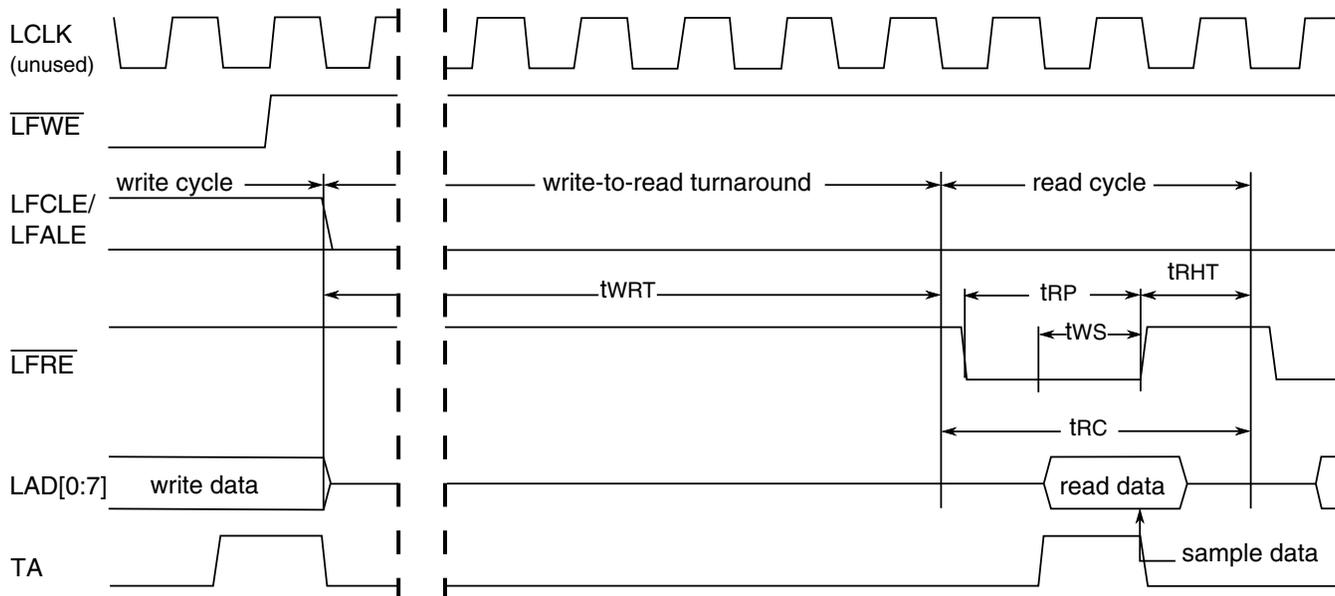
12.4.3.3.4 FCM read data timing

The timing for read data transfers is shown in the figure below.

Upon assertion of LFRE_B, the Flash device will enable its output drivers and drive valid read data while LFRE_B is held low.

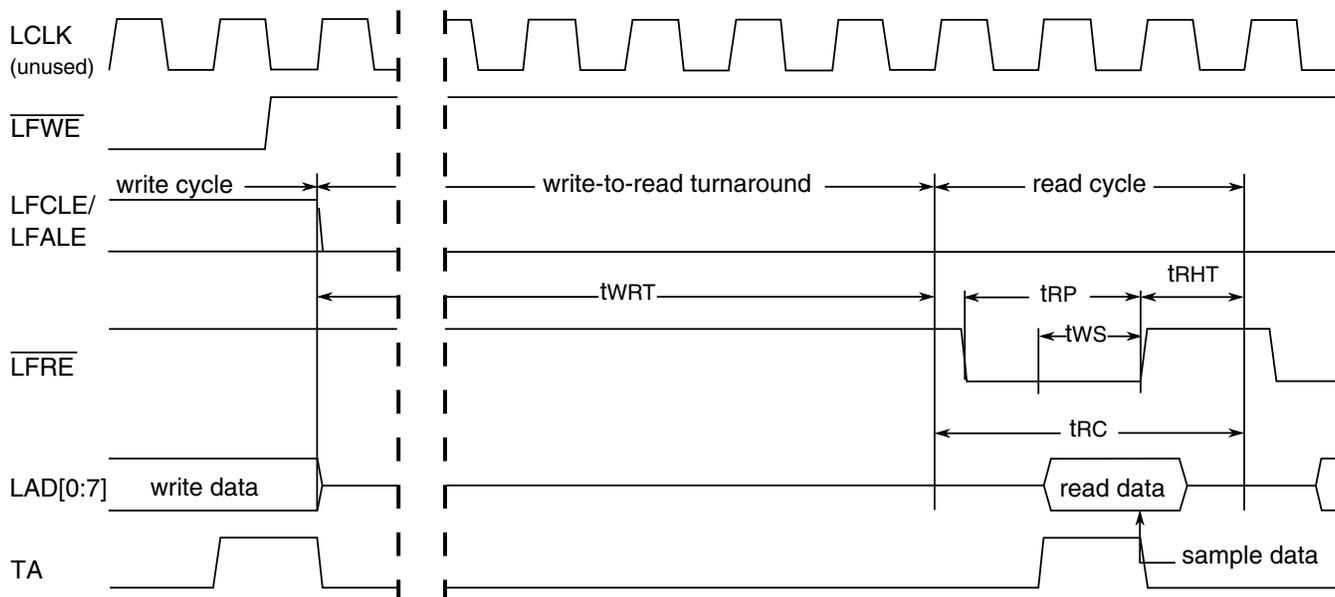
FCM samples read data on the rising edge of LFRE_B, which follows an optional number of wait states. Note that FCM will delay the first read if a RBW or RSW instruction is issued, in which case LFRB_B sample timing takes effect (see [FCM ready/busy timing](#)).

eLBC functional description



Notes: t_{RP} = \overline{LFRE} pulse time, read period. t_{WS} = Read wait state time.
 t_{RHT} = \overline{LFRE} hold time. t_{RC} = Read data cycle time.
 t_{WRT} = Write to read turnaround time.

Figure 12-111. FCM read data timing (for $TRLX = 0$, $RST = 0$, $SCY = 1$)



Notes: t_{RP} = \overline{LFRE} pulse time, read period. t_{WS} = Read wait state time.
 t_{RHT} = \overline{LFRE} hold time. t_{RC} = Read data cycle time.
 t_{WRT} = Write to read turnaround time.

Figure 12-112. FCM read data timing (for $TRLX = 0$, $RST = 0$, $SCY = 1$)

Table 12-197. FCM read data timing parameters

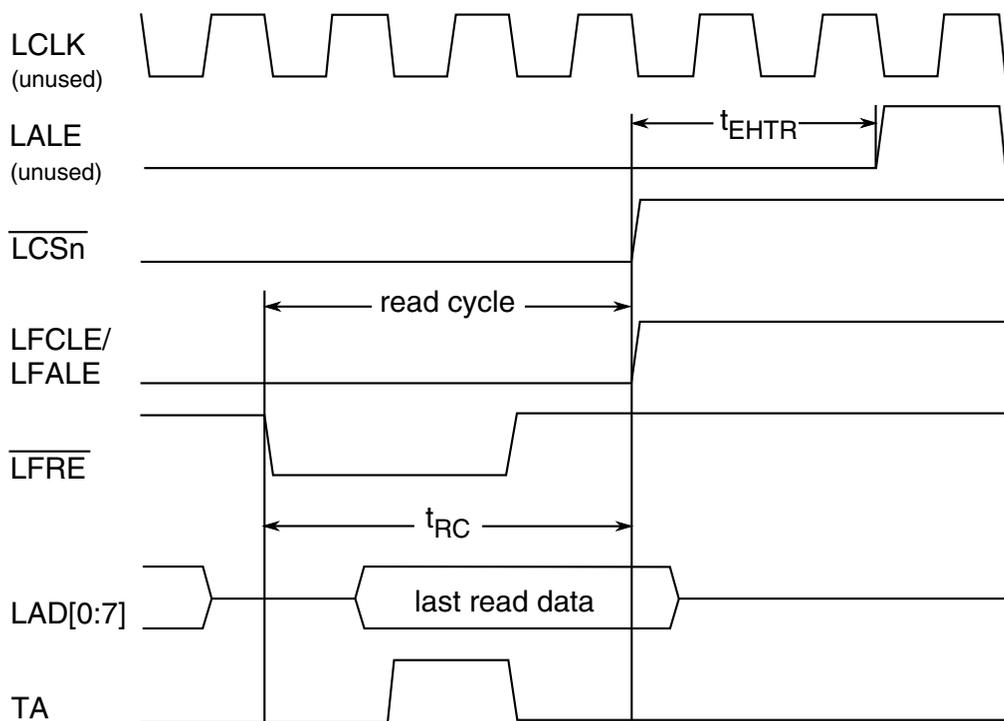
Option Register Attributes		Timing Parameter (LCLK Clock Cycles) ¹				
TRLX	RST	t_{RP}	t_{RHT}	t_{WS}	t_{RC}	t_{WRT}
0	0	$\frac{3}{4} + SCY$	1	SCY	$2 + SCY$	$4 \times (2 + SCY)$
0	1	$1 + SCY$	1	SCY	$2 + SCY$	$4 \times (2 + SCY)$
1	0	$\frac{1}{2} + 2 \times SCY$	2	$2 \times SCY$	$3 + 2 \times SCY$	$8 \times (2 + SCY)$
1	1	$1 + 2 \times SCY$	2	$2 \times SCY$	$3 + 2 \times SCY$	$8 \times (2 + SCY)$

1. In the parameters, SCY refers to a delay of ORn[SCY] clock cycles.

12.4.3.3.5 FCM extended read hold timing

Allowance for slow output driver turn-off when reading NAND Flash EEPROMs is made via setting of ORn[EHTR] and ORn[TRLX].

The extended read data hold time, shown at t_{EHTR} in [Figure 12-99](#) [Figure 12-98](#) and [Figure 12-114](#) [Figure 12-113](#), is a delay inserted by FCM between the last data read and another eLBC memory access (requiring LALE assertion). $LCSn_B$ is negated during t_{EHTR} to allow external devices and bus transceivers time to disable their drivers.



Notes: t_{RC} = Read data cycle time.
 t_{EHTR} = Extended read data hold time.

Figure 12-113. FCM read data timing with extended hold time (for TRLX = 0, EHTR = 1, RST = 1, SCY = 1)

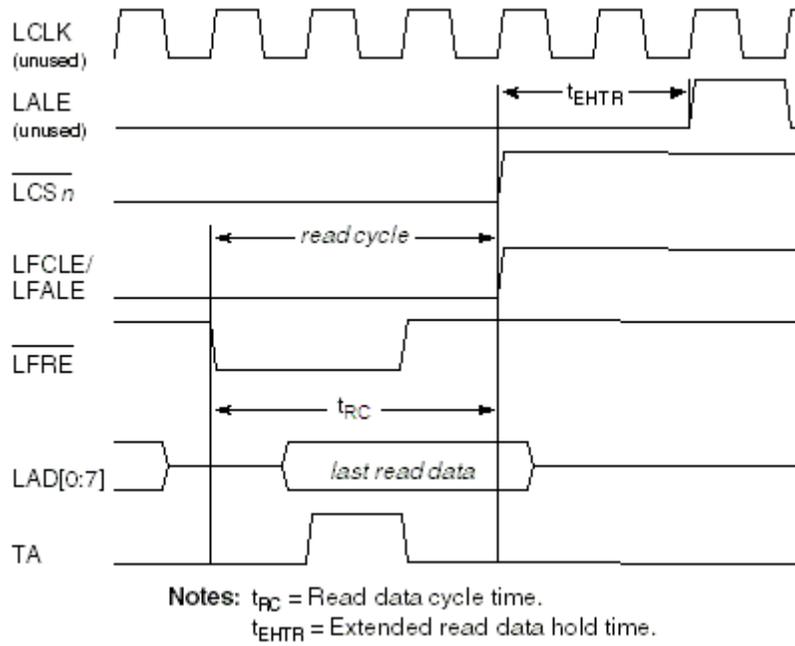


Figure 12-114. FCM read data timing with extended hold time (for $TRLX = 0$, $EHTR = 1$, $RST = 1$, $SCY = 1$)

12.4.3.4 FCM boot chip-select operation

Boot chip-select operation allows address decoding for a boot ROM before system initialization.

LCS0_B is the boot chip-select output; its operation differs from other external chip-select outputs after a system reset. When the core begins accessing memory after system reset, LCS0_B is asserted initially to load a 4-Kbyte boot block into the FCM buffer RAM, but core instruction fetches occur from the buffer RAM.

12.4.3.4.1 FCM bank 0 reset initialization

The boot chip-select also provides a programmable port size, which is configured during reset.

The boot chip-select does not provide write protection. LCS0_B operates this way until the first write to OR0 and it can be used as any other chip-select register after the preferred address range is loaded into BR0. After the first write to OR0, the boot chip-select can be restarted only with a hardware reset.

Table 12-198. Boot bank field values after reset for FCM as boot controller

Register	Field	Setting
BR0	BA	0000_0000_0000_0000_0
	PS	From cfg_rom_loc
	DECC	From cfg_elbc_ecc
	WP	0
	MSEL	001
	V	1
OR0	AM	0000_0000_0000_0000_0
	BCTLD	0
	PGS	From cfg_rom_loc
	CSCT	1
	CST	1
	CHT	1
	RST	1
	SCY	010
	TRLX	1
	EHTR	1

12.4.3.4.2 Boot block loading into the FCM buffer RAM

If FCM is selected as the boot ROM controller from power-on-reset configuration, eLBC will automatically load from bank 0 a single 4 Kbyte page of boot code into the FCM buffer RAM during HRESET_B.

The CPU can execute boot code directly from the FCM buffer RAM, but must ensure that any further data read from the NAND Flash EEPROM is transferred under software control in order to continue the bootstrap process.

Since OR0[AM] is initially cleared during reset, all CPU fetches to eLBC will access the FCM buffer RAM, which appears in the memory map as a 4-Kbyte RAM. No NAND Flash spare regions are mapped during boot, therefore only 4 Kbytes of contiguous, main region data, loaded from the first pages of the boot block, are accessible in eLBC bank 0.

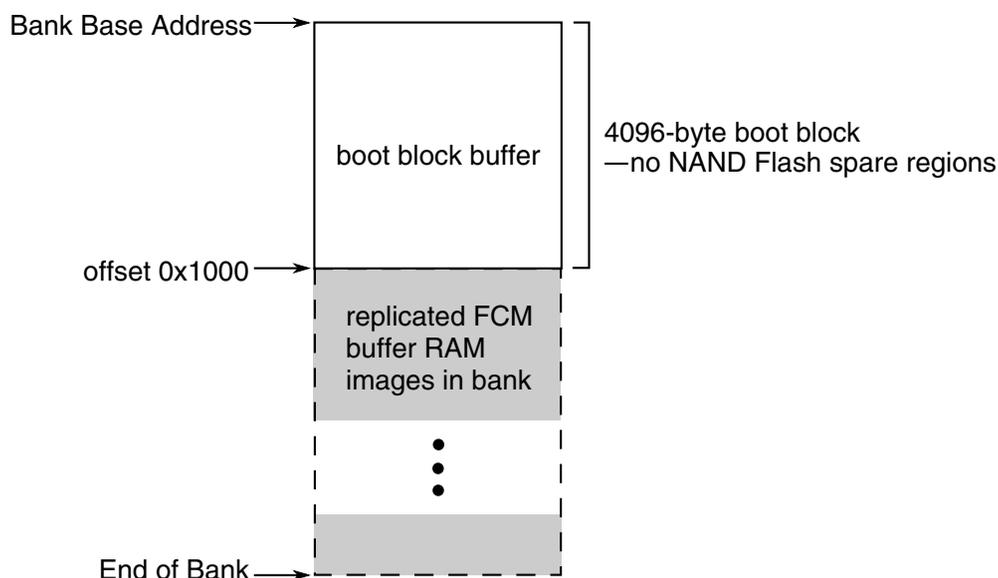


Figure 12-115. FCM buffer RAM memory map during boot loading

The process for booting is as follows:

1. Following negation of HRESET_B , eLBC is released from reset and commences automatic boot block loading if FCM is selected as the boot ROM location. Small-page or large-page, 8-bit NAND Flash devices can be used for boot loading when enabled with LCS0_B. eLBC drives LFWP_B low during boot accesses to prevent accidental erasure of the NAND Flash boot ROM.
2. FCM starts searching for a valid boot block at block index 0.
3. FCM reads the spare regions of the first two pages of the current block, checking the bad block indication (BI) bytes to validate the block for reading. BI bytes must all hold the value 0xFF for the page to be considered readable.
 - For small-page devices, BI is a single byte read from spare region byte offset 5.
 - For large-page devices, BI is a single byte read from spare region byte offset 0.

If either of the first two pages of the current block are marked invalid, then the boot block index is incremented by 1, and FCM repeats step 3. eLBC will continue searching for a bootable block indefinitely, therefore at least one block must be marked valid for boot loading to proceed. At the conclusion of the boot block search, the value of FBAR[BLK] points to the boot block.

4. The FCM optionally performs ECC checking at boot time depending on the configuration selected during reset. If ECC checking is enabled, the FCM recovers from the spare region the stored ECC for each 512-byte block of boot data. The boot block must be prepared with ECC protection. During ECC generation, software should use FMR[ECCM] = 0 for small-page devices, and FMR[ECCM] = 1 for large-page devices.

5. FCM performs a sequence of random-access page reads, reading entire pages from the boot block until 4 Kbytes have been saved to the FCM buffer RAM. If ECC checking is enabled, the ECC of each 512-byte region is verified and single-bit errors are corrected if possible. If FCM is unable to correct ECC errors, eLBC halts the boot process and signals an unrecoverable error by asserting the *hreset_req_B* signal.
6. The CPU now commences fetching instructions, in random order, from the FCM buffer RAM. This first-level boot loader typically copies a secondary boot loader into system memory, and continues booting from there. Boot software must clear FMR[BOOT] to enable normal operation of FCM.

12.4.4 User-programmable machines (UPMs)

UPMs are flexible interfaces that connect to a wide range of memory devices.

At the heart of each UPM is an internal RAM array that specifies the logical value driven on the external memory control signals (*LCS_n_B*, *LBS_B*[0:1] and *LGPL*[0:5]) for a given clock cycle. Each word in the RAM array provides bits that allow a memory access to be controlled with a resolution of up to one quarter of the external bus clock period on the byte-select and chip-select lines. A gap of 2 dead LCLK cycles is present on the UPM interface between UPM transactions.

NOTE

If the *LGPL4/ LGTA_B/ LFRB_B/ LUPWAIT/ LPBSE* signal is used as both an input and an output, a weak pull-up is required. See the device hardware specifications for details regarding termination options.

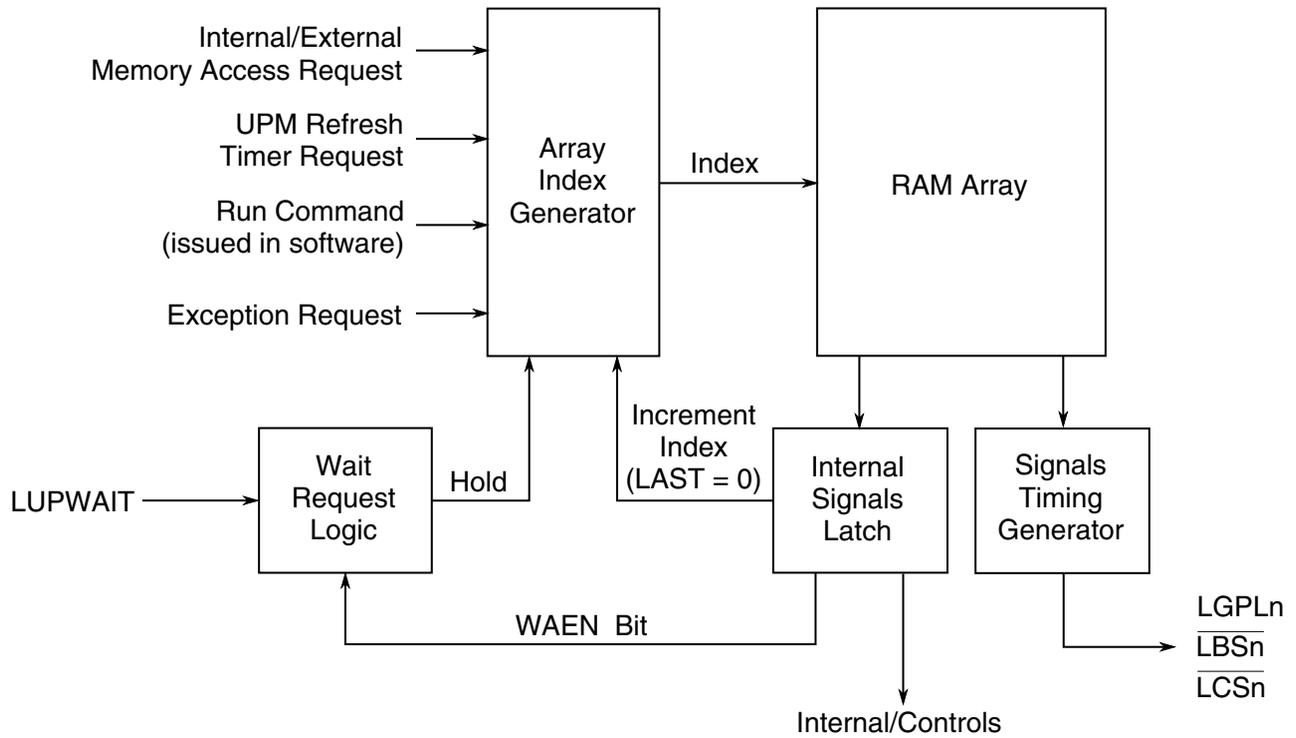


Figure 12-116. User-programmable machine functional block diagram

The following events initiate a UPM cycle:

- Any internal device requests an external memory access to an address space mapped to a chip-select serviced by the UPM
- A UPM refresh timer expires and requests a transaction, such as a DRAM refresh
- A bus monitor time-out error during a normal UPM cycle redirects the UPM to execute an exception sequence

The RAM array contains 64 words of 32-bits each. The signal timing generator loads the RAM word from the RAM array to drive the general-purpose lines, byte-selects, and chip-selects. If the UPM reads a RAM word with WAEN set, the external LUPWAIT signal is sampled and synchronized by the memory controller and the current request is frozen.

12.4.4.1 UPM requests

A special pattern location in the RAM array is associated with each of the possible UPM requests.

An internal device's request for a memory access initiates one of the following patterns (MxMR[OP] = 00):

- Read single-beat pattern (RSS)

- Read burst cycle pattern (RBS)
- Write single-beat pattern (WSS)
- Write burst cycle pattern (WBS)

A UPM refresh timer request pattern initiates a refresh timer pattern (RTS).

An exception (caused by a bus monitor time-out error) occurring while another UPM pattern is running initiates an exception condition pattern (EXS).

The figure and table below show the start addresses of these patterns in the UPM RAM, according to cycle type. RUN commands (MxMR[OP] = 11), however, can initiate patterns starting at any of the 64 UPM RAM words.

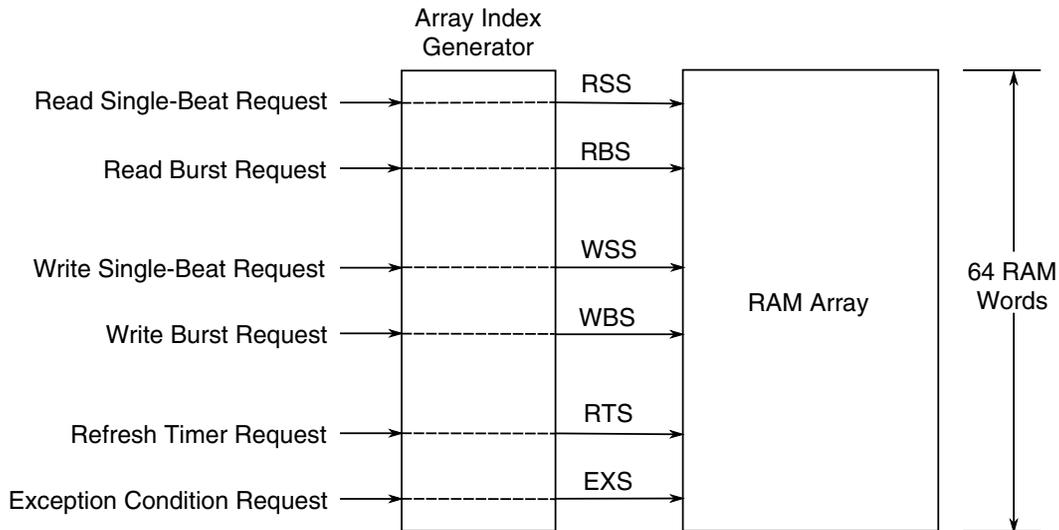


Figure 12-117. RAM array indexing

Table 12-199. UPM routines start addresses

UPM Routine	Routine Start Address
Read single-beat (RSS)	0x00
Read burst (RBS)	0x08
Write single-beat (WSS)	0x18
Write burst (WBS)	0x20
Refresh timer (RTS)	0x30
Exception condition (EXS)	0x3C

12.4.4.1.1 Memory access requests

The user must ensure that the UPM is appropriately initialized before a request occurs.

The UPM supports two types of memory reads and writes:

- A single-beat transfer transfers one operand consisting of up to a single word (dependent on port size). A single-beat cycle starts with one transfer start and ends with one transfer acknowledge.
- A burst transfer transfers exactly 4 double words regardless of port size. For 32-bit accesses, the burst cycle starts with one transfer start but ends after eight transfer acknowledges, whereas an 8-bit device requires 32 transfer acknowledges.

The user must ensure that patterns for single-beat transfers contain one, and only one, transfer acknowledge (UTA bit in RAM word set high) and for a burst transfer, contain the exact number of transfer acknowledges required.

Any transfers that do not naturally fit single or burst transfers are synthesized as a series of single transfers. These accesses are treated by the UPM as back-to-back, single-beat transfers. Burst transfers can also be inhibited by setting $OR_n[BI]$. Burst performance can be achieved by ensuring that UPM transactions are 32-byte aligned with a transaction size being some multiple of 32-bytes, which is a natural fit for cache-line transfers, for example.

12.4.4.1.2 UPM refresh timer requests

Each UPM contains a refresh timer that can be programmed to generate refresh service requests of a particular pattern in the RAM array.

The figure below shows the clock division hardware associated with memory refresh timer request generation. The UPM refresh timer register (LURT) defines the period for the timers associated with all three UPMs.

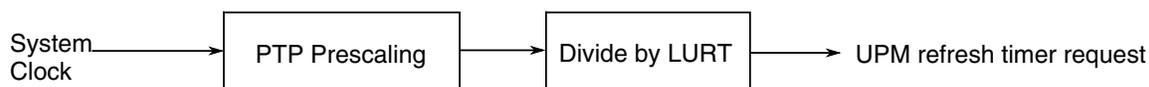


Figure 12-118. Memory refresh timer request block diagram

By default, all local bus refreshes are performed using the refresh pattern of UPMA. This means that if refresh is required, $MAMR[RFEN]$ must be set. It also means that only one refresh routine should be programmed and be placed in UPMA, which serves as the refresh executor. Any banks assigned to a UPM are provided with the common UPMA refresh pattern if the $RFEN$ bit of the corresponding UPM is set, concurrently. UPMA assigned banks, therefore, always receive refresh services when $MAMR[RFEN]$ is set, while UPMB and UPMC assigned banks also receive (the same) refresh services if the corresponding $MxMR[RFEN]$ bits are set. In this scenario, more than one chip select may assert at the same time, as refresh pattern runs for all banks assigned to UPM with $RFEN$ bit set.

12.4.4.1.3 Software requests-RUN command

Software can start a request to the UPM by issuing a RUN command to the UPM.

Some memory devices have their own signal handshaking protocol to put them into special modes, such as self-refresh mode.

For these special cycles, the user creates a special RAM pattern that can be stored in any unused areas in the UPM RAM. Then a RUN command is used to run the cycle. The UPM runs the pattern beginning at the specified RAM location until it encounters a RAM word with its LAST bit set. The RUN command is issued by setting $MxMR[OP] = 11$ and accessing UPM_n memory region with any write transaction that hits the corresponding UPM machine. $MxMR[MAD]$ determines the starting address in the RAM array for the pattern.

Note that transfer acknowledges (UTA bit in the RAM word) are ignored for software (RUN command) requests, and hence the LAD signals remain high-impedance unless the normal initial LALE occurs or the RUN pattern causes assertion of LALE to occur on changes to the RAM word AMX field.

12.4.4.1.4 Exception requests

When the eLBC under UPM control initiates an access to a memory device and an exception occurs (bus monitor time-out), the UPM provides a mechanism by which memory control signals can meet the device's timing requirements without losing data.

The mechanism is the exception pattern that defines how the UPM negates its signals in a controlled manner.

12.4.4.2 Programming the UPMs

The UPM is a micro sequencer that requires microinstructions or RAM words to generate signal timings for different memory cycles.

Follow these steps to program the UPMs:

1. Set up BR_n and OR_n registers.
2. Write patterns into the RAM array.
3. Program MRTPR, LURT and $MAMR[RFEN]$ if refresh is required.
4. Program $MxMR$.

Patterns are written to the RAM array by setting $MxMR[OP] = 01$ and accessing the UPM with any write transaction that hits the relevant chip select. The entire array is thus programmed by an alternating series of writes: to MDR (RAM word to be written) each

time followed by a read from MDR and then followed by a (dummy) write transaction to the relevant UPM assigned bank. A read from MDR is required to ensure that the MDR update has occurred prior to the (dummy) write transaction.

RAM array contents may also be read for debug purposes, for example, by alternating dummy read transactions, each time followed by reads of MDR (when $MxMR[OP] = 10$).

NOTE

$MxMR$ /MDR registers should not be updated while dummy read/write access is still in progress. If the $MxMR[MAD]$ is incremented then the previous dummy transaction is already completed.

In order to enforce proper ordering between updates to the $MxMR$ /MDR register and the dummy accesses to the UPM memory region, two rules must be followed:

- Since the result of any update to the $MxMR$ /MDR register must be in effect before the dummy read or write to the UPM region, a write to $MxMR$ /MDR should be followed immediately by a read of $MxMR$ /MDR.
- The UPM memory region should have the same MMU settings as the memory region containing the $MxMR$ configuration register; both should be mapped by the MMU as cache-inhibited and guarded. This prevents the CPU from re-ordering a read of the UPM memory around the read of $MxMR$. Once the programming of the UPM array is complete the MMU setting for the associated address range can be set to the proper mode for normal operation, such as cacheable and copyback.

For proper signalling, the following guidelines must be followed while programming UPM RAM words:

- For UPM reads, program UTA and LAST in the same or consecutive RAM words.
- For UPM burst reads, program last UTA and LAST in the same or consecutive RAM words.
- For UPM writes, program UTA and LAST in the same RAM word.
- For UPM burst writes, program last UTA and LAST in the same RAM word.

12.4.4.2.1 UPM programming example (two sequential writes to the RAM array)

The following example further illustrates the steps required to perform two writes to the RAM array at non-sequential addresses assuming that the relevant BR_n and OR_n registers have been previously set up:

1. Program MxMR for the first write (with the desired RAM array address).
2. Write pattern/data to MDR to ensure that the MxMR has already been updated with the desired configuration.
3. Read MDR to ensure that the MDR has already been updated with the desired pattern. (Or, read MxMR register if step 2 is not performed.)
4. Perform a dummy write transaction.
5. Read/check MxMR[MAD]. If incremented, the previous dummy write transaction is completed; proceed to step 6. Repeat step 5 until incremented.
6. Program MxMR for the second write with the desired RAM array address.
7. Write pattern/data to MDR to ensure that the MxMR has already been updated with the desired configuration.
8. Read MDR to ensure that the MDR has already been updated with the desired pattern.
9. Perform a dummy write transaction.
10. Read/check MxMR[MAD]. If incremented, the previous dummy write transaction is completed.

Note that if steps 1 and 2 or steps 6 and 7 are reversed, step 3 or 8 (as appropriate) is replaced by the following:

- Read MxMR to ensure that the MxMR has already been updated with the desired configuration.

12.4.4.2.2 UPM programming example (two sequential reads from the RAM array)

RAM array contents may also be read for debug purposes, for example, by alternating dummy read transactions, each time followed by reads of MDR ($MxMR[OP] = 0b10$).

The following example further illustrates the steps required to perform two reads from the RAM array at non-sequential addresses assuming that the relevant BR_n and OR_n registers have been previously set up:

1. Program MxMR for the first read with the desired RAM array address.
2. Read MxMR to ensure that the MxMR has already been updated with the desired configuration, such as RAM array address.
3. Perform a dummy read transaction.

4. Read/check MxMR[MAD]. If incremented, the previous dummy read transaction is completed; proceed to step 5. Repeat step 4 until incremented.
5. Read MDR.
6. Program MxMR for the second read with the desired RAM array address.
7. Read MxMR to ensure that the MxMR has already been updated with the desired configuration, such as RAM array address.
8. Perform a dummy read transaction.
9. Read/check MxMR[MAD]. If incremented, the previous dummy read transaction is completed; proceed to step 10. Repeat step 9 until incremented.
10. Read MDR.

12.4.4.3 UPM signal timing

RAM word fields specify the value of the various external signals at a granularity of up to four values for each bus clock cycle.

The signal timing generator causes external signals to behave according to timing specified in the current RAM word. Each bit in the RAM word relating to LCS_n_B and LBS_B timing specifies the value of the corresponding external signal at each quarter phase of the bus clock.

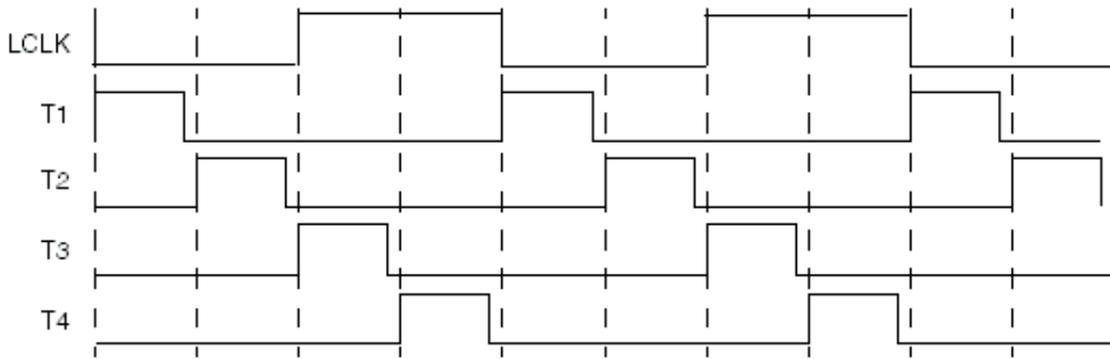


Figure 12-119. UPM clock scheme for LCRR[CLKDIV] = 4 and higher order ratios (PLL bypass)

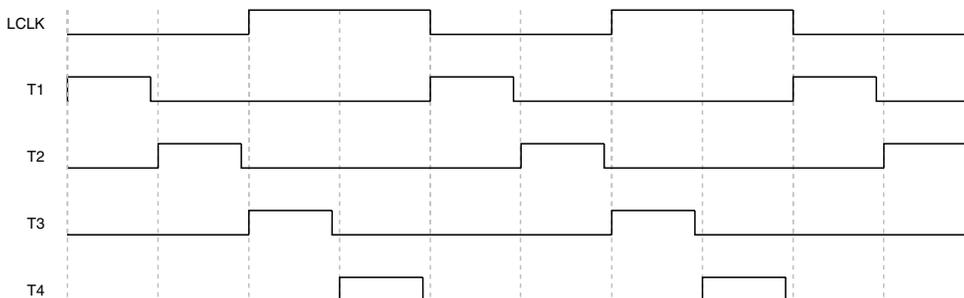


Figure 12-120. UPM clock scheme for LCRR[CLKDIV] = 4 and higher order ratios (PLL enable)

12.4.4.4 RAM array

The RAM array for each UPM is 64 locations deep and 32 bits wide.

The signals at the bottom of the figure are UPM outputs. The selected LCS_n_B is for the bank that matches the current address. The selected LBS_B is for the byte lanes read or written by the access.

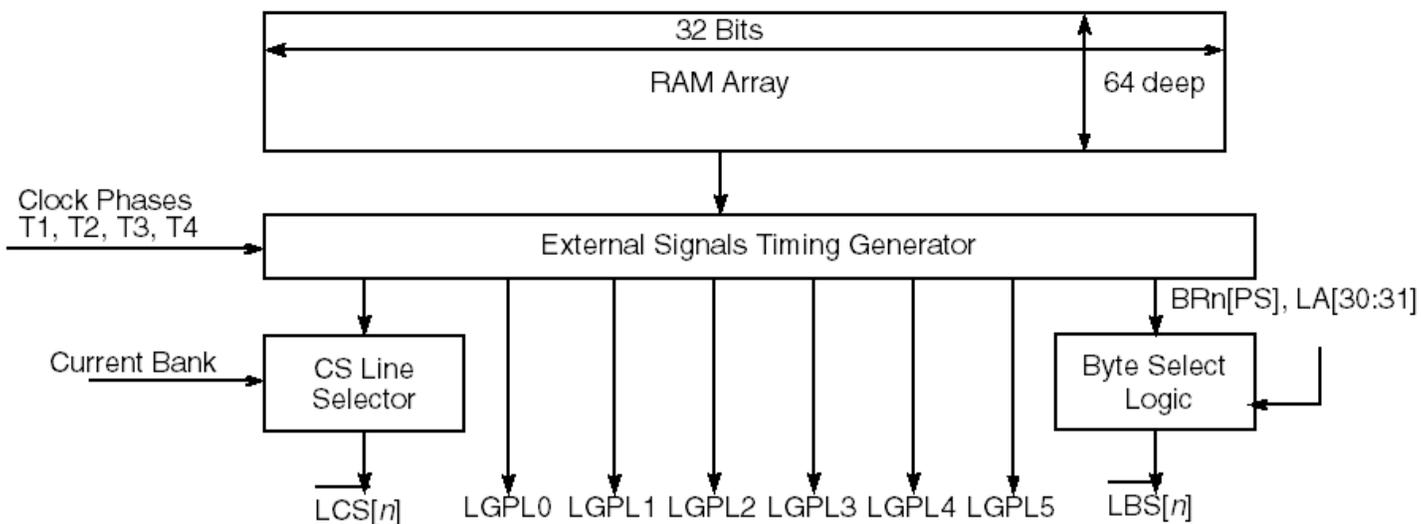


Figure 12-121. RAM array and signal generation

12.4.4.4.1 RAM words

The RAM word is a 32-bit microinstruction stored in one of 64 locations in the RAM array.

It specifies timing for external signals controlled by the UPM. The CST_n and BST_n bits determine the state of UPM signals LCS_n_B and $LBS_B[0:1]$ at each quarter phase of the bus clock.

Table 12-200. RAM word fields

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	CST1	CST2	CST3	CST4	BST1	BST2	BST3	BST4	G0L		G0H		G1T1	G1T3	G2T1	G2T3
W																
Resett	All zeros															
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	G3T1	G3T3	G4T1/ DLT3	G4T3/ WAEN	G5T1	G5T3	REDO		LOO P	EXE N	AMX		NA	UTA	TODT	LAS T
W																
Resett	All zeros															

Table 12-201. RAM word field descriptions

Bits	Name	Description
0	CST1	Chip select timing 1. Defines the state (0 or 1) of LCS_n_B during bus clock quarter phase 1.
1	CST2	Chip select timing 2. Defines the state (0 or 1) of LCS_n_B during bus clock quarter phase 2.
2	CST3	Chip select timing 3. Defines the state (0 or 1) of LCS_n_B during bus clock quarter phase 3.
3	CST4	Chip select timing 4. Defines the state (0 or 1) of LCS_n_B during bus clock quarter phase 4.
4	BST1	Byte select timing 1. Defines the state (0 or 1) of LBS_B during bus clock quarter phase 1.
5	BST2	Byte select timing 2. Defines the state (0 or 1) of LBS_B during bus clock quarter phase 2.
6	BST3	Byte select timing 3. Defines the state (0 or 1) of LBS_B during bus clock quarter phase 3.
7	BST4	Byte select timing 4. Defines the state (0 or 1) of LBS_B during bus clock quarter phase 4.
8-9	G0L	General purpose line 0 lower. Defines the state of $LGPL0$ during the bus clock quarter phases 1 and 2 (first half phase). 00 Value defined by $MxMR[G0CL]$ 01 Reserved 10 0 11 1

Table continues on the next page...

Table 12-201. RAM word field descriptions (continued)

Bits	Name	Description
10-11	G0H	General purpose line 0 higher. Defines the state of LGPL0 during the bus clock quarter phases 3 and 4 (second half phase). 00 Value defined by MxMR[GOCL] 01 Reserved 10 0 11 1
12	G1T1	General purpose line 1 timing 1. Defines the state (0 or 1) of LGPL1 during bus clock quarter phases 1 and 2 (first half phase).
13	G1T3	General purpose line 1 timing 3. Defines the state (0 or 1) of LGPL1 during bus clock quarter phases 3 and 4 (second half phase)
14	G2T1	General purpose line 2 timing 1. Defines state (0 or 1) of LGPL2 during bus clock quarter phases 1 and 2 (first half phase).
15	G2T3	General purpose line 2 timing 3. Defines the state (0 or 1) of LGPL2 during bus clock quarter phases 3 and 4 (second half phase).
16	G3T1	General purpose line 3 timing 1. Defines the state (0 or 1) of LGPL3 during bus clock quarter phases 1 and 2 (first half phase).
17	G3T3	General purpose line 3 timing 3. Defines the state (0 or 1) of LGPL3 during bus clock quarter phases 3 and 4 (second half phase).
18	G4T1/DLT3	General purpose line 4 timing 1/delay time 3. The function of this bit is determined by MxMR[GPL4]. If MxMR[GPL4] = 0 and LGPL4/LUPWAIT pin functions as an output (LGPL4), G4T1/DLT3 defines the state (0 or 1) of LGPL4 during bus clock quarter phases 1 and 2 (first half phase). If MxMR[GPL4] = 1 and LGPL4/LUPWAIT functions as an input (LUPWAIT), if a read burst or single read is executed, G4T1/DLT3 defines the sampling of the data bus as follows: 0 In the current word, the data bus should be sampled at the start of bus clock quarter phase 1 of the next bus clock cycle. 1 In the current word, the data bus should be sampled at the start of bus clock quarter phase 3 of the current bus clock cycle.
19	G4T3/WAEN	General purpose line 4 timing 3/wait enable. Bit function is determined by MxMR[GPL4]. If MxMR[GPL4] = 0 and LGPL4/LUPWAIT pin functions as an output (LGPL4), G4T3/WAEN defines the state (0 or 1) of LGPL4 during bus clock quarter phases 3 and 4 (second half phase). If MxMR[GPL4] = 1 and LGPL4/LUPWAIT functions as an input (LUPWAIT), G4T3/WAEN is used to enable the wait mechanism: 0 LUPWAIT detection is disabled. 1 LUPWAIT is enabled. If LUPWAIT is detected as being asserted, a freeze in the external signals logical values occurs until LUPWAIT is detected as being negated.
20	G5T1	General purpose line 5 timing 1. Defines the state (0 or 1) of LGPL5 during bus clock quarter phases 1 and 2 (first half phase).
21	G5T3	General purpose line 5 timing 3. Defines the state (0 or 1) of LGPL5 during bus clock quarter phases 3 and 4 (second half phase).

Table continues on the next page...

Table 12-201. RAM word field descriptions (continued)

Bits	Name	Description
22-23	REDO	<p>Redo current RAM word. Defines the number of times to execute the current RAM word.</p> <p>00 Once (normal operation)</p> <p>01 Twice</p> <p>10 Three times</p> <p>11 Four times</p>
24	LOOP	<p>Loop start/end. The first RAM word in the RAM array where LOOP is 1 is recognized as the loop start word. The next RAM word where LOOP is 1 is the loop end word. RAM words between, and including the start and end words, are defined as part of the loop. The number of times the UPM executes this loop is defined in the corresponding loop fields of the MxMR.</p> <p>0 The current RAM word is not the loop start word or loop end word.</p> <p>1 The current RAM word is the start or end of a loop.</p> <p>NOTE: AMX must not change values in any RAM word which begins a loop</p>
25	EXEN	<p>Exception enable. Allows branching to an exception pattern at the exception start address (EXS). When an internal bus monitor time-out exception is recognized and EXEN in the RAM word is set, the UPM branches to the special exception start address (EXS) and begins operating as the pattern defined there specifies.</p> <p>The user should provide an exception pattern to negate signals controlled by the UPM in a controlled fashion. For DRAM control, a handler should negate RAS and CAS to prevent data corruption. If EXEN = 0, exceptions are ignored by UPM (but not by local bus) and execution continues. After the UPM branches to the exception start address, it continues reading until the LAST bit is set in the RAM word.</p> <p>0 The UPM continues executing the remaining RAM words, ignoring any internal bus monitor time-out.</p> <p>1 The current RAM word allows a branch to the exception pattern after the current cycle if an exception condition is detected.</p>
26-27	AMX	<p>Address multiplexing. Determines the source of LAD during an LALE phase. Any change in the AMX field initiates a new LALE (address) phase.</p> <p>00 LAD (and/or in conjunction with LA) is the non-multiplexed address. For example, column address.</p> <p>01 Reserved</p> <p>10 LAD (and/or in conjunction with LA) is driven with the multiplexed address according to MxMR[AM]. For example, row address. See Address multiplexing (AMX) for more information.</p> <p>11 LAD (and/or in conjunction with LA) is driven with the contents of MAR. Used, for example, to initialize a mode.</p> <p>NOTE: Source ID debug mode is only supported for the AMX = 00 setting.</p> <p>NOTE: AMX must not change values in any RAM word which begins a loop.</p>

Table continues on the next page...

Table 12-201. RAM word field descriptions (continued)

Bits	Name	Description
28	NA	<p>Next burst address. Determines when the address is incremented during a burst access.</p> <p>0 The address increment function is disabled.</p> <p>1 The address is incremented in the next cycle. In conjunction with the $BRn[PS]$, the increment value of LAN is 1 or 2 for port sizes of 8 and 16 bits, respectively.</p>
29	UTA	<p>UPM transfer acknowledge. Indicates assertion of transfer acknowledge in the current cycle.</p> <p>0 Transfer acknowledge is not asserted in the current cycle.</p> <p>1 Transfer acknowledge is asserted in the current cycle.</p> <p>In case of UPM writes, program UTA and LAST in same RAM word.</p> <p>In case of UPM reads, program UTA and LAST in consecutive or same RAM words.</p>
30	TODT	<p>Turn-on disable timer. The disable timer associated with each UPM allows a minimum time to be guaranteed between two successive accesses to the same memory bank. This feature is critical when DRAM requires a RAS precharge time. TODT turns the timer on to prevent another UPM access to the same bank until the timer expires. The disable timer period is determined in $MxMR[DSn]$. The disable timer does not affect memory accesses to different banks. Note that TODT must be set together with LAST, otherwise it is ignored.</p> <p>0 The disable timer is turned off.</p> <p>1 The disable timer for the current bank is activated preventing a new access to the same bank (when controlled by the UPMs) until the disable timer expires. For example, precharge time.</p>
31	LAST	<p>Last word. When LAST is read in a RAM word, the current UPM pattern terminates and control signal timing set in the RAM word is applied to the current (and last) cycle. However, if the disable timer is activated and the next access is to the same bank, execution of the next UPM pattern is held off and the control signal values specified in the last word are extended in duration for the number of clock cycles specified in $MxMR[DSn]$.</p> <p>Note that UPM continue to execute RAM words until it finds LAST, irrespective of the assigned region for various patterns like RSS.</p> <p>0 The UPM continues executing RAM words.</p> <p>1 Indicates the last RAM word in the program. The service to the UPM request is done after this cycle concludes.</p> <p>In case of UPM writes, program UTA and LAST in same RAM word.</p> <p>In case of UPM reads, program UTA and LAST in consecutive or same RAM words.</p>

12.4.4.4.2 Chip-select signal timing (CST_n)

If BR_n[MSEL] of the accessed bank selects a UPM on the currently requested cycle, the UPM manipulates the LCS_B _n for that bank with timing as specified in the UPM RAM word CST_n fields.

The selected UPM affects only the assertion and negation of the appropriate LCS_n_B signal. The state of the selected LCS_n_B signal of the corresponding bank depends on the value of each CST_n bit.

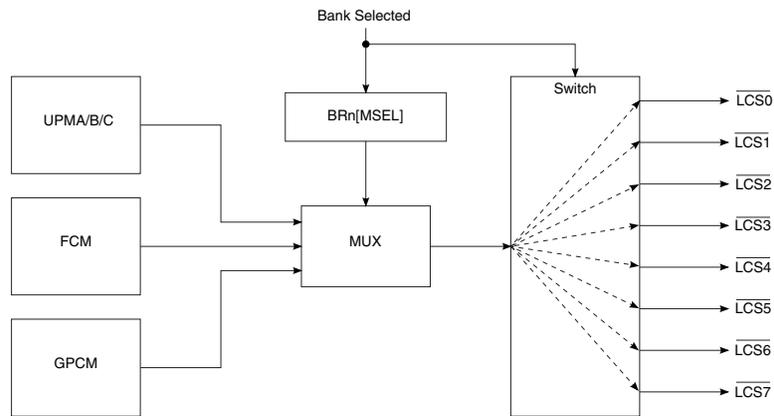


Figure 12-122. LCS_n_B Signal selection

12.4.4.4.3 Byte select signal timing (BST_n)

If BR_n[MSEL] of the accessed memory bank selects a UPM on the currently requested cycle, the selected UPM affects the assertion and negation of the appropriate LBS_B[0:1] signal.

The timing of both byte-select signals is specified in the RAM word. However, LBS_B[0:1] are also controlled by the port size of the accessed bank, the number of bytes to transfer, and the address accessed.

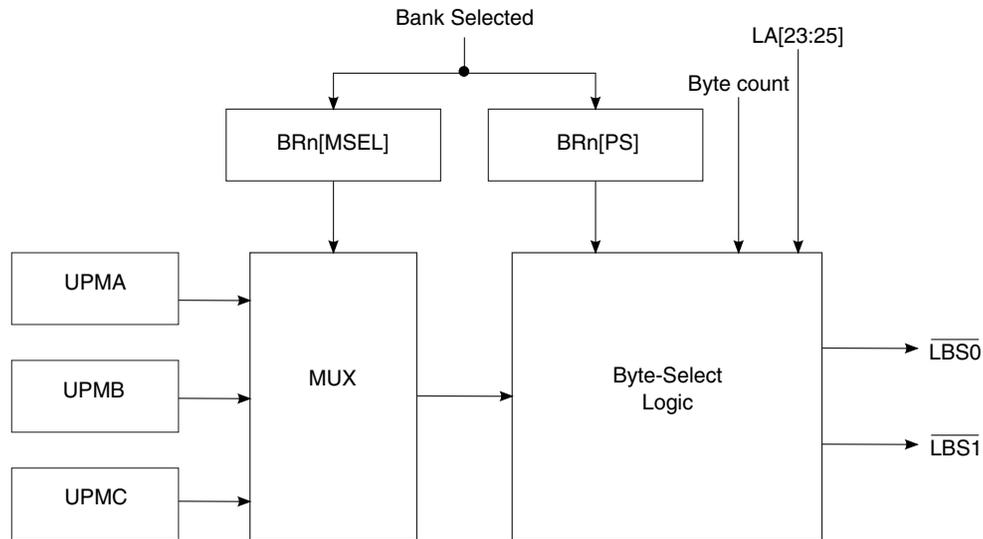


Figure 12-123. LBS_B Signal selection

The uppermost byte select ($LBS0_B$), when asserted, indicates that $LAD[0:7]$ contains valid data during a cycle. Likewise, $LBS1_B$ indicates that $LAD[8:15]$ contain valid data. For a UPM refresh timer request, all $LBS_B[0:1]$ signals are asserted/negated by the UPM according to the refresh pattern only. Following any internal bus monitor exception, the $LBS_B[0:1]$ signals are negated regardless of the exception handling provided by any UPM exception pattern to prevent spurious writes to external RAM.

12.4.4.4.4 General-purpose signals ($GnTn$, GOn)

The general-purpose signals ($LGPL[0:5]$) each have two bits in the RAM word that define the logical value of the signal to be changed at the rising edge of the bus clock and/or at the falling edge of the bus clock.

$LGPL0$ offers enhancements beyond the other $LGPLn$ lines.

$LGPL0$ can be controlled by an address line specified in $MxMR[G0CL]$. To use this feature, $G0H$ and $G0L$ should be set in the RAM word. For example, for a SIMM with multiple banks, this address line can be used to switch between internal memory device banks.

12.4.4.4.5 Loop control (LOOP)

The LOOP bit in the RAM word specifies the beginning and end of a set of UPM RAM words that are to be repeated.

The first time LOOP = 1, the memory controller recognizes it as a loop start word and loads the memory loop counter with the corresponding contents of the loop field shown in the table below. The next RAM word for which LOOP = 1 is recognized as a loop end word. When it is reached, the loop counter is decremented by one.

Continued loop execution depends on the loop counter. If the counter is not zero, the next RAM word executed is the loop start word. Otherwise, the next RAM word executed is the one after the loop end word. Loops can be executed sequentially but cannot be nested. Also, special care must be taken:

- LAST and LOOP must not be set together.
- Loop start word should not have an AMX change with regard to the previous word.

Table 12-202. MxMR loop field use

Request Serviced	Loop Field
Read single-beat cycle	RLF
Read burst cycle	RLF
Write single-beat cycle	WLF
Write burst cycle	WLF
Refresh timer expired	TLF
RUN command	RLF

12.4.4.4.6 Repeat execution of current RAM word (REDO)

The REDO function is useful for wait-state insertion in a long UPM routine that would otherwise need too many RAM words.

Setting the REDO bits of the RAM word to a nonzero value causes the UPM to re-execute the current RAM word up to three more times, as defined in the REDO field of the current RAM word.

Special care must be taken in the following cases:

- When UTA and REDO are set together, TA is asserted the number of times specified by the REDO function.
- When NA and REDO are set together, the address is incremented the number of times specified by the REDO function.
- When LOOP and REDO are set together, the loop mechanism works as usual and the line is repeated according to the REDO function.
- LAST and REDO must not be set together.
- REDO should not be used within the exception routine.

Table 12-203. UPM address multiplexing (continued)

	msb		Internal Transaction Address																												lsb			
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
AMX = 00 (Col)																					2	2	2	2	2	2	2	2	2	2	2	3	3	
AMX = 10 MxMR[AM] = 100 (Row)				1	1	1	1	2	2	2	2	2	2	2	2	2	2	2	3	3														
AMX = 00 (Col)																					2	2	2	2	2	2	2	2	2	2	3	3		
AMX = 10 MxMR[AM] = 101 (Row)			1	1	1	1	2	2	2	2	2	2	2	2	2	2	2	3	3															
AMX = 00 (Col)																					1	2	2	2	2	2	2	2	2	2	3	3		
AMX = 10 MxMR[AM] = 110	Reserved																																	
AMX = 10 MxMR[AM] = 111	Reserved																																	

Note that any change to the AMX field from one RAM word to the next RAM word executed results in an address phase on the {LADn, LAN} bus with the assertion of LALE for the number of cycles set for LALE in the ORn and LCRR registers. The LGPL[0:5] signals maintain the value specified in the RAM word during the LALE phase.

NOTE

AMX must not change values in any RAM word which begins a loop.

12.4.4.4.8 Data valid and data sample control (UTA)

When a read access is handled by the UPM, and the UTA bit is 1 (data is to be sampled by the eLBC), the value of the DLT3 bit in the same RAM word, in conjunction with MxMR[GPL4], determines when the data input is sampled by the eLBC as follows:

- If MxMR[GPL4] = 1 (G4T4/DLT3 functions as DLT3) and DLT3 = 1 in the RAM word, data is latched on the falling edge of the bus clock instead of the rising edge. The eLBC samples the data on the next falling edge of the bus clock, which is during

the middle of the current bus cycle. This feature should be used only in systems without external synchronous bus devices that require mid-cycle sampling.

- If $MxMR[GPL4] = 0$ (G4T4/DLT3 functions as G4T4), or if $MxMR[GPL4] = 1$ but $DLT3 = 0$ in the RAM word, data is latched on the rising edge of the bus clock, which occurs at the end of the current bus clock cycle (normal operation).

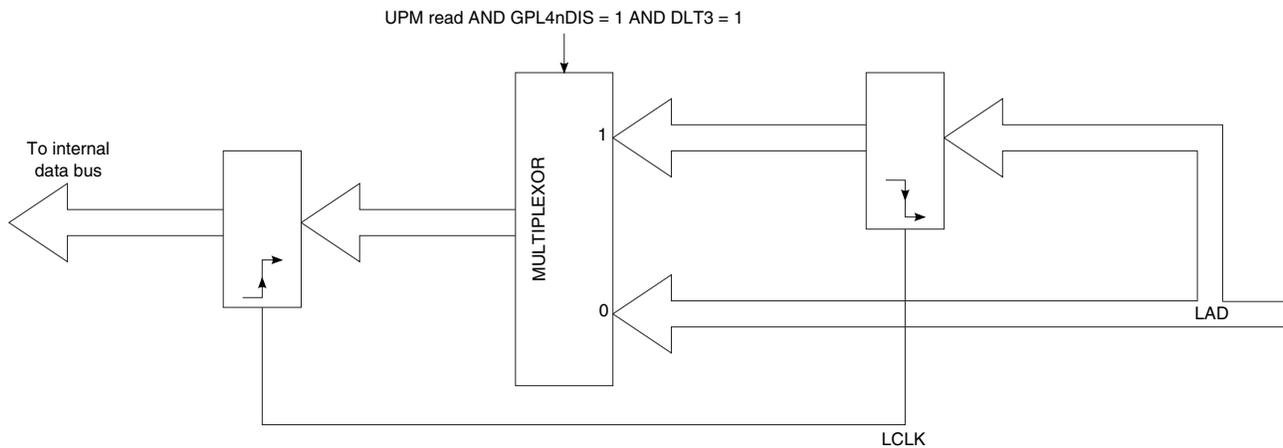


Figure 12-124. UPM read access data sampling

12.4.4.4.9 LGPL[0:5] signal negation (LAST)

When the LAST bit is read in a RAM word, the current UPM pattern is terminated at the end of the current cycle.

On the next cycle (following LAST) all the UPM signals are negated unconditionally (driven to logic 1), unless there is a back-to-back UPM request pending. In this case, the signal values for the cycle following the one in which the LAST bit was set are taken from the first RAM word of the pending UPM routine.

In case of UPM writes, program UTA and LAST in same RAM word. In case of UPM reads, program UTA and LAST in consecutive or same RAM words.

12.4.4.4.10 Wait mechanism (WAEN)

The WAEN bit in the RAM array word can be used to enable the UPM wait mechanism in selected UPM RAM words.

If the UPM reads a RAM word with WAEN set, the external LUPWAIT signal is sampled and synchronized by the memory controller as if it were an asynchronous signal. The WAEN bit is ignored if $LAST = 1$ in the same RAM word.

Synchronization of LUPWAIT starts at the rising/falling edge of the LCLK and takes at least 1 LCLK cycle to get affected. If LUPWAIT is asserted and WAEN = 1 in the current UPM word, the UPM is frozen until LUPWAIT is negated. The value of external signals driven by the UPM remains as indicated in the previous RAM word. When LUPWAIT is negated, the UPM continues normal functions. Note that during WAIT cycles, the UPM does not handle data.

The figure below shows how the WAEN bit in the word read by the UPM and the LUPWAIT signal are used to hold the UPM in a particular state until LUPWAIT is negated. As the example shows, the LCS_n_B and $LGPL1$ states and the WAEN value are frozen until LUPWAIT is recognized as negated. WAEN is typically set before the line that contains $UTA = 1$. Note that if WAEN and NA are both set in the same RAM word, NA causes the burst address to increment once as normal regardless of whether the UPM freezes.

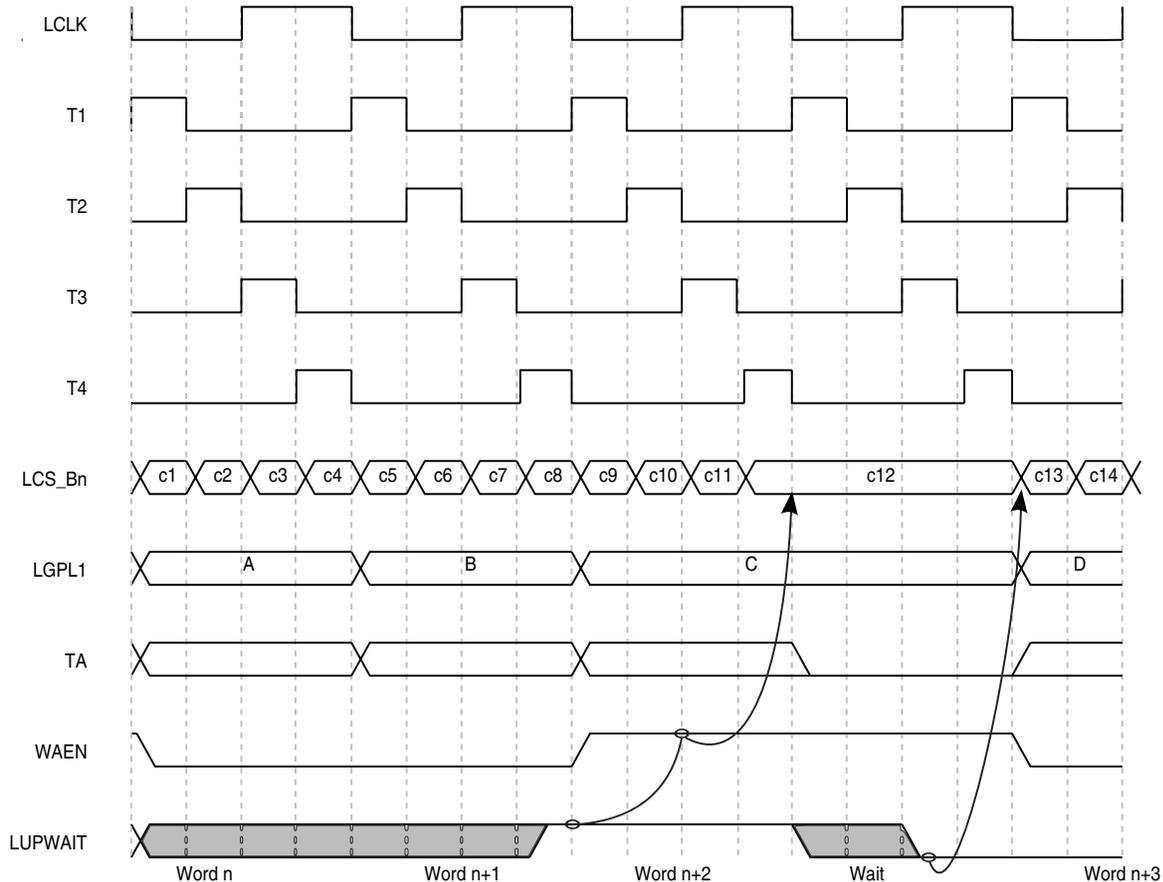


Figure 12-125. Effect of LUPWAIT signal with PLL bypass

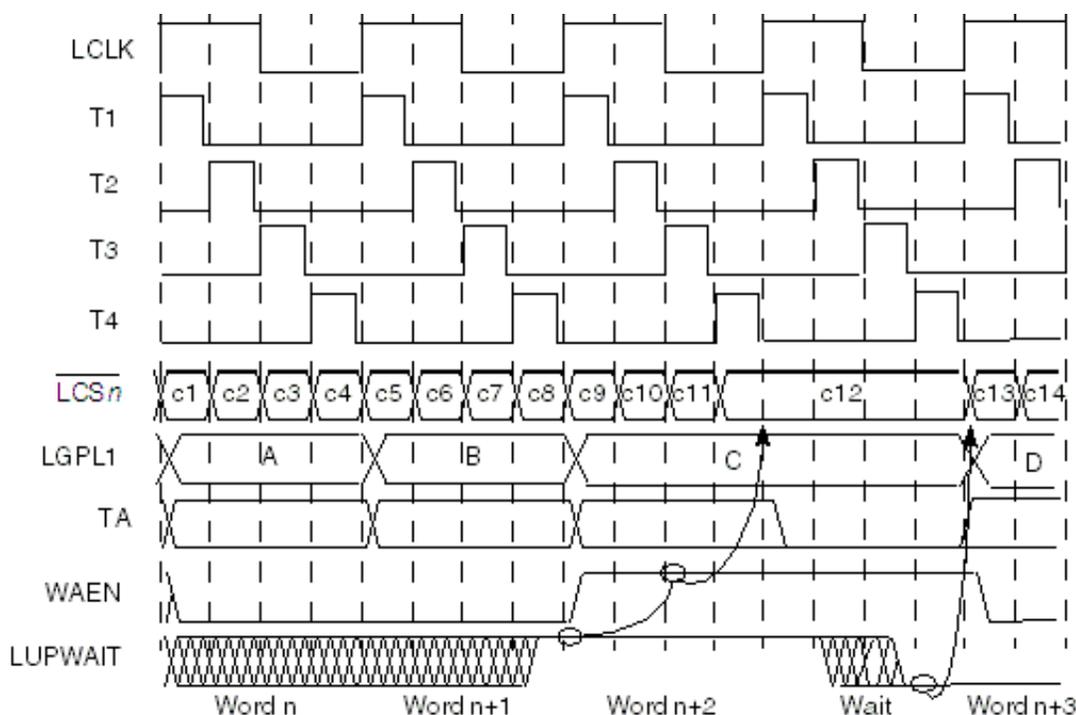


Figure 12-126. Effect of LUPWAIT signal with PLL enabled

If LUPWAIT is to be considered an asynchronous signal, which can be asserted/negated at any time, no UPM RAM word must contain both WAEN = 1 and UTA = 1 simultaneously.

12.4.4.5 Synchronous sampling of LUPWAIT for early transfer acknowledge

However, programming WAEN = 1 and UTA = 1 in the same RAM word, under certain conditions, allows the UPM to treat LUPWAIT as a synchronous signal, which must meet set-up and hold times in relation to the rising edge of the bus clock.

The conditions are as follows:

- The PLL must be enabled, that is, LCRR[PBYP] = 0.
- DLT3 bit must be cleared in the same RAM word to avoid mid-sampling of read data.
- LBCR[LPBSE] = 0 and MXMR[GPL4] = 1
- The combination WAEN=1 and UTA=1 should be in the RAM word next to the word which gets frozen by LUPWAIT assertion. This condition limits the use of this mode to cases where the exact cycle of LUPWAIT assertion is predictable.

In this mode, as soon as UPM samples LUPWAIT negated on the rising edge of the bus clock, it immediately generates an internal transfer acknowledge, which allows a data transfer one bus clock cycle later. The generation of transfer acknowledge is early because LUPWAIT is not re-synchronized. The acknowledge occurs early or normally depending on whether the UPM was already frozen in WAIT cycles or not. This feature allows the synchronous negation of LUPWAIT to affect a data transfer, even if UTA, WAEN, and LAST are set simultaneously.

12.4.4.6 Extended hold time on read accesses-UPM

Slow memory devices that take a long time to turn off their data bus drivers on read accesses should choose some non-zero combination of $OR_n[TRLX]$ and $OR_n[EHTR]$.

The next accesses after a read access to the slow memory device is delayed by the number of clock cycles specified in the OR_n register in addition to any existing bus turnaround cycle.

12.5 eLBC initialization/application information

This section provides information about the following:

- [Interfacing to peripherals in different address modes](#)
- [Bus turnaround](#)
- [Interface to different port-size devices](#)
- [Command sequence examples for NAND flash EEPROM](#)
- [Interfacing to fast-page mode DRAM using UPM](#)
- [Interfacing to ZBT SRAM using UPM](#)

Note that the data paths for configuration of eLBC and data transfers are different so there should be a delay of 200 clock cycles between register access and data access to avoid synchronization issues.

12.5.1 Interfacing to peripherals in different address modes

This section provides guidelines for interfacing to peripherals.

12.5.1.1 Multiplexed address/data bus for 32-bit addressing

In order to reduce pins on the local bus, address and data signals are multiplexed. To build the address, an external latch is used to demultiplex and reconstruct the original address.

Because the LALE signal provides the correct timing to control a standard logic latch, no external intelligence is needed. To pass data, the LAD signals can be directly connected to the data signals of the memory/peripheral.

Transactions on the local bus begin with an address phase. The eLBC drives the transaction address on the LAD signals and asserts the LALE signal to latch the address. This assertion causes address bits A[16:31] to appear on LAD[0:15]. The eLBC can then continue on into the data phase.

The eLBC supports port sizes of 8 bits. When there is an access larger than the port size, the eLBC breaks up the access into smaller transactions using the non-multiplexed address signals LA_n .

In addition, the eLBC supports burst transfers in the UPM machine. To minimize the amount of address phases needed on the local bus and to optimize the throughput, LA_n are driven separately. All other address bits, A[0:15], must be reconstructed through the latch.

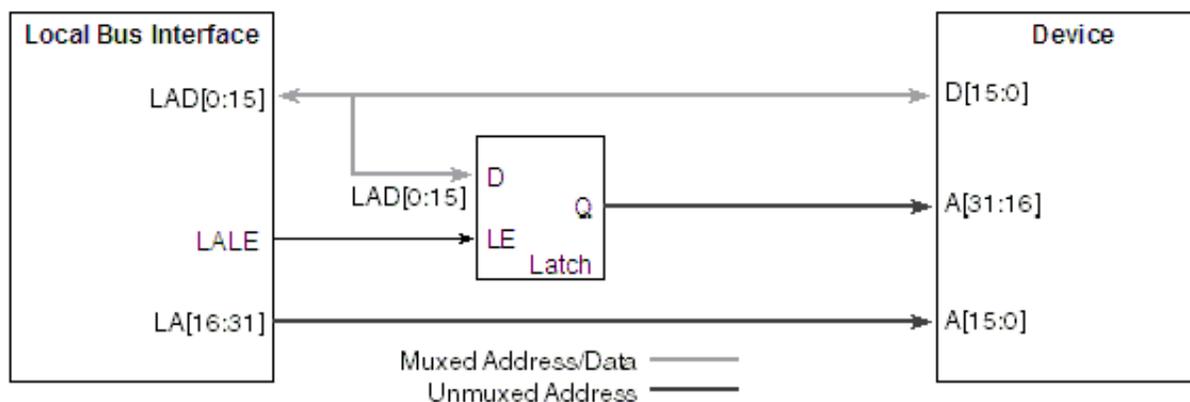


Figure 12-127. Multiplexed address/data bus for 32-bit addressing

12.5.1.2 Non-multiplexed address and data buses

For small address space applications the address latch may be eliminated entirely if the local bus address is taken entirely from $LA[16:31]$, in which case addresses driven onto LAD during address phases are simply ignored. The connection is illustrated in the figure

below. In non-multiplexed mode, the waveforms remain the same except for a few things, such as ASHIFT parameter, LAD bus turnaround time, LALE timings, that can be ignored.

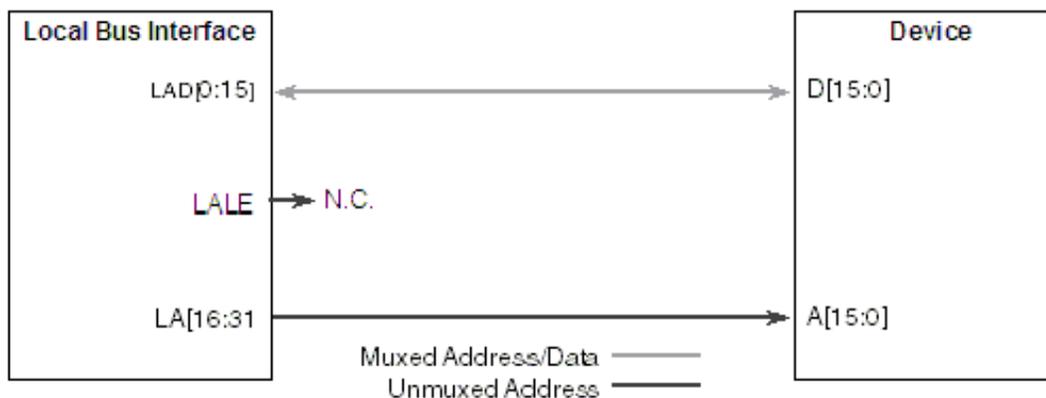


Figure 12-128. Non-multiplexed address and data buses

12.5.1.3 Peripheral hierarchy on the local bus for high bus speeds

To achieve the highest possible bus speeds on the local bus, it is recommended to reduce the number of devices connected directly to the bus.

For best results, only one bank of synchronous SRAMs should have a direct connection, and a bus demultiplexor should be used to replace separate latch and separate bus transceiver combinations. The figure below shows an example of such a hierarchy. This section is only a guideline, and the board designer must simulate the electric characteristics of the scenario to determine the maximum operating frequency.

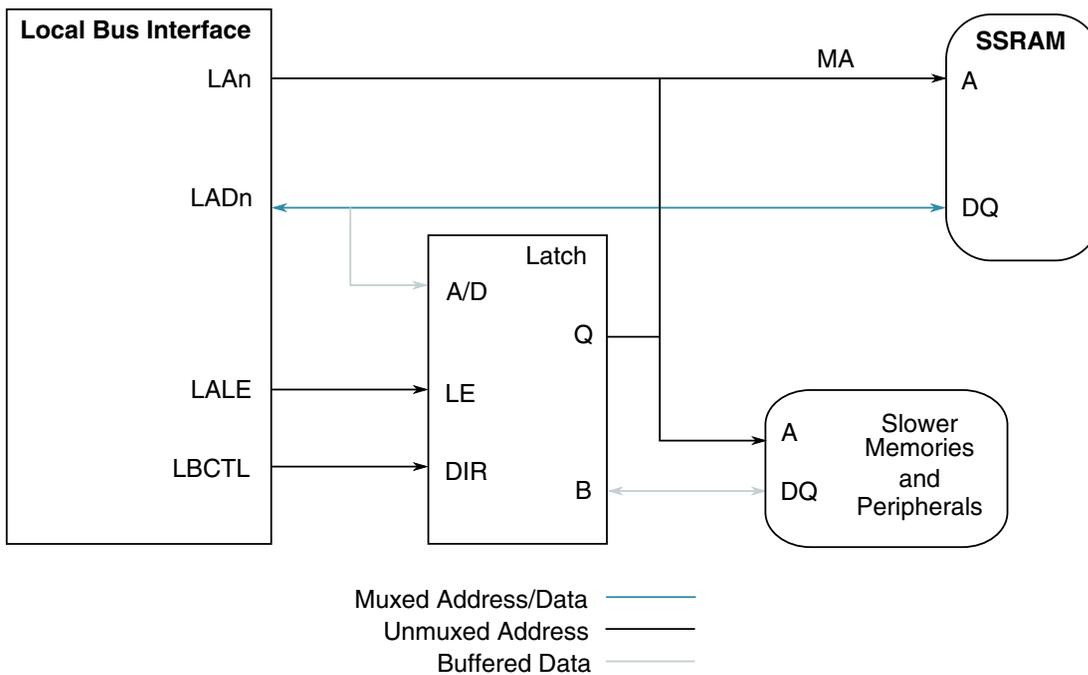


Figure 12-129. Local bus peripheral hierarchy for high bus speeds

12.5.1.4 GPCM timings

In case a system contains a memory hierarchy with high speed synchronous memories (synchronous SRAM) and lower speed asynchronous memories (for example, FLASH EPROM and peripherals) the GPCM-controlled memories should be decoupled by buffers to reduce capacitive loading on the bus.

Those buffers have to be taken into account for the timing calculations.

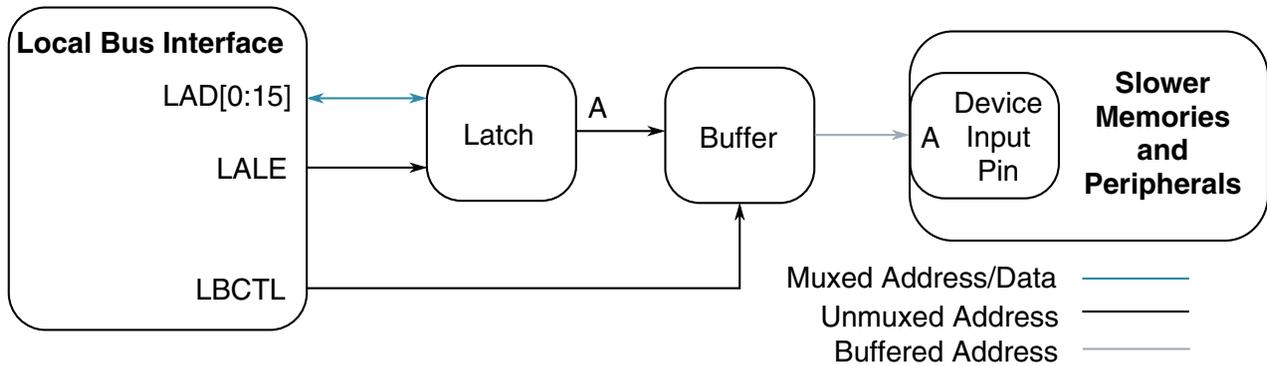


Figure 12-130. GPCM address timings

To calculate address setup timing for a slower peripheral/memory device, several parameters have to be added: propagation delay for the address latch, propagation delay for the buffer and the address setup for the actual peripheral. Typical values for the two propagation delays are in the order of 3 to 6 ns, so for a 100-MHz bus frequency, LCS_B should arrive on the order of 2 to 3 bus clocks later.

For data timings, only the propagation delay of one buffer plus the actual data setup time has to be considered.

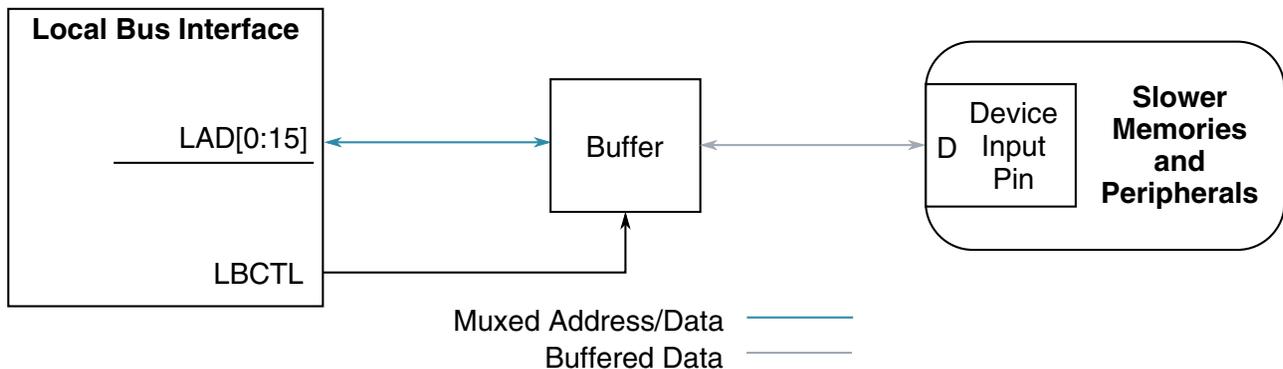


Figure 12-131. GPCM data timings

12.5.2 Bus turnaround

Because the local bus uses multiplexed address and data, special consideration must be given to avoid bus contention at bus turnaround.

The following cases must be examined:

- Address phase after previous read
- Read data phase after address phase
- Read-modify-write cycle for parity protected memory banks
- UPM cycles with additional address phases

The bus does not change direction for the following cases so they need no special attention:

- Continued burst after the first beat
- Write data phase after address phase
- Address phase after previous write

12.5.2.1 Address phase after previous read

During a read cycle, the memory/peripheral drives the bus and the bus transceiver drives LAD.

After the data has been sampled, the output drivers of the external device must be disabled. This can take some time; for slow devices the EHTR feature of the GPCM or the programmability of the UPM should be used to guarantee that those devices have stopped driving the bus when the eLBC memory controller ends the bus cycle.

In this case, after the previous cycle ends, LBCTL goes high and changes the direction of the bus transceiver. The eLBC then inserts a bus turnaround cycle to avoid contention. The external device has now already placed its data signals in high impedance and no bus contention will occur.

12.5.2.2 Read data phase after address phase

During the address phase, LAD actively drives the address and LBCTL is high, driving the bus transceivers in the same direction as during a write.

After the end of the address phase, LBCTL goes low and changes the direction of the bus transceiver. The eLBC places the LAD signals in high impedance after its $t_{dis}(LB)$. The LBCTL will have its new state after $t_{en}(LB)$ and, because this is an asynchronous input, the transceiver starts to drive those signals after its $t_{en}(\text{transceiver})$ time. The system designer has to ensure, that $[t_{en}(LB) + t_{en}(\text{transceiver})]$ is larger than $t_{dis}(LB)$ to avoid bus contention.

12.5.2.3 Read-modify-write cycle for parity protected memory banks

Principally, a read-modify-write cycle is a read cycle immediately followed by a write cycle.

Because the write cycle will have a new address phase in any case, this basically is the same case as an address phase after a previous read.

12.5.2.4 UPM cycles with additional address phases

The flexibility of the UPM allows the user to insert additional address phases during read cycles by changing the AMX field, therefore turning around the bus during one pattern.

The eLBC automatically inserts a single bus turnaround cycle if the bus (LAD) was previously high impedance for any reason, such as a read, before LALE is driven and LAD is driven with the new address. The turnaround cycle is not inserted on a write, because the bus was already driven to begin with.

However, bus contention could potentially still occur on the far side of a bus transceiver. It is the responsibility of the designer of the UPM pattern to guarantee that enough idle cycles are inserted in the UPM pattern to avoid this.

12.5.3 Interface to different port-size devices

The eLBC supports 8- and 16-bit data port sizes.

However, the bus requires that the portion of the data bus used for a transfer to or from a particular port size be fixed. A 16-bit port must reside on LAD[0-15], and an 8-bit port must reside on LAD[0-7]. The local bus always tries to transfer the maximum amount of data on all bus cycles

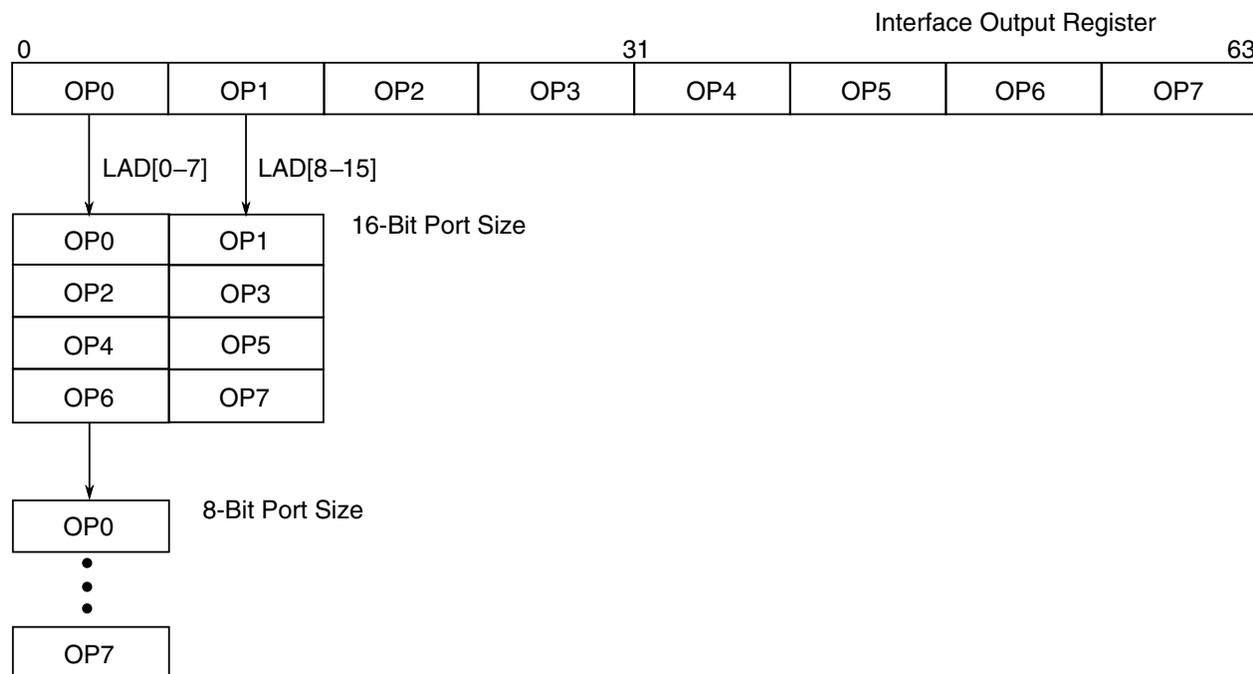


Figure 12-132. Interface to different port-size devices

Table 12-204. Data bus drive requirements for read cycles

Transfer Size	Address State ¹ 3 lsbs	Port Size/LAD Data Bus Assignments							
		16-Bit				8-Bit			
		0-7	8-15	16-23	24-31	0-7	8-15	16-23	24-31
Byte	000	OP0	-			OP0			
	001	-	OP1			OP1			
	010	OP2	-			OP2			
	011	-	OP3			OP3			
	100	OP4	-			OP4			
	101	-	OP5			OP5			
	110	OP6	-			OP6			
	111	-	OP7			OP7			
Half Word	000	OP0	OP1			OP0			
	001	-	OP1			OP1			
	010	OP2	OP3			OP2			
	100	OP4	OP5			OP4			
	101	-	OP5			OP5			
	110	OP6	OP7			OP6			
Word	000	OP0	OP1			OP0			
	100	OP4	OP5			OP4			

1. Address state is the calculated address for port size.

12.5.4 Command sequence examples for NAND flash EEPROM

To program the eLBC and FCM for executing NAND Flash command sequences, command codes and pause states should be obtained from the relevant NAND Flash device data sheet and programmed into FCM configuration registers.

This section illustrates some common sequences for large-page, multi-gigabit NAND Flash EEPROMs; however, details should be verified against manufacturers' specific programming data.

Throughout these examples it is assumed that one or more banks of eLBC has been configured under FCM control ($BR_n[MSEL] = 001$), with base address, port size, ECC mode, and timing parameters configured in accordance with the device hardware specifications.

12.5.4.1 NAND flash soft reset command sequence example

An example of configuring FCM to execute a soft reset command to large-page NAND Flash is shown in the table below.

This sequence does not require use of the shared FCM buffer RAM.

The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled.

Table 12-205. FCM register settings for soft reset (ORn[PGS] = 1)

Register	Initial Contents	Description
FCR	0xFF00_0000	CMD0 = 0xFF = reset command; other commands unused
FBAR	-	Unused
FPAR	-	Unused
FBCR	-	Unused
MDR	-	Unused
FIR	0x4000_0000	OP0 = CM0 = command 0; OP1-OP7 = NOP

12.5.4.2 NAND flash read status command sequence example

An example of configuring FCM to execute a status read command to large-page NAND Flash is shown in the table below.

This sequence does not require use of the shared FCM buffer RAM, but reads the NAND Flash status into register MDR[AS0] .

The sequence is initiated by writing FMR[OP] = 10 and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled.

Table 12-206. FCM register settings for status read (ORn[PGS] = 1)

Register	Initial Contents	Description
FCR	0x7000_0000	CMD0 = 0x70 = read status command; other commands unused
FBAR	-	Unused
FPAR	-	Unused
FBCR	-	Unused
MDR	-	Status returned in AS0

Table continues on the next page...

Table 12-206. FCM register settings for status read (ORn[PGS] = 1) (continued)

Register	Initial Contents	Description
FIR	0x4B00_0000	OP0 = CM0 = command 0; OP1 = RS = read status to MDR; OP2-OP7 = NOP

12.5.4.3 NAND flash read identification command sequence example

An example of configuring FCM to execute a status ID command to large-page NAND Flash is shown in the table below.

This sequence does not require use of the shared FCM buffer RAM, but uses MDR to set up a dummy address prior to the sequence, and then to receive the first 4 bytes of ID during the sequence.

The sequence is initiated by writing FMR[OP] = 10, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled. MDR[AS3-AS0] then can be read to obtain the first 4 bytes of NAND Flash ID.

Table 12-207. FCM register settings for ID read (ORn[PGS] = 1)

Register	Initial Contents	Description
FCR	0x9000_0000	CMD0 = 0x90 = read ID command; other commands unused
FBAR	-	Unused
FPAR	-	Unused
FBCR	-	Unused
MDR	All zeros	AS0 = 0x00 = dummy address for read ID command; AS0-AS3 return with first 4 bytes of ID code
FIR	0x43BB_BB0	OP0 = CM0 = command 0; OP1 = UA = user address from MDR; OP2-OP6 = RS = read 4 bytes ID into MDR[AS3-AS0]; OP7 = NOP
FMR	0x0000_F002	
LSOR	0x0000_0001	Assumes NAND flash is on eLBC CS1

12.5.4.4 NAND flash page read command sequence example

An example of configuring FCM to execute a random page read command to large-page NAND Flash is shown in the table below.

This sequence reads an entire page (main and spare region) into the shared FCM buffer RAM, checking ECC as it proceeds.

The sequence is initiated by writing FMR[OP] = 11, and issuing a special operation to the bank. A few cycles before completion itself, FECCn gets updated with the ECC bytes for the main region validated by FECCn[0]. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled. Once the sequence has completed, the shared buffer (buffer 1 for page index 5) and transfer error registers (LTECCR that reports the 512 blocks with unibit /multibit errors if any) are valid.

Table 12-208. FCM register settings for page read (ORn[PGS] = 1)

Register	Initial Contents	Description
FCR	0x0030_0000	CMD0 = 0x00 = random read address entry; CMD1 = 0x30 = read page
FBAR	block index (for example block 0x0001_0AB4)	BLK locates index of 128-Kbyte block
FPAR	page offset (for example 0x0000_5000 locates page 5, buffer 1)	PI locates page index in BLK; PI mod 2 indexes FCM buffer RAM; MS = 0 and CI = 0
FBCR	All zeros	BC = 0 to read entire 2112-byte page with ECC check
MDR	-	Unused
FIR	0x4125_E000	OP0 = CM0 = command 0; OP1 = CA = column address; OP2 = PA = page address; OP3 = CM1 = command 1; OP4 = RBW = wait on Flash ready and read data into FCM buffer; OP5-OP7 = NOP

12.5.4.5 NAND flash block erase command sequence example

An example of configuring FCM to execute a block erase command to large-page NAND Flash is shown in the table below.

This sequence does not require use of the shared FCM buffer RAM, but returns with the erase status in MDR[AS0].

The sequence is initiated by writing $FMR[OP] = 11$, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTESR[CC]) if interrupts are enabled.

Note that operations specified by OP3 and OP4 (status read) should never be skipped while erasing a NAND Flash device, because, in case that happens, contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP3 and OP4 operations are skipped, it may also happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.

Table 12-209. FCM register settings for block erase (ORn[PGS] = 1)

Register	Initial Contents	Description
FCR	0x6070_D000	CMD0 = 0x60 = block address entry; CMD1 = 0x70 = read status CMD2 = 0xD0 = erase block;
FBAR	Block Index (for example block 0x0001_0AB4)	BLK locates index of 128-Kbyte block
FPAR	All zeros	PI = 0 to locate block boundary
FBCR	-	Unused
MDR	-	Returns with AS0 holding erase status
FIR	0x426D_B000	OP0 = CM0 = command 0; OP1 = PA = page address; OP2 = CM2 = command 2; OP3 = CW1 = wait on Flash ready and issue command 1; OP4 = RS = read erase status into MDR[AS0]; OP5-OP7 = NOP

12.5.4.6 NAND flash program command sequence example

An example of configuring FCM to execute a program command to large-page NAND Flash is shown in the table below.

This sequence writes an entire page (main and spare region) from the shared FCM buffer RAM, generating ECC as it proceeds.

The shared buffer (buffer 1 for page index 5) must be initialized by software prior to starting the sequence. The sequence is initiated by writing FMR[OP] = 11, and issuing a special operation to the bank. At the conclusion of the sequence, eLBC will issue a command complete interrupt (LTER[CC]) if interrupts are enabled. The status of the programming operation is returned in MDR[AS0].

Note that operations specified by OP5 and OP6 (status read) should never be skipped while programming a NAND Flash device, because, in case that happens, contention may arise on LGPL4. A possible case is that the next transaction from eLBC may try to use that pin as an output and since the NAND Flash device might already be driving it, contention will occur. In case OP5 and OP6 operations are skipped, it may also happen that a new command is issued to the NAND Flash device even when the device has not yet finished processing the previous request. This may also result in unpredictable behavior.

Table 12-210. FCM register settings for page program (ORn[PGS] = 1)

Register	Initial Contents	Description
FCR	0x8070_1000	CMD0 = 0x80 = page address and data entry; CMD1 = 0x70 = read status CMD2 = 0x10 = program page;
FBAR	block index (for example block 0x0001_0AB4)	BLK locates index of 128-Kbyte block
FPAR	page offset (for example 0x0000_5000 locates page 5, buffer 1)	PI locates page index in BLK; PI mod 2 indexes FCM buffer RAM; MS = 0 and CI = 0
FBCR	All zeros	BC = 0 to write entire 2112-Byte page with ECC generation
MDR	-	Returns with AS0 holding program status
FIR	0x4128_6DB0	OP0 = CM0 = command 0; OP1 = CA = column address; OP2 = PA = page address; OP3 = WB = write data from buffer; OP4 = CM2 = command 2; OP5 = CW1 = wait on Flash ready and issue command 1; OP6 = RS = read erase status into MDR[AS0]; OP7 = NOP

12.5.5 Interfacing to fast-page mode DRAM using UPM

Connecting the local bus UPM controller to a DRAM device requires a detailed examination of the timing diagrams representing the possible memory cycles that must be performed when accessing this device.

This section describes timing diagrams for various UPM configurations for fast-page mode DRAM, with LCRR[CLKDIV] = 4 (clock ratio of 8) or 8 (clock ratio of 16). The illustrative examples shown in [Table 12-211](#)-[Table 12-215](#) may not represent the timing necessary for any specific device used with the eLBC. Here, LGPL1 is programmed to drive R/W of the DRAM, although any LGPL_n signal may be used for this purpose.

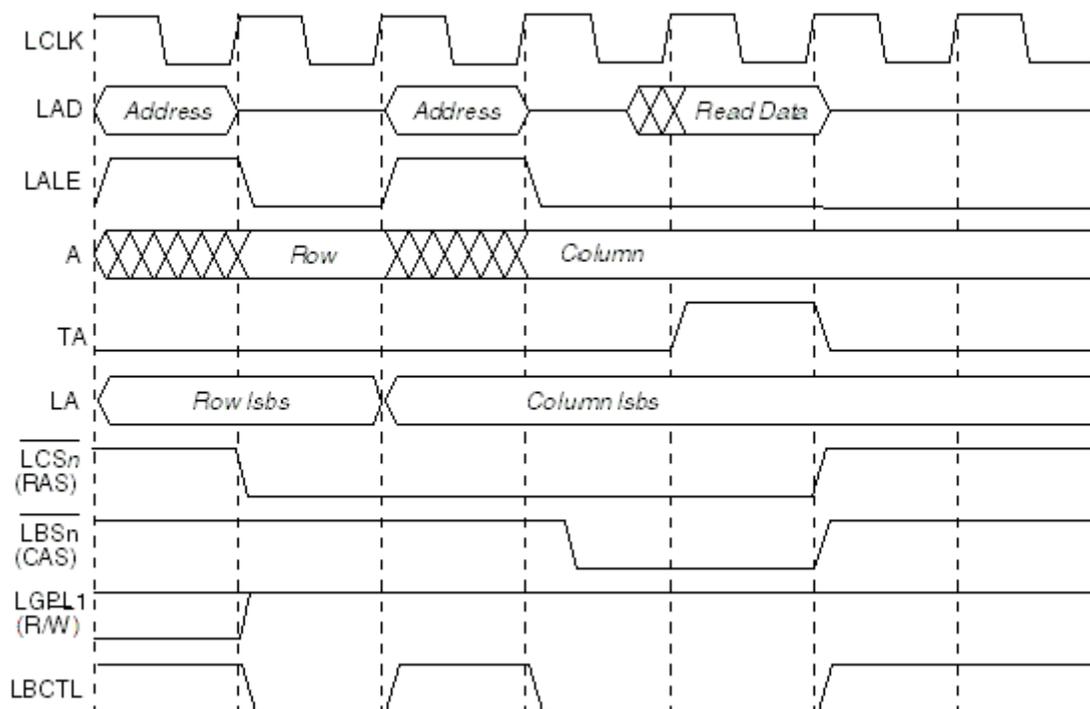


Figure 12-133. Single-beat read access to FPM DRAM

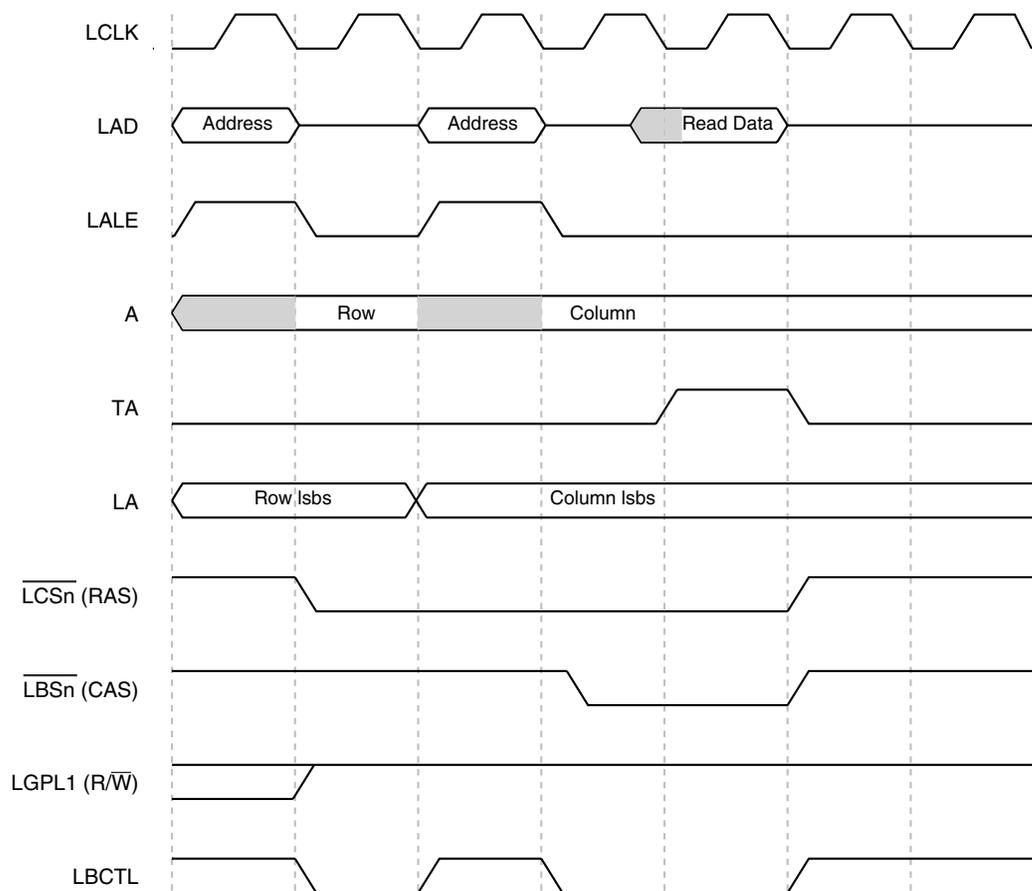


Figure 12-134. Single-beat read access to FPM DRAM

Table 12-211. Single-beat read access to FPM DRAM

cst1	0	LALE pause (due to change in AMX)	0	0	Bit 0
cst2	0		0	0	Bit 1
cst3	0		0	0	Bit 2
cst4	0		0	0	Bit 3
bst1	1		1	0	Bit 4
bst2	1		0	0	Bit 5
bst3	1		0	0	Bit 6
bst4	1		0	0	Bit 7
g0l0					Bit 8
g0l1					Bit 9
g0h0					Bit 10
g0h1					Bit 11
g1t1	1		1	1	Bit 12
g1t3	1		1	1	Bit 13
g2t1					Bit 14
g2t3					Bit 15
g3t1					Bit 16
g3t3					Bit 17
g4t1					Bit 18
g4t3					Bit 19
g5t1					Bit 20
g5t3					Bit 21
redo[0]					Bit 22
redo[1]					Bit 23
loop	0		0	0	Bit 24
exen	0		0	0	Bit 25
amx0	1		0	0	Bit 26
amx1	0		0	0	Bit 27
na	0		0	0	Bit 28
uta	0		0	1	Bit 29
todt	0		0	1	Bit 30
last	0		0	1	Bit 31
	RSS	RSS + 1	RSS + 1	RSS + 2	

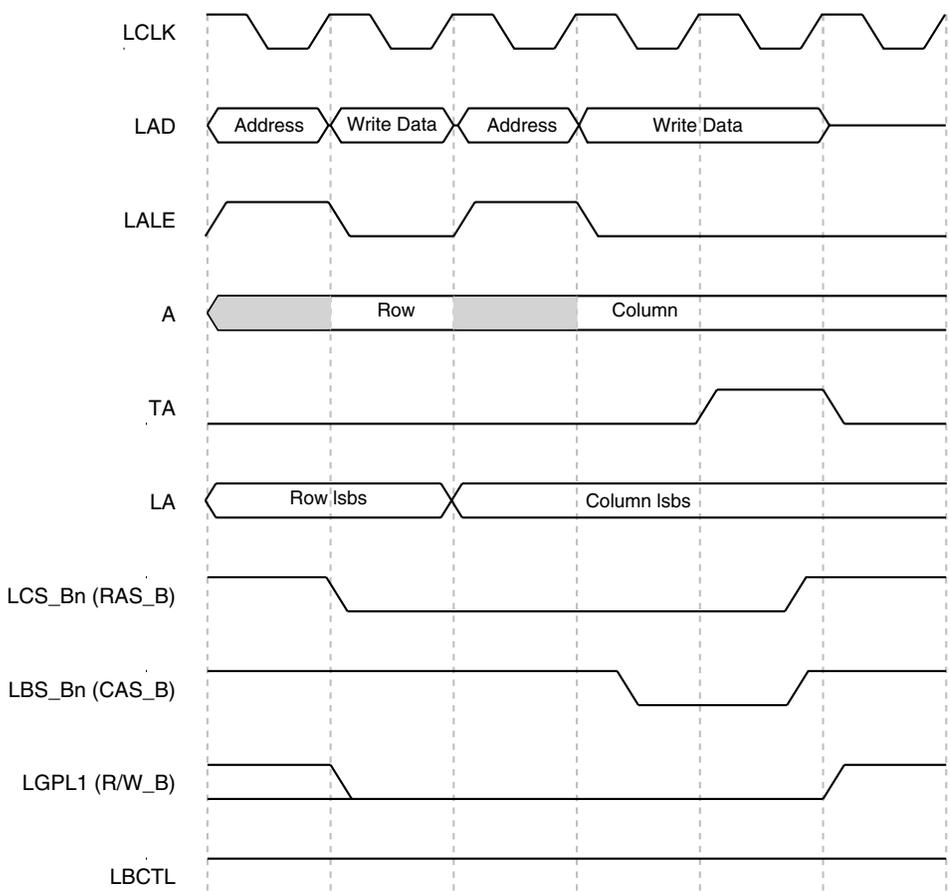


Figure 12-135. Single-beat write access to FPM DRAM

Table 12-212. Single-beat write access to FPM DRAM

cst1	0	LALE pause (due to change in AMX)	0	0	Bit 0
cst2	0		0	0	Bit 1
cst3	0		0	0	Bit 2
cst4	0		0	1	Bit 3
bst1	1		1	0	Bit 4
bst2	1		1	0	Bit 5
bst3	1		0	0	Bit 6
bst4	1		0	1	Bit 7
g0l0					Bit 8
g0l1					Bit 9
g0h0					Bit 10
g0h1					Bit 11
g1t1	0		0	0	Bit 12
g1t3	0		0	0	Bit 13
g2t1					Bit 14
g2t3					Bit 15
g3t1					Bit 16
g3t3					Bit 17
g4t1					Bit 18
g4t3					Bit 19
g5t1					Bit 20
g5t3					Bit 21
redo[0]					Bit 22
redo[1]					Bit 23
loop	0		0	0	Bit 24
exen	0		0	0	Bit 25
amx0	1		0	0	Bit 26
amx1	0		0	0	Bit 27
na	0		0	0	Bit 28
uta	0		0	1	Bit 29
todt	0		0	1	Bit 30
last	0		0	1	Bit 31
	WSS	WSS + 1	WSS + 1	WSS + 2	

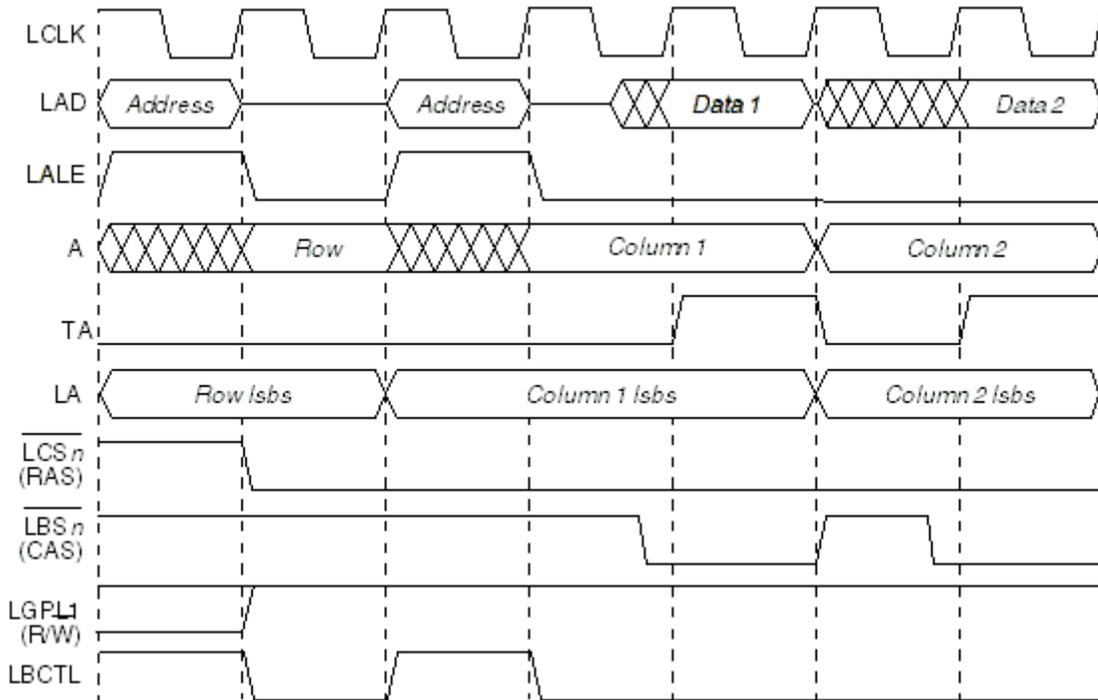


Figure 12-136. Burst read access to FPM DRAM using LOOP (two beats shown)

Table 12-213. Burst read access to FPM DRAM using LOOP (two beats shown)

cst1	0	LALE pause (due to change in AMX)	0	0	1	Bit 0
cst2	0		0	0	1	Bit 1
cst3	0		0	0	1	Bit 2
cst4	0		0	0	1	Bit 3
bst1	1		1	0	1	Bit 4
bst2	1		1	0	1	Bit 5
bst3	1		1	0	1	Bit 6
bst4	1		0	0	1	Bit 7
g0l0						Bit 8
g0l1						Bit 9
g0h0						Bit 10
g0h1						Bit 11
g1t1	1		1	1	1	Bit 12
g1t3	1		1	1	1	Bit 13
g2t1						Bit 14
g2t3						Bit 15
g3t1						Bit 16
g3t3						Bit 17
g4t1						Bit 18
g4t3						Bit 19
g5t1						Bit 20
g5t3						Bit 21
redo[0]						Bit 22
redo[1]						Bit 23
loop	0		1	1	0	Bit 24
exen	0		0	1	0	Bit 25
amx0	1		0	0	0	Bit 26
amx1	0		0	0	0	Bit 27
na	0		0	1	0	Bit 28
uta	0		0	1	0	Bit 29
todt	0		0	0	1	Bit 30
last	0		0	0	1	Bit 31
	RBS		RBS + 1	RBS + 2	RBS + 3	

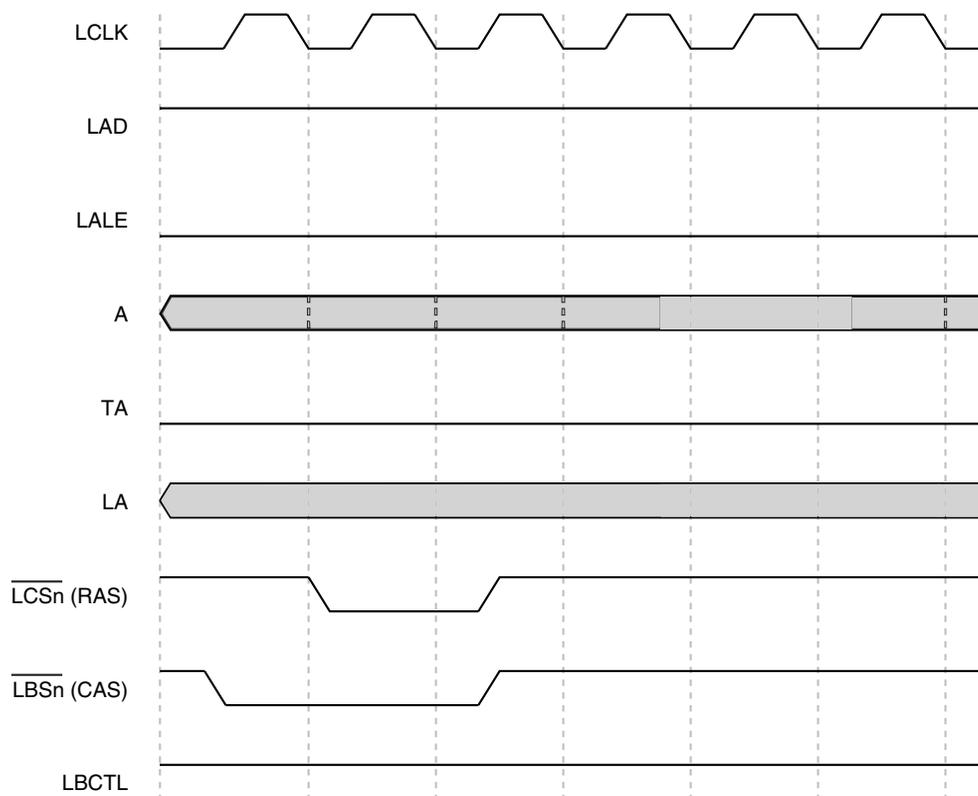


Figure 12-137. Refresh cycle (CBR) to FPM DRAM

Table 12-214. Refresh cycle (CBR) to FPM DRAM

cst1	1	0	0	Bit 0
cst2	1	0	0	Bit 1
cst3	1	0	1	Bit 2
cst4	1	0	1	Bit 3
bst1	1	0	0	Bit 4
bst2	0	0	0	Bit 5
bst3	0	0	1	Bit 6
bst4	0	0	1	Bit 7
g0l0				Bit 8
g0l1				Bit 9
g0h0				Bit 10
g0h1				Bit 11
g1t1				Bit 12
g1t3				Bit 13
g2t1				Bit 14
g2t3				Bit 15

Table continues on the next page...

Table 12-214. Refresh cycle (CBR) to FPM DRAM (continued)

g3t1				Bit 16
g3t3				Bit 17
g4t1				Bit 18
g4t3				Bit 19
g5t1				Bit 20
g5t3				Bit 21
redo[0]				Bit 22
redo[1]				Bit 23
loop	0	0	0	Bit 24
exen	0	0	0	Bit 25
amx0	0	0	0	Bit 26
amx1	0	0	0	Bit 27
na	0	0	0	Bit 28
uta	0	0	0	Bit 29
todt	0	0	1	Bit 30
last	0	0	1	Bit 31
	PTS	PTS + 1	PTS + 2	

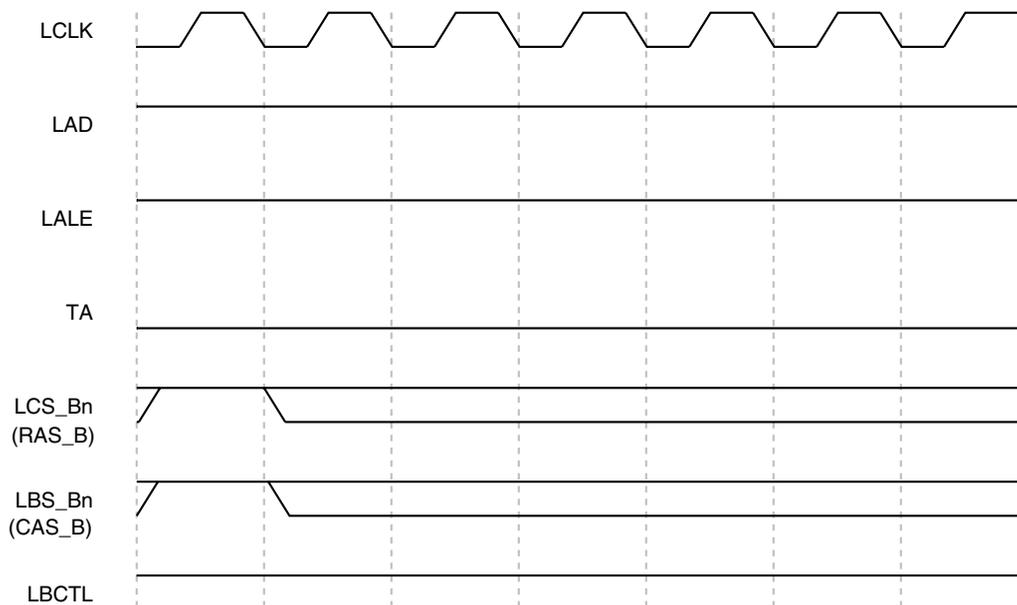


Figure 12-138. Exception cycle

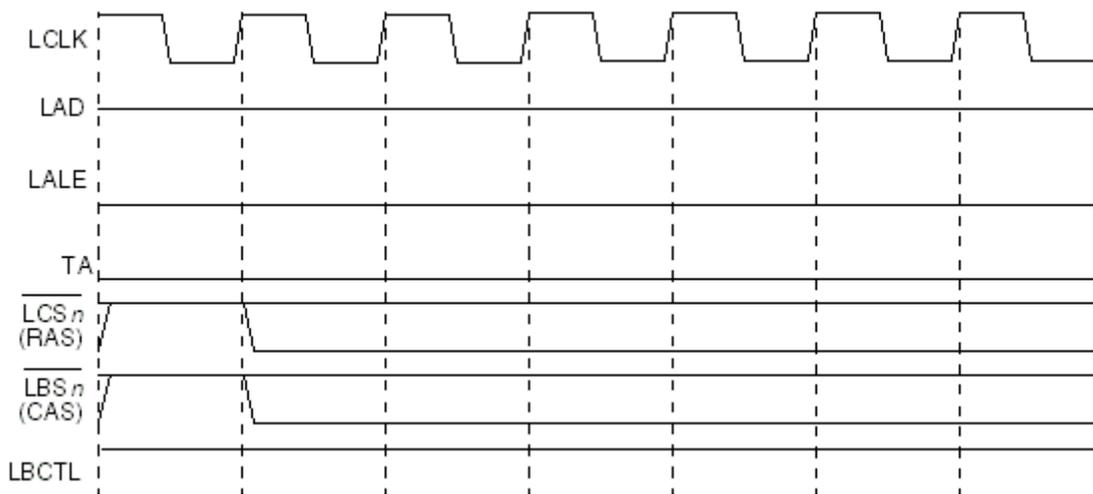


Figure 12-139. Exception cycle

Table 12-215. Exception cycle

cst1	1	Bit 0
cst2	1	Bit 1
cst3	1	Bit 2
cst4	1	Bit 3
bst1	1	Bit 4
bst2	1	Bit 5
bst3	1	Bit 6
bst4	1	Bit 7
g0l0		Bit 8
g0l1		Bit 9
g0h0		Bit 10
g0h1		Bit 11
g1t1		Bit 12
g1t3		Bit 13
g2t1		Bit 14
g2t3		Bit 15
g3t1		Bit 16
g3t3		Bit 17
g4t1		Bit 18
g4t3		Bit 19
g5t1		Bit 20
g5t3		Bit 21
redo[0]		Bit 22
redo[1]		Bit 23
loop	0	Bit 24
exen	0	Bit 25

Table continues on the next page...

Table 12-215. Exception cycle (continued)

amx0	0	Bit 26
amx1	0	Bit 27
na	0	Bit 28
uta	0	Bit 29
todt	1	Bit 30
last	1	Bit 31
	EXS	

12.5.6 Interfacing to ZBT SRAM using UPM

ZBT SRAMs have been designed to optimize the performance of table access in networking applications.

This section describes how to interface to ZBT SRAMs. The figure below shows the connections. The UPM is used to generate control signals. The same interfacing is used for pipelined and flow-through versions of ZBT SRAMs. However different UPM patterns must be generated for those cases. As ZBT SRAMs are mostly used by performance-critical applications, it is assumed here that, typically, the maximum width of the local bus of 16 bits will be used.

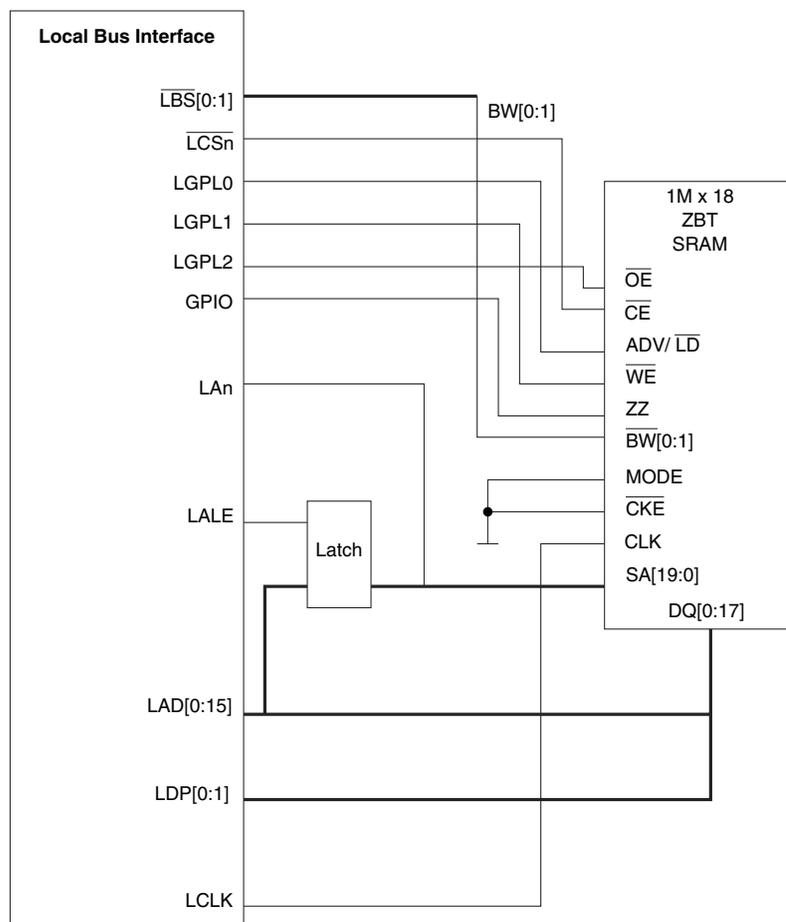


Figure 12-140. Interface to ZBT SRAM

ZBT SRAMs allow different configurations. For the local bus, the burst order should be set to linear burst order by tying the mode pin to GND. CKE_B should also be tied to ground.

ZBT SRAMs perform four-beat bursts. Because the eLBC generates sixteen-beat transactions (for 16-bit ports) the UPM breaks down each burst into four consecutive four-beat bursts. The internal address generator of the eLBC generates the new LA{A27, A28} for the second, third, and fourth burst. In other words, because linear burst is used on the SRAM, the device itself bursts with the burst addresses of [0:1:2:3]. The local bus always generates linear bursts and expects [0:1:2:3:4:5:6:7:....:15]. Therefore, four consecutive linear bursts of the ZBT SRAM with {A27, A28} = {0,0} for the first burst, {A27, A28} = {0,1} for the second burst, {A27, A28} = {1,0} for the third burst, and {A27, A28} = {1,1} for the fourth burst give the desired burst pattern.

Chapter 13

DMA Controller

The DMA controller transfers blocks of data between the many interface and functional modules of this chip, independent of the core or external hosts.

13.1 DMA overview

This chapter describes the DMA controller offered on this chip.

It describes a single DMA controller, which is instantiated on this chip. As such, all functionality is identical between the two controllers with the exception of the register offsets, as noted in the DMA register summary table , and the external signals, as noted in DMA signals table.

The DMA controller transfers blocks of data between the many interface and functional modules of this chip, independent of the cores or external hosts. It has four high-speed DMA channels. Both the cores and external devices can initiate DMA transfers. The channels are capable of complex data movement and advanced transaction chaining. Operations such as descriptor fetches and block transfers are initiated by each channel. A channel is selected by the arbitration logic and information is passed to the source and destination control blocks for processing. The source and destination blocks generate read and write requests to the address tenure engine, which manages the DMA master port address interface. After a transaction is accepted by the master port, control is transferred to the data tenure engine that manages the read and write data transfers. A channel remains active in the shared resources for the duration of the data transfer unless the allotted bandwidth per channel is reached.

This figure shows the block diagram of the DMA controller.

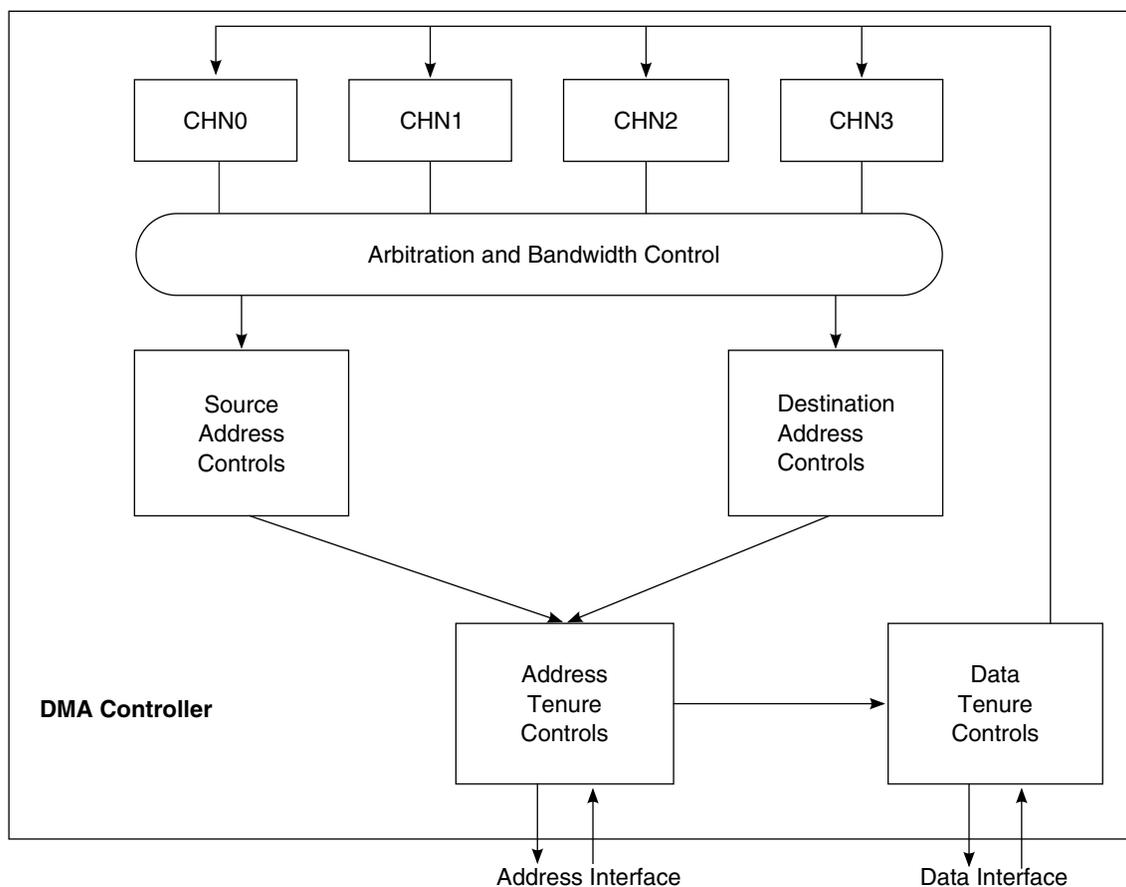


Figure 13-1. DMA block diagram

13.1.1 DMA features summary

The DMA controller offers the following features:

- Four high-speed/high-bandwidth channels accessible by local and remote masters
- Basic DMA operation modes (direct, simple chaining)
- Extended DMA operation modes (advanced chaining and stride capability)
- Cascading descriptor chains
- Misaligned transfers
- Programmable bandwidth control between channels
- Up to 256 bytes for DMA sub-block transfers to maximize performance
- Interrupt on error and completed segment, list, or link
- Externally-controlled transfer using DMA_DREQ_B, DMA_DACK_B, and DMA_DDONE_B

13.1.2 DMA modes of operation

The DMA module has two modes of operation: basic and extended.

Basic mode is the DMA legacy mode, which does not support advanced features. Extended mode supports advanced features such as striding and flexible descriptor structures.

These two basic modes allow users to initiate and end DMA transfers in various ways. [Table 13-1](#) summarizes the relationship between the modes and the following features:

- Direct mode-No descriptors are involved. Software must initialize the required fields as described in [Table 13-2](#) before starting a transfer.
- Chaining mode-Software must initialize descriptors in memory and the required fields as described in [Table 13-2](#) before starting a transfer.
- Single-write start mode-The DMA process can be started using a single-write command to either the descriptor address register in one of the chaining modes or the source/destination address registers in one of the direct modes.
- External control capability-This allows an external agent to start, pause, and check the status of a DMA transfer that has already been initialized.
- Channel continue capability-The channel continue capability allows software the flexibility of having the DMA controller start with descriptors that have already been programmed while software continues to build more descriptors in memory.
- Channel abort capability-The software can abort a previously initiated transfer by setting the bit $MRn[CA]$. The DMA controller terminates all outstanding transfers initiated by the channel without generating any errors before entering an idle state.

Table 13-1. Relationship of modes and features

Mode	Mode with one additional feature	Mode with two additional features
B (Basic)	BD (basic direct)	BDS (BD single-write start)
		BDE (BD external control)
	BC (basic chaining)	BCE (BC external control)
		BCS (BC single-write start)
Ext (Extended)	ExtD (extended direct)	ExtDS (ExtD single-write start)
		ExtDE (ExtD external control)
	ExtC (extended chaining)	ExtCE (ExtC external control)
		ExtCS (ExtC single-write start)

The following table describes bit settings required for each DMA mode of operation.

Table 13-2. DMA mode bit settings

Modes with features	$MRn[XFE]$	$MRn[CTM]$	$MRn[SRW]$	$MRn[CDSM_SWSM]$	$MRn[EMS_EN]$
Basic direct modes					

Table continues on the next page...

Table 13-2. DMA mode bit settings (continued)

Modes with features	MRn[XFE]	MRn[CTM]	MRn[SRW]	MRn[CDSM_SWSM]	MRn[EMS_EN]
Basic direct	0	1	0	0	0
Basic direct external control	0	1	0	0	1
Basic direct single-write start	0	1	1	1 or 0	0
Basic chaining modes					
Basic chaining	0	0	Reserved	0	0
Basic chaining external control	0	0	Reserved	0	1
Basic chaining single-write start	0	0	Reserved	1	0
Extended direct modes					
Extended direct	1	1	0	0	0
Extended direct external control	1	1	0	0	1
Extended direct single-write start	1	1	1	1 or 0	0
Extended chaining modes					
Extended chaining	1	0	Reserved	0	0
Extended chaining external control	1	0	Reserved	0	1
Extended chaining single-write start	1	0	Reserved	1	0

See [DMA functional description](#) for details on these modes.

This figure shows the general DMA operational flow chart.

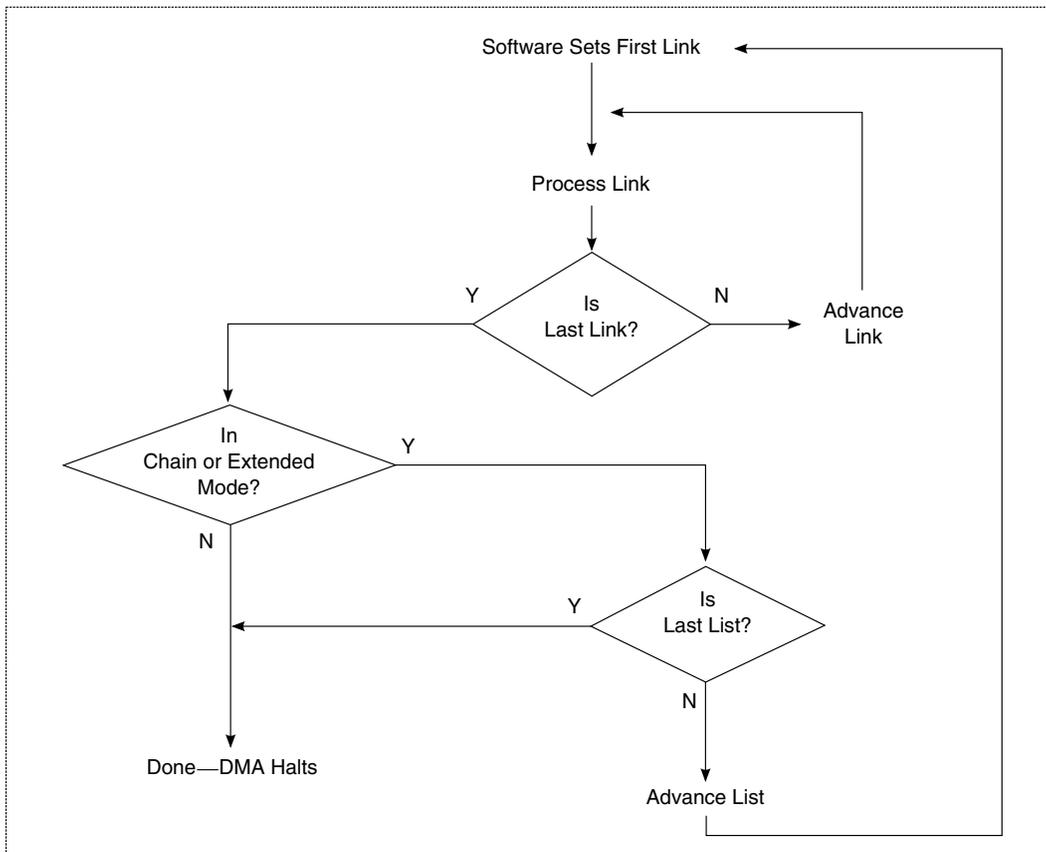


Figure 13-2. DMA operational flow chart

13.2 DMA external signal description

This section describes the external DMA signals.

13.2.1 Signal overview

This figure summarizes the DMA controller signals.

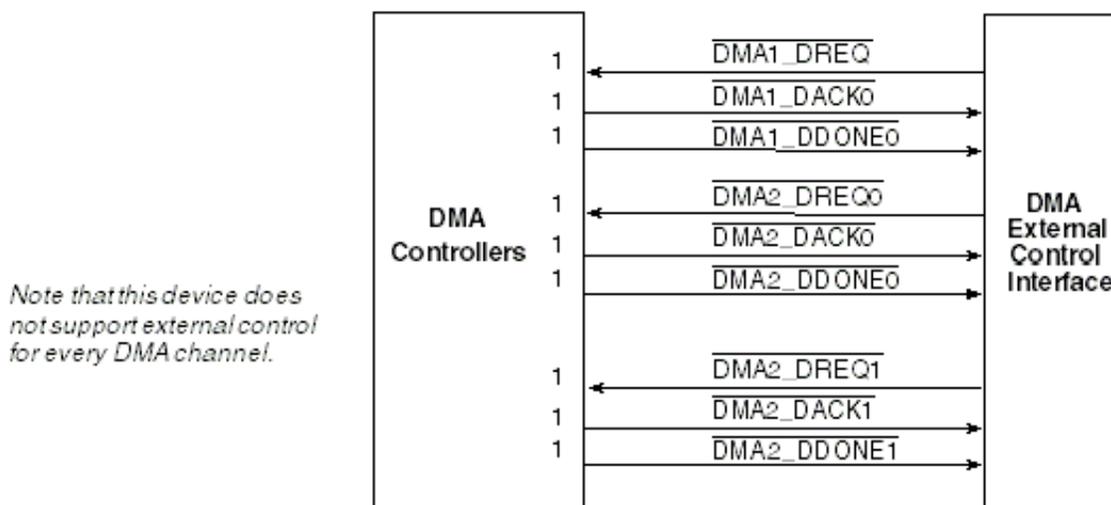


Figure 13-3. DMA signal summary

Note that the three DMA signals for DMA 2 channel 0 are multiplexed with the LCS_B[5:7] signals on the device. These functions are mutually exclusive and the active function is specified in the PMUXCR register of the global utilities block as described in [Alternate function signal multiplex control \(GUTS_PMUXCR\)](#).

13.2.2 DMA signal descriptions

This table describes the external DMA control signals. These signals are only utilized when operating in external master mode.

See [External control mode transfer](#).

Table 13-3. DMA signals-detailed signal descriptions

Signal	I / O	Description
DMA_DREQ_B n	I	DMA request. Indicates the start of a DMA transfer or a restart from a paused request. Assertion of DMA_DREQ_B n causes MRn[CS] to be set, thereby activating the corresponding DMA channel.
		State Meaning Asserted-Assertion of DMA_DREQ_B n while DMA_DACK_B n is negated causes a new transfer to start OR resumes a paused transfer if the EMP_EN bit is set. Assertion while DMA_DACK_B n is asserted results in an illegal condition. Negated-Negation while DMA_DACK_B n is asserted has no effect. Negation before the assertion of DMA_DACK_B n results in an illegal condition.
		Timing Assertion-Can be asserted asynchronously Negation- Must remain asserted at least until the assertion of the corresponding DMA_DACK_B n

Table continues on the next page...

Table 13-3. DMA signals-detailed signal descriptions (continued)

Signal	I / O	Description
DMA_DACK_B <i>n</i>	O	DMA acknowledge. Indicates that a DMA transfer is currently in progress
	State Meaning	Asserted-Indicates that a DMA transfer is currently in progress. Asserted after the assertion of DMA_DREQ_B <i>n</i> to indicate the start of a transfer Negated-Negated after finishing a complete transfer or after entering a paused state if MRn[EMP_EN] is set
	Timing	Assertion-Asynchronous assertion; asserted for more than three system clocks Negation-Asynchronous negation; negated for more than three system clocks
DMA_DDONE_B <i>n</i>	O	DMA done. Indicates that a DMA transfer is complete
	State Meaning	Asserted-Indicates transfer completion. SRn[CB] is clear. Note, however, that write data may still be queued at the target interface or in the process of transfer on an external interface. Negated-Indicates that the current transfer is in process
	Timing	Assertion-Always asserts asynchronously after the negation of the final DMA_DACK_B <i>n</i> to indicate completion of a transfer. For a paused transfer, DMA_DDONE_B <i>n</i> is asserted asynchronously after the negation of the final DMA_DACK_B <i>n</i> . Negation-Negated asynchronously after the assertion of DMA_DREQ_B <i>n</i> for the next transfer

13.3 DMA controller memory map

This section provides a detailed description of all accessible DMA memory and registers. The descriptions include individual, bit-level descriptions and reset states of each register. Undefined 4-byte address spaces within the map are reserved.

This table lists the DMA registers and their offsets. Note that the full register address is comprised of the programmable CCSRBAR together with the fixed DMA block base address and offset listed in table below.

DMA memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_1100	DMA n mode register (DMA1_MR0)	32	R/W	0800_0000h	13.3.1/693
2_1104	DMA n status register (DMA1_SR0)	32	R/W	0000_0000h	13.3.2/697
2_1108	DMA n current link descriptor extended address register (DMA1_ECLNDAR0)	32	R/W	0000_0000h	13.3.3/698
2_110C	DMA n current link descriptor address register (DMA1_CLNDAR0)	32	R/W	0000_0000h	13.3.4/699

Table continues on the next page...

DMA memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_1110	DMA n source attributes register (DMA1_SATR0)	32	R/W	0000_0000h	13.3.5/701
2_1114	DMA n source address register (DMA1_SAR0)	32	R/W	0000_0000h	13.3.6/702
2_1118	DMA n destination attributes register (DMA1_DATR0)	32	R/W	0000_0000h	13.3.7/703
2_111C	DMA n destination address register (DMA1_DAR0)	32	R/W	0000_0000h	13.3.8/704
2_1120	DMA n byte count register (DMA1_BCR0)	32	R/W	0000_0000h	13.3.9/705
2_1124	DMA n extended next link descriptor address register (DMA1_ENLNDAR0)	32	R/W	0000_0000h	13.3.10/705
2_1128	DMA n next link descriptor address register (DMA1_NLNDAR0)	32	R/W	0000_0000h	13.3.11/706
2_1130	DMA n extended current list descriptor address register (DMA1_ECLSDAR0)	32	R/W	0000_0000h	13.3.12/707
2_1134	DMA n current list descriptor address register (DMA1_CLSDAR0)	32	R/W	0000_0000h	13.3.13/708
2_1138	DMA n extended next list descriptor address register (DMA1_ENLS DAR0)	32	R/W	0000_0000h	13.3.14/709
2_113C	DMA n next list descriptor address register (DMA1_NLS DAR0)	32	R/W	0000_0000h	13.3.15/709
2_1140	DMA n source stride register (DMA1_SSR0)	32	R/W	0000_0000h	13.3.16/710
2_1144	DMA n destination stride register (DMA1_DSR0)	32	R/W	0000_0000h	13.3.17/711
2_1180	DMA n mode register (DMA1_MR1)	32	R/W	0800_0000h	13.3.1/693
2_1184	DMA n status register (DMA1_SR1)	32	R/W	0000_0000h	13.3.2/697
2_1188	DMA n current link descriptor extended address register (DMA1_ECLNDAR1)	32	R/W	0000_0000h	13.3.3/698
2_118C	DMA n current link descriptor address register (DMA1_CLNDAR1)	32	R/W	0000_0000h	13.3.4/699
2_1190	DMA n source attributes register (DMA1_SATR1)	32	R/W	0000_0000h	13.3.5/701
2_1194	DMA n source address register (DMA1_SAR1)	32	R/W	0000_0000h	13.3.6/702
2_1198	DMA n destination attributes register (DMA1_DATR1)	32	R/W	0000_0000h	13.3.7/703
2_119C	DMA n destination address register (DMA1_DAR1)	32	R/W	0000_0000h	13.3.8/704
2_11A0	DMA n byte count register (DMA1_BCR1)	32	R/W	0000_0000h	13.3.9/705
2_11A4	DMA n extended next link descriptor address register (DMA1_ENLNDAR1)	32	R/W	0000_0000h	13.3.10/705
2_11A8	DMA n next link descriptor address register (DMA1_NLNDAR1)	32	R/W	0000_0000h	13.3.11/706
2_11B0	DMA n extended current list descriptor address register (DMA1_ECLSDAR1)	32	R/W	0000_0000h	13.3.12/707
2_11B4	DMA n current list descriptor address register (DMA1_CLSDAR1)	32	R/W	0000_0000h	13.3.13/708
2_11B8	DMA n extended next list descriptor address register (DMA1_ENLS DAR1)	32	R/W	0000_0000h	13.3.14/709

Table continues on the next page...

DMA memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_11BC	DMA n next list descriptor address register (DMA1_NLSDAR1)	32	R/W	0000_0000h	13.3.15/709
2_11C0	DMA n source stride register (DMA1_SSR1)	32	R/W	0000_0000h	13.3.16/710
2_11C4	DMA n destination stride register (DMA1_DSR1)	32	R/W	0000_0000h	13.3.17/711
2_1200	DMA n mode register (DMA1_MR2)	32	R/W	0800_0000h	13.3.1/693
2_1204	DMA n status register (DMA1_SR2)	32	R/W	0000_0000h	13.3.2/697
2_1208	DMA n current link descriptor extended address register (DMA1_ECLNDAR2)	32	R/W	0000_0000h	13.3.3/698
2_120C	DMA n current link descriptor address register (DMA1_CLNDAR2)	32	R/W	0000_0000h	13.3.4/699
2_1210	DMA n source attributes register (DMA1_SATR2)	32	R/W	0000_0000h	13.3.5/701
2_1214	DMA n source address register (DMA1_SAR2)	32	R/W	0000_0000h	13.3.6/702
2_1218	DMA n destination attributes register (DMA1_DATR2)	32	R/W	0000_0000h	13.3.7/703
2_121C	DMA n destination address register (DMA1_DAR2)	32	R/W	0000_0000h	13.3.8/704
2_1220	DMA n byte count register (DMA1_BCR2)	32	R/W	0000_0000h	13.3.9/705
2_1224	DMA n extended next link descriptor address register (DMA1_ENLNDAR2)	32	R/W	0000_0000h	13.3.10/705
2_1228	DMA n next link descriptor address register (DMA1_NLNDAR2)	32	R/W	0000_0000h	13.3.11/706
2_1230	DMA n extended current list descriptor address register (DMA1_ECLSDAR2)	32	R/W	0000_0000h	13.3.12/707
2_1234	DMA n current list descriptor address register (DMA1_CLSDAR2)	32	R/W	0000_0000h	13.3.13/708
2_1238	DMA n extended next list descriptor address register (DMA1_ENLSDAR2)	32	R/W	0000_0000h	13.3.14/709
2_123C	DMA n next list descriptor address register (DMA1_NLSDAR2)	32	R/W	0000_0000h	13.3.15/709
2_1240	DMA n source stride register (DMA1_SSR2)	32	R/W	0000_0000h	13.3.16/710
2_1244	DMA n destination stride register (DMA1_DSR2)	32	R/W	0000_0000h	13.3.17/711
2_1280	DMA n mode register (DMA1_MR3)	32	R/W	0800_0000h	13.3.1/693
2_1284	DMA n status register (DMA1_SR3)	32	R/W	0000_0000h	13.3.2/697
2_1288	DMA n current link descriptor extended address register (DMA1_ECLNDAR3)	32	R/W	0000_0000h	13.3.3/698
2_128C	DMA n current link descriptor address register (DMA1_CLNDAR3)	32	R/W	0000_0000h	13.3.4/699
2_1290	DMA n source attributes register (DMA1_SATR3)	32	R/W	0000_0000h	13.3.5/701
2_1294	DMA n source address register (DMA1_SAR3)	32	R/W	0000_0000h	13.3.6/702
2_1298	DMA n destination attributes register (DMA1_DATR3)	32	R/W	0000_0000h	13.3.7/703

Table continues on the next page...

DMA memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_129C	DMA n destination address register (DMA1_DAR3)	32	R/W	0000_0000h	13.3.8/704
2_12A0	DMA n byte count register (DMA1_BCR3)	32	R/W	0000_0000h	13.3.9/705
2_12A4	DMA n extended next link descriptor address register (DMA1_ENLDAR3)	32	R/W	0000_0000h	13.3.10/705
2_12A8	DMA n next link descriptor address register (DMA1_NLDAR3)	32	R/W	0000_0000h	13.3.11/706
2_12B0	DMA n extended current list descriptor address register (DMA1_ECLSDAR3)	32	R/W	0000_0000h	13.3.12/707
2_12B4	DMA n current list descriptor address register (DMA1_CLSDAR3)	32	R/W	0000_0000h	13.3.13/708
2_12B8	DMA n extended next list descriptor address register (DMA1_ENLSDAR3)	32	R/W	0000_0000h	13.3.14/709
2_12BC	DMA n next list descriptor address register (DMA1_NLSDAR3)	32	R/W	0000_0000h	13.3.15/709
2_12C0	DMA n source stride register (DMA1_SSR3)	32	R/W	0000_0000h	13.3.16/710
2_12C4	DMA n destination stride register (DMA1_DSR3)	32	R/W	0000_0000h	13.3.17/711
2_1300	DMA general status register (DMA1_DGSR)	32	R	0000_0000h	13.3.18/712
C100	DMA n mode register (DMA2_MR0)	32	R/W	0800_0000h	13.3.1/693
C104	DMA n status register (DMA2_SR0)	32	R/W	0000_0000h	13.3.2/697
C108	DMA n current link descriptor extended address register (DMA2_ECLNDAR0)	32	R/W	0000_0000h	13.3.3/698
C10C	DMA n current link descriptor address register (DMA2_CLNDAR0)	32	R/W	0000_0000h	13.3.4/699
C110	DMA n source attributes register (DMA2_SATR0)	32	R/W	0000_0000h	13.3.5/701
C114	DMA n source address register (DMA2_SAR0)	32	R/W	0000_0000h	13.3.6/702
C118	DMA n destination attributes register (DMA2_DATR0)	32	R/W	0000_0000h	13.3.7/703
C11C	DMA n destination address register (DMA2_DAR0)	32	R/W	0000_0000h	13.3.8/704
C120	DMA n byte count register (DMA2_BCR0)	32	R/W	0000_0000h	13.3.9/705
C124	DMA n extended next link descriptor address register (DMA2_ENLDAR0)	32	R/W	0000_0000h	13.3.10/705
C128	DMA n next link descriptor address register (DMA2_NLDAR0)	32	R/W	0000_0000h	13.3.11/706
C130	DMA n extended current list descriptor address register (DMA2_ECLSDAR0)	32	R/W	0000_0000h	13.3.12/707
C134	DMA n current list descriptor address register (DMA2_CLSDAR0)	32	R/W	0000_0000h	13.3.13/708
C138	DMA n extended next list descriptor address register (DMA2_ENLSDAR0)	32	R/W	0000_0000h	13.3.14/709
C13C	DMA n next list descriptor address register (DMA2_NLSDAR0)	32	R/W	0000_0000h	13.3.15/709

Table continues on the next page...

DMA memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C140	DMA n source stride register (DMA2_SSR0)	32	R/W	0000_0000h	13.3.16/710
C144	DMA n destination stride register (DMA2_DSR0)	32	R/W	0000_0000h	13.3.17/711
C180	DMA n mode register (DMA2_MR1)	32	R/W	0800_0000h	13.3.1/693
C184	DMA n status register (DMA2_SR1)	32	R/W	0000_0000h	13.3.2/697
C188	DMA n current link descriptor extended address register (DMA2_ECLNDAR1)	32	R/W	0000_0000h	13.3.3/698
C18C	DMA n current link descriptor address register (DMA2_CLNDAR1)	32	R/W	0000_0000h	13.3.4/699
C190	DMA n source attributes register (DMA2_SATR1)	32	R/W	0000_0000h	13.3.5/701
C194	DMA n source address register (DMA2_SAR1)	32	R/W	0000_0000h	13.3.6/702
C198	DMA n destination attributes register (DMA2_DATR1)	32	R/W	0000_0000h	13.3.7/703
C19C	DMA n destination address register (DMA2_DAR1)	32	R/W	0000_0000h	13.3.8/704
C1A0	DMA n byte count register (DMA2_BCR1)	32	R/W	0000_0000h	13.3.9/705
C1A4	DMA n extended next link descriptor address register (DMA2_ENLNDAR1)	32	R/W	0000_0000h	13.3.10/705
C1A8	DMA n next link descriptor address register (DMA2_NLNDAR1)	32	R/W	0000_0000h	13.3.11/706
C1B0	DMA n extended current list descriptor address register (DMA2_ECLSDAR1)	32	R/W	0000_0000h	13.3.12/707
C1B4	DMA n current list descriptor address register (DMA2_CLSDAR1)	32	R/W	0000_0000h	13.3.13/708
C1B8	DMA n extended next list descriptor address register (DMA2_ENLSDAR1)	32	R/W	0000_0000h	13.3.14/709
C1BC	DMA n next list descriptor address register (DMA2_NLSDAR1)	32	R/W	0000_0000h	13.3.15/709
C1C0	DMA n source stride register (DMA2_SSR1)	32	R/W	0000_0000h	13.3.16/710
C1C4	DMA n destination stride register (DMA2_DSR1)	32	R/W	0000_0000h	13.3.17/711
C200	DMA n mode register (DMA2_MR2)	32	R/W	0800_0000h	13.3.1/693
C204	DMA n status register (DMA2_SR2)	32	R/W	0000_0000h	13.3.2/697
C208	DMA n current link descriptor extended address register (DMA2_ECLNDAR2)	32	R/W	0000_0000h	13.3.3/698
C20C	DMA n current link descriptor address register (DMA2_CLNDAR2)	32	R/W	0000_0000h	13.3.4/699
C210	DMA n source attributes register (DMA2_SATR2)	32	R/W	0000_0000h	13.3.5/701
C214	DMA n source address register (DMA2_SAR2)	32	R/W	0000_0000h	13.3.6/702
C218	DMA n destination attributes register (DMA2_DATR2)	32	R/W	0000_0000h	13.3.7/703
C21C	DMA n destination address register (DMA2_DAR2)	32	R/W	0000_0000h	13.3.8/704
C220	DMA n byte count register (DMA2_BCR2)	32	R/W	0000_0000h	13.3.9/705

Table continues on the next page...

DMA memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C224	DMA n extended next link descriptor address register (DMA2_ENLNDAR2)	32	R/W	0000_0000h	13.3.10/705
C228	DMA n next link descriptor address register (DMA2_NLNDAR2)	32	R/W	0000_0000h	13.3.11/706
C230	DMA n extended current list descriptor address register (DMA2_ECLSDAR2)	32	R/W	0000_0000h	13.3.12/707
C234	DMA n current list descriptor address register (DMA2_CLSDAR2)	32	R/W	0000_0000h	13.3.13/708
C238	DMA n extended next list descriptor address register (DMA2_ENLSDAR2)	32	R/W	0000_0000h	13.3.14/709
C23C	DMA n next list descriptor address register (DMA2_NLSDAR2)	32	R/W	0000_0000h	13.3.15/709
C240	DMA n source stride register (DMA2_SSR2)	32	R/W	0000_0000h	13.3.16/710
C244	DMA n destination stride register (DMA2_DSR2)	32	R/W	0000_0000h	13.3.17/711
C280	DMA n mode register (DMA2_MR3)	32	R/W	0800_0000h	13.3.1/693
C284	DMA n status register (DMA2_SR3)	32	R/W	0000_0000h	13.3.2/697
C288	DMA n current link descriptor extended address register (DMA2_ECLNDAR3)	32	R/W	0000_0000h	13.3.3/698
C28C	DMA n current link descriptor address register (DMA2_CLNDAR3)	32	R/W	0000_0000h	13.3.4/699
C290	DMA n source attributes register (DMA2_SATR3)	32	R/W	0000_0000h	13.3.5/701
C294	DMA n source address register (DMA2_SAR3)	32	R/W	0000_0000h	13.3.6/702
C298	DMA n destination attributes register (DMA2_DATR3)	32	R/W	0000_0000h	13.3.7/703
C29C	DMA n destination address register (DMA2_DAR3)	32	R/W	0000_0000h	13.3.8/704
C2A0	DMA n byte count register (DMA2_BCR3)	32	R/W	0000_0000h	13.3.9/705
C2A4	DMA n extended next link descriptor address register (DMA2_ENLNDAR3)	32	R/W	0000_0000h	13.3.10/705
C2A8	DMA n next link descriptor address register (DMA2_NLNDAR3)	32	R/W	0000_0000h	13.3.11/706
C2B0	DMA n extended current list descriptor address register (DMA2_ECLSDAR3)	32	R/W	0000_0000h	13.3.12/707
C2B4	DMA n current list descriptor address register (DMA2_CLSDAR3)	32	R/W	0000_0000h	13.3.13/708
C2B8	DMA n extended next list descriptor address register (DMA2_ENLSDAR3)	32	R/W	0000_0000h	13.3.14/709
C2BC	DMA n next list descriptor address register (DMA2_NLSDAR3)	32	R/W	0000_0000h	13.3.15/709
C2C0	DMA n source stride register (DMA2_SSR3)	32	R/W	0000_0000h	13.3.16/710
C2C4	DMA n destination stride register (DMA2_DSR3)	32	R/W	0000_0000h	13.3.17/711

Table continues on the next page...

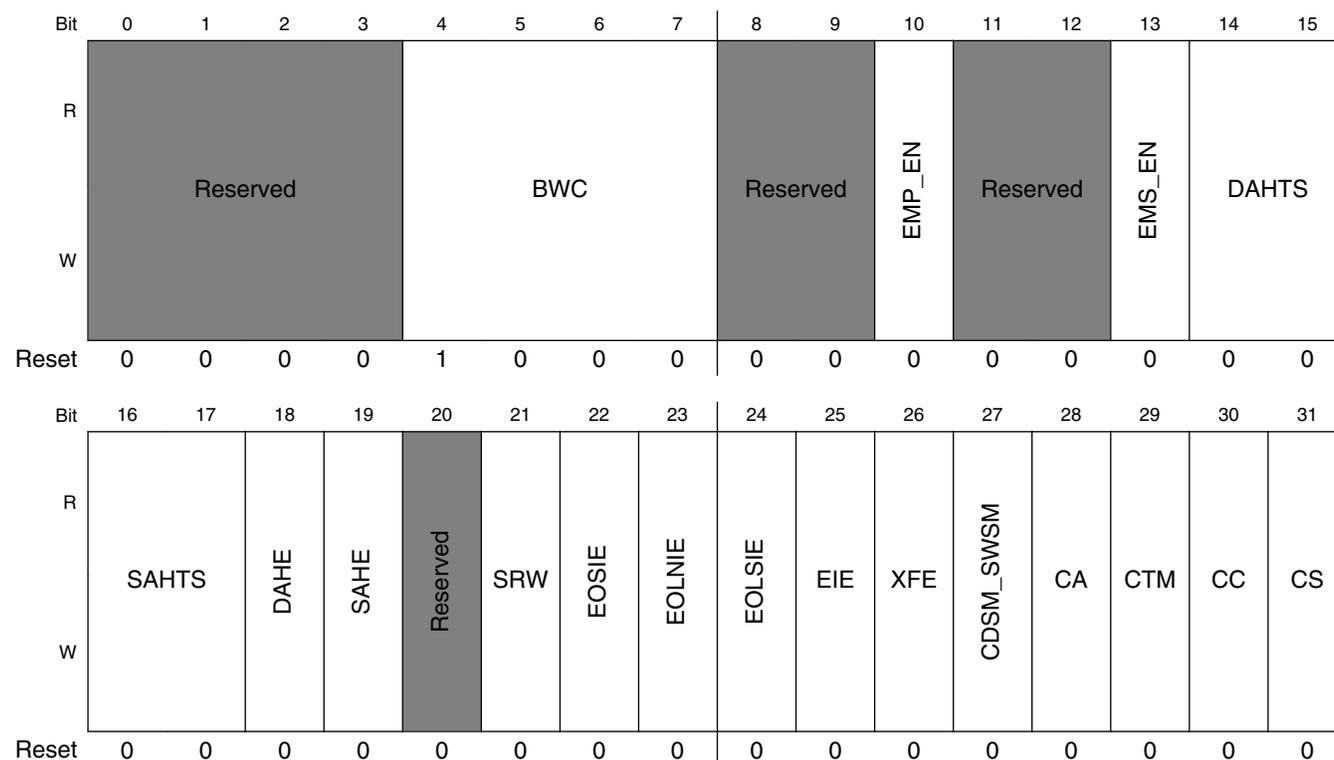
DMA memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C300	DMA general status register (DMA2_DGSR)	32	R	0000_0000h	13.3.18/ 712

13.3.1 DMA n mode register (DMAx_MRn)

The mode register allows software to start a DMA transfer and to control various DMA transfer characteristics.

Address: Base address + 100h offset + (128d × i), where i=0d to 3d



DMAx_MRn field descriptions

Field	Description
0–3 -	This field is reserved. Reserved
4–7 BWC	Bandwidth/pause control.

Table continues on the next page...

DMAx_MRn field descriptions (continued)

Field	Description
	<p>If multiple channels are executing transfers concurrently the value of MRn[BWC] determines how many bytes a given channel is allowed to transfer before the DMA engine pauses the current channel and switches to the next channel.</p> <p>If only one channel is executing transfers the value of MRn[BWC] dictates how many bytes are allowed to transfer before pausing the channel, after which a new assertion of DREQ resumes channel operation.</p> <p>0000 1 byte 0001 2 bytes 0010 4 bytes 0011 8 bytes 0100 16 bytes 0101 32 bytes 0110 64 bytes 0111 128 bytes 1000 256 bytes 1001 512 bytes 1010 1024 bytes 1011-1110 Reserved 1111 Disable bandwidth sharing to allow uninterrupted transfers from each channel.</p>
8–9 -	This field is reserved. Reserved
10 EMP_EN	External master pause enable. Valid only if MRn[EMS_EN] is set. 0 Disable the external master pause feature. 1 Enable the external master pause feature. Channel is paused as described by MRn[BWC].
11–12 -	This field is reserved. Reserved
13 EMS_EN	External master start enable. This bit does not apply to single-write start modes (direct or chaining). 0 Disable the channel from being started by an external DMA start pin. 1 Enable the channel to be started by an external DMA start pin, which sets MRn[CS].
14–15 DAHTS	Destination address hold transfer size. Indicates the transfer size used for each transaction while MRn[DAHE] is set. The byte count register must be in multiples of the size and the destination address register must be aligned based on the size. The transfer size assigned to MRn[DAHTS] must be equal to or smaller than that assigned to MRn[BWC] to avoid undefined behavior. 00 1 byte 01 2 bytes 10 4 bytes 11 8 bytes
16–17 SAHTS	Source address hold transfer size. Indicates the transfer size used for each transaction while MRn[SAHE] is set. The byte count register must be in multiples of the size and the source address register must be aligned based on the size. The transfer size assigned to MRn[SAHTS] must be equal to or smaller than that assigned to MRn[BWC] to avoid undefined behavior. 00 1 byte 01 2 bytes 10 4 bytes 11 8 bytes

Table continues on the next page...

DMAx_MRn field descriptions (continued)

Field	Description
18 DAHE	Destination address hold enable 0 Disable destination address hold 1 Enable the DMA controller to hold the destination address of a transfer to the size specified by MRn[DAHTS]. Hardware only supports aligned transfers for this feature.
19 SAHE	Source address hold enable 0 Disable source address hold 1 Enable the DMA controller to hold the source address of a transfer to the size specified by MRn[SAHTS]. Hardware only supports aligned transfers for this feature.
20 -	This field is reserved. Reserved
21 SRW	Single register write (Direct mode only; reserved for chaining mode.) 0 Normal operation 1 Enable a write to the source address register to simultaneously set MRn[CS], starting a DMA transfer, when MRn[CDSM_SWSM] is also set. Setting this bit and clearing CDSM_SWSM causes a write to the destination address register to simultaneously set MRn[CS], starting a DMA transfer.
22 EOSIE	End-of-segments interrupt enable 0 Do not generate an interrupt at the completion of a data transfer. CLNDARn[EOSIE] overrides this bit on a link descriptor basis. 1 Generate an interrupt at the completion of a data transfer (That is, SRn[EOSI] is set). This bit overrides the CLNDARn[EOSIE].
23 EOLNIE	End-of-links interrupt enable 0 Do not generate an interrupt at the completion of a list of DMA transfers. 1 Generate an interrupt at the completion of a list of DMA transfers (That is, NLNDARn[EOLND] is set).
24 EOLSIE	End-of-lists interrupt enable 0 Do not generate an interrupt at the completion of all DMA transfers. 1 Generate an interrupt at the completion of all DMA transfers (That is, NLNDARn[EOLND] and NLSDARn[EOLSD] are set).
25 EIE	Error interrupt enable 0 Do not generate an interrupt if a programming or transfer error is detected. 1 Generate an interrupt if a programming or transfer error is detected.
26 XFE	Extended features enable 0 Disable striding feature in direct mode or disable list chaining feature in chaining mode. 1 Disable striding feature in direct mode or disable list chaining feature in chaining mode.
27 CDSM_SWSM	In chaining mode, current descriptor start mode/single-write start mode is as follows: <ul style="list-style-type: none"> • In basic mode (MRn[XFE] is cleared), setting this bit causes a write to the current link descriptor address register to simultaneously set MRn[CS], starting a DMA transfer. • In extended chaining mode (MRn[XFE] is set), setting this bit causes a write to the current list descriptor address register to simultaneously set MRn[CS], starting a DMA transfer. In direct mode, setting this bit and MRn[SRW] causes a write to the source address register to simultaneously set MRn[CS], starting a DMA transfer. Clearing this bit and setting MRn[SRW] causes a write to the destination address register to simultaneously set MRn[CS], starting a DMA transfer. This bit must be cleared when MRn[SRW] is cleared.

Table continues on the next page...

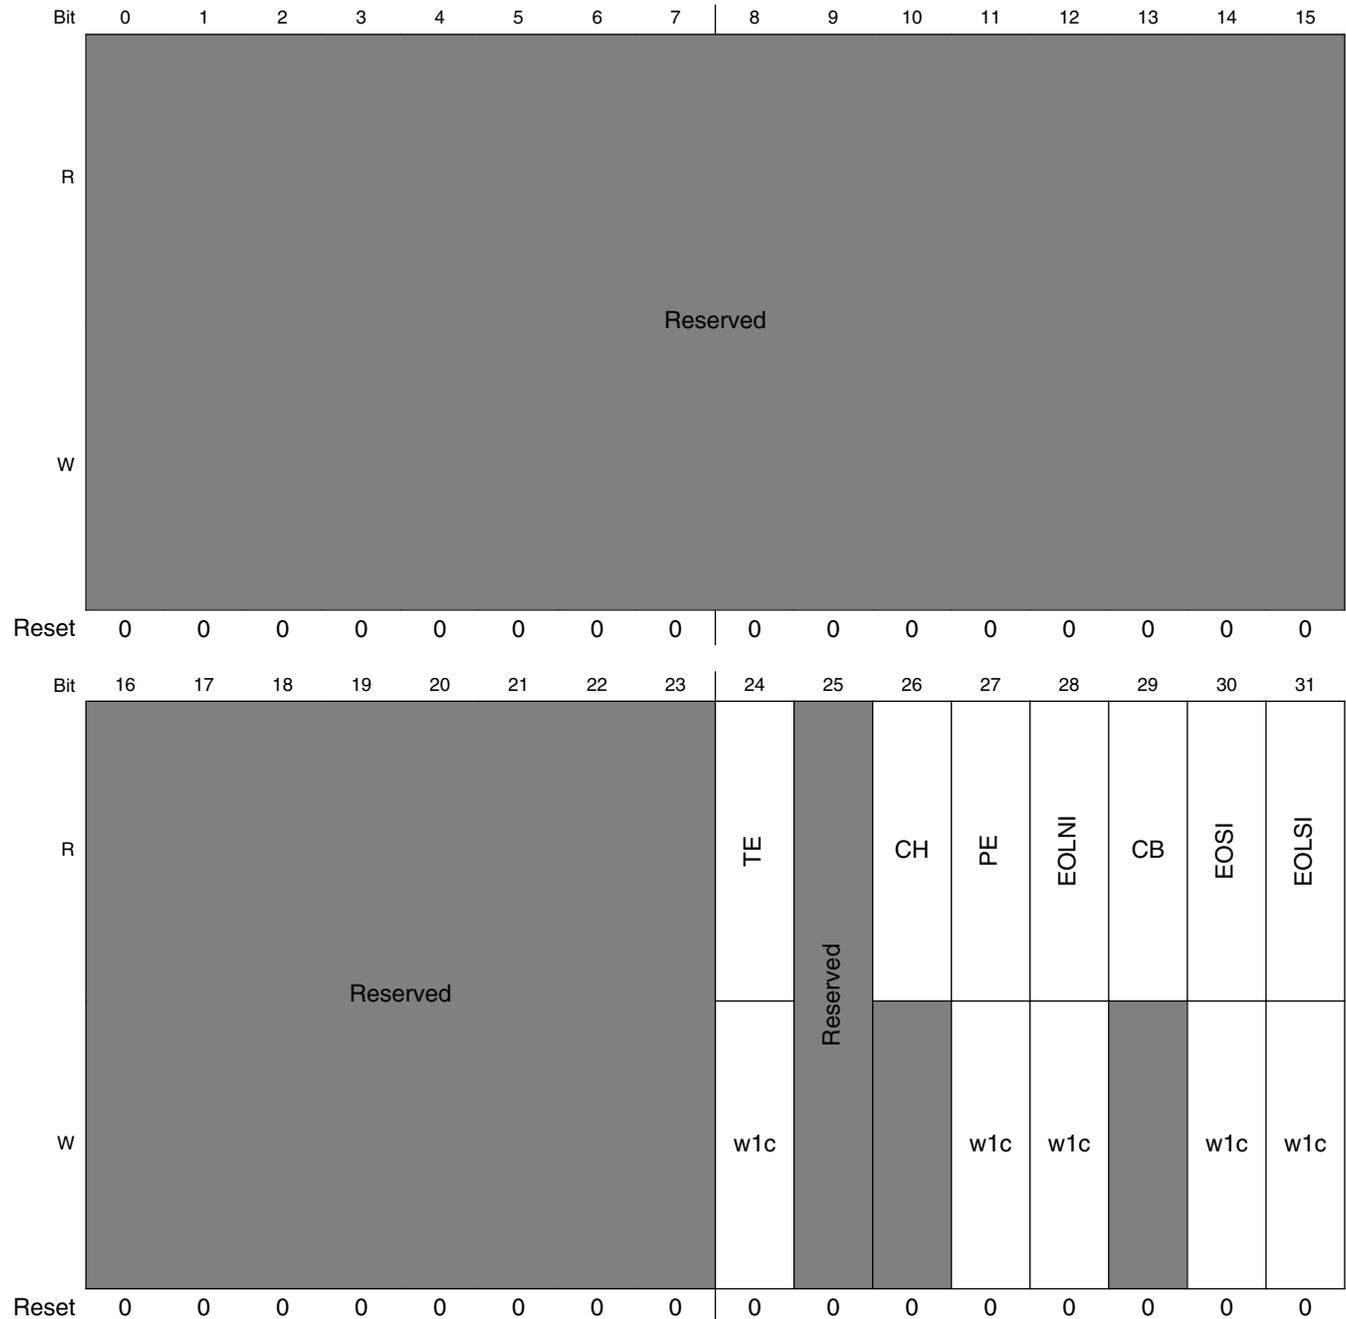
DMAx_MRn field descriptions (continued)

Field	Description
28 CA	Channel abort 0 No effect 1 Cause the current transfer to be aborted and SRn[CB] to be cleared if the channel is busy. The channel remains in the idle state until a new transfer is programmed.
29 CTM	Channel transfer mode 0 Configure the channel in chaining mode. 1 Configure the channel into direct mode. This means that software is responsible for placing all the required parameters into necessary registers to start the DMA process.
30 CC	Channel continue. This bit applies only to chaining mode and is cleared by hardware after the first descriptor read when continuing a transfer. This bit is reserved for external master mode. 0 No effect 1 The DMA transfer restarts the transferring process starting at the current descriptor address.
31 CS	Channel start. This bit is also automatically set by hardware during single-write start mode and external master start enable mode. Note that in external control mode, deasserting DMA_DREQ does NOT clear this bit. 0 Halt the DMA process if channel is busy (SRn[CB] is set). No effect if the channel is not busy. 1 Start the DMA process if channel is not busy (CB is cleared). If the channel was halted (CS = 0 and CB = 1), the transfer continues from the point at which it was halted.

13.3.2 DMA n status register (DMAx_SRn)

The status registers report various DMA conditions during and after a DMA transfer.

Address: Base address + 104h offset + (128d × i), where i=0d to 3d



DMAx_SRn field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24 TE	Transfer error (Bit reset, write 1 to clear) 0 No error condition during the DMA transfer 1 Error condition during the DMA transfer. See DMA errors , for additional information.
25 -	This field is reserved. Reserved
26 CH	Channel halted. Cleared automatically by hardware if MR n [CS] is set again for resuming a halted transfer 0 Channel is not halted. If software attempts to halt an idle channel (SRn[CB] is cleared), this bit remains 0. 1 DMA transfer was successfully halted by software and can be resumed.
27 PE	Programming error (bit reset, write 1 to clear) 0 No programming error detected 1 A programming error is detected that prevents the DMA transfer from occurring.
28 EOLNI	End-of-links interrupt. After transferring the last block of data in the last link descriptor, if MRn[EOLNIE] is set, then this bit is set and an interrupt is generated. (Bit reset, write 1 to clear)
29 CB	Channel busy 0 DMA transfer is finished, an error occurred, or a channel abort occurred. 1 DMA transfer is currently in progress.
30 EOSI	End-of-segment interrupt. In chaining mode, after finishing a data transfer, if MRn[EOSIE] is set or if CLNDARn[EOSIE] is set, this bit gets set and an interrupt is generated. In direct mode, if MRn[EOSIE] is set, this bit gets set and an interrupt is generated. (Bit reset, write 1 to clear)
31 EOLSI	End-of-list interrupt. After transferring the last block of data in the last list descriptor, if MRn[EOLSIE] is set, then this bit is set and an interrupt is generated. (Bit reset, write 1 to clear)

13.3.3 DMA n current link descriptor extended address register (DMAx_ECLNDARn)

These registers contain the upper 4 bits of the address of the current link descriptor. See [DMA n current link descriptor address register \(DMA_CLNDARn\)](#) for a description of basic chaining mode.

Address: Base address + 108h offset + (128d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															ECLNDA																
W	Reserved															ECLNDA																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

DMA_x_ECLNDAR_n field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 ECLNDA	Current link descriptor extended address (upper 4 bits of 36-bit address)

13.3.4 DMA *n* current link descriptor address register (DMA_x_CLNDAR_n)

Current link descriptor address registers contain the address of the current link descriptor. For devices with 36-bit addressing, the upper 4 bits of the address are described in [DMA *n* current link descriptor extended address register \(DMA_x_ECLNDAR_n\)](#).

In basic chaining mode, shown in the figure below, software must initialize these registers to point to the first link descriptors in memory.

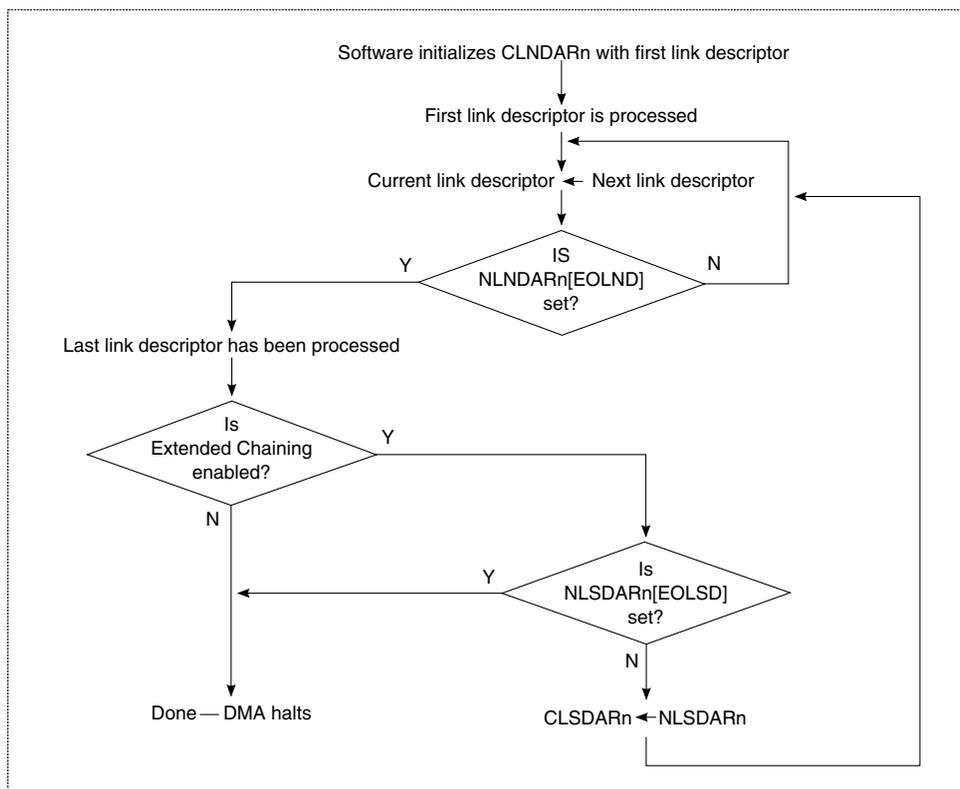


Figure 13-102. Basic Chaining Mode Flow Chart

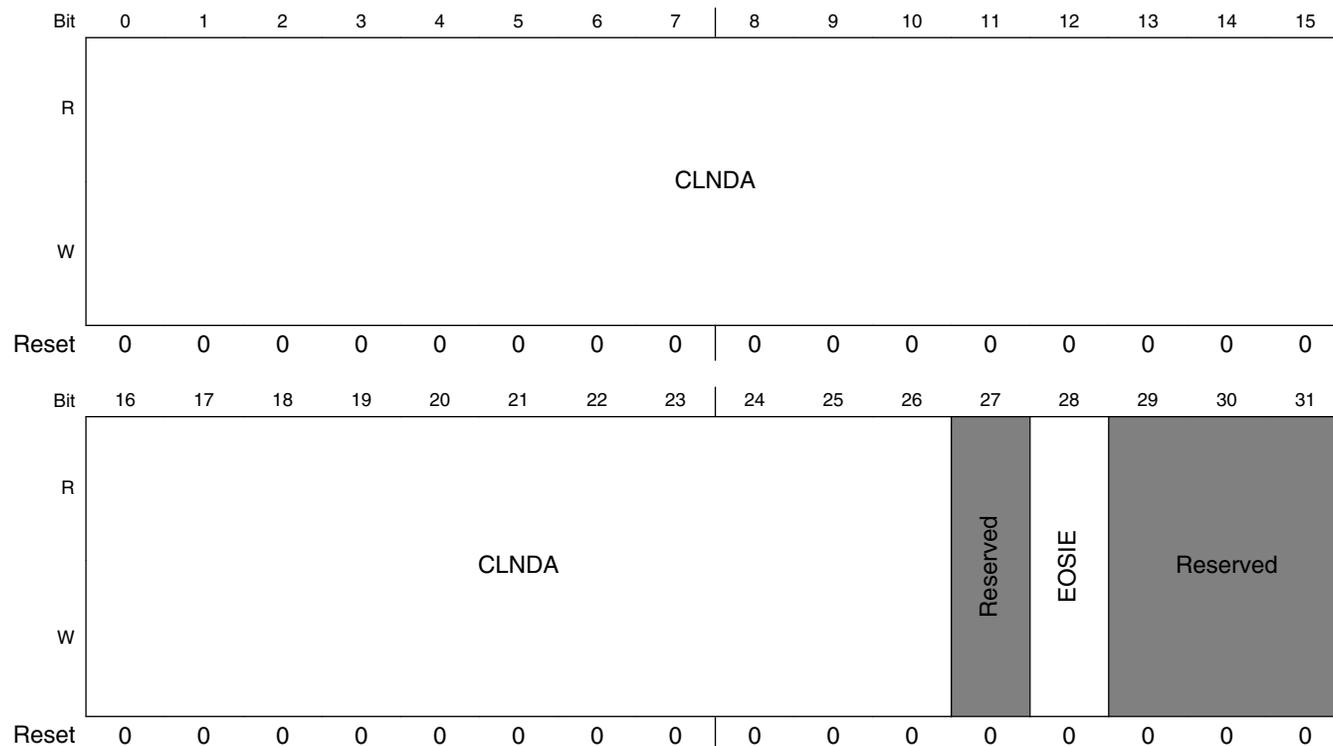
After the current descriptor is processed, the current link descriptor address register is loaded from the next link descriptor address registers and NLNDAR_n [EOLND] in the next link descriptor address register is examined. If EOLND is zero, the DMA controller

DMA controller memory map

reads in the new current link descriptor for processing. If EOLND is set, the last descriptor of the list was just completed. If extended chaining mode is not enabled, all DMA transfers are complete and the DMA controller halts.

If extended chaining mode is enabled, the DMA controller examines the state of NLSDAR n [EOLSD] in the next list descriptor address register. If EOLSD is clear, the controller loads the contents of the next list descriptor address register into the current list descriptor address register and reads the new list descriptor from memory. If EOLSD is set, all DMA transfers are complete and the DMA controller halts.

Address: Base address + 10Ch offset + (128d × i), where i=0d to 3d



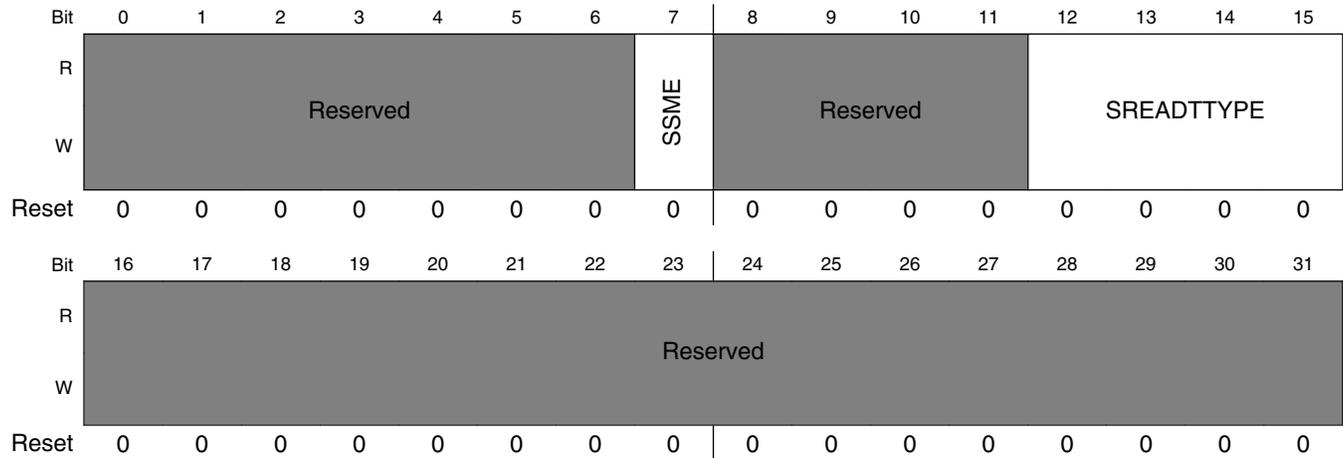
DMAx_CLNDARn field descriptions

Field	Description
0–26 CLNDA	Current link descriptor address. Contains the current descriptor address of the buffer descriptor in memory. The descriptor must be aligned to a 32-byte boundary. (This is the lower portion of the 36-bit address formed by CLNDARn[CLNDA] and ECLNDARn[ECLNDA].)
27 -	This field is reserved. Reserved
28 EOSIE	End-of-segment interrupt enable 0 Do not generate an interrupt upon completion of the current DMA transfer for the current link descriptor. 1 Generate an interrupt upon completion of the current DMA transfer for the current link descriptor.
29–31 -	This field is reserved. Reserved

13.3.5 DMA n source attributes register (DMAx_SATRn)

The source attributes registers contain the transaction attributes to be used for the DMA operation. Stride mode is enabled by setting SATR *n* [SSME].

Address: Base address + 110h offset + (128d × i), where i=0d to 3d



DMAx_SATRn field descriptions

Field	Description
0–6 -	This field is reserved. Reserved
7 SSME	Source stride mode enable Ignored in basic mode (MRn[XFE] is cleared). Striding on the source address can be accomplished by enabling SATRn[SSME] and setting the desired stride size and distance in the SSRn. 0 Stride mode disabled 1 Stride mode enabled
8–11 -	This field is reserved. Reserved
12–15 SREADTTYPE	DMA source transaction type. Reserved values result in a programming error being detected and logged in SR[PE]. Transaction type to run on local address space . 0000-0011 Reserved 0100 Read, do not snoop local processor 0101 Read, snoop local processor 0110-1111 Reserved
16–31 -	This field is reserved. Reserved

13.3.6 DMA n source address register (DMAx_SARn)

The source address registers contain the address from which the DMA controller reads data. In direct mode, if MR n [CDSM_SWSM] and MR n [SRW] are set, a write to this register simultaneously sets MR n [CS], starting a DMA transfer. Software must ensure that this is a valid address.

Address: Base address + 114h offset + (128d × i), where i=0d to 3d



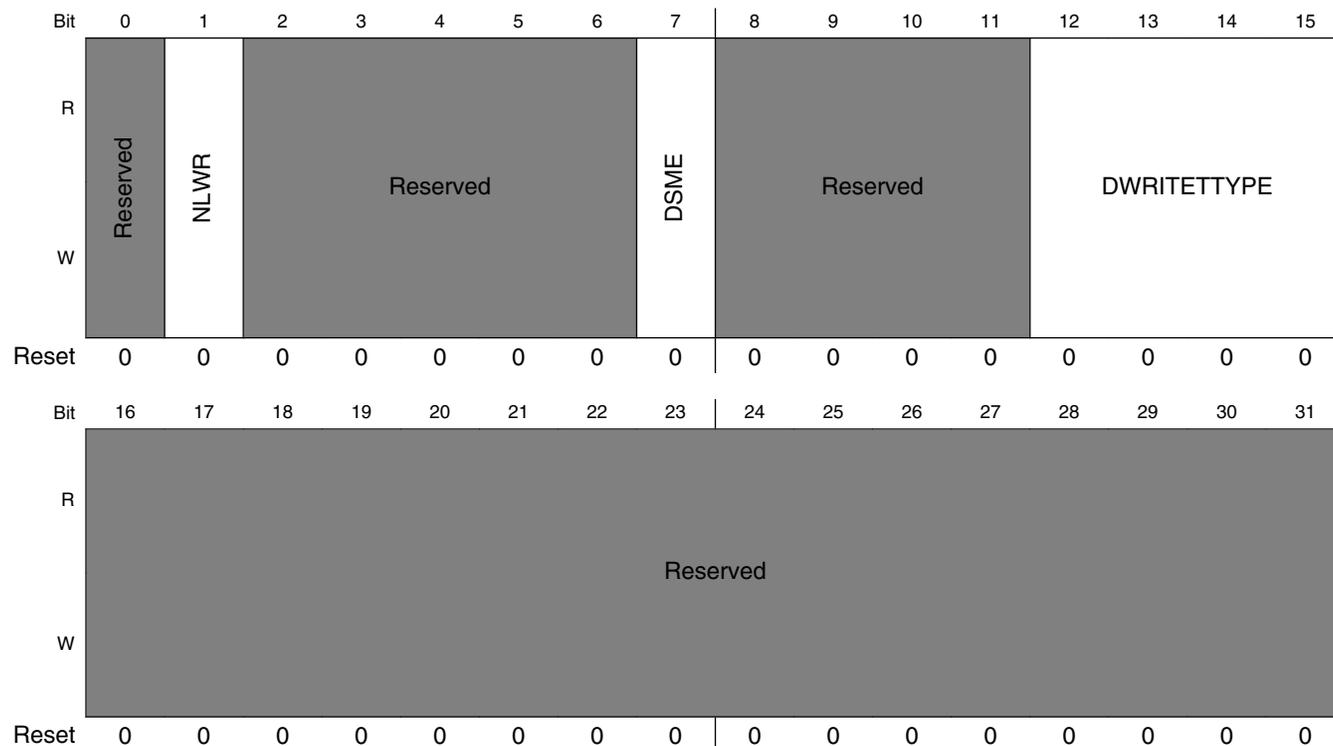
DMAx_SARn field descriptions

Field	Description
0–31 SAD	Source address. This register contains the low-order bits of the 36-bit source address of the DMA transfer. The contents are updated after every DMA write operation unless the final stride of a striding operation is less than the stride size, in which case it remains equal to the address from which the last stride began.

13.3.7 DMA n destination attributes register (DMAx_DATRn)

The destination attributes registers contain the transaction attributes for the DMA operation. Stride mode is enabled by setting DATR n [DSME].

Address: Base address + 118h offset + (128d × i), where i=0d to 3d



DMAx_DATRn field descriptions

Field	Description
0 -	This field is reserved. Reserved
1 NLWR	No last write with response. Performance impact: <ul style="list-style-type: none"> When NLWR is cleared only one "write-with-response" transaction can be processed/in-flight at a time, which prevents pipelining of "write-with-response" transactions. When NLWR is set "write-without-response" transactions can be issued in a pipelined manner not being stalled by a write response. Error impact: <ul style="list-style-type: none"> When NLWR is cleared the returned response includes error status which is used to update the SR[TE] field. When NLWR is set, since no response is returned, the SR[TE] field is not updated with write transaction error status.

Table continues on the next page...

DMAx_DATRn field descriptions (continued)

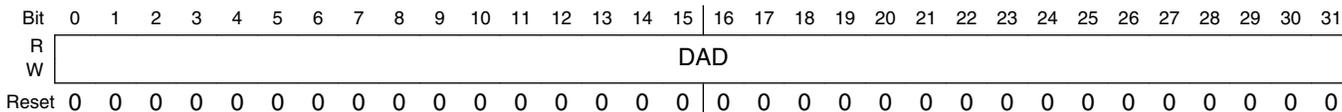
Field	Description
	<p>0 "Write-with-response" transactions are generated for the last transaction in a multi-transaction transfer. Single writes will be NWRITE_R. A response is returned from the module targeted by the write transfer (ex. DDR controller, eLBC, PCI-Express,...) informing the DMA of write completion and error status.</p> <p>1 "Write-without-response" transactions are generated for the last transaction in a multi-transaction transfer. The DMA considers a transfer to be complete when the last write transaction is issued (not necessarily when the transaction has reached the target).</p>
2–6 -	This field is reserved. Reserved
7 DSME	<p>Destination stride mode enable</p> <p>Ignored in basic mode (MRn[XFE] is cleared). Striding on the destination address can be accomplished by setting DSME and setting the desired stride size and distance in DSRn.</p> <p>0 Stride mode disabled 1 Stride mode enabled</p>
8–11 -	This field is reserved. Reserved
12–15 DWRITETYPE	<p>DMA destination transaction type. Reserved values result in a programming error being detected and logged in SR[PE].</p> <p>Transaction type to run on local address space</p> <p>0000-0011 Reserved 0100 Write, do not snoop local processor 0101 Write, snoop local processor 0110 Reserved 0111 Reserved 1000-1111 Reserved</p>
16–31 -	This field is reserved. Reserved

13.3.8 DMA n destination address register (DMAx_DARn)

The destination address registers contain the addresses to which the DMA controller writes data.

In direct mode, if MR n [SRW] is set and MR n [CDSM_SWSM] is cleared, a write to this register simultaneously sets MR n [CS], starting a DMA transfer. Software must ensure that this is a valid address.

Address: Base address + 11Ch offset + (128d × i), where i=0d to 3d



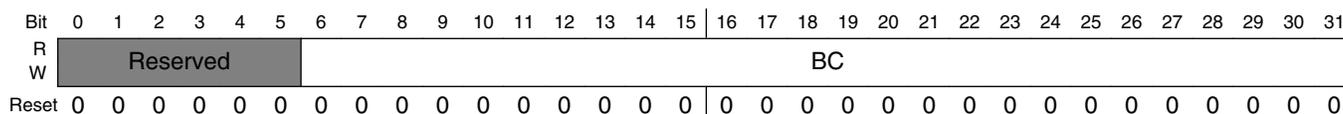
DMA_x_DAR_n field descriptions

Field	Description
0–31 DAD	Destination address. This register contains the low-order bits of the 36-bit destination address of the DMA transfer. The contents are updated after every DMA write operation unless the final stride of a striding operation is less than the stride size, in which case it remains equal to the address from which the last stride began.

13.3.9 DMA n byte count register (DMA_x_BCR_n)

The byte count register contains the number of bytes to transfer.

Address: Base address + 120h offset + (128d × i), where i=0d to 3d



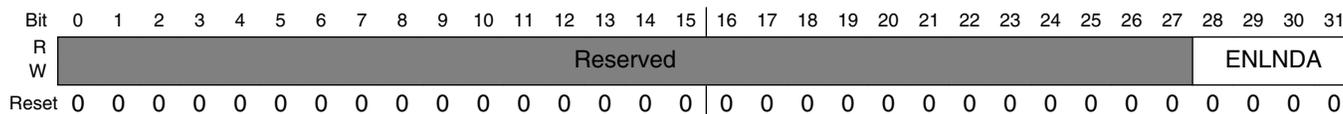
DMA_x_BCR_n field descriptions

Field	Description
0–5 -	This field is reserved. Reserved
6–31 BC	Byte count. Contains the number of bytes to transfer. The value in this register is decremented after each DMA read operation. The maximum transfer size is (2 ²⁶) - 1 bytes.

13.3.10 DMA n extended next link descriptor address register (DMA_x_ENLNDAR_n)

These registers contain the upper 4 bits of the address of the next link descriptor in memory. See [DMA n next link descriptor address register \(DMA_NLNDAR_n\)](#).

Address: Base address + 124h offset + (128d × i), where i=0d to 3d



DMAx_ENLNDARn field descriptions

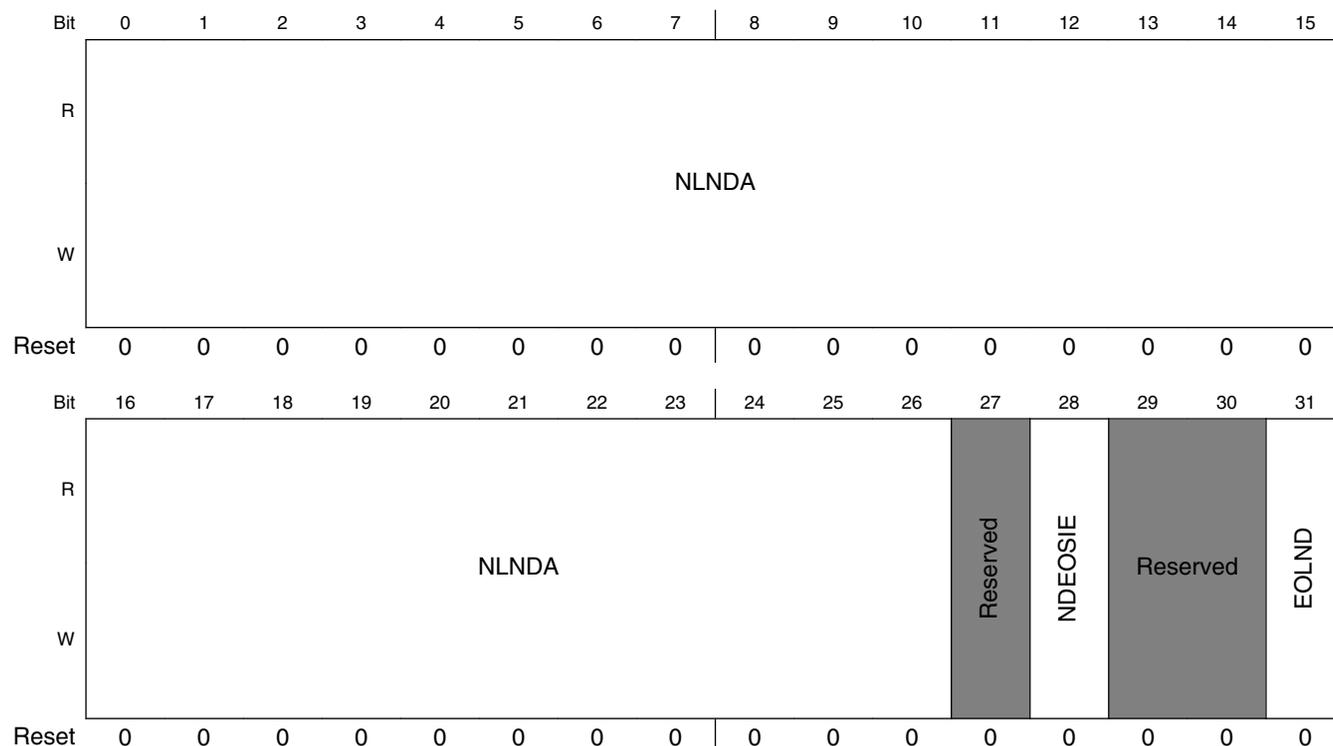
Field	Description
0–27 -	This field is reserved. Reserved
28–31 ENLNDAR	Next link descriptor extended address bits (upper 4 bits of 36-bit address)

13.3.11 DMA n next link descriptor address register (DMAx_NLNDARn)

Current link descriptor address registers contain the address for the next link descriptor in memory. For devices with 36-bit addressing, the upper 4 bits of the address are described in [DMA n extended next link descriptor address register \(DMAx_ENLNDARn\)](#).

Contents transferred to the current descriptor address registers become effective for the current transfer in basic and extended chaining modes.

Address: Base address + 128h offset + (128d × i), where i=0d to 3d



DMAx_NLNDARn field descriptions

Field	Description
0–26 NLNDA	Next link descriptor address. Contains the next link descriptor address in memory. The descriptor must be aligned to a 32-byte boundary.

Table continues on the next page...

DMAx_NLNDARn field descriptions (continued)

Field	Description
27 -	This field is reserved. Reserved
28 NDEOSIE	Next descriptor end-of-segment interrupt enable 0 Do not generate an interrupt if the current DMA transfer for the current descriptor is finished. 1 Generate an interrupt if the current DMA transfer for the current descriptor is finished.
29–30 -	This field is reserved. Reserved
31 EOLND	End-of-links descriptor. This bit is ignored in direct mode. 0 This descriptor is not the last link descriptor in memory for this list. 1 This descriptor is the last link descriptor in memory for this list. If this bit is set, the DMA controller advances to the next list descriptor in memory if NLSDARn[EOLSD] is also set in extended mode.

13.3.12 DMA n extended current list descriptor address register (DMAx_ECLSDARn)

These registers contain the upper 4 bits of the current address of the list descriptor in memory in extended chaining mode. See [DMA n current list descriptor address register \(DMA_CLSDARn\)](#).

Address: Base address + 130h offset + (128d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															ECLSDA																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

DMAx_ECLSDARn field descriptions

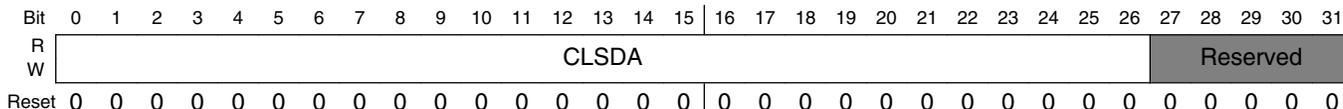
Field	Description
0–27 -	This field is reserved. Reserved
28–31 ECLSDA	Current list descriptor extended address bits (upper 4 bits of 36-bit address)

13.3.13 DMA n current list descriptor address register (DMAx_CLSDARn)

The current list descriptor address registers contain the current address of the list descriptor in memory in extended chaining mode. For chips with 36-bit addressing, the upper 4 bits of the address are described in [DMA n extended current list descriptor address register \(DMA_ECLSDARn\)](#).

In extended chaining mode, software must initialize CLS DAR *n* and ECLSDAR *n* to point to the first list descriptor in memory. After finishing the last link descriptor in the current list, the DMA controller loads the contents of the next list descriptor address register into the current list descriptor address register. If NLSDAR *n* [EOLSD] in the next list descriptor address register is clear, the DMA controller reads the new current list descriptor from memory to process that list. If EOLSD in the next list descriptor address register is set and the last link in the current list is finished, all DMA transfers are complete.

Address: Base address + 134h offset + (128d × i), where i=0d to 3d



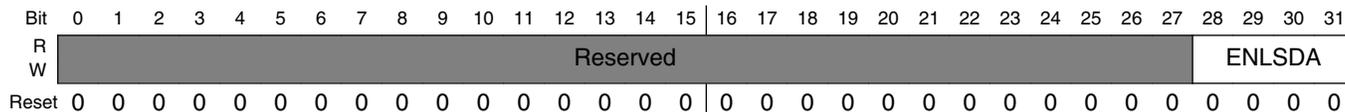
DMAx_CLSDARn field descriptions

Field	Description
0–26 CLSDA	Current list descriptor address. Contains the low-order bits of the 36-bit current list descriptor address of the buffer descriptor in memory in extended chaining mode. The descriptor must be aligned to a 32-byte boundary.
27–31 -	This field is reserved. Reserved

13.3.14 DMA n extended next list descriptor address register (DMAx_ENLSDARn)

These registers contain the upper 4 bits of the address of the next list descriptor in memory. See [DMA n next list descriptor address register \(DMA_NLSDARn\)](#).

Address: Base address + 138h offset + (128d × i), where i=0d to 3d



DMAx_ENLSDARn field descriptions

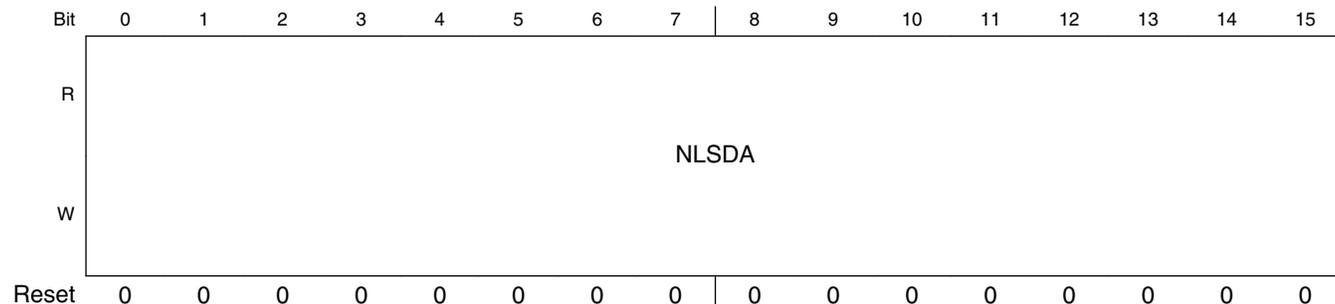
Field	Description
0–27 -	This field is reserved. Reserved
28–31 ENLSDA	Next list descriptor extended address bits (upper 4 bits of 36-bit address)

13.3.15 DMA n next list descriptor address register (DMAx_NLSDARn)

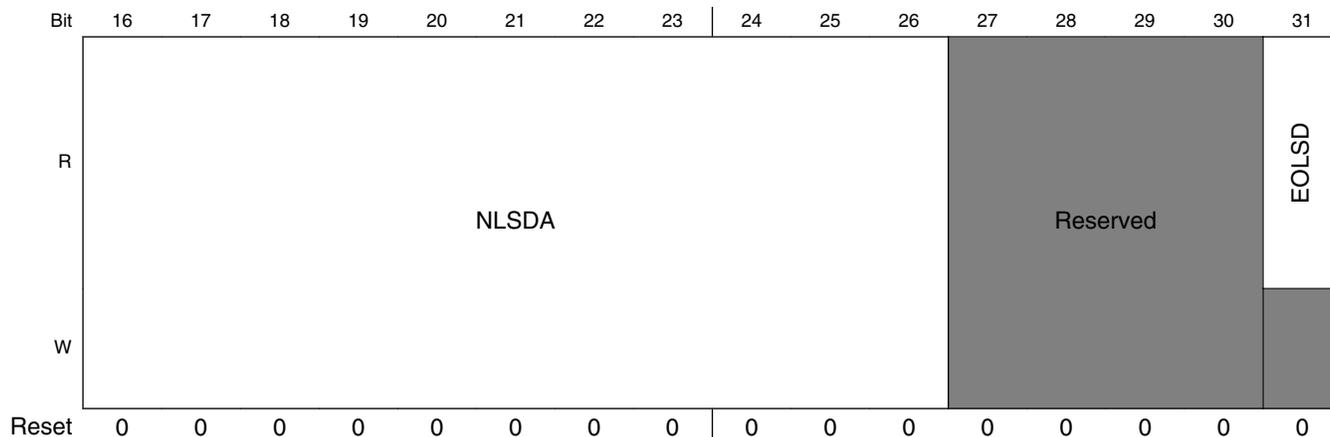
The next list descriptor address registers contain the address for the next list descriptor in memory. For devices with 36-bit addressing, the upper 4 bits of the address are described in [DMA n extended next list descriptor address register \(DMA_ENLSDARn\)](#).

If the contents are transferred to the current list descriptor address register, they become effective for the current transfer in extended chaining mode.

Address: Base address + 13Ch offset + (128d × i), where i=0d to 3d



DMA controller memory map



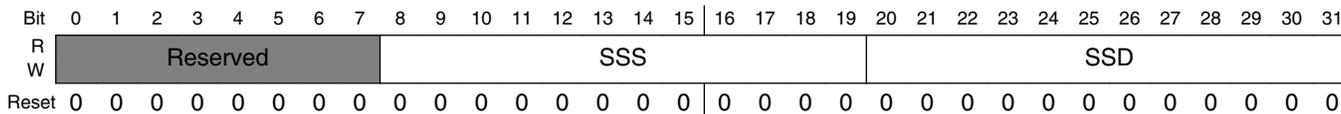
DMAx_NLSDARn field descriptions

Field	Description
0–26 NLSDA	Next list descriptor address. Contains the low-order bits of the 36-bit next descriptor address of the buffer descriptor in memory. The descriptor must be aligned on a 32-byte boundary.
27–30 -	This field is reserved. Reserved
31 EOLSD	End-of-lists descriptor. This bit is ignored in direct mode. 0 This list descriptor is not the last list descriptor in memory. 1 This list descriptor is the last list descriptor in memory. If this bit is set, then the DMA controller halts after the last link descriptor transaction is finished.

13.3.16 DMA n source stride register (DMAx_SSRn)

The source stride register contains the stride size and distance. Note that the source stride information is loaded when a new list descriptor is read from memory. Therefore, the source stride register is applicable for all link descriptors in the new list. Changing the source stride information for a link requires that a new list be generated.

Address: Base address + 140h offset + (128d × i), where i=0d to 3d



DMAx_SSRn field descriptions

Field	Description
0–7 -	This field is reserved. Reserved

Table continues on the next page...

DMA_x_SSR_n field descriptions (continued)

Field	Description
8–19 SSS	Source stride size. Number of bytes to transfer before jumping to the next address as specified in the source stride distance field.
20–31 SSD	Source stride distance. The source stride distance in bytes from start byte to start byte.

13.3.17 DMA n destination stride register (DMA_x_DSR_n)

The destination stride register contains the stride size, and distance. Note that the destination stride information is loaded when a new list descriptor is read from memory. Therefore, the destination stride register is applicable for all link descriptors in the new list. Changing the destination stride information for a link requires that a new list be generated.

Address: Base address + 144h offset + (128d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

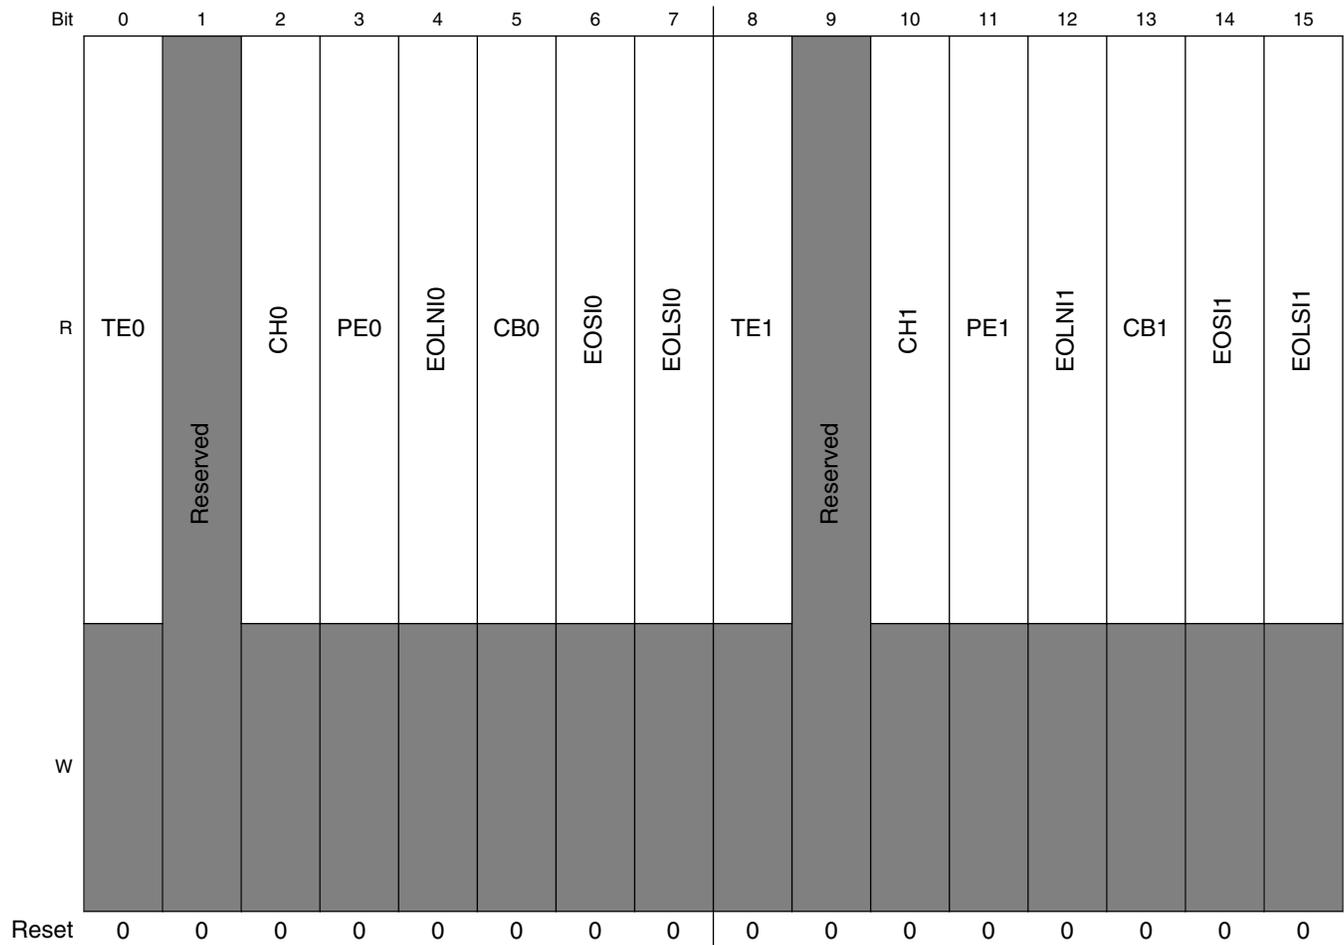
DMA_x_DSR_n field descriptions

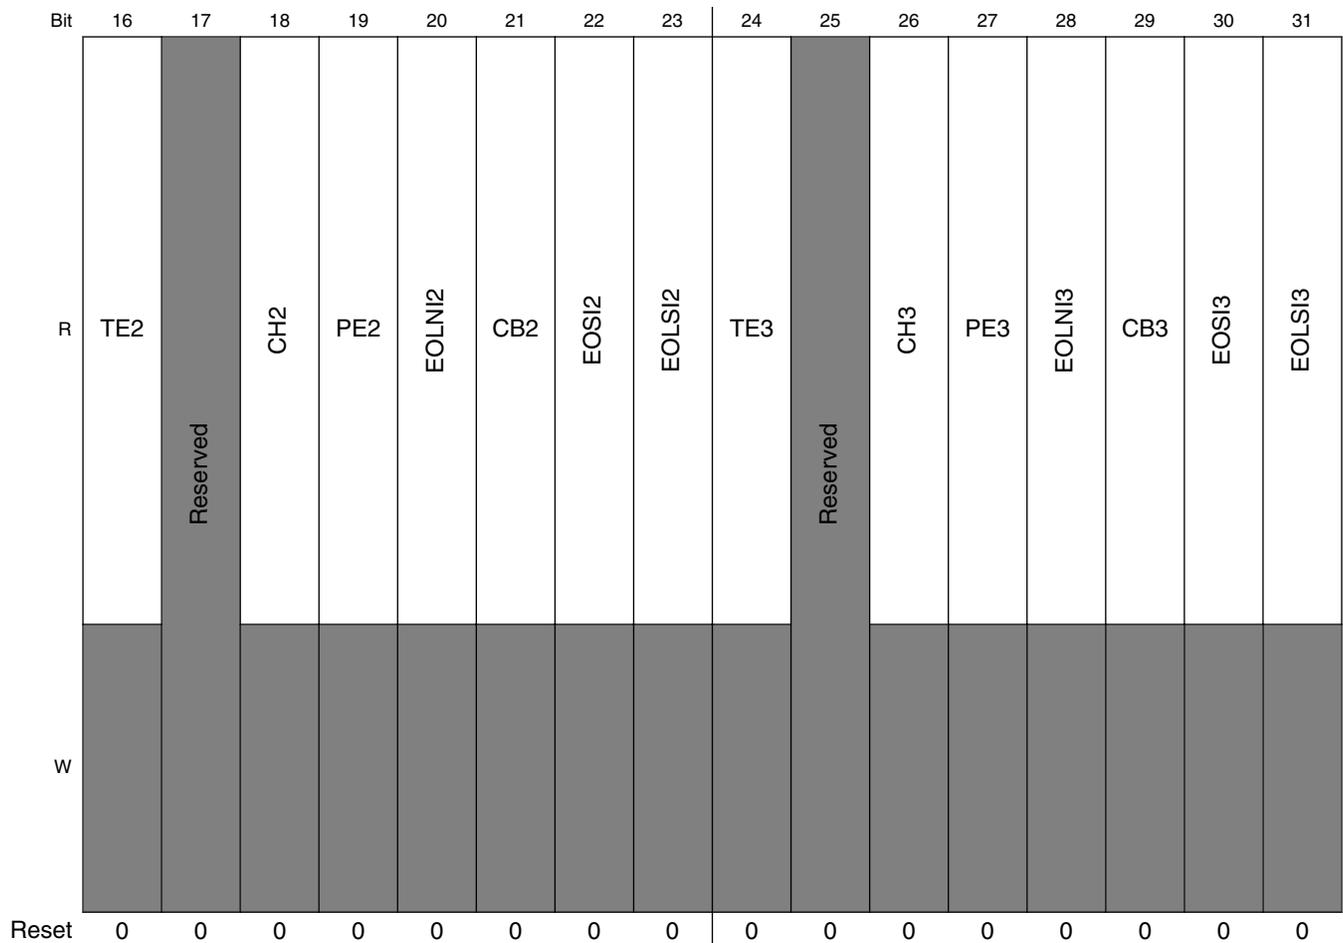
Field	Description
0–7 -	This field is reserved. Reserved
8–19 DSS	Destination stride size. Number of bytes to transfer before jumping to the next address as specified in the destination stride distance field.
20–31 DSD	Destination stride distance. The destination stride distance in bytes from start byte to start byte.

13.3.18 DMA general status register (DMAx_DGSR)

The DMA general status register combines all of the status bits each channel into one register. This register is read-only.

Address: Base address + 300h offset





DMAx_DGSR field descriptions

Field	Description
0 TE0	Transfer error, channel 0 0 Normal operation 1 An error condition occurred during the DMA transfer.
1 -	This field is reserved. Reserved
2 CH0	Channel halted, channel 0
3 PE0	Programming error, channel 0
4 EOLNI0	End-of-links interrupt, channel 0
5 CB0	Channel busy, channel 0
6 EOSI0	End-of-segment interrupt, channel 0
7 EOLSI0	End-of-lists/direct interrupt, channel 0

Table continues on the next page...

DMAx_DGSR field descriptions (continued)

Field	Description
8 TE1	Transfer error, channel 1 0 Normal operation 1 An error condition occurred during the DMA transfer.
9 -	This field is reserved. Reserved
10 CH1	Channel halted, channel 1
11 PE1	Programming error, channel 1
12 EOLNI1	End-of-links interrupt, channel 1
13 CB1	Channel busy, channel 1
14 EOSI1	End-of-segment interrupt, channel 1
15 EOLSI1	End-of-lists/direct interrupt, channel 1
16 TE2	Transfer error, channel 2 0 Normal operation 1 An error condition occurred during the DMA transfer.
17 -	This field is reserved. Reserved
18 CH2	Channel halted, channel 2
19 PE2	Programming error, channel 2
20 EOLNI2	End-of-links interrupt, channel 2
21 CB2	Channel busy, channel 2
22 EOSI2	End-of-segment interrupt, channel 2
23 EOLSI2	End-of-lists/direct interrupt, channel 2
24 TE3	Transfer error, channel 3 0 Normal operation 1 An error condition occurred during the DMA transfer.
25 -	This field is reserved. Reserved
26 CH3	Channel halted, channel 3
27 PE3	Programming error, channel 3

Table continues on the next page...

DMAx_DGSR field descriptions (continued)

Field	Description
28 EOLNI3	End-of-links interrupt, channel 3
29 CB3	Channel busy, channel 3
30 EOSI3	End-of-segment interrupt, channel 3
31 EOLSI3	End-of-lists/direct interrupt, channel 3

13.4 DMA functional description

This section describes the function of the DMA controller.

13.4.1 DMA channel operation

All DMA channels support two different modes of operation: a basic mode ($MRn[XFE]$ is cleared) and an extended mode ($MRn[XFE]$ is set). In both modes, a channel can be activated by clearing and setting $MRn[CS]$, or through the single-write start mode using $MRn[CDSM_SWSM]$ and $MRn[SRW]$, or through an external control mode using $MRn[EMS_EN]$ and the external DMA_DREQ (DMA1_DREQ, DMA2_DREQ0, and DMA2_DREQ1) signal.

In basic mode, the channel can be programmed in basic direct mode or basic chaining mode. In extended mode, the channel can be programmed in extended direct mode or extended chaining mode. Extended mode provides more capabilities, such as extended descriptor chaining, striding capabilities, and a more flexible descriptor structure.

The DMA controller supports misaligned transfers for both the source and destination addresses. In order to maximize performance, the source and destination engines issue one or more transactions to reach the desired alignment based on the rules described in [Source/destination transaction size calculations](#). The DMA always reads/writes the maximum number of bytes for a given transfer as described by the capability inputs of the DMA controller except for globally coherent transactions that use the size of the cache coherence granule as described by the mode select input. Using 256 bytes over the RapidIO interface reduces packet overhead which translates to increased bandwidth utilization through the interface.

The DMA controller supports bandwidth control, which prevents a channel from consuming all the data bandwidth in the controller. Each channel is allowed to consume the bandwidth of the shared resources as specified by the bandwidth control value. After the channel uses its allotted bandwidth, the arbiter grants the next channel access to the shared resources. The arbitration is round robin between the channels. This feature is also used to implement the external control pause feature. If the external control start and pause are enabled in the MR_n , the channel enters a paused state after transferring the data described in the bandwidth control. External control can restart the channel from a paused state.

The DMA programming model permits software to program each DMA engine independently to interrupt on completed segment, chain, or error. It also provides the capability for software to resume the DMA engine from a hardware halted condition by setting the channel continue bit, $MR_n[CC]$. See [Table 13-262](#) for more complete descriptions of the channel states and state transitions.

13.4.1.1 Source/destination transaction size calculations

The DMA controller may issue smaller transactions from the source and destination address engines in an effort to reach alignment for improved performance.

The flow chart below shows the decision points made in determining the transaction size. The starting *Txfer_Size* is determined by $\min(BCR[BC], MR[BWC])$, *Stride_En* is determined by $SATR_n[SSME]$ or $DATR_n[DSME]$, and *Stride_Size* is determined by $SSR_n[SSS]$ or $DSR_n[DSS]$.

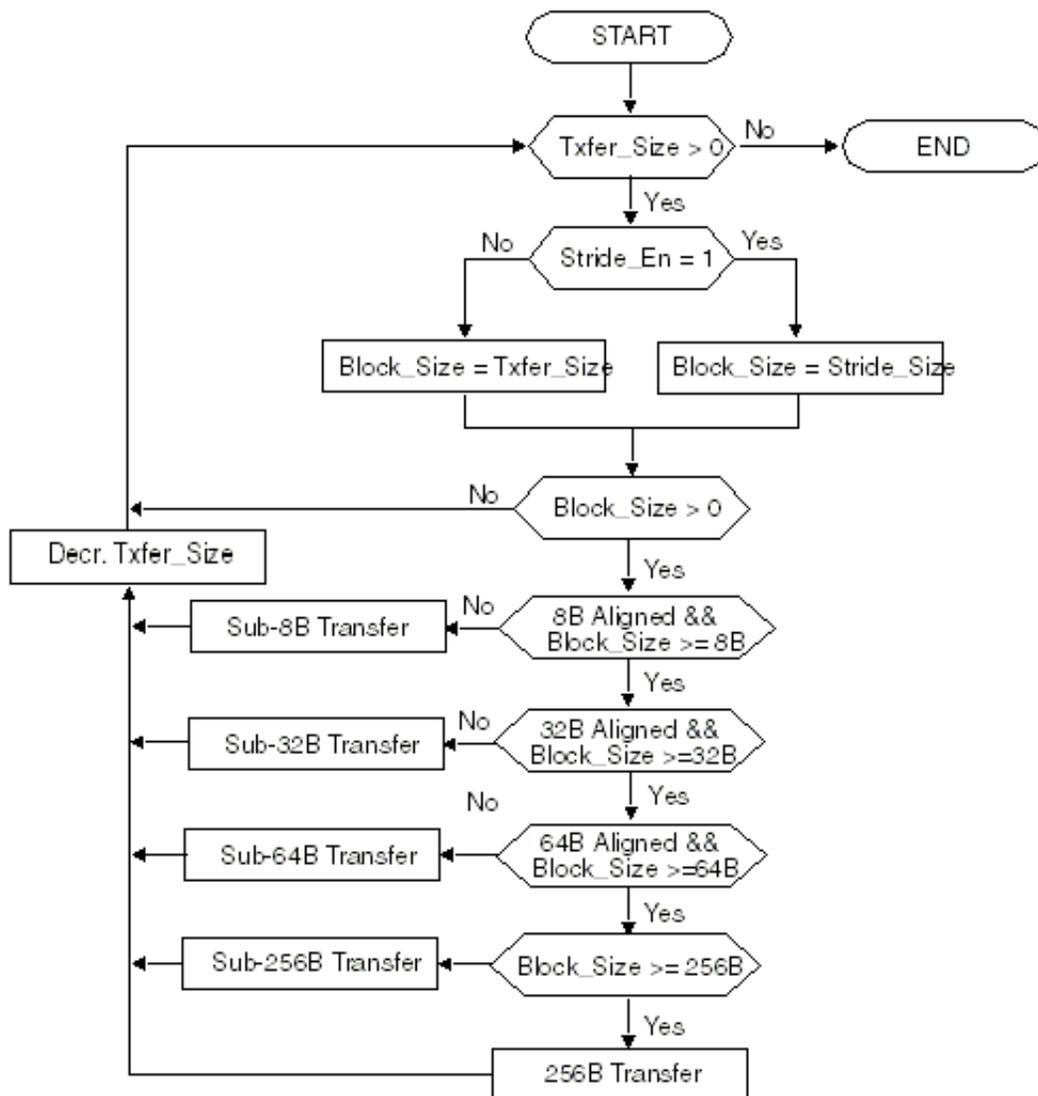


Figure 13-277. Source/destination engine transaction size flow chart

For example, if BCR[BC]=512 bytes and MR[BWC]=256 bytes, reading from starting address 0x5D0 will result in the following transaction sizes:

```

-- Channel arbitration --
0x5D0 - 16 bytes
0x5E0 - 32 bytes
0x600 - 128 bytes
0x680 - 64 bytes
0x6C0 - 16 bytes
-- Channel arbitration --
0x6D0 - 16 bytes
0x6E0 - 32 bytes
0x700 - 128 bytes
0x780 - 64 bytes
0x7C0 - 16 bytes
    
```

In the example above, the bandwidth control limits the channel from ever reaching the maximum transaction size of 256 bytes. Software should align addresses or increase the available bandwidth for best performance.

13.4.1.2 Basic DMA mode transfer

This mode is primarily included for backward compatibility with existing DMA controllers which use a simple programming model. This is the default mode out of reset.

The different modes of operation under the basic mode are explained in the following sections.

13.4.1.2.1 Basic direct mode

In basic direct mode, the DMA controller does not read descriptors from memory, but instead uses the current parameters programmed in the DMA registers to start the DMA transfer.

Software is responsible for initializing SAR_n , $SATR_n$, DAR_n , $DATR_n$, and BCR_n registers. The DMA transfer is started when $MR_n[CS]$ is set. Software is expected to program all the appropriate registers before setting $MR_n[CS]$ to a 1. The transfer is finished after all the bytes specified in the byte count register have been transferred or if an error condition occurs. The sequence of events to start and complete a transfer in basic direct mode is as follows:

1. Poll the channel state (see [Table 13-262](#)) to confirm that the specific DMA channel is idle.
2. Initialize SAR_n , $SATR_n$, DAR_n , $DATR_n$ and BCR_n .
3. Set the mode register channel transfer mode bit, $MR_n[CTM]$, to indicate direct mode. Other control parameters may also be initialized in the mode register.
4. Clear, then set the mode register channel start bit, $MR_n[CS]$, to start the DMA transfer.
5. $SR_n[CB]$ is set by the DMA controller to indicate the DMA transfer is in progress.
6. $SR_n[CB]$ is automatically cleared by the DMA controller after the transfer is finished, or if the transfer is aborted ($MR_n[CA]$ transitions from a 0 to 1), or if a transfer error occurs.
7. End of segment interrupt is generated if $MR_n[EOSIE]$ is set.

13.4.1.2.2 Basic, direct, single-write start mode

In basic direct single-write start mode, the DMA controller does not read descriptors from memory, but instead uses the current parameters programmed in the DMA registers to start the DMA transfer.

Software is responsible for initializing the $SATR_n$, $DATR_n$, and BCR_n registers. Setting $MR_n[SRW]$ configures the DMA controller to begin the DMA transfer either when SAR_n is written or when DAR_n is written, determined by the state of $MR_n[CDSM_SWSM]$. Writing to SAR_n initiates the DMA transfer if $MR_n[CDSM_SWSM]$ is set. Writing to DAR_n initiates the DMA transfer if $MR_n[CDSM_SWSM]$ is cleared. The DMA controller automatically sets the channel start bit, $MR_n[CS]$. Software is expected to program all the appropriate registers before writing the source or destination address registers. The transfer is finished after all the bytes specified in the byte count register have been transferred or if an error condition occurs. The sequence of events to start and complete a transfer in single-write start basic direct mode is as follows:

1. Poll the channel state (see [Table 13-262](#)) to confirm that the specific DMA channel is idle.
2. Initialize the source attributes ($SATR_n$), $DATR_n$, and BCR_n registers.
3. Set the mode register channel transfer mode bit, $MR_n[CTM]$, and the single-write start direct mode bit, $MR_n[SRW]$. Other control parameters may also be initialized in the mode register. Set $MR_n[CDSM_SWSM]$ for transfers started using SAR_n . Clear $MR_n[CDSM_SWSM]$ for transfers started using the DAR_n .
4. A write to the source or destination address register starts the DMA transfer and automatically sets $MR_n[CS]$.
5. $SR_n[CB]$ is set by the DMA controller to indicate the DMA transfer is in progress.
6. $SR_n[CB]$ is automatically cleared by the DMA controller after the transfer is finished, if the transfer is aborted ($MR_n[CA]$ transitions from a 0 to 1), or if a transfer error occurs.
7. End of segment interrupt is generated if $MR_n[EOSIE]$ is set.

13.4.1.2.3 Basic chaining mode

In basic chaining mode, software must first build link descriptor segments in memory.

Then the current link descriptor address register must be initialized to point to the first descriptor in memory. The DMA controller loads descriptors from memory prior to a DMA transfer. The DMA controller begins the transfer according to the link descriptor information loaded for the segment. After the current segment is finished, the DMA controller reads the next link descriptor from memory and begins another DMA transfer.

The transfer is finished if the current link descriptor is the last one in memory or if an error condition occurs. The sequence of events to start and complete a transfer in chaining mode is as follows:

1. Build link descriptor segments in memory.
2. Poll the channel state (see [Table 13-262](#)) to confirm that the specific DMA channel is idle.
3. Initialize CLNDAR_n and ECLNDAR_n to point to the first link descriptor in memory.
4. Clear the mode register channel transfer mode bit, MR_n[CTM], as well as MR_n[XFE], to indicate basic chaining mode. Other control parameters may also be initialized in the mode register.
5. Clear and then set the mode register channel start bit, MR_n[CS], to start the DMA transfer.
6. SR_n[CB] is set by the DMA controller to indicate the DMA transfer is in progress.
7. SR_n[CB] is automatically cleared by the DMA controller after finishing the transfer of the last descriptor segment, if the transfer is aborted (MR_n[CA] transitions from a 0 to 1), or if an error occurs during any of the transfers.

13.4.1.2.4 Basic chaining, single-write start mode

Basic chaining, single-write start mode allows a chain to be started by writing the current link descriptor address register (CLNDAR_n).

(Note that ECLNDAR_n must be written *first* so that the full 36-bit descriptor address is present when the chain starts.) Setting MR_n[CDSM_SWSM] in the mode register causes MR_n[CS] to be automatically set when the current link descriptor address register is written. The sequence of events to start and complete a chain using single-write start mode is as follows:

1. Set the mode register current descriptor start mode bit, MR_n[CDSM_SWSM], and clear the extended features enable bit MR_n[XFE]. Also, clear the channel transfer mode bit, MR_n[CTM]. This initialization indicates basic chaining and single-write start mode. Also other control parameters may be initialized in the mode register.
2. Build link descriptor segments in memory.
3. Poll the channel state (see [Table 13-262](#)) to confirm that the specific DMA channel is idle.
4. Initialize CLNDAR_n and ECLNDAR_n to point to the first descriptor segment in memory. This write automatically causes the DMA controller to begin the link descriptor fetch and set MR_n[CS].
5. SR_n[CB] is set by the DMA controller to indicate the DMA transfer is in progress.
6. SR_n[CB] is automatically cleared by the DMA controller after finishing the transfer of the last descriptor segment, if the transfer is aborted (MR_n[CA] transitions from a 0 to 1), or if an error occurs during any of the transfers.

13.4.1.3 Extended DMA mode transfer

The extended DMA mode also operates in chaining and direct mode.

It offers additional capability over the basic mode by supporting striding and a more flexible descriptor structure. This additional functionality also requires a new and more complex programming model. The extended DMA mode is activated by setting $MR_n[XFE]$.

13.4.1.3.1 Extended direct mode

Extended direct mode has the same functionality as basic direct mode with the addition of stride capabilities.

The bit settings are the same as in direct mode with the exception of the $MR_n[XFE]$ being set. Striding on the source address can be accomplished by setting $SATR_n[SSME]$ and setting the desired stride size and distance in SSR_n . Striding on the destination address can be accomplished by setting $DATR_n[DSME]$ and setting the desired stride size and distance in DSR_n .

13.4.1.3.2 Extended direct, single-write start mode

Extended direct, single-write start mode has the same functionality as the basic direct single-write start mode with the addition of stride capabilities.

The bit settings are also the same with the exception of $MR_n[XFE]$ being set. Striding on the source address can be accomplished by setting $SATR_n[SSME]$ and setting the desired stride size and distance in SSR_n . Striding on the destination address can be accomplished by setting $DATR_n[DSME]$ and setting the desired stride size and distance in DSR_n .

13.4.1.3.3 Extended chaining mode

In extended chaining mode, the software must first build list and link descriptor segments in memory.

Then $CLSDAR_n$ and $ECLSDAR_n$ must be initialized to point to the first list descriptor in memory. The DMA controller loads list descriptors and link descriptors from memory prior to a DMA transfer. The DMA controller begins the transfer according to the link descriptor information loaded. Once the current link descriptor is finished, the DMA controller reads the next link descriptor from memory and begins another DMA transfer. If the current link descriptor is the last in the list, the DMA controller reads the next list

descriptor in memory. The transfer is finished if the current link descriptor is the last one in the last list in memory or if an error condition occurs. The sequence of events to start and complete a transfer in extended chaining mode is as follows:

1. Build link and list descriptor segments in memory.
2. Poll the channel state (see [Table 13-262](#)) to confirm that the specific DMA channel is idle.
3. Initialize CLSDAR_n and ECLSDAR_n to point to the first list descriptor in memory.
4. Clear the mode register channel transfer mode bit, MR_n[CTM], to indicate chaining mode. MR_n[XFE] must be set to indicate extended DMA mode. Other control parameters may also be initialized in the mode register.
5. Clear and then set the mode register channel start bit, MR_n[CS], to start the DMA transfer.
6. SR_n[CB] is set by the DMA controller to indicate the DMA transfer is in progress.
7. SR_n[CB] is automatically cleared by the DMA controller after finishing the transfer of the last descriptor segment, if the transfer is aborted (MR_n[CA] transitions from a 0 to 1), or if an error occurs during any of the transfers.

13.4.1.3.4 Extended chaining, single-write start mode

In the extended mode, the single-write start feature allows a chain to be started by writing the current list descriptor pointer.

Setting MR_n[CDSM_SWSM] causes MR_n[CS] to be set automatically when CLSDAR_n is written. (Note that ECLSDAR_n must be written *first* so that the full 36-bit descriptor address is present when the chain starts.) The sequence of events to start and complete an extended chain using single-write start mode is as follows:

1. Set MR_n[CDSM_SWSM] and MR_n[XFE] and clear MR_n[CTM] to indicate extended chaining and single-write start mode. Also other control parameters may be initialized in the mode register.
2. Build list and link descriptor segments in local memory.
3. Poll the channel state (see [Table 13-262](#)) to confirm that the specific DMA channel is idle.
4. Initialize the current list descriptor address register to point to the first list descriptor segment in memory. This write automatically causes the DMA controller to begin the list descriptor fetch and set MR_n[CS].
5. SR_n[CB] is set by the DMA controller to indicate the DMA transfer is in progress.
6. SR_n[CB] is automatically cleared by the DMA controller after finishing the transfer of the last descriptor segment, or if the transfer is aborted (MR_n[CA] transitions from a 0 to 1), or if an error occurs during any of the transfers.

13.4.1.4 External control mode transfer

An external control can be used to control all DMA channels by setting $MR_n[EMS_EN]$.

The external control can direct the DMA channel in the following transfer modes:

- Basic direct
- Basic chaining
- Extended direct
- Extended chaining

The external control and the DMA controller use a well defined protocol to communicate. The external control can start or restart a paused DMA transfer. The DMA controller acknowledges a DMA transfer in progress and also indicates a transfer completion. Note that external control cannot cause a channel to enter a paused state.

The pause feature can be enabled by setting $MR_n[EMP_EN]$. $MR_n[BWC]$ specifies how much data to allow a specific channel to transfer before entering a paused state by clearing $MR_n[CS]$. Note however, that write data for a paused transfer may not have reached the target interface when so indicated. The channel can be restarted from a paused state by the asserted edge of \overline{DREQ} as driven by an external master. In chaining modes, the channel does not pause for descriptor fetch transfer; it only pauses during the actual data transfer.

Note that when operating the DMA in chaining mode the register byte count field, $BCR[BC]$, must be initialized to zero before enabling the pause feature. In chaining modes, the channel does not pause for descriptor fetch transfer.

The following signals are defined for the external control interface:

- $\overline{DMA_DREQ}$ -Asserting edge triggers a DMA transfer start or restart from a pause request. Sets $MR_n[CS]$. (Note that negating $\overline{DMA_DREQ}$ does NOT clear $MR_n[CS]$.)
- $\overline{DMA_DACK}$ -Indicates a DMA transfer currently in progress. $SR_n[CB]$ is set.
- $\overline{DMA_DDONE}$ -Indicates the completion of the DMA controller's involvement in the transfer and the readiness to accept a new DMA command. $SR_n[CB]$ is clear. Note however, that write data may still be queued at the target interface or in the process of transfer on an external interface.

Detailed descriptions of the external control interface are in [Table 13-3](#). The timing diagram of the external control interface is shown in this figure.

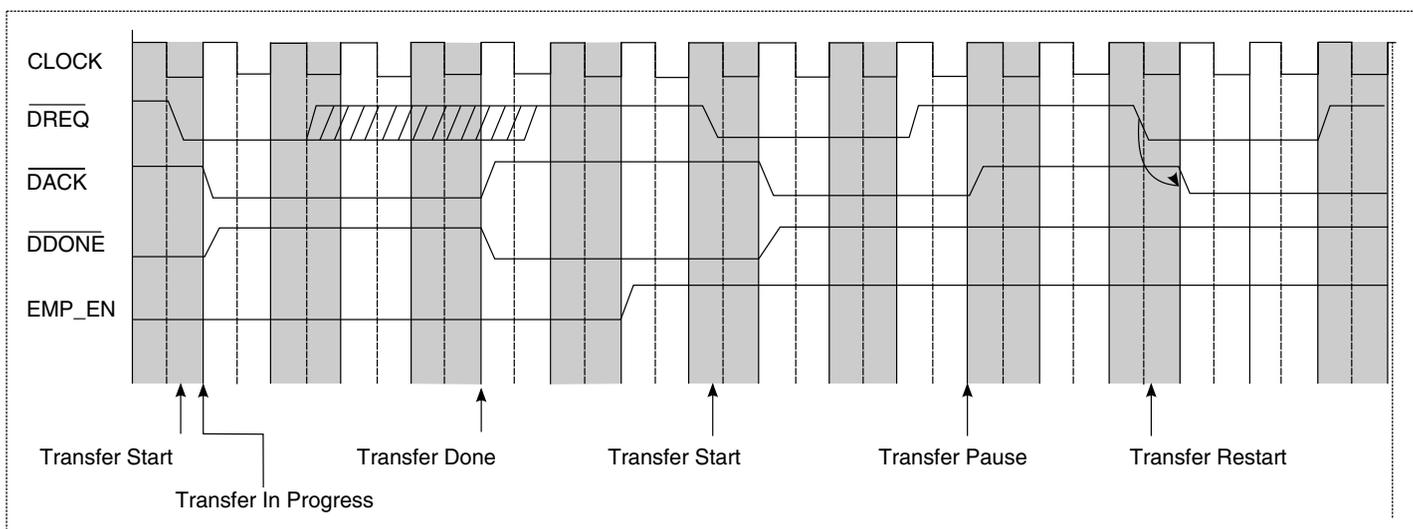


Figure 13-278. External control interface timing

13.4.1.5 Channel continue mode for cascading transfer chains

The channel continue mode (enabled when $MR_n[CC]$ is set) offers software the flexibility of having the DMA controller start on descriptors that have already been programmed while software continues to build more descriptors in memory.

Software can set the end-of-links descriptor (EOLND) in basic mode, or end-of-lists descriptor (EOLSD) in extended mode, to cause the channel to go into a halted state while software continues to build other descriptors in memory. Software can then set CC to force hardware to continue where it left off. Channel continue is only meaningful for chaining modes, not direct mode.

If CC is set by software while the channel is busy with a transfer, the DMA controller finishes all transfers until it reaches the EOLND in basic mode or EOLSD in extended mode. The DMA controller then refetches the last link descriptor in basic mode or the last list descriptor in extended mode and clears the channel continue bit. If EOLND or EOLSD is still set for their respective modes, the DMA controller remains in the idle state. If EOLND or EOLSD is not set, the DMA controller continues the transfer by refetching the new descriptor. The channel busy ($SR_n[CB]$) bit is cleared when the DMA controller reaches EOLND/EOLSD and is set again when it initiates the refetch of the link or list descriptor.

If CC is set by software while the channel is not busy with a transfer, the DMA controller refetches the last link descriptor in basic mode, or the last list descriptor in extended mode and clears the channel continue bit. If EOLND or EOLSD is still set for their

respective modes, the DMA controller remains in the idle state. If the EOLND or EOLSD bits are not set, the DMA controller continues the transfer by refetching the new descriptor.

13.4.1.5.1 Basic mode

On a channel continue, the descriptor at the current link descriptor address registers (CLNDAR n and ECLNDAR n) is refetched to get the next link descriptor address field as updated by software.

The channel halts if NLNDAR n [EOLND] is still set. If EOLND is zero, the next link descriptor address is copied into CLNDAR n and ECLNDAR n and the channel continues with another descriptor fetch of the current link descriptor address. As a result, two link descriptor fetches always exist after channel continue before starting the first transfer.

13.4.1.5.2 Extended mode

On a channel continue, the descriptor at the current list descriptor (CLSDAR n and ECLSDAR n) address register is refetched to get the next list descriptor address field as updated by software.

The channel halts if NLSDAR n [EOLSD] is still set. If not, the next list descriptor address is copied into the CLSDAR n and ECLSDAR n registers and the channel continues with another descriptor fetch of the current list descriptor address. As a result, two list descriptor fetches always exist after channel continue before the first link descriptor fetch and the first transfer.

13.4.1.6 Channel abort

Software can abort a previously initiated transfer by setting MR n [CA].

Once the DMA channel controller detects a zero-to-one transition of MR n [CA], it finishes the current sub-block transfer and halts all further activity. The controller then waits for all previously initiated transfers from the specified channel to drain and clears SR n [CB]. Successful completion of a software initiated abort request can be recognized by MR n [CA] being set and SR n [CB] being cleared. Obviously, if the controller was already halted because of an error condition (SR n [TE] is set), or the channel has completed all transfers, then SR n [CB] being cleared may not signify that the controller entered a halt state due to the abort request.

13.4.1.7 Bandwidth control

MRn[BWC] specifies how much data to allow a specific channel to transfer before allowing the next channel to use the shared data transfer hardware.

This promotes equitable bandwidth allocation between channels. However, if only one channel is busy, hardware overrides the specified bandwidth control size value. The DMA controller allows a channel to transfer up to 1 Kbyte at a time when no other channel is active.

The maximum performance can be achieved on a single channel by disabling bandwidth control. This is recommended especially if only a single channel is utilized.

13.4.1.8 Channel state

This table defines the state of a channel based on the values of the channel start (MRn[CS]), channel busy (SRn[CB]), transfer error (SRn[TE]), and channel continue (MRn[CC]) bits.

Table 13-262. Channel state table

MRn[CS]	SRn[CB]	SRn[TE]	MRn[CC]	Channel state
0	0	0	0	Idle state. This is the state of the bits out of reset.
0	0	0	1	Channel continue unexpected. Channel remains idle
0	0	1	0	Error occurred after software halted the channel.
0	0	1	1	Channel continue unexpected. Channel remains in error halt state
0	1	0	0	Software halted channel. The channel was busy and software cleared MRn[CS].
0	1	0	1	Channel remains in halt state.
-	1	1	-	The channel has encountered an error condition and it is trying to halt.
1	0	0	0	Ready to start a transfer, or transfer completed
1	0	0	1	Continue transfer (only meaningful in chaining mode, not direct mode). In direct mode, the channel continue has no effect.
1	0	1	0	Error occurred during transfer
1	0	1	1	Channel remains in error halt state
1	1	0	0	Transfer in progress
1	1	0	1	Continue after reaching the end of list/link, or the first descriptor fetch after channel continue

13.4.1.9 Illustration of stride size and stride distance

If operating in stride mode, the stride size defines the amount of data to transfer before jumping to the next quantity of data as specified by the stride distance.

The stride distance is added to the current base address to point to the next quantity of data to be transferred. This figure illustrates the stride size and distance parameters.

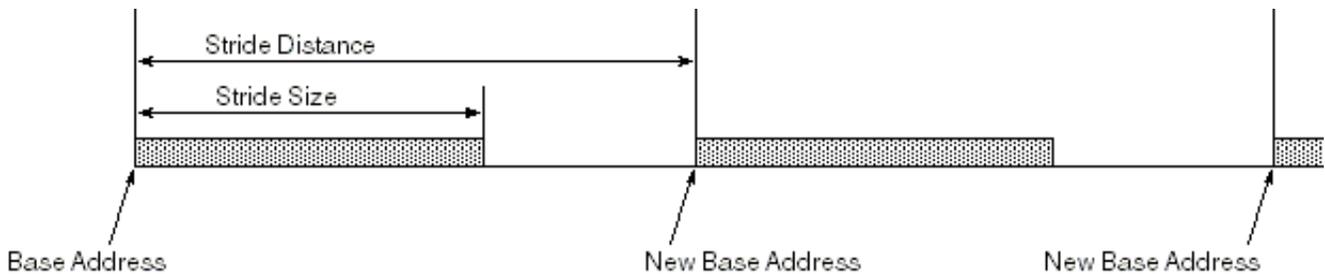


Figure 13-279. Stride size and stride distance

As shown, each time the stride distance is added to the base address, the resulting address becomes the new base address. This sequence repeats until the amount of data transferred equals the transfer size.

13.4.2 DMA transfer interfaces

The DMA can be used to achieve data transfers across the entire memory map.

Note that a single DMA transfer in any of the direct or chaining modes must not cross a 16 GB (34-bit) address boundary.

13.4.3 DMA errors

On a transfer error (for example, non-correctable ECC errors on memory accesses, parity errors on local bus flash controller, address mapping errors, and so on), the DMA halts by setting $SR_n[TE]$ and generates an interrupt if $MR_n[EIE]$ is set.

On a programming error, the DMA sets $SR_n[PE]$ and generates an interrupt if $MR_n[EIE]$ is set. The DMA controller detects the following programming errors:

- Transfer started with a byte count of zero
- Stride transfer started with a stride size of zero
- Transfer started with a priority of three
- Illegal type, defined by $SATR_n[SREADTTYPE]$ and $DATR_n[DWRITETTYPE]$, used for the transfer.

13.4.4 DMA descriptors

The DMA engine recognizes the following types of descriptors:

- List descriptors connect lists of link descriptors.
- Link descriptors describe the DMA activity that is to take place.

DMA descriptors are built in either local or remote memory and are connected by the next descriptor fields. Only link descriptors contain information for the DMA controller to transfer data. Software must ensure that each descriptor is 32-byte aligned. The last link descriptor in the last list in memory sets the NLNDAR_n[EOLND] bit in the next link descriptor and NLSDAR_n[EOLSD] in the next list descriptor fields indicating that these are the last descriptors in memory. Software initializes the current list descriptor address register to point to the first list descriptor in memory. The DMA controller traverses through the descriptor lists until the last link descriptor is met. For each link descriptor in the chain, the DMA controller starts a new DMA transfer with the control parameters specified by that descriptor. Link and list descriptor fetches always snoop the local memory space.

NOTE

Software must ensure that each descriptor is aligned on a 32-byte boundary.

This table summarizes the DMA list descriptors.

Table 13-263. List DMA descriptor summary

Descriptor field	Description
Next list descriptor extended address	Points to the next list descriptor in memory. After the DMA controller reads the descriptor from memory, this field is loaded into the next list descriptor extended address registers.
Next list descriptor address	Points to the next list descriptor in memory. After the DMA controller reads the descriptor from memory, this field is loaded into the next list descriptor address registers.
First link descriptor extended address	Points to the first link descriptor in memory for this list. After the DMA controller reads the descriptor from memory, this field is loaded into the current link descriptor extended address registers.
First link descriptor address	Points to the first link descriptor in memory for this list. After the DMA controller reads the descriptor from memory, this field is loaded into the current link descriptor address registers.
Source stride	Contains the stride information used for the data source if striding is enabled for a link in the list
Destination stride	Contains the stride information used for the data destination if striding is enabled for a link in the list

This table summarizes the DMA link descriptors.

Table 13-264. Link DMA descriptor summary

Descriptor field	Description
Source attributes register	Contains source transaction attributes
Source address	Contains the source address of the DMA transfer. After the DMA controller reads the descriptor from memory, this field is loaded into the Source address register.
Destination attributes register	Contains destination transaction attributes

Table continues on the next page...

Table 13-264. Link DMA descriptor summary (continued)

Descriptor field	Description
Destination address	Contains the destination address of the DMA transfer. After the DMA controller reads the descriptor from memory, this field is loaded into the destination address register.
Next link descriptor extended address	Points to the next link descriptor in memory. After the DMA controller reads the link descriptor from memory, this field is loaded into the extended next link descriptor address registers
Next link descriptor address	Points to the next link descriptor in memory. After the DMA controller reads the link descriptor from memory, this field is loaded into the next link descriptor address registers.
Byte count	Contains the number of bytes to transfer. After the DMA controller reads the descriptor from memory, this field is loaded into the byte count register.

This figure describes the DMA transaction flow.

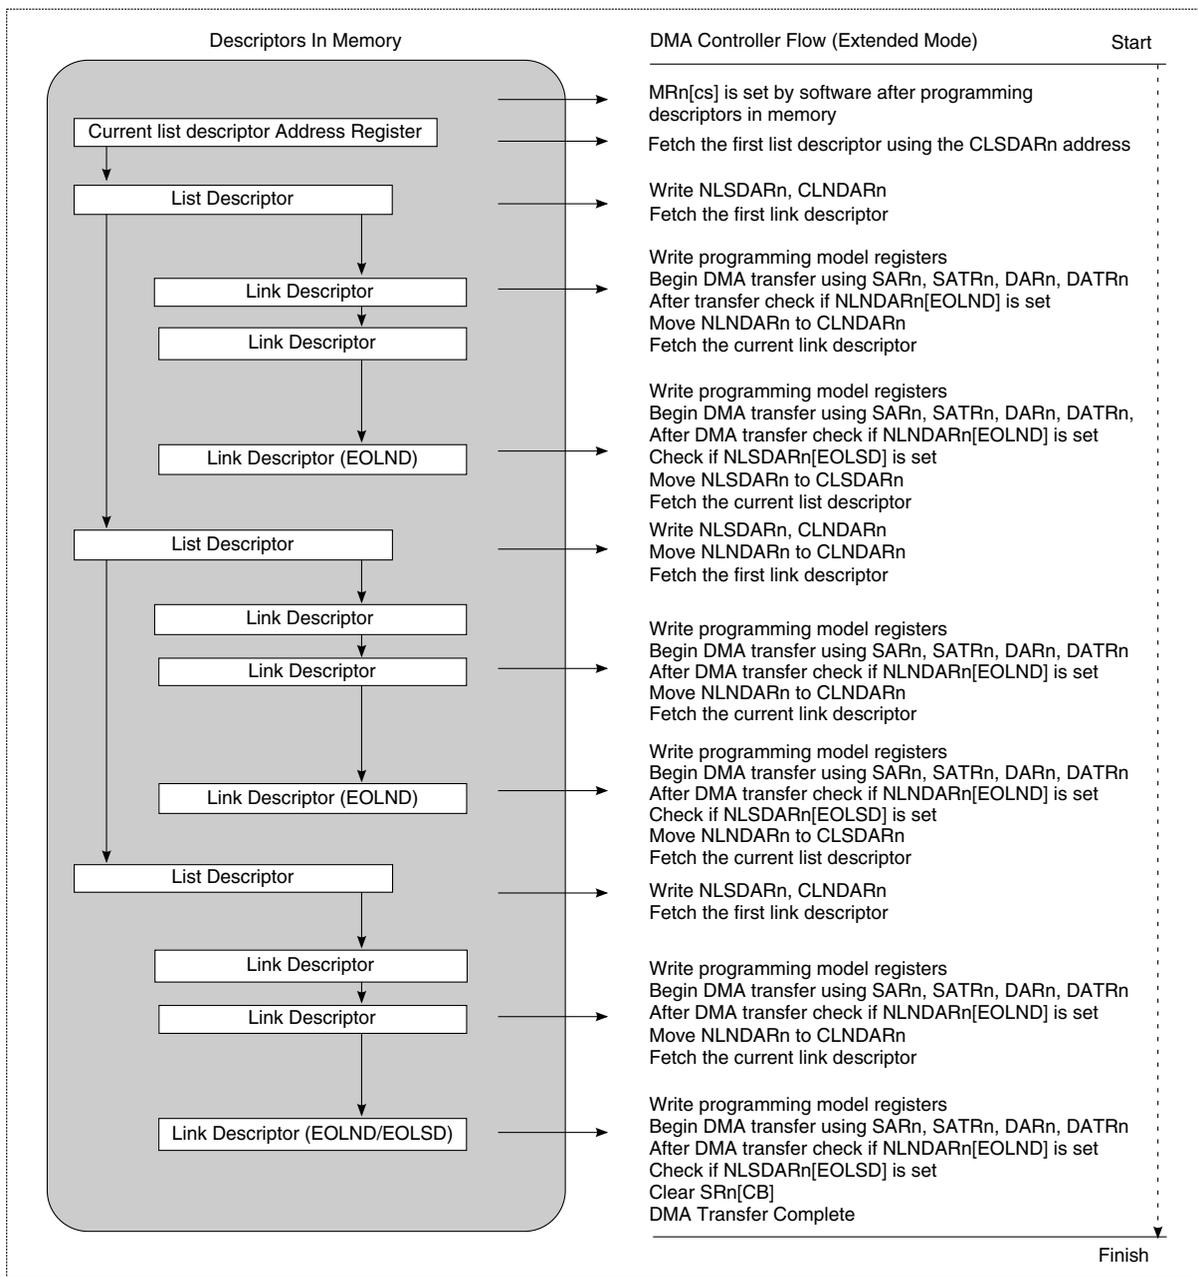


Figure 13-280. DMA transaction flow with DMA descriptors

This table describes the format of the list descriptors.

Offset	
0x00	Next List Descriptor Extended Address
0x04	Next List Descriptor Address
0x08	First Link Descriptor Extended Address
0x0C	First Link Descriptor Address
0x10	Source Stride
0x14	Destination Stride
0x18	Reserved
0x1C	Reserved

Figure 13-281. List descriptor format

This table describes the format of the link descriptors.

Offset	
0x00	Source Attributes
0x04	Source Address
0x08	Destination Attributes
0x0C	Destination Address
0x10	Next Link Descriptor Extended Address
0x14	Next Link Descriptor Address
0x18	Byte Count
0x1C	Reserved

Figure 13-282. Link descriptor format

13.4.5 DMA controller limitations and restrictions

This section is intended to help software maximize the DMA performance and avoid DMA programming errors.

The limitations of the DMA controller are the following:

- Due to the limited number of buffers that the DMA controller can use, avoid stride sizes less than 64 bytes. Maximum utilization is obtained from strides greater than or equal to 256 bytes. However, small stride sizes can be used for scatter-gather functions.
- Coherent reads or writes are broken up into cache line accesses in the DMA.

The DMA controller restrictions are as follows:

- All interface capabilities from where descriptors are being fetched must support read sizes of 32 bytes or greater.
- If MR_n[SAHE] is set, the source interface transfer size capability must be greater than or equal to MR_n[SAHTS]. The source address must be aligned to a size specified by SAHTS.
- If MR_n[DAHE] is set, the destination interface transfer size capability must be greater than or equal to MR_n[DAHTS]. The destination address must be aligned to the size specified by DAHTS.
- Destination striding is not supported if MR_n[DAHE] is set and source striding is not supported if MR_n[SAHE] is set.
- A single DMA transfer in any of the direct or chaining modes must not cross a 16 GB (34-bit) address boundary.

13.5 DMA system considerations

This section provides information about how to make most effective use of the DMA channels.

The following table lists all of the DMA controllers (both captive and general-purpose) on the device and the most likely DMA targets on and off-chip. The bus controllers themselves cannot initiate DMA transfers; rather, they operate transparently with respect to both on- and off-chip DMA controllers. The codes in the table cells have the following meaning:

- Y-Supported
- NR-Possible, but not recommended. Inefficient use of system resources
- NS-Possible, but not supported. Resulting system behavior is not defined.

Table 13-265. P2020 DMA paths

DMA Controllers		On-Chip Targets		Off-Chip Targets				
		L2	Configurati on Registers	DDR	Local Bus	DUART FIFO	Serial RapidIO	PCI Express
On-chip	Ethernet	Y	NS	Y	Y	Y	Y	Y
	eSDHC	Y	NS	Y	Y	Y	Y	Y
	USB	Y	NS	Y	Y	Y	Y	Y
	4 Channel	Y	NR	Y	Y	Y	Y	Y

Table continues on the next page...

Table 13-265. P2020 DMA paths (continued)

DMA Controllers		On-Chip Targets		Off-Chip Targets				
		L2	Configurati on Registers	DDR	Local Bus	DUART FIFO	Serial RapidIO	PCI Express
Off-chip	Serial RapidIO	Y	Y	Y	Y	Y	-	Y
	PCI Express	Y	Y	Y	Y	Y	Y	-
NOTE: On-chip target configuration registers include I ² C data register. NOTE: On-chip Ethernet captive resource. Not available to external masters. NOTE: On-chip 4-channel controller can serve external masters.								

13.5.1 Unusual DMA scenarios

The following is a description of unusual DMA paths including explanations of why some functional modules cannot serve as DMA targets.

The following topics are addressed:

- DMA transaction initiators (masters)
- DMA targets, that is, data sources or destinations
- Transparency of the bus controllers to DMA transactions
- What is useful as opposed to what is possible. For example, any register can be addressed through an internal control bus, which means configuration and control registers can be DMA targets.

13.5.1.1 DMA to core

The L1 cache cannot be a direct DMA target because it cannot be directly addressed by software.

However, DMA access into the L1 cache occurs indirectly if a block of memory that is cached in the L1 is specified as the DMA target. This effect is deterministic if the target memory block was locked into the L1 with cache locking instructions.

13.5.1.2 DMA to configuration, control, and status registers (CCSR)

Because any internal register can be addressed with the four-channel DMA controller, configuration, control, and status registers throughout the device are valid DMA targets.

However, the primary purpose of DMA, to reduce processor load by moving large blocks of data, is not served by DMA transfers of configuration data. For example, while it is possible to DMA into the I²C controller or programmable interrupt controller (PIC), doing so is extremely inefficient and is seldom beneficial in normal operation. The overhead of creating DMA descriptors far exceeds any savings in CPU cycles.

13.5.1.3 DMA to I²C

The I²C controller is not transparent to DMA transfers.

Observe the caveats listed in [DMA to configuration, control, and status registers \(CCSR\)](#) when accessing any I²C register, including the data register (I2CDR).

13.5.1.4 DMA to DUART

The DUART provides complete and sophisticated DMA support, which is described in [FIFO Mode](#).

Chapter 14

Enhanced three-speed Ethernet controllers (eTSEC)

The enhanced three-speed Ethernet controllers (eTSEC) provides the flexibility to accelerate the identification and retrieval of standard and non-standard protocols carried over Ethernet, including both IP versions 4 and 6 and TCP/UDP.

14.1 eTSEC overview

The enhanced three-speed Ethernet controllers (eTSECs) of the device interface to 10 Mbps, 100 Mbps, and 1 Gbps Ethernet/IEEE Std 802.3 networks .

For Ethernet, an external PHY or SerDes device is required to complete the interface to the media. Each eTSEC supports multiple standard media-independent interfaces. Multiple eTSECs are available, providing flexible options for connectivity and control access at different speeds.

The eTSEC provides the flexibility to accelerate the identification and retrieval of standard and non-standard protocols carried over Ethernet, including both IP versions 4 and 6 and TCP/UDP. CPU-intensive parsing and checksum operations can be optionally off-loaded to an eTSEC to accelerate existing TCP/IP stacks. On transmission, varying fractions of link bandwidth can be allocated to each of multiple transmit queues through a modified weighted round-robin scheduler. On receive, an arbitrary set of queue selection rules can be programmed into each eTSEC to implement flexible quality of service or firewall strategies based on high-level protocol identification. Without enabling these advanced features, each eTSEC emulates a PowerQUICC III TSEC, allowing existing driver software to be re-used with minimal change. Each eTSEC is organized as shown in the figure below.

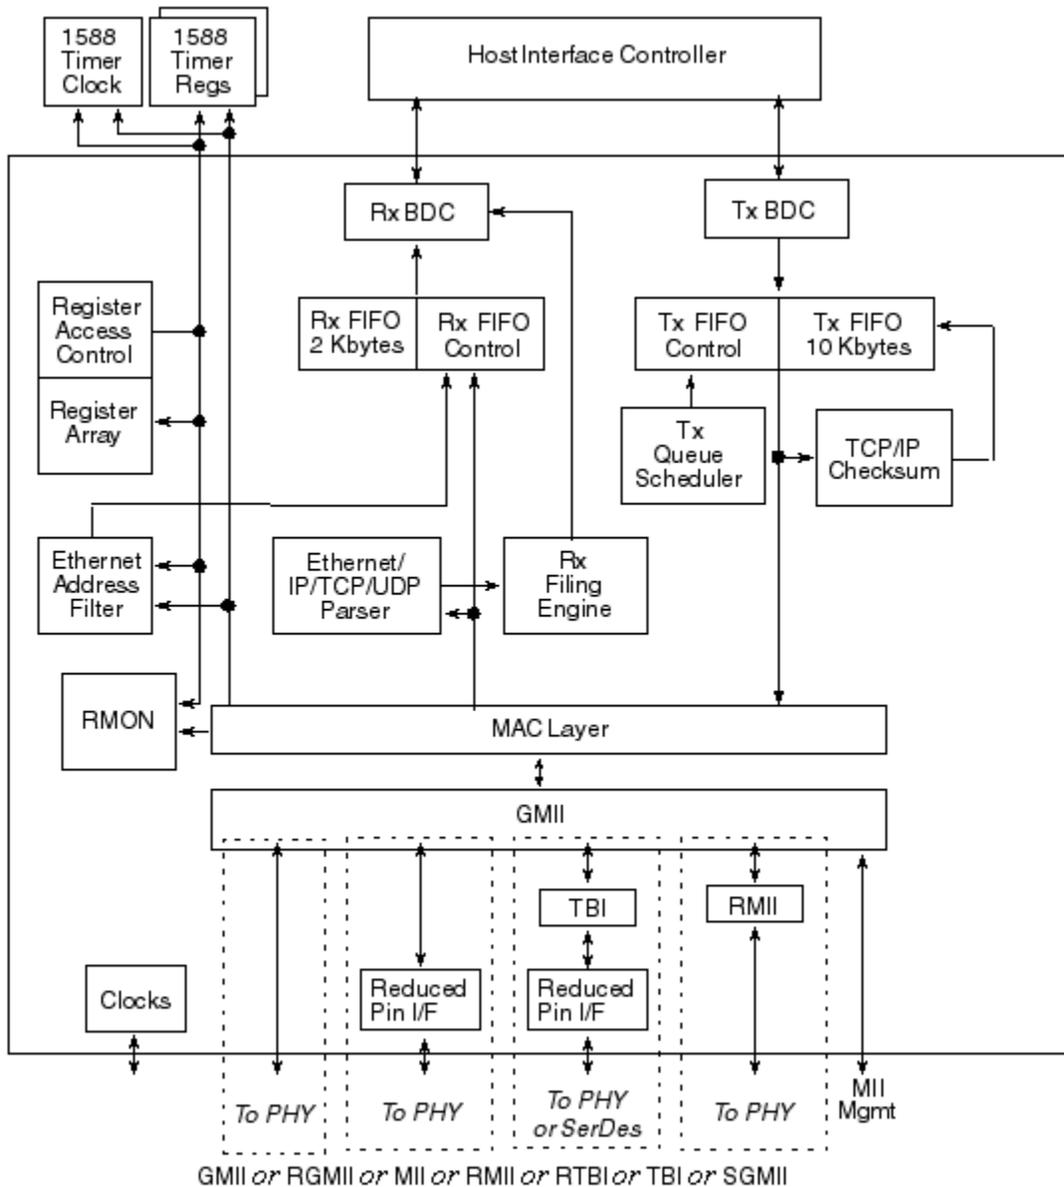


Figure 14-1. eTSEC block diagram

14.2 eTSEC features

The eTSECs of the device include these distinctive features:

- IEEE 802.3, 802.3u, 802.3x, 802.3z, 802.3ac, 802.3ab compatible
- Support for different Ethernet physical interfaces:
 - 10/100 Mbps IEEE 802.3 GMII
 - 1000 Mbps full-duplex IEEE 802.3 GMII
 - 10/100 Mbps IEEE 802.3 MII and RMII

- 10/100 Mbps RGMII
- 1000 Mbps full-duplex RGMII and RTBI (carrier extend symbols in full duplex mode are not supported)
- 10/100 Mbps SGMII
- 1000 Mbps full-duplex SGMII (carrier extend symbols in full duplex mode are not supported)
- 1000 Mbps IEEE 802.3z ten-bit interface (TBI) (carrier extend symbols in full duplex mode are not supported)
- Single-clock TBI
- TCP/IP off-load
 - IP v4 and IP v6 header recognition on receive
 - IP v4 header checksum verification and generation
 - TCP and UDP checksum verification and generation
 - Per-packet configurable off-load
 - Recognition of VLAN, stacked-VLAN, 802.2, PPPoE session, MPLS stacks, ARP, and ESP/AH IP-Security headers
- Quality of service (QoS) support
 - Transmission from up to eight queues
 - Priority-based queue selection
 - Modified weighted round-robin queue selection with fair bandwidth allocation
 - Reception to up to eight physical queues
 - 64 virtual receive queues overlaid on 8 physical buffer descriptor rings
 - Table-oriented queue filing strategy based on 16 header fields or flags
 - Frame rejection support for filtering applications
 - Filing based on Ethernet, IP, and TCP/UDP properties, including VLAN fields, Ether-type, IP protocol type, IP TOS or differentiated services, IP source and destination addresses, TCP/UDP port numbers
- Interrupt coalescing
 - Packet-count-based thresholds for both receive and transmit
 - Timer-based thresholds
- Full- and half-duplex Ethernet support (1000 Mbps supports only full duplex):
 - IEEE 802.3 full-duplex flow control (automatic PAUSE frame generation or software programmed PAUSE frame generation and recognition)
 - Programmable maximum frame length supports jumbo frames (up to 9.6 Kbytes) and IEEE 802.1 virtual local area network (VLAN) tags and priority
 - VLAN insertion and deletion
 - Per-frame VLAN control word or default VLAN for each eTSEC
 - Extracted VLAN control word passed to software separately
 - Programmable VLAN tag to support metropolitan bridging
 - Retransmission following a collision

- Support for CRC generation and verification of inbound/outbound packets
- Programmable Ethernet preamble insertion and extraction of up to 7 bytes
- MAC address recognition:
 - Exact match on primary and virtual 48-bit unicast addresses
 - VRRP and HSRP support for seamless router fail-over
 - In addition to primary station address, up to fifteen additional exact-match MAC addresses supported
 - Broadcast address (accept/reject)
 - Hash table match on up to 256 unicast/multicast or 512 multicast-only addresses
 - Promiscuous mode
- Remote network monitoring (RMON) statistics support
 - 32-bit byte counters
 - Carry/Overflow of counter interrupts
- Backward compatibility with TSEC
 - PowerQUICC III buffer descriptor (BD) format and rings supported
 - Common register memory map, with specific exceptions:
 - Out-of-sequence transmit BD not supported
 - Internal DMA BD pointers and data counts not visible
 - MINFLR register not supported
 - Reset state of eTSEC defaults to common PowerQUICC III TSEC subset
 - TSEC_ID register permits TSEC versus enhanced TSEC differentiation
- Hardware assist for 1588-compliant timestamping (1588 not supported in conjunction with SGMII 10/100)
 - Per packet timestamp tag for Receive
 - Programmable timestamp capture for Transmit
 - Recognition of PTP packet
 - Periodic Pulse Generation
 - Self-correcting precision timer with nano-second resolution
 - Phase aligned adjustable (divide by N) clock output
 - Two 64-bit alarm (future time) registers for future time comparison

14.3 eTSEC modes of operation

The eTSEC's primary operational modes are the following:

- Full- and half-duplex operation

This is determined by the MACCFG2 register's full-duplex bit (MACCFG2[Full Duplex]). Full-duplex mode is intended for use on point-to-point links between switches or end node to switch. Half-duplex mode is used in connections between an end node and a repeater or between repeaters.

If configured in half-duplex mode (10- and 100-Mbps operation; MACCFG2[Full Duplex] is cleared), the MAC complies with the IEEE CSMA/CD access method.

If configured in full-duplex mode (10/100/1000 Mbps operation; MACCFG2[Full Duplex] is set), the MAC supports flow control. If flow control is enabled, it allows the MAC to receive or send PAUSE frames.

- 10- and 100-Mbps MII interface operation

The MAC-PHY interface operates in MII mode by setting MACCFG2[I/F Mode] = 01. The MII is the media-independent interface defined by the 802.3 standard for 10/100 Mbps operation. The speed of operation is determined by the TSECn_TX_CLK and TSECn_RX_CLK signals, which are driven by the transceiver. The transceiver either auto-negotiates the speed, or it may be controlled by software using the serial management interface (MDC/MDIO signals) to the transceiver.

Clause 22.2.4 of the IEEE 802.3 specification describes the MII management interface.

- 10- and 100-Mbps RMII interface operation

The RMII is the reduced media-independent interface defined by the RMII Consortium (March 1998) for 10/100 Mbps operation. The speed of operation is determined by the TSECn_TX_CLK signal, which is driven by the transceiver.

- 1000 Mbps GMII and TBI interface operation

The MAC-PHY interface operates in GMII mode by setting MACCFG2[I/F Mode] = 10. The GMII is the gigabit media-independent interface defined by the 802.3 standard for 1000-Mbps operation.

Independently, the MAC-PHY interface can also operate in TBI mode. Note that either the TBI or GMII interface is chosen, not both at the same time. TBI is the 10-bit interface which contains PCS functions (10-bit encoding/decoding) as defined by the 802.3 standard.

In reduced-pin count mode (RGMII or RTBI), the MAC remains configured in GMII or TBI but the eTSEC muxes and decodes the input signals and provides the MAC with the expected interface.

eTSEC provides the TSECn_GTX_CLK to the PHY in either GMII or TBI mode of operation.

- MAC address recognition options

The options supported are promiscuous, broadcast, exact unicast address match, exact unicast virtual address match to support router redundancy, and multicast hash match. For detailed descriptions refer to [Frame recognition](#).

eTSEC supports automatic LAN-initiated wake-up during power management through the AMD Magic Packet protocol, as described in [Magic Packet mode](#).

- Receive frame parsing options

Frame parsing options are to disable parsing (no TCP/IP off-load), IP header parsing, and TCP or UDP parsing. Parsing must be enabled to make use of receive queue filing algorithms. The options are detailed in [TCP/IP off-load](#).

- Receive queue selection options

Received frames are by default sent to a single buffer descriptor ring. If multiple receive queues are enabled, a receive queue filer can be programmed with selection criteria to differentiate received frames and file them to different buffer descriptor rings. See [Quality of service \(QoS\) provision](#), for detailed descriptions.

- TCP/IP transmit options

Frames for transmission may be sent as-is, with IP header processing, or TCP header processing. The transmit buffer descriptors, described in [Transmit data buffer descriptors \(TxBD\)](#), enable these options and operate with parameters prepended to frame buffers, as described in [TCP/IP off-load](#).

- Transmit queue selection options

The options supported are single transmit queue, priority-based queue selection, and modified weighted round-robin queueing. These options are described further in [Transmit control register \(ETSEC_TCTRL\)](#).

- RMON support

Standard Ethernet interface management information base (MIBs) can be generated through the RMON MIB counters.

- Internal loop back supported for all interfaces except when configured for half-duplex operation. Note that MAC loop back is not supported for 10-bit modes (RTBI or SGMII).

Internal loop back mode is selected through the loop back bit in the MACCFG1 register. See [Interface mode configuration](#), for details.

14.4 eTSEC external signals description

This section defines the eTSEC interface signals. The buses are described using the bus convention used in IEEE 802.3 because the PHY follows this same convention.

(That is, TxD[7 :0] means 0 is the lsb.) Note that except for external physical interfaces the buses and registers follow a big-endian format, where 0 denotes the msb.

Each eTSEC network interface supports multiple options:

- The MII option requires 18 I/O signals (including the MDIO and MDC MII management interface) and supports both a data and a management interface to the PHY (transceiver) device. The MII option supports both 10- and 100-Mbps Ethernet rates.
- The GMII option is a superset of the MII signals and supports a 1000-Mbps Ethernet rate.
- The TBI interface shares signals with the GMII interface signals.
- The RGMII, RTBI, and RMII options are reduced-pin implementations of the GMII, TBI, and MII interfaces, respectively.
- SGMII interfaces are offered via the SerDes interface signals.
- 1588 timer signals

Table 14-1. eTSEC_n network interface signal properties

Signal name	Function	Reset State
TSECn_COL	MII-collision, input	-
TSECn_CRS	MII-carrier sense, input TBI-signal detect, input	-
TSECn_GTX_CLK	RTBI, RGMII-inverted transmit clock feedback, output TBI-continuous transmit clock feedback, output GMII, MII, RMII-transmit clock feedback when transmission is enabled, zero otherwise, output	0
EC_GTX_CLK125	Oscillator source for GMII, TBI, RGMII, RTBI transmit clock, input, shared by all eTSECs	-
EC_MDC	Management clock, output.	0
EC_MDIO	Management data, bidirectional.	Hi-Z (input)
TSECn_RX_CLK	GMII, MII, RGMII-receive clock, input TBI-PMA receive clock 0, input	-

Table continues on the next page...

Table 14-1. eTSECn network interface signal properties (continued)

Signal name	Function	Reset State
TSECn_RX_DV	GMII, MII-receive data valid, input TBI-receive code group (RCG) bit 8, input RGMII (RX_CLK rising)-receive data valid, input RGMII (RX_CLK falling)-receive error, input RTBI (RX_CLK rising)-receive code group (RCG) bit 4, input RTBI (RX_CLK falling)-receive code group (RCG) bit 9, input RMII-CRS_DV carrier sense/data valid, input	-
TSECn_RXD[7:4]	GMII-receive data bits 7:4 input TBI-RCG bits 7:4, input MII, RGMII, RTBI, RMII-unused	-
TSECn_RXD[3:0]	GMII, MII-Receive data bits 3:0, input TBI-RCG bits 3:0, input RGMII (RX_CLK rising) -Receive data bits 3:0, input RGMII (RX_CLK falling)-Receive data bits 7:4, input RTBI (RX_CLK rising)-RCG bits 3:0, input RTBI (RX_CLK falling)-RCG bits 8:5, input RMII-RXD[1:0] receive data bits, input RMII-RXD[3:2] are unused	-
TSECn_RX_ER	GMII, MII, RMII-Receive error, input TBI-RCG bit 9, input RGMII, RTBI-Unused	-
TSECn_TX_CLK	MII-transmit clock, input TBI-PMA receive clock 1, input RMII-reference transmit and receive clock, input RGMII, RTBI-unused	-
TSECn_TXD[7:4]	GMII-transmit data bit 7:4, output TBI-transmit code group (TCG) bit 7:4, output MII, RGMII, RTBI, RMII-unused	0000
TSECn_TXD[3:0]	GMII, MII-Transmit data bits 3:0, output TBI-TCG bits 3:0, output RGMII (TX_CLK rising)-Transmit data bits 3:0, output RGMII (TX_CLK falling)-Transmit data bits 7:4, output RTBI (TX_CLK rising)-TCG bits 3:0, output RTBI (TX_CLK falling)-TCG bits 8:5, output RMII-TXD[1:0] transmit data bits, output RMII-TXD[3:2] unused	0000

Table continues on the next page...

Table 14-1. eTSEC_n network interface signal properties (continued)

Signal name	Function	Reset State
TSEC _n _TX_ER	GMII, MII-transmit error, output RGMII, RTBI, RMII-unused, output driven zero TBI-TCG bit 9, output	0
TSEC _n _TX_EN	GMII, MII, RMII-Transmit data valid, output TBI-TCG bit 8, output RGMII (TX_CLK rising)-Transmit data enabled, output RGMII (TX_CLK falling)-Transmit error, output RTBI (TX_CLK rising)-TCG bit 4, output RTBI (TX_CLK falling)-TCG bit 9, output	0
TSEC_1588_CLK_IN	1588-Clock input External high precision timer reference clock input (chip external input pin).	-
TSEC_1588_CLK_OUT	1588-Clock output Phase aligned timer clock divider output (chip external output pin).	0
TSEC_1588_TRIG_IN1	1588-Trigger in 1 External timer trigger input 1 . This is an asynchronous general purpose input (chip external input pin).	-
TSEC_1588_TRIG_IN2	1588-Trigger in 2 External timer trigger input 2 . This is an asynchronous general purpose input (chip external input pin).	-
TSEC_1588_PULSE_OUT1	1588-Pulse out 1 Timer pulse per period 1. It is phase aligned with 1588 timer clock (chip external output pin).	0
TSEC_1588_PULSE_OUT2	1588-Pulse out 2 Timer pulse per period 2. It is phase aligned with 1588 timer clock (chip external output pin).	0
TSEC_1588_ALARM_OUT1	1588-Timer alarm 1 Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_1588_ALARM _n _H/L register to deactivate this output (chip external output pin).	0
TSEC_1588_ALARM_OUT2	1588-Timer alarm 2 Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_1588_ALARM _n _H/L register to deactivate this output (chip external output pin).	0
SD_TX[<i>n</i> -1] SD_TX[<i>n</i> -1]	SGMII transmit data (and complement)	-
SD_RX[<i>n</i> -1] SD_RX[<i>n</i> -1]	SGMII receive data (and complement)	-
SD_REF_CLK SD_REF_CLK	SGMII SerDes PLL reference clock (and complement)	-
SD_IMP_CAL_TX	Transmitter impedance calibration	-
SD_IMP_CAL_RX	Receiver impedance calibration	-
SD_PLL_TPA	PLL test point analog	-
SD_PLL_TPD	PLL test point digital	-

14.4.1 eTSEC detailed signal descriptions

Below is a description of the eTSEC interface signals.

For RGMII mode details refer to the Hewlett-Packard reduced gigabit media-independent interface (RGMII) specification version 1.2a, dated 9/22/2000. RMII mode details follow the RMII Consortium Specification, dated 3/20/1998. All other modes follow the IEEE 802.3 standard, 2000 Edition. Input signals not used are internally disabled. Except for TSECn_GTX_CLK, output signals not used are driven low.

Table 14-2. eTSEC signals-detailed signal descriptions

Signal	I/O	Description
TSECn_COL	I	Collision input. The behavior of this signal is not specified while in full-duplex mode.
		<p>State meaning</p> <p>Asserted/Negated-In MII mode, this signal is asserted upon detection of a collision, and must remain asserted while the collision persists.</p> <p>This signal is not used in the following modes:</p> <ul style="list-style-type: none"> • RMII • GMII • TBI • RTBI • RGMII
		<p>Timing</p> <p>Asserted/Negated-This signal is not required to transition synchronously with TSECn_TX_CLK or TSECn_RX_CLK.</p>
TSECn_CRS	I	Carrier sense input. In TBI and RTBI modes, this signal is used as SDET (signal detect). In TBI mode SDET must be tied high externally on the board. In RTBI mode SDET is tied high internally.
		<p>This signal is not used in the following modes:</p> <ul style="list-style-type: none"> • RMII • GMII • RGMII
		<p>State meaning</p> <p>Asserted/Negated-In MII mode, TSECn_CRS is asserted while the transmit or receive medium is not idle. In the event of a collision, TSECn_CRS must remain asserted for the duration of the collision.</p>
		<p>Timing</p> <p>Asserted/Negated-This signal is not required to transition synchronously with TSECn_TX_CLK or TSECn_RX_CLK.</p>
TSECn_GTX_CLK	O	<p>Gigabit transmit clock. This signal is an output from the eTSEC into the PHY. TSECn_GTX_CLK is a 125-MHz clock that provides a timing reference for TX_EN, TXD, and TX_ER in the following modes:</p> <ul style="list-style-type: none"> • GMII • TBI • RTBI <p>In RGMII mode, TSECn_GTX_CLK becomes the transmit clock and provides timing reference during 1000Base-T (125 MHz), 100Base-T (25 MHz) and 10Base-T (2.5 MHz) transmissions.</p> <p>This signal feeds back the uninverted transmit clock in MII mode, but feeds back an inverted transmit clock in RTBI or RGMII modes.</p> <p>This signal is driven low unless transmission is enabled, or the eTSEC is in TBI mode.</p>

Table continues on the next page...

Table 14-2. eTSEC signals-detailed signal descriptions (continued)

Signal	I/O	Description		
EC_GTX_CLK125	I	<p>Gigabit transmit 125-MHz source. This signal must be generated externally with a crystal or oscillator, or is sometimes provided by the PHY. EC_GTX_CLK125 is a 125-MHz input into the eTSEC and is used to generate all 125-MHz related signals and clocks in the following modes:</p> <ul style="list-style-type: none"> • GMII • TBI • RTBI • RGMII <p>This input is not used in these modes:</p> <ul style="list-style-type: none"> • RMII • SGMII • MII 		
EC_MDC	O	<p>Management data clock.</p> <p>This signal is a clock (typically 2.5 MHz) supplied by the MAC (IEEE set minimum period of 400 ns or a frequency of 2.5 MHz, but the device may be configured up to 12.5 MHz if supported by the PHY at that speed.) The frequency can be modified by writing to MIIMCFG[28:31] of the eTSEC1 controller.</p>		
EC_MDIO	I/O	Management data input/output.		
		<table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-EC_MDIO is a bidirectional signal to input PHY-supplied status during management read cycles and output control during MII management write cycles. Addressed using eTSEC1 memory-mapped registers.</td> </tr> <tr> <td>Timing</td> <td>Asserted/Negated-This signal is required to be synchronous with the EC_MDC signal.</td> </tr> </table>	State Meaning	Asserted/Negated-EC_MDIO is a bidirectional signal to input PHY-supplied status during management read cycles and output control during MII management write cycles. Addressed using eTSEC1 memory-mapped registers.
State Meaning	Asserted/Negated-EC_MDIO is a bidirectional signal to input PHY-supplied status during management read cycles and output control during MII management write cycles. Addressed using eTSEC1 memory-mapped registers.			
Timing	Asserted/Negated-This signal is required to be synchronous with the EC_MDC signal.			
TSEC _n _RX_CLK	I	<p>Receive clock. In GMII, MII, or RGMII mode, the receive clock TSEC_n_RX_CLK is a continuous clock (2.5, 25, or 125 MHz) that provides a timing reference for TSEC_n_RX_DV, TSEC_n_RXD, and TSEC_n_RX_ER.</p> <p>In TBI mode, TSEC_n_RX_CLK is the input for a 62.5 MHz PMA receive clock, 0 split phase with PMA_RX_CLK1 and is supplied by the SerDes.</p> <p>In RTBI mode it is a 125-MHz receive clock.</p> <p>In RMII mode this clock is not used for the receive clock, as RMII uses a shared reference clock.</p>		
TSEC _n _RX_DV	I	<p>Receive data valid. In GMII or MII mode, if TSEC_n_RX_DV is asserted, the PHY is indicating that valid data is present on the GMII and MII interfaces.</p> <p>In RGMII mode, TSEC_n_RX_DV becomes RX_CTL. The RX_DV and RX_ERR are received on this signal on the rising and falling edges of TSEC_n_RX_CLK.</p> <p>In TBI mode, TSEC_n_RX_DV represents receive code group (RCG) bit 8. Together, with RCG[9] and RCG[7:0], they represent the 10-bit encoded symbol of GMII receive signals.</p> <p>In RTBI mode, TSEC_n_RX_DV represents receive code group (RCG) bit 4 and 9. On the positive edge of the TSEC_n_RX_CLK, RCG[4] and RCG[3:0] represent the first half of the 10-bit encoded symbol. On the negative edge of the TSEC_n_RX_CLK, RCG[9] and RCG[8:5] represent the second half of the 10-bit encoded symbol.</p> <p>In RMII mode the PHY asserts TSEC_n_RX_DV (CRS_DV) when the receive medium is non-idle. This signal asserts asynchronously with respect to the RMII reference clock, but negates synchronously to indicate loss of carrier.</p>		

Table continues on the next page...

Table 14-2. eTSEC signals-detailed signal descriptions (continued)

Signal	I/O	Description		
TSEc _n _RXD[7:0]	I	<p>Receive data in. In GMII mode, TSEc_n_RXD[7:4] with TSEc_n_RXD[3:0], represent one complete octet of data to be transferred from the PHY to the MAC when TSEc_n_RX_DV is asserted.</p> <p>In TBI mode, TSEc_n_RXD[7:4] represents RCG[7:4]. Together, with RCG[9:8] and RCG[3:0], they represent the 10-bit encoded symbol of GMII receive signals.</p> <p>In GMII or MII mode, TSEc_n_RXD[3:0] represents a nibble of data to be transferred from the PHY to the MAC when TSEc_n_RX_DV is asserted. A completely-formed SFD must be passed across the MII. While TSEc_n_RX_DV is not asserted, TSEc_n_RXD has no meaning.</p> <p>In RGMII mode, data bits 3:0 are received on the rising edge of TSEc_n_RX_CLK.</p> <p>In RTBI mode, TSEc_n_RXD[3:0] represents RCG[3:0] on the rising edge of TSEc_n_RX_CLK and RCG[8:5] are received on the falling edge of TSEc_n_RX_CLK.</p> <p>In TBI mode, TSEc_n_RXD[3:0] represents RCG[3:0]. Together, with RCG[9:4], they represent the 10-bit encoded symbol of GMII receive signals.</p> <p>In RMII mode TSEc_n_RXD[1:0] represents RXD[1:0], which is considered valid when TSEc_n_RX_DV (CRS_DV) is asserted, or invalid otherwise.</p>		
TSEc _n _RX_ER	I	<p>Receive error</p> <table border="1"> <tr> <td>State Meaning</td> <td>Asserted/Negated-In GMII, MII, or RMII mode, if TSEc_n_RX_ER and TSEc_n_RX_DV are asserted, the PHY has detected an error in the current frame. In TBI mode, this signal represents RCG[9]. Together, with RCG[8:0], they represent the 10-bit encoded symbol of GMII receive signals. This signal is not used in the RTBI or RGMII modes.</td> </tr> </table>	State Meaning	Asserted/Negated-In GMII, MII, or RMII mode, if TSEc _n _RX_ER and TSEc _n _RX_DV are asserted, the PHY has detected an error in the current frame. In TBI mode, this signal represents RCG[9]. Together, with RCG[8:0], they represent the 10-bit encoded symbol of GMII receive signals. This signal is not used in the RTBI or RGMII modes.
State Meaning	Asserted/Negated-In GMII, MII, or RMII mode, if TSEc _n _RX_ER and TSEc _n _RX_DV are asserted, the PHY has detected an error in the current frame. In TBI mode, this signal represents RCG[9]. Together, with RCG[8:0], they represent the 10-bit encoded symbol of GMII receive signals. This signal is not used in the RTBI or RGMII modes.			
TSEc _n _TX_CLK	I	<p>Transmit clock in. In MII mode, TSEc_n_TX_CLK is a continuous clock (2.5 or 25 MHz) that provides a timing reference for the TSEc_n_TX_EN, TSEc_n_TXD, and TSEc_n_TX_ER signals.</p> <p>In GMII mode, this signal provides the 2.5 or 25 MHz timing reference during 10Base-T and 100Base-T and comes from the PHY. In 1000Base-T this clock is not used and TSEc_n_GTX_CLK (125 MHz) becomes the timing reference. The TSEc_n_GTX_CLK is generated in the eTSEC and provided to the PHY and the MAC. The TSEc_n_TX_CLK is generated in the PHY and provided to the MAC.</p> <p>In TBI mode, this signal is PMA receive clock 1 at 62.5 MHz, split phase with PMA_RX_CLK0, and is supplied by the SerDes.</p> <p>In RMII mode this signal is the reference clock shared between transmit and receive, and is supplied by the PHY.</p> <p>This signal is not used in the eTSEC RTBI or RGMII modes.</p>		

Table continues on the next page...

Table 14-2. eTSEC signals-detailed signal descriptions (continued)

Signal	I/O	Description
TSEC _n _TXD[7:0]	O	<p>Transmit data out. In GMII mode, TSEC_n_TXD[7:0] represents one complete octet of data to be sent from the MAC to the PHY when TSEC_TX_DV is asserted and has no meaning while TSEC_n_TX_EN is negated.</p> <p>In TBI mode, TSEC_n_TXD[7:4] represents transmit code group (TCG) bits 7:4. Together, with TCG[9:8] and TCG[3:0], they represent the 10-bit encoded symbol.</p> <p>In GMII or MII mode, TSEC_n_TXD[3:0] represent a nibble of data to be sent from the MAC to the PHY when TSEC_n_TX_EN is asserted and have no meaning while TSEC_n_TX_EN is negated.</p> <p>In RGMII or RTBI mode, data bits 3:0 are transmitted on the rising edge of TSEC_n_TX_CLK, and data bits 7:4 are transmitted on the falling edge of TSEC_n_TX_CLK.</p> <p>In TBI mode, TSEC_n_TXD[3:0] represents TCG[3:0]. Together, with TCG[9:4], they represent the 10-bit encoded symbol.</p> <p>In RMII mode, TSEC_n_TXD[1:0] represents TXD[1:0], which is valid data sent to the PHY when TSEC_n_TX_EN is asserted, or undefined otherwise.</p> <p>Note that some of these signals are also used during reset to configure the eTSEC interface mode.</p>
TSEC _n _TX_EN	O	<p>Transmit data valid. In GMII, MII, or RMII mode, if TSEC_n_TX_EN is asserted, the MAC is indicating that valid data is present on the GMII's or the MII's TSEC_n_TXD signals.</p> <p>In RGMII mode, TSEC_n_TX_EN becomes TX_CTL. TX_EN and TX_ERR are asserted on this signal on rising and falling edges of the TSEC_n_GTX_CLK, respectively.</p> <p>In TBI mode, TSEC_n_TX_EN represents TCG[8]. Together, with TCG[9] and TCG[7:0], they represent the 10-bit encoded symbol.</p> <p>In RTBI mode, TSEC_n_TX_EN represents TCG[4] on the rising edge and TCG[9] on the falling edge of TSEC_n_GTX_CLK, respectively. Together with TCG[3:0] and TCG[8:5], they represent the 10-bit encoded symbol.</p>
TSEC _n _TX_ER	O	<p>Transmit error. In GMII or MII mode, assertion of TSEC_n_TX_ER for one or more clock cycles while TSEC_n_TX_EN is asserted causes the PHY to transmit one or more illegal symbols. Asserting TSEC_n_TX_ER has no effect while operating at 10 Mbps or while TSEC_n_TX_EN is negated. This signal transitions synchronously with respect to TSEC_n_TX_CLK.</p> <p>In TBI mode, TSEC_n_TX_ER represents TCG[9]. Together, with TCG[8:0], they represent the 10-bit encoded symbol.</p> <p>This signal is not used in the eTSEC RMII, RTBI, or RGMII modes and is driven low.</p>
TSEC_1588_CLK_IN	I	1588 clock in. External high precision timer reference clock input (chip external input pin).
TSEC_1588_CLK_OUT	O	1588 clock out. Phase aligned timer clock divider output (chip external output pin).
TSEC_1588_TRIG_IN1	I	1588 trigger in 1 . External timer trigger input 1 .This is an asynchronous general purpose input (chip external input pin).
TSEC_1588_TRIG_IN2	I	1588 trigger in 2 . External timer trigger input 2 .This is an asynchronous general purpose input (chip external input pin).
TSEC_1588_PULSE_OUT1	O	1588 pulse out 1. Timer pulse per period 1. It is phase aligned with 1588 timer clock (chip external output pin)
TSEC_1588_PULSE_OUT2	O	1588 pulse out 2. Timer pulse per period 2. It is phase aligned with 1588 timer clock (chip external output pin)

Table continues on the next page...

Table 14-2. eTSEC signals-detailed signal descriptions (continued)

Signal	I/O	Description
TSEC_1588_ALARM_OUT1	O	1588 timer alarm 1 . Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_ 1588_ALARMn_H/L register to deactivate this output (chip external output pin)
TSEC_1588_ALARM_OUT2	O	1588 timer alarm 2 . Timer current time is equal to or greater than alarm time comparator register. User reprograms the TSEC_ 1588_ALARMn_H/L register to deactivate this output (chip external output pin)
SD_TX[n-1] SD_TX[n-1]	O	SGMII transmit data (and complement) When in SGMII interface mode: <ul style="list-style-type: none"> • eTSEC2 utilizes SD_TX[2] and $\overline{\text{SD_TX}}[2]$ • eTSEC3 utilizes SD_TX[3] and $\overline{\text{SD_TX}}[3]$
SD_RX[n-1] SD_RX[n-1]	I	SGMII receive data (and complement) When in SGMII interface mode: <ul style="list-style-type: none"> • eTSEC2 utilizes SD_RX[2] and $\overline{\text{SD_RX}}[2]$ • eTSEC3 utilizes SD_RX[3] and $\overline{\text{SD_RX}}[3]$
SD_REF_CLK SD_REF_CLK	I	SGMII SerDes PLL reference clock (and complement)

14.5 eTSEC memory map/register definition

14.5.1 eTSEC top-level module memory map

Each of the eTSECs is allocated 4 Kbytes of memory-mapped space.

The space for each eTSEC is divided as indicated in the table below.

Table 14-3. Module memory map summary

Address offset	Function
000-0FF	eTSEC general control/status registers
100-2FF	eTSEC transmit control/status registers
300-4FF	eTSEC receive control/status registers
500-5FF	MAC registers
600-7FF	RMON MIB registers
800-8FF	Hash table registers
900-AFF	-
B00-BFF	DMA system registers
C00-C3F	Lossless flow control registers
C40-DFF	-
E00-EFF	1588 Hardware assist

14.5.2 eTSEC Memory Map/Register Definition

This section provides a detailed description of all the eTSEC registers. Because all of the eTSEC registers are 32 bits wide, only 32-bit register accesses are supported.

The eTSEC memory mapped registers are accessed by reading and writing to an address comprised of the base address (specified in CCSRBAR as defined in your reference manual's chapter "Memory Map") plus the block base address, plus the offset of the specific register to be accessed. Note that all memory-mapped registers must only be accessed as 32-bit quantities.

The table below lists the offset, name, and a cross-reference to the complete description of each register. The offsets to the memory map table are applicable to each eTSEC. Block base addresses are as follows:

- eTSEC1 starts at 0x2_4000 address offset
- eTSEC2 starts at 0x2_5000 address offset
- eTSEC3 starts at 0x2_6000 address offset

ETSEC memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4000	Controller ID register (ETSEC1_TSEC_ID)	32	R	0124_0000h	14.52.1/776
2_4004	Controller ID register (ETSEC1_TSEC_ID2)	32	R	0030_00F0h	14.52.2/776
2_4010	Interrupt event register (ETSEC1_IEVENT)	32	w1c	0000_0000h	14.52.3/777
2_4014	Interrupt mask register (ETSEC1_IMASK)	32	R/W	0000_0000h	14.52.4/783
2_4018	Error disabled register (ETSEC1_EDIS)	32	R/W	0000_0000h	14.52.5/785
2_4020	Ethernet control register (ETSEC1_ECCTRL)	32	R/W	See section	14.52.6/787
2_4028	Pause time value register (ETSEC1_PTV)	32	R/W	0000_0000h	14.52.7/791
2_402C	DMA control register (ETSEC1_DMACTRL)	32	R/W	0000_0000h	14.52.8/791
2_4030	TBI PHY address register (ETSEC1_TBIPA)	32	R/W	0000_0000h	14.52.9/793
2_4100	Transmit control register (ETSEC1_TCTRL)	32	R/W	0000_0000h	14.52.10/795

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4104	Transmit status register (ETSEC1_TSTAT)	32	w1c	0000_0000h	14.52.11/ 798
2_4108	Default VLAN control word (ETSEC1_DFVLAN)	32	R/W	8100_0000h	14.52.12/ 802
2_4110	Transmit interrupt coalescing register (ETSEC1_TXIC)	32	R/W	0000_0000h	14.52.13/ 803
2_4114	Transmit queue control register (ETSEC1_TQUEUE)	32	R/W	0000_8000h	14.52.14/ 804
2_4140	TxBD Rings 0-3 round-robin weightings (ETSEC1_TR03WT)	32	R/W	0000_0000h	14.52.15/ 805
2_4144	TxBD Rings 4-7 round-robin weightings (ETSEC1_TR47WT)	32	R/W	0000_0000h	14.52.16/ 806
2_4180	Tx data buffer pointer high bits (ETSEC1_TBDBPH)	32	R/W	0000_0000h	14.52.17/ 807
2_4184	TxBD pointer for ring n (ETSEC1_TBPTR0)	32	R/W	0000_0000h	14.52.18/ 807
2_418C	TxBD pointer for ring n (ETSEC1_TBPTR1)	32	R/W	0000_0000h	14.52.18/ 807
2_4194	TxBD pointer for ring n (ETSEC1_TBPTR2)	32	R/W	0000_0000h	14.52.18/ 807
2_419C	TxBD pointer for ring n (ETSEC1_TBPTR3)	32	R/W	0000_0000h	14.52.18/ 807
2_41A4	TxBD pointer for ring n (ETSEC1_TBPTR4)	32	R/W	0000_0000h	14.52.18/ 807
2_41AC	TxBD pointer for ring n (ETSEC1_TBPTR5)	32	R/W	0000_0000h	14.52.18/ 807
2_41B4	TxBD pointer for ring n (ETSEC1_TBPTR6)	32	R/W	0000_0000h	14.52.18/ 807
2_41BC	TxBD pointer for ring n (ETSEC1_TBPTR7)	32	R/W	0000_0000h	14.52.18/ 807
2_4200	TxBD base address high bits (ETSEC1_TBASEH)	32	R/W	0000_0000h	14.52.19/ 808
2_4204	TxBD base address of ring n (ETSEC1_TBASE0)	32	R/W	0000_0000h	14.52.20/ 809
2_420C	TxBD base address of ring n (ETSEC1_TBASE1)	32	R/W	0000_0000h	14.52.20/ 809
2_4214	TxBD base address of ring n (ETSEC1_TBASE2)	32	R/W	0000_0000h	14.52.20/ 809
2_421C	TxBD base address of ring n (ETSEC1_TBASE3)	32	R/W	0000_0000h	14.52.20/ 809
2_4224	TxBD base address of ring n (ETSEC1_TBASE4)	32	R/W	0000_0000h	14.52.20/ 809
2_422C	TxBD base address of ring n (ETSEC1_TBASE5)	32	R/W	0000_0000h	14.52.20/ 809

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4234	TxBD base address of ring n (ETSEC1_TBASE6)	32	R/W	0000_0000h	14.52.20/809
2_423C	TxBD base address of ring n (ETSEC1_TBASE7)	32	R/W	0000_0000h	14.52.20/809
2_4280	Tx time stamp identification tag (set n) register (ETSEC1_TMR_TXTS1_ID)	32	R	0000_0000h	14.52.21/809
2_4284	Tx time stamp identification tag (set n) register (ETSEC1_TMR_TXTS2_ID)	32	R	0000_0000h	14.52.21/809
2_42C0	Tx time stamp high (ETSEC1_TMR_TXTS1_H)	32	R	0000_0000h	14.52.22/810
2_42C4	Tx time stamp low (ETSEC1_TMR_TXTS1_L)	32	R	0000_0000h	14.52.23/810
2_42C8	Tx time stamp high (ETSEC1_TMR_TXTS2_H)	32	R	0000_0000h	14.52.22/810
2_42CC	Tx time stamp low (ETSEC1_TMR_TXTS2_L)	32	R	0000_0000h	14.52.23/810
2_4300	Receive control register (ETSEC1_RCTRL)	32	R/W	0000_0000h	14.52.24/811
2_4304	Receive status register (ETSEC1_RSTAT)	32	w1c	0000_0000h	14.52.25/814
2_4310	Receive interrupt coalescing register (ETSEC1_RXIC)	32	R/W	0000_0000h	14.52.26/816
2_4314	Receive queue control register. (ETSEC1_RQUEUE)	32	R/W	0080_0080h	14.52.27/817
2_4330	Receive bit field extract control register (ETSEC1_RBIFX)	32	R/W	0000_0000h	14.52.28/819
2_4334	Receive queue filing table address register (ETSEC1_RQFAR)	32	R/W	0000_0000h	14.52.29/821
2_4338	Receive queue filing table control register (ETSEC1_RQFCR)	32	R/W	See section	14.52.30/822
2_433C	Receive queue filing table property register (ETSEC1_RQFPR)	32	R/W	0000_0000h	14.52.31/824
2_4340	Maximum receive buffer length register (ETSEC1_MRBLR)	32	R/W	0000_0000h	14.52.32/830
2_4380	Rx data buffer pointer high bits (ETSEC1_RBDBPH)	32	R/W	0000_0000h	14.52.33/831
2_4384	RxBD pointer for ring n (ETSEC1_RBPTR0)	32	R/W	0000_0000h	14.52.34/831
2_438C	RxBD pointer for ring n (ETSEC1_RBPTR1)	32	R/W	0000_0000h	14.52.34/831
2_4394	RxBD pointer for ring n (ETSEC1_RBPTR2)	32	R/W	0000_0000h	14.52.34/831
2_439C	RxBD pointer for ring n (ETSEC1_RBPTR3)	32	R/W	0000_0000h	14.52.34/831

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_43A4	RxBD pointer for ring n (ETSEC1_RBPTR4)	32	R/W	0000_0000h	14.52.34/831
2_43AC	RxBD pointer for ring n (ETSEC1_RBPTR5)	32	R/W	0000_0000h	14.52.34/831
2_43B4	RxBD pointer for ring n (ETSEC1_RBPTR6)	32	R/W	0000_0000h	14.52.34/831
2_43BC	RxBD pointer for ring n (ETSEC1_RBPTR7)	32	R/W	0000_0000h	14.52.34/831
2_4400	RxBD base address high bits (ETSEC1_RBSEH)	32	R/W	0000_0000h	14.52.35/832
2_4404	RxBD base address of ring n (ETSEC1_RBSE0)	32	R/W	0000_0000h	14.52.36/833
2_440C	RxBD base address of ring n (ETSEC1_RBSE1)	32	R/W	0000_0000h	14.52.36/833
2_4414	RxBD base address of ring n (ETSEC1_RBSE2)	32	R/W	0000_0000h	14.52.36/833
2_441C	RxBD base address of ring n (ETSEC1_RBSE3)	32	R/W	0000_0000h	14.52.36/833
2_4424	RxBD base address of ring n (ETSEC1_RBSE4)	32	R/W	0000_0000h	14.52.36/833
2_442C	RxBD base address of ring n (ETSEC1_RBSE5)	32	R/W	0000_0000h	14.52.36/833
2_4434	RxBD base address of ring n (ETSEC1_RBSE6)	32	R/W	0000_0000h	14.52.36/833
2_443C	RxBD base address of ring n (ETSEC1_RBSE7)	32	R/W	0000_0000h	14.52.36/833
2_44C0	Rx timer time stamp register high (ETSEC1_TMR_RXTS_H)	32	R	0000_0000h	14.52.37/833
2_44C4	Rx timer time stamp register low (ETSEC1_TMR_RXTS_L)	32	R	0000_0000h	14.52.38/834
2_4500	MAC configuration register 1 (ETSEC1_MACCFG1)	32	R/W	0000_0000h	14.52.39/834
2_4504	MAC configuration register 2 (ETSEC1_MACCFG2)	32	R/W	0000_7000h	14.52.40/837
2_4508	Inter-packet/inter-frame gap register (ETSEC1_IPGIFG)	32	R/W	4060_5060h	14.52.41/839
2_450C	Half-duplex control (ETSEC1_HAFDUP)	32	R/W	00A1_F037h	14.52.42/840
2_4510	Maximum frame length (ETSEC1_MAXFRM)	32	R/W	0000_0600h	14.52.43/841
2_4520	MII management configuration (ETSEC1_MIIMCFG)	32	R/W	0000_0007h	14.52.44/842
2_4524	MII management command (ETSEC1_MIIMCOM)	32	R/W	0000_0000h	14.52.45/844

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4528	MII management address (ETSEC1_MIIMADD)	32	R/W	0000_0000h	14.52.46/ 845
2_452C	MII management control (ETSEC1_MIIMCON)	32	W	0000_0000h	14.52.47/ 845
2_4530	MII management status (ETSEC1_MIIMSTAT)	32	R	0000_0000h	14.52.48/ 846
2_4534	MII management indicator (ETSEC1_MIIMIND)	32	R	0000_0000h	14.52.49/ 846
2_453C	Interface status (ETSEC1_IFSTAT)	32	R	0000_0000h	14.52.50/ 847
2_4540	MAC station address register 1 (ETSEC1_MACSTNADDR1)	32	R/W	0000_0000h	14.52.51/ 849
2_4544	MAC station address register 2 (ETSEC1_MACSTNADDR2)	32	R/W	0000_0000h	14.52.52/ 849
2_4548	MAC exact match address n, part 1 (ETSEC1_MAC01ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_454C	MAC exact match address n, part 2 (ETSEC1_MAC01ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_4550	MAC exact match address n, part 1 (ETSEC1_MAC02ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_4554	MAC exact match address n, part 2 (ETSEC1_MAC02ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_4558	MAC exact match address n, part 1 (ETSEC1_MAC03ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_455C	MAC exact match address n, part 2 (ETSEC1_MAC03ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_4560	MAC exact match address n, part 1 (ETSEC1_MAC04ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_4564	MAC exact match address n, part 2 (ETSEC1_MAC04ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_4568	MAC exact match address n, part 1 (ETSEC1_MAC05ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_456C	MAC exact match address n, part 2 (ETSEC1_MAC05ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_4570	MAC exact match address n, part 1 (ETSEC1_MAC06ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_4574	MAC exact match address n, part 2 (ETSEC1_MAC06ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_4578	MAC exact match address n, part 1 (ETSEC1_MAC07ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_457C	MAC exact match address n, part 2 (ETSEC1_MAC07ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_4580	MAC exact match address n, part 1 (ETSEC1_MAC08ADDR1)	32	R/W	0000_0000h	14.52.53/ 850

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4584	MAC exact match address n, part 2 (ETSEC1_MAC08ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_4588	MAC exact match address n, part 1 (ETSEC1_MAC09ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_458C	MAC exact match address n, part 2 (ETSEC1_MAC09ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_4590	MAC exact match address n, part 1 (ETSEC1_MAC10ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_4594	MAC exact match address n, part 2 (ETSEC1_MAC10ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_4598	MAC exact match address n, part 1 (ETSEC1_MAC11ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_459C	MAC exact match address n, part 2 (ETSEC1_MAC11ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_45A0	MAC exact match address n, part 1 (ETSEC1_MAC12ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_45A4	MAC exact match address n, part 2 (ETSEC1_MAC12ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_45A8	MAC exact match address n, part 1 (ETSEC1_MAC13ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_45AC	MAC exact match address n, part 2 (ETSEC1_MAC13ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_45B0	MAC exact match address n, part 1 (ETSEC1_MAC14ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_45B4	MAC exact match address n, part 2 (ETSEC1_MAC14ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_45B8	MAC exact match address n, part 1 (ETSEC1_MAC15ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_45BC	MAC exact match address n, part 2 (ETSEC1_MAC15ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_4680	Transmit and receive 64-byte frame counter (ETSEC1_TR64)	32	R/W	0000_0000h	14.52.55/851
2_4684	Transmit and receive 65- to 127-byte frame counter (ETSEC1_TR127)	32	R/W	0000_0000h	14.52.56/852
2_4688	Transmit and receive 128- to 255-byte frame counter (ETSEC1_TR255)	32	R/W	0000_0000h	14.52.57/852
2_468C	Transmit and receive 256- to 511-byte frame counter (ETSEC1_TR511)	32	R/W	0000_0000h	14.52.58/853
2_4690	Transmit and receive 512- to 1023-byte frame counter (ETSEC1_TR1K)	32	R/W	0000_0000h	14.52.59/853
2_4694	Transmit and receive 1024- to 1518-byte frame counter (ETSEC1_TRMAX)	32	R/W	0000_0000h	14.52.60/853
2_4698	Transmit and receive 1519- to 1522-byte good VLAN frame count (ETSEC1_TRMGV)	32	R/W	0000_0000h	14.52.61/854

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_469C	Receive byte counter (ETSEC1_RBYT)	32	R/W	0000_0000h	14.52.62/854
2_46A0	Receive packet counter (ETSEC1_RPKT)	32	R/W	0000_0000h	14.52.63/855
2_46A4	Receive FCS error counter (ETSEC1_RFCS)	32	R/W	0000_0000h	14.52.64/855
2_46A8	Receive multicast packet counter (ETSEC1_RMCA)	32	R/W	0000_0000h	14.52.65/855
2_46AC	Receive broadcast packet counter (ETSEC1_RBCA)	32	R/W	0000_0000h	14.52.66/856
2_46B0	Receive control frame packet counter (ETSEC1_RXCF)	32	R/W	0000_0000h	14.52.67/856
2_46B4	Receive PAUSE frame packet counter (ETSEC1_RXPF)	32	R/W	0000_0000h	14.52.68/857
2_46B8	Receive unknown OP code counter (ETSEC1_RXUO)	32	R/W	0000_0000h	14.52.69/857
2_46BC	Receive alignment error counter (ETSEC1_RALN)	32	R/W	0000_0000h	14.52.70/857
2_46C0	Receive frame length error counter (ETSEC1_RFLR)	32	R/W	0000_0000h	14.52.71/858
2_46C4	Receive code error counter (ETSEC1_RCDE)	32	R/W	0000_0000h	14.52.72/858
2_46C8	Receive carrier sense error counter (ETSEC1_RCSE)	32	R/W	0000_0000h	14.52.73/859
2_46CC	Receive undersize packet counter (ETSEC1_RUND)	32	R/W	0000_0000h	14.52.74/859
2_46D0	Receive oversize packet counter (ETSEC1_ROVR)	32	R/W	0000_0000h	14.52.75/859
2_46D4	Receive fragments counter (ETSEC1_RFRG)	32	R/W	0000_0000h	14.52.76/860
2_46D8	Receive jabber counter (ETSEC1_RJBR)	32	R/W	0000_0000h	14.52.77/860
2_46DC	Receive drop counter (ETSEC1_RDRP)	32	R/W	0000_0000h	14.52.78/861
2_46E0	Transmit byte counter (ETSEC1_TBYT)	32	R/W	0000_0000h	14.52.79/861
2_46E4	Transmit packet counter (ETSEC1_TPKT)	32	R/W	0000_0000h	14.52.80/861
2_46E8	Transmit multicast packet counter (ETSEC1_TMCA)	32	R/W	0000_0000h	14.52.81/862
2_46EC	Transmit broadcast packet counter (ETSEC1_TBCA)	32	R/W	0000_0000h	14.52.82/862
2_46F0	Transmit PAUSE control frame counter (ETSEC1_TXPF)	32	R/W	0000_0000h	14.52.83/863

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_46F4	Transmit deferral packet counter (ETSEC1_TDFR)	32	R/W	0000_0000h	14.52.84/ 863
2_46F8	Transmit excessive deferral packet counter (ETSEC1_TEDF)	32	R/W	0000_0000h	14.52.85/ 863
2_46FC	Transmit single collision packet counter (ETSEC1_TSCL)	32	R/W	0000_0000h	14.52.86/ 864
2_4700	Transmit multiple collision packet counter (ETSEC1_TMCL)	32	R/W	0000_0000h	14.52.87/ 864
2_4704	Transmit late collision packet counter (ETSEC1_TLCL)	32	R/W	0000_0000h	14.52.88/ 865
2_4708	Transmit excessive collision packet counter (ETSEC1_TXCL)	32	R/W	0000_0000h	14.52.89/ 865
2_470C	Transmit total collision counter (ETSEC1_TNCL)	32	R/W	0000_0000h	14.52.90/ 865
2_4714	Transmit drop frame counter (ETSEC1_TDRP)	32	R/W	0000_0000h	14.52.91/ 866
2_4718	Transmit jabber frame counter (ETSEC1_TJBR)	32	R/W	0000_0000h	14.52.92/ 866
2_471C	Transmit FCS error counter (ETSEC1_TFCS)	32	R/W	0000_0000h	14.52.93/ 867
2_4720	Transmit control frame counter (ETSEC1_TXCF)	32	R/W	0000_0000h	14.52.94/ 867
2_4724	Transmit oversize frame counter (ETSEC1_TOVR)	32	R/W	0000_0000h	14.52.95/ 867
2_4728	Transmit undersize frame counter (ETSEC1_TUND)	32	R/W	0000_0000h	14.52.96/ 868
2_472C	Transmit fragments frame counter (ETSEC1_TFRG)	32	R/W	0000_0000h	14.52.97/ 868
2_4730	Carry register one registerCleared on read. (ETSEC1_CAR1)	32	w1c	0000_0000h	14.52.98/ 869
2_4734	Carry register two register 3 (ETSEC1_CAR2)	32	R/W	0000_0000h	14.52.99/ 871
2_4738	Carry register one mask register (ETSEC1_CAM1)	32	R/W	FE03_FFFFh	14.52.100/ 873
2_473C	Carry register two mask register (ETSEC1_CAM2)	32	R/W	000F_FFFDh	14.52.101/ 875
2_4740	Receive filer rejected packet counter (ETSEC1_RREJ)	32	R/W	0000_0000h	14.52.102/ 876
2_4800	Individual/group address register n (ETSEC1_IGADDR0)	32	R/W	0000_0000h	14.52.103/ 877
2_4804	Individual/group address register n (ETSEC1_IGADDR1)	32	R/W	0000_0000h	14.52.103/ 877
2_4808	Individual/group address register n (ETSEC1_IGADDR2)	32	R/W	0000_0000h	14.52.103/ 877

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_480C	Individual/group address register n (ETSEC1_IGADDR3)	32	R/W	0000_0000h	14.52.103/877
2_4810	Individual/group address register n (ETSEC1_IGADDR4)	32	R/W	0000_0000h	14.52.103/877
2_4814	Individual/group address register n (ETSEC1_IGADDR5)	32	R/W	0000_0000h	14.52.103/877
2_4818	Individual/group address register n (ETSEC1_IGADDR6)	32	R/W	0000_0000h	14.52.103/877
2_481C	Individual/group address register n (ETSEC1_IGADDR7)	32	R/W	0000_0000h	14.52.103/877
2_4880	Group address register n (ETSEC1_GADDR0)	32	R/W	0000_0000h	14.52.104/878
2_4884	Group address register n (ETSEC1_GADDR1)	32	R/W	0000_0000h	14.52.104/878
2_4888	Group address register n (ETSEC1_GADDR2)	32	R/W	0000_0000h	14.52.104/878
2_488C	Group address register n (ETSEC1_GADDR3)	32	R/W	0000_0000h	14.52.104/878
2_4890	Group address register n (ETSEC1_GADDR4)	32	R/W	0000_0000h	14.52.104/878
2_4894	Group address register n (ETSEC1_GADDR5)	32	R/W	0000_0000h	14.52.104/878
2_4898	Group address register n (ETSEC1_GADDR6)	32	R/W	0000_0000h	14.52.104/878
2_489C	Group address register n (ETSEC1_GADDR7)	32	R/W	0000_0000h	14.52.104/878
2_4BF8	Attribute register (ETSEC1_ATTR)	32	R/W	0000_0000h	14.52.105/879
2_4BFC	Attribute extract length and extract index register (ETSEC1_ATTRELI)	32	R/W	0000_0000h	14.52.106/881
2_4C00	Receive Queue Parameters register n (ETSEC1_RQPRM0)	32	R/W	0000_0000h	14.52.107/882
2_4C04	Receive Queue Parameters register n (ETSEC1_RQPRM1)	32	R/W	0000_0000h	14.52.107/882
2_4C08	Receive Queue Parameters register n (ETSEC1_RQPRM2)	32	R/W	0000_0000h	14.52.107/882
2_4C0C	Receive Queue Parameters register n (ETSEC1_RQPRM3)	32	R/W	0000_0000h	14.52.107/882
2_4C10	Receive Queue Parameters register n (ETSEC1_RQPRM4)	32	R/W	0000_0000h	14.52.107/882
2_4C14	Receive Queue Parameters register n (ETSEC1_RQPRM5)	32	R/W	0000_0000h	14.52.107/882
2_4C18	Receive Queue Parameters register n (ETSEC1_RQPRM6)	32	R/W	0000_0000h	14.52.107/882

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4C1C	Receive Queue Parameters register n (ETSEC1_RQPRM7)	32	R/W	0000_0000h	14.52.107/882
2_4C44	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR0)	32	R/W	0000_0000h	14.52.108/883
2_4C4C	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR1)	32	R/W	0000_0000h	14.52.108/883
2_4C54	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR2)	32	R/W	0000_0000h	14.52.108/883
2_4C5C	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR3)	32	R/W	0000_0000h	14.52.108/883
2_4C64	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR4)	32	R/W	0000_0000h	14.52.108/883
2_4C6C	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR5)	32	R/W	0000_0000h	14.52.108/883
2_4C74	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR6)	32	R/W	0000_0000h	14.52.108/883
2_4C7C	Last Free RxBD pointer for ring n (ETSEC1_RFBPTR7)	32	R/W	0000_0000h	14.52.108/883
2_5000	Controller ID register (ETSEC2_TSEC_ID)	32	R	0124_0000h	14.52.1/776
2_5004	Controller ID register (ETSEC2_TSEC_ID2)	32	R	0030_00F0h	14.52.2/776
2_5010	Interrupt event register (ETSEC2_IEVENT)	32	w1c	0000_0000h	14.52.3/777
2_5014	Interrupt mask register (ETSEC2_IMASK)	32	R/W	0000_0000h	14.52.4/783
2_5018	Error disabled register (ETSEC2_EDIS)	32	R/W	0000_0000h	14.52.5/785
2_5020	Ethernet control register (ETSEC2_ECCTRL)	32	R/W	See section	14.52.6/787
2_5028	Pause time value register (ETSEC2_PTV)	32	R/W	0000_0000h	14.52.7/791
2_502C	DMA control register (ETSEC2_DMACTRL)	32	R/W	0000_0000h	14.52.8/791
2_5030	TBI PHY address register (ETSEC2_TBIPA)	32	R/W	0000_0000h	14.52.9/793
2_5100	Transmit control register (ETSEC2_TCTRL)	32	R/W	0000_0000h	14.52.10/795
2_5104	Transmit status register (ETSEC2_TSTAT)	32	w1c	0000_0000h	14.52.11/798
2_5108	Default VLAN control word (ETSEC2_DFVLAN)	32	R/W	8100_0000h	14.52.12/802
2_5110	Transmit interrupt coalescing register (ETSEC2_TXIC)	32	R/W	0000_0000h	14.52.13/803

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_5114	Transmit queue control register (ETSEC2_TQUEUE)	32	R/W	0000_8000h	14.52.14/804
2_5140	TxBD Rings 0-3 round-robin weightings (ETSEC2_TR03WT)	32	R/W	0000_0000h	14.52.15/805
2_5144	TxBD Rings 4-7 round-robin weightings (ETSEC2_TR47WT)	32	R/W	0000_0000h	14.52.16/806
2_5180	Tx data buffer pointer high bits (ETSEC2_TBDBPH)	32	R/W	0000_0000h	14.52.17/807
2_5184	TxBD pointer for ring n (ETSEC2_TBPTR0)	32	R/W	0000_0000h	14.52.18/807
2_518C	TxBD pointer for ring n (ETSEC2_TBPTR1)	32	R/W	0000_0000h	14.52.18/807
2_5194	TxBD pointer for ring n (ETSEC2_TBPTR2)	32	R/W	0000_0000h	14.52.18/807
2_519C	TxBD pointer for ring n (ETSEC2_TBPTR3)	32	R/W	0000_0000h	14.52.18/807
2_51A4	TxBD pointer for ring n (ETSEC2_TBPTR4)	32	R/W	0000_0000h	14.52.18/807
2_51AC	TxBD pointer for ring n (ETSEC2_TBPTR5)	32	R/W	0000_0000h	14.52.18/807
2_51B4	TxBD pointer for ring n (ETSEC2_TBPTR6)	32	R/W	0000_0000h	14.52.18/807
2_51BC	TxBD pointer for ring n (ETSEC2_TBPTR7)	32	R/W	0000_0000h	14.52.18/807
2_5200	TxBD base address high bits (ETSEC2_TBASEH)	32	R/W	0000_0000h	14.52.19/808
2_5204	TxBD base address of ring n (ETSEC2_TBASE0)	32	R/W	0000_0000h	14.52.20/809
2_520C	TxBD base address of ring n (ETSEC2_TBASE1)	32	R/W	0000_0000h	14.52.20/809
2_5214	TxBD base address of ring n (ETSEC2_TBASE2)	32	R/W	0000_0000h	14.52.20/809
2_521C	TxBD base address of ring n (ETSEC2_TBASE3)	32	R/W	0000_0000h	14.52.20/809
2_5224	TxBD base address of ring n (ETSEC2_TBASE4)	32	R/W	0000_0000h	14.52.20/809
2_522C	TxBD base address of ring n (ETSEC2_TBASE5)	32	R/W	0000_0000h	14.52.20/809
2_5234	TxBD base address of ring n (ETSEC2_TBASE6)	32	R/W	0000_0000h	14.52.20/809
2_523C	TxBD base address of ring n (ETSEC2_TBASE7)	32	R/W	0000_0000h	14.52.20/809
2_5280	Tx time stamp identification tag (set n) register (ETSEC2_TMR_TXTS1_ID)	32	R	0000_0000h	14.52.21/809

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_5284	Tx time stamp identification tag (set n) register (ETSEC2_TMR_TXTS2_ID)	32	R	0000_0000h	14.52.21/809
2_52C0	Tx time stamp high (ETSEC2_TMR_TXTS1_H)	32	R	0000_0000h	14.52.22/810
2_52C4	Tx time stamp low (ETSEC2_TMR_TXTS1_L)	32	R	0000_0000h	14.52.23/810
2_52C8	Tx time stamp high (ETSEC2_TMR_TXTS2_H)	32	R	0000_0000h	14.52.22/810
2_52CC	Tx time stamp low (ETSEC2_TMR_TXTS2_L)	32	R	0000_0000h	14.52.23/810
2_5300	Receive control register (ETSEC2_RCTRL)	32	R/W	0000_0000h	14.52.24/811
2_5304	Receive status register (ETSEC2_RSTAT)	32	w1c	0000_0000h	14.52.25/814
2_5310	Receive interrupt coalescing register (ETSEC2_RXIC)	32	R/W	0000_0000h	14.52.26/816
2_5314	Receive queue control register. (ETSEC2_RQUEUE)	32	R/W	0080_0080h	14.52.27/817
2_5330	Receive bit field extract control register (ETSEC2_RBIFX)	32	R/W	0000_0000h	14.52.28/819
2_5334	Receive queue filing table address register (ETSEC2_RQFAR)	32	R/W	0000_0000h	14.52.29/821
2_5338	Receive queue filing table control register (ETSEC2_RQFCR)	32	R/W	See section	14.52.30/822
2_533C	Receive queue filing table property register (ETSEC2_RQFPR)	32	R/W	0000_0000h	14.52.31/824
2_5340	Maximum receive buffer length register (ETSEC2_MRBLR)	32	R/W	0000_0000h	14.52.32/830
2_5380	Rx data buffer pointer high bits (ETSEC2_RBDBPH)	32	R/W	0000_0000h	14.52.33/831
2_5384	RxBD pointer for ring n (ETSEC2_RBPTR0)	32	R/W	0000_0000h	14.52.34/831
2_538C	RxBD pointer for ring n (ETSEC2_RBPTR1)	32	R/W	0000_0000h	14.52.34/831
2_5394	RxBD pointer for ring n (ETSEC2_RBPTR2)	32	R/W	0000_0000h	14.52.34/831
2_539C	RxBD pointer for ring n (ETSEC2_RBPTR3)	32	R/W	0000_0000h	14.52.34/831
2_53A4	RxBD pointer for ring n (ETSEC2_RBPTR4)	32	R/W	0000_0000h	14.52.34/831
2_53AC	RxBD pointer for ring n (ETSEC2_RBPTR5)	32	R/W	0000_0000h	14.52.34/831
2_53B4	RxBD pointer for ring n (ETSEC2_RBPTR6)	32	R/W	0000_0000h	14.52.34/831

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_53BC	RxBD pointer for ring n (ETSEC2_RBPTR7)	32	R/W	0000_0000h	14.52.34/831
2_5400	RxBD base address high bits (ETSEC2_RBSEH)	32	R/W	0000_0000h	14.52.35/832
2_5404	RxBD base address of ring n (ETSEC2_RBSE0)	32	R/W	0000_0000h	14.52.36/833
2_540C	RxBD base address of ring n (ETSEC2_RBSE1)	32	R/W	0000_0000h	14.52.36/833
2_5414	RxBD base address of ring n (ETSEC2_RBSE2)	32	R/W	0000_0000h	14.52.36/833
2_541C	RxBD base address of ring n (ETSEC2_RBSE3)	32	R/W	0000_0000h	14.52.36/833
2_5424	RxBD base address of ring n (ETSEC2_RBSE4)	32	R/W	0000_0000h	14.52.36/833
2_542C	RxBD base address of ring n (ETSEC2_RBSE5)	32	R/W	0000_0000h	14.52.36/833
2_5434	RxBD base address of ring n (ETSEC2_RBSE6)	32	R/W	0000_0000h	14.52.36/833
2_543C	RxBD base address of ring n (ETSEC2_RBSE7)	32	R/W	0000_0000h	14.52.36/833
2_54C0	Rx timer time stamp register high (ETSEC2_TMR_RXTS_H)	32	R	0000_0000h	14.52.37/833
2_54C4	Rx timer time stamp register low (ETSEC2_TMR_RXTS_L)	32	R	0000_0000h	14.52.38/834
2_5500	MAC configuration register 1 (ETSEC2_MACCFG1)	32	R/W	0000_0000h	14.52.39/834
2_5504	MAC configuration register 2 (ETSEC2_MACCFG2)	32	R/W	0000_7000h	14.52.40/837
2_5508	Inter-packet/inter-frame gap register (ETSEC2_IPGIFG)	32	R/W	4060_5060h	14.52.41/839
2_550C	Half-duplex control (ETSEC2_HAFDUP)	32	R/W	00A1_F037h	14.52.42/840
2_5510	Maximum frame length (ETSEC2_MAXFRM)	32	R/W	0000_0600h	14.52.43/841
2_5520	MII management configuration (ETSEC2_MIIMCFG)	32	R/W	0000_0007h	14.52.44/842
2_5524	MII management command (ETSEC2_MIIMCOM)	32	R/W	0000_0000h	14.52.45/844
2_5528	MII management address (ETSEC2_MIIMADD)	32	R/W	0000_0000h	14.52.46/845
2_552C	MII management control (ETSEC2_MIIMCON)	32	W	0000_0000h	14.52.47/845
2_5530	MII management status (ETSEC2_MIIMSTAT)	32	R	0000_0000h	14.52.48/846

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_5534	MII management indicator (ETSEC2_MIIMIND)	32	R	0000_0000h	14.52.49/ 846
2_553C	Interface status (ETSEC2_IFSTAT)	32	R	0000_0000h	14.52.50/ 847
2_5540	MAC station address register 1 (ETSEC2_MACSTNADDR1)	32	R/W	0000_0000h	14.52.51/ 849
2_5544	MAC station address register 2 (ETSEC2_MACSTNADDR2)	32	R/W	0000_0000h	14.52.52/ 849
2_5548	MAC exact match address n, part 1 (ETSEC2_MAC01ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_554C	MAC exact match address n, part 2 (ETSEC2_MAC01ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5550	MAC exact match address n, part 1 (ETSEC2_MAC02ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_5554	MAC exact match address n, part 2 (ETSEC2_MAC02ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5558	MAC exact match address n, part 1 (ETSEC2_MAC03ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_555C	MAC exact match address n, part 2 (ETSEC2_MAC03ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5560	MAC exact match address n, part 1 (ETSEC2_MAC04ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_5564	MAC exact match address n, part 2 (ETSEC2_MAC04ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5568	MAC exact match address n, part 1 (ETSEC2_MAC05ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_556C	MAC exact match address n, part 2 (ETSEC2_MAC05ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5570	MAC exact match address n, part 1 (ETSEC2_MAC06ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_5574	MAC exact match address n, part 2 (ETSEC2_MAC06ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5578	MAC exact match address n, part 1 (ETSEC2_MAC07ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_557C	MAC exact match address n, part 2 (ETSEC2_MAC07ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5580	MAC exact match address n, part 1 (ETSEC2_MAC08ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_5584	MAC exact match address n, part 2 (ETSEC2_MAC08ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5588	MAC exact match address n, part 1 (ETSEC2_MAC09ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_558C	MAC exact match address n, part 2 (ETSEC2_MAC09ADDR2)	32	R/W	0000_0000h	14.52.54/ 851

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_5590	MAC exact match address n, part 1 (ETSEC2_MAC10ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_5594	MAC exact match address n, part 2 (ETSEC2_MAC10ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5598	MAC exact match address n, part 1 (ETSEC2_MAC11ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_559C	MAC exact match address n, part 2 (ETSEC2_MAC11ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_55A0	MAC exact match address n, part 1 (ETSEC2_MAC12ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_55A4	MAC exact match address n, part 2 (ETSEC2_MAC12ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_55A8	MAC exact match address n, part 1 (ETSEC2_MAC13ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_55AC	MAC exact match address n, part 2 (ETSEC2_MAC13ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_55B0	MAC exact match address n, part 1 (ETSEC2_MAC14ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_55B4	MAC exact match address n, part 2 (ETSEC2_MAC14ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_55B8	MAC exact match address n, part 1 (ETSEC2_MAC15ADDR1)	32	R/W	0000_0000h	14.52.53/ 850
2_55BC	MAC exact match address n, part 2 (ETSEC2_MAC15ADDR2)	32	R/W	0000_0000h	14.52.54/ 851
2_5680	Transmit and receive 64-byte frame counter (ETSEC2_TR64)	32	R/W	0000_0000h	14.52.55/ 851
2_5684	Transmit and receive 65- to 127-byte frame counter (ETSEC2_TR127)	32	R/W	0000_0000h	14.52.56/ 852
2_5688	Transmit and receive 128- to 255-byte frame counter (ETSEC2_TR255)	32	R/W	0000_0000h	14.52.57/ 852
2_568C	Transmit and receive 256- to 511-byte frame counter (ETSEC2_TR511)	32	R/W	0000_0000h	14.52.58/ 853
2_5690	Transmit and receive 512- to 1023-byte frame counter (ETSEC2_TR1K)	32	R/W	0000_0000h	14.52.59/ 853
2_5694	Transmit and receive 1024- to 1518-byte frame counter (ETSEC2_TRMAX)	32	R/W	0000_0000h	14.52.60/ 853
2_5698	Transmit and receive 1519- to 1522-byte good VLAN frame count (ETSEC2_TRMGV)	32	R/W	0000_0000h	14.52.61/ 854
2_569C	Receive byte counter (ETSEC2_RBYT)	32	R/W	0000_0000h	14.52.62/ 854
2_56A0	Receive packet counter (ETSEC2_RPKT)	32	R/W	0000_0000h	14.52.63/ 855
2_56A4	Receive FCS error counter (ETSEC2_RFCS)	32	R/W	0000_0000h	14.52.64/ 855

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_56A8	Receive multicast packet counter (ETSEC2_RMCA)	32	R/W	0000_0000h	14.52.65/ 855
2_56AC	Receive broadcast packet counter (ETSEC2_RBCA)	32	R/W	0000_0000h	14.52.66/ 856
2_56B0	Receive control frame packet counter (ETSEC2_RXCF)	32	R/W	0000_0000h	14.52.67/ 856
2_56B4	Receive PAUSE frame packet counter (ETSEC2_RXPF)	32	R/W	0000_0000h	14.52.68/ 857
2_56B8	Receive unknown OP code counter (ETSEC2_RXUO)	32	R/W	0000_0000h	14.52.69/ 857
2_56BC	Receive alignment error counter (ETSEC2_RALN)	32	R/W	0000_0000h	14.52.70/ 857
2_56C0	Receive frame length error counter (ETSEC2_RFLR)	32	R/W	0000_0000h	14.52.71/ 858
2_56C4	Receive code error counter (ETSEC2_RCDE)	32	R/W	0000_0000h	14.52.72/ 858
2_56C8	Receive carrier sense error counter (ETSEC2_RCSE)	32	R/W	0000_0000h	14.52.73/ 859
2_56CC	Receive undersize packet counter (ETSEC2_RUND)	32	R/W	0000_0000h	14.52.74/ 859
2_56D0	Receive oversize packet counter (ETSEC2_ROVR)	32	R/W	0000_0000h	14.52.75/ 859
2_56D4	Receive fragments counter (ETSEC2_RFRG)	32	R/W	0000_0000h	14.52.76/ 860
2_56D8	Receive jabber counter (ETSEC2_RJBR)	32	R/W	0000_0000h	14.52.77/ 860
2_56DC	Receive drop counter (ETSEC2_RDRP)	32	R/W	0000_0000h	14.52.78/ 861
2_56E0	Transmit byte counter (ETSEC2_TBYT)	32	R/W	0000_0000h	14.52.79/ 861
2_56E4	Transmit packet counter (ETSEC2_TPKT)	32	R/W	0000_0000h	14.52.80/ 861
2_56E8	Transmit multicast packet counter (ETSEC2_TMCA)	32	R/W	0000_0000h	14.52.81/ 862
2_56EC	Transmit broadcast packet counter (ETSEC2_TBCA)	32	R/W	0000_0000h	14.52.82/ 862
2_56F0	Transmit PAUSE control frame counter (ETSEC2_TXPF)	32	R/W	0000_0000h	14.52.83/ 863
2_56F4	Transmit deferral packet counter (ETSEC2_TDFR)	32	R/W	0000_0000h	14.52.84/ 863
2_56F8	Transmit excessive deferral packet counter (ETSEC2_TEDF)	32	R/W	0000_0000h	14.52.85/ 863
2_56FC	Transmit single collision packet counter (ETSEC2_TSCL)	32	R/W	0000_0000h	14.52.86/ 864

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_5700	Transmit multiple collision packet counter (ETSEC2_TMCL)	32	R/W	0000_0000h	14.52.87/ 864
2_5704	Transmit late collision packet counter (ETSEC2_TLCL)	32	R/W	0000_0000h	14.52.88/ 865
2_5708	Transmit excessive collision packet counter (ETSEC2_TXCL)	32	R/W	0000_0000h	14.52.89/ 865
2_570C	Transmit total collision counter (ETSEC2_TNCL)	32	R/W	0000_0000h	14.52.90/ 865
2_5714	Transmit drop frame counter (ETSEC2_TDRP)	32	R/W	0000_0000h	14.52.91/ 866
2_5718	Transmit jabber frame counter (ETSEC2_TJBR)	32	R/W	0000_0000h	14.52.92/ 866
2_571C	Transmit FCS error counter (ETSEC2_TFCS)	32	R/W	0000_0000h	14.52.93/ 867
2_5720	Transmit control frame counter (ETSEC2_TXCF)	32	R/W	0000_0000h	14.52.94/ 867
2_5724	Transmit oversize frame counter (ETSEC2_TOVR)	32	R/W	0000_0000h	14.52.95/ 867
2_5728	Transmit undersize frame counter (ETSEC2_TUND)	32	R/W	0000_0000h	14.52.96/ 868
2_572C	Transmit fragments frame counter (ETSEC2_TFRG)	32	R/W	0000_0000h	14.52.97/ 868
2_5730	Carry register one registerCleared on read. (ETSEC2_CAR1)	32	w1c	0000_0000h	14.52.98/ 869
2_5734	Carry register two register 3 (ETSEC2_CAR2)	32	R/W	0000_0000h	14.52.99/ 871
2_5738	Carry register one mask register (ETSEC2_CAM1)	32	R/W	FE03_FFFFh	14.52.100/ 873
2_573C	Carry register two mask register (ETSEC2_CAM2)	32	R/W	000F_FFFDh	14.52.101/ 875
2_5740	Receive filer rejected packet counter (ETSEC2_RREJ)	32	R/W	0000_0000h	14.52.102/ 876
2_5800	Individual/group address register n (ETSEC2_IGADDR0)	32	R/W	0000_0000h	14.52.103/ 877
2_5804	Individual/group address register n (ETSEC2_IGADDR1)	32	R/W	0000_0000h	14.52.103/ 877
2_5808	Individual/group address register n (ETSEC2_IGADDR2)	32	R/W	0000_0000h	14.52.103/ 877
2_580C	Individual/group address register n (ETSEC2_IGADDR3)	32	R/W	0000_0000h	14.52.103/ 877
2_5810	Individual/group address register n (ETSEC2_IGADDR4)	32	R/W	0000_0000h	14.52.103/ 877
2_5814	Individual/group address register n (ETSEC2_IGADDR5)	32	R/W	0000_0000h	14.52.103/ 877

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_5818	Individual/group address register n (ETSEC2_IGADDR6)	32	R/W	0000_0000h	14.52.103/877
2_581C	Individual/group address register n (ETSEC2_IGADDR7)	32	R/W	0000_0000h	14.52.103/877
2_5880	Group address register n (ETSEC2_GADDR0)	32	R/W	0000_0000h	14.52.104/878
2_5884	Group address register n (ETSEC2_GADDR1)	32	R/W	0000_0000h	14.52.104/878
2_5888	Group address register n (ETSEC2_GADDR2)	32	R/W	0000_0000h	14.52.104/878
2_588C	Group address register n (ETSEC2_GADDR3)	32	R/W	0000_0000h	14.52.104/878
2_5890	Group address register n (ETSEC2_GADDR4)	32	R/W	0000_0000h	14.52.104/878
2_5894	Group address register n (ETSEC2_GADDR5)	32	R/W	0000_0000h	14.52.104/878
2_5898	Group address register n (ETSEC2_GADDR6)	32	R/W	0000_0000h	14.52.104/878
2_589C	Group address register n (ETSEC2_GADDR7)	32	R/W	0000_0000h	14.52.104/878
2_5BF8	Attribute register (ETSEC2_ATTR)	32	R/W	0000_0000h	14.52.105/879
2_5BFC	Attribute extract length and extract index register (ETSEC2_ATTRELI)	32	R/W	0000_0000h	14.52.106/881
2_5C00	Receive Queue Parameters register n (ETSEC2_RQPRM0)	32	R/W	0000_0000h	14.52.107/882
2_5C04	Receive Queue Parameters register n (ETSEC2_RQPRM1)	32	R/W	0000_0000h	14.52.107/882
2_5C08	Receive Queue Parameters register n (ETSEC2_RQPRM2)	32	R/W	0000_0000h	14.52.107/882
2_5C0C	Receive Queue Parameters register n (ETSEC2_RQPRM3)	32	R/W	0000_0000h	14.52.107/882
2_5C10	Receive Queue Parameters register n (ETSEC2_RQPRM4)	32	R/W	0000_0000h	14.52.107/882
2_5C14	Receive Queue Parameters register n (ETSEC2_RQPRM5)	32	R/W	0000_0000h	14.52.107/882
2_5C18	Receive Queue Parameters register n (ETSEC2_RQPRM6)	32	R/W	0000_0000h	14.52.107/882
2_5C1C	Receive Queue Parameters register n (ETSEC2_RQPRM7)	32	R/W	0000_0000h	14.52.107/882
2_5C44	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR0)	32	R/W	0000_0000h	14.52.108/883
2_5C4C	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR1)	32	R/W	0000_0000h	14.52.108/883

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_5C54	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR2)	32	R/W	0000_0000h	14.52.108/883
2_5C5C	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR3)	32	R/W	0000_0000h	14.52.108/883
2_5C64	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR4)	32	R/W	0000_0000h	14.52.108/883
2_5C6C	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR5)	32	R/W	0000_0000h	14.52.108/883
2_5C74	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR6)	32	R/W	0000_0000h	14.52.108/883
2_5C7C	Last Free RxBD pointer for ring n (ETSEC2_RFBPTR7)	32	R/W	0000_0000h	14.52.108/883
2_6000	Controller ID register (ETSEC3_TSEC_ID)	32	R	0124_0000h	14.52.1/776
2_6004	Controller ID register (ETSEC3_TSEC_ID2)	32	R	0030_00F0h	14.52.2/776
2_6010	Interrupt event register (ETSEC3_IEVENT)	32	w1c	0000_0000h	14.52.3/777
2_6014	Interrupt mask register (ETSEC3_IMASK)	32	R/W	0000_0000h	14.52.4/783
2_6018	Error disabled register (ETSEC3_EDIS)	32	R/W	0000_0000h	14.52.5/785
2_6020	Ethernet control register (ETSEC3_ECCTRL)	32	R/W	See section	14.52.6/787
2_6028	Pause time value register (ETSEC3_PTV)	32	R/W	0000_0000h	14.52.7/791
2_602C	DMA control register (ETSEC3_DMACTRL)	32	R/W	0000_0000h	14.52.8/791
2_6030	TBI PHY address register (ETSEC3_TBIPA)	32	R/W	0000_0000h	14.52.9/793
2_6100	Transmit control register (ETSEC3_TCTRL)	32	R/W	0000_0000h	14.52.10/795
2_6104	Transmit status register (ETSEC3_TSTAT)	32	w1c	0000_0000h	14.52.11/798
2_6108	Default VLAN control word (ETSEC3_DFVLAN)	32	R/W	8100_0000h	14.52.12/802
2_6110	Transmit interrupt coalescing register (ETSEC3_TXIC)	32	R/W	0000_0000h	14.52.13/803
2_6114	Transmit queue control register (ETSEC3_TQUEUE)	32	R/W	0000_8000h	14.52.14/804
2_6140	TxBD Rings 0-3 round-robin weightings (ETSEC3_TR03WT)	32	R/W	0000_0000h	14.52.15/805
2_6144	TxBD Rings 4-7 round-robin weightings (ETSEC3_TR47WT)	32	R/W	0000_0000h	14.52.16/806

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_6180	Tx data buffer pointer high bits (ETSEC3_TBDBPH)	32	R/W	0000_0000h	14.52.17/807
2_6184	TxBD pointer for ring n (ETSEC3_TBPTR0)	32	R/W	0000_0000h	14.52.18/807
2_618C	TxBD pointer for ring n (ETSEC3_TBPTR1)	32	R/W	0000_0000h	14.52.18/807
2_6194	TxBD pointer for ring n (ETSEC3_TBPTR2)	32	R/W	0000_0000h	14.52.18/807
2_619C	TxBD pointer for ring n (ETSEC3_TBPTR3)	32	R/W	0000_0000h	14.52.18/807
2_61A4	TxBD pointer for ring n (ETSEC3_TBPTR4)	32	R/W	0000_0000h	14.52.18/807
2_61AC	TxBD pointer for ring n (ETSEC3_TBPTR5)	32	R/W	0000_0000h	14.52.18/807
2_61B4	TxBD pointer for ring n (ETSEC3_TBPTR6)	32	R/W	0000_0000h	14.52.18/807
2_61BC	TxBD pointer for ring n (ETSEC3_TBPTR7)	32	R/W	0000_0000h	14.52.18/807
2_6200	TxBD base address high bits (ETSEC3_TBASEH)	32	R/W	0000_0000h	14.52.19/808
2_6204	TxBD base address of ring n (ETSEC3_TBASE0)	32	R/W	0000_0000h	14.52.20/809
2_620C	TxBD base address of ring n (ETSEC3_TBASE1)	32	R/W	0000_0000h	14.52.20/809
2_6214	TxBD base address of ring n (ETSEC3_TBASE2)	32	R/W	0000_0000h	14.52.20/809
2_621C	TxBD base address of ring n (ETSEC3_TBASE3)	32	R/W	0000_0000h	14.52.20/809
2_6224	TxBD base address of ring n (ETSEC3_TBASE4)	32	R/W	0000_0000h	14.52.20/809
2_622C	TxBD base address of ring n (ETSEC3_TBASE5)	32	R/W	0000_0000h	14.52.20/809
2_6234	TxBD base address of ring n (ETSEC3_TBASE6)	32	R/W	0000_0000h	14.52.20/809
2_623C	TxBD base address of ring n (ETSEC3_TBASE7)	32	R/W	0000_0000h	14.52.20/809
2_6280	Tx time stamp identification tag (set n) register (ETSEC3_TMR_TXTS1_ID)	32	R	0000_0000h	14.52.21/809
2_6284	Tx time stamp identification tag (set n) register (ETSEC3_TMR_TXTS2_ID)	32	R	0000_0000h	14.52.21/809
2_62C0	Tx time stamp high (ETSEC3_TMR_TXTS1_H)	32	R	0000_0000h	14.52.22/810
2_62C4	Tx time stamp low (ETSEC3_TMR_TXTS1_L)	32	R	0000_0000h	14.52.23/810

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_62C8	Tx time stamp high (ETSEC3_TMR_TXTS2_H)	32	R	0000_0000h	14.52.22/810
2_62CC	Tx time stamp low (ETSEC3_TMR_TXTS2_L)	32	R	0000_0000h	14.52.23/810
2_6300	Receive control register (ETSEC3_RCTRL)	32	R/W	0000_0000h	14.52.24/811
2_6304	Receive status register (ETSEC3_RSTAT)	32	w1c	0000_0000h	14.52.25/814
2_6310	Receive interrupt coalescing register (ETSEC3_RXIC)	32	R/W	0000_0000h	14.52.26/816
2_6314	Receive queue control register. (ETSEC3_RQUEUE)	32	R/W	0080_0080h	14.52.27/817
2_6330	Receive bit field extract control register (ETSEC3_RBIFX)	32	R/W	0000_0000h	14.52.28/819
2_6334	Receive queue filing table address register (ETSEC3_RQFAR)	32	R/W	0000_0000h	14.52.29/821
2_6338	Receive queue filing table control register (ETSEC3_RQFCR)	32	R/W	See section	14.52.30/822
2_633C	Receive queue filing table property register (ETSEC3_RQFPR)	32	R/W	0000_0000h	14.52.31/824
2_6340	Maximum receive buffer length register (ETSEC3_MRBLR)	32	R/W	0000_0000h	14.52.32/830
2_6380	Rx data buffer pointer high bits (ETSEC3_RBDBPH)	32	R/W	0000_0000h	14.52.33/831
2_6384	RxBD pointer for ring n (ETSEC3_RBPTR0)	32	R/W	0000_0000h	14.52.34/831
2_638C	RxBD pointer for ring n (ETSEC3_RBPTR1)	32	R/W	0000_0000h	14.52.34/831
2_6394	RxBD pointer for ring n (ETSEC3_RBPTR2)	32	R/W	0000_0000h	14.52.34/831
2_639C	RxBD pointer for ring n (ETSEC3_RBPTR3)	32	R/W	0000_0000h	14.52.34/831
2_63A4	RxBD pointer for ring n (ETSEC3_RBPTR4)	32	R/W	0000_0000h	14.52.34/831
2_63AC	RxBD pointer for ring n (ETSEC3_RBPTR5)	32	R/W	0000_0000h	14.52.34/831
2_63B4	RxBD pointer for ring n (ETSEC3_RBPTR6)	32	R/W	0000_0000h	14.52.34/831
2_63BC	RxBD pointer for ring n (ETSEC3_RBPTR7)	32	R/W	0000_0000h	14.52.34/831
2_6400	RxBD base address high bits (ETSEC3_RBASEH)	32	R/W	0000_0000h	14.52.35/832
2_6404	RxBD base address of ring n (ETSEC3_RBASE0)	32	R/W	0000_0000h	14.52.36/833

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_640C	RxBD base address of ring n (ETSEC3_RBASE1)	32	R/W	0000_0000h	14.52.36/833
2_6414	RxBD base address of ring n (ETSEC3_RBASE2)	32	R/W	0000_0000h	14.52.36/833
2_641C	RxBD base address of ring n (ETSEC3_RBASE3)	32	R/W	0000_0000h	14.52.36/833
2_6424	RxBD base address of ring n (ETSEC3_RBASE4)	32	R/W	0000_0000h	14.52.36/833
2_642C	RxBD base address of ring n (ETSEC3_RBASE5)	32	R/W	0000_0000h	14.52.36/833
2_6434	RxBD base address of ring n (ETSEC3_RBASE6)	32	R/W	0000_0000h	14.52.36/833
2_643C	RxBD base address of ring n (ETSEC3_RBASE7)	32	R/W	0000_0000h	14.52.36/833
2_64C0	Rx timer time stamp register high (ETSEC3_TMR_RXTS_H)	32	R	0000_0000h	14.52.37/833
2_64C4	Rx timer time stamp register low (ETSEC3_TMR_RXTS_L)	32	R	0000_0000h	14.52.38/834
2_6500	MAC configuration register 1 (ETSEC3_MACCFG1)	32	R/W	0000_0000h	14.52.39/834
2_6504	MAC configuration register 2 (ETSEC3_MACCFG2)	32	R/W	0000_7000h	14.52.40/837
2_6508	Inter-packet/inter-frame gap register (ETSEC3_IPGIFG)	32	R/W	4060_5060h	14.52.41/839
2_650C	Half-duplex control (ETSEC3_HAFDUP)	32	R/W	00A1_F037h	14.52.42/840
2_6510	Maximum frame length (ETSEC3_MAXFRM)	32	R/W	0000_0600h	14.52.43/841
2_6520	MII management configuration (ETSEC3_MIIMCFG)	32	R/W	0000_0007h	14.52.44/842
2_6524	MII management command (ETSEC3_MIIMCOM)	32	R/W	0000_0000h	14.52.45/844
2_6528	MII management address (ETSEC3_MIIMADD)	32	R/W	0000_0000h	14.52.46/845
2_652C	MII management control (ETSEC3_MIIMCON)	32	W	0000_0000h	14.52.47/845
2_6530	MII management status (ETSEC3_MIIMSTAT)	32	R	0000_0000h	14.52.48/846
2_6534	MII management indicator (ETSEC3_MIIMIND)	32	R	0000_0000h	14.52.49/846
2_653C	Interface status (ETSEC3_IFSTAT)	32	R	0000_0000h	14.52.50/847
2_6540	MAC station address register 1 (ETSEC3_MACSTNADDR1)	32	R/W	0000_0000h	14.52.51/849

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_6544	MAC station address register 2 (ETSEC3_MACSTNADDR2)	32	R/W	0000_0000h	14.52.52/849
2_6548	MAC exact match address n, part 1 (ETSEC3_MAC01ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_654C	MAC exact match address n, part 2 (ETSEC3_MAC01ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6550	MAC exact match address n, part 1 (ETSEC3_MAC02ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_6554	MAC exact match address n, part 2 (ETSEC3_MAC02ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6558	MAC exact match address n, part 1 (ETSEC3_MAC03ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_655C	MAC exact match address n, part 2 (ETSEC3_MAC03ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6560	MAC exact match address n, part 1 (ETSEC3_MAC04ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_6564	MAC exact match address n, part 2 (ETSEC3_MAC04ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6568	MAC exact match address n, part 1 (ETSEC3_MAC05ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_656C	MAC exact match address n, part 2 (ETSEC3_MAC05ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6570	MAC exact match address n, part 1 (ETSEC3_MAC06ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_6574	MAC exact match address n, part 2 (ETSEC3_MAC06ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6578	MAC exact match address n, part 1 (ETSEC3_MAC07ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_657C	MAC exact match address n, part 2 (ETSEC3_MAC07ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6580	MAC exact match address n, part 1 (ETSEC3_MAC08ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_6584	MAC exact match address n, part 2 (ETSEC3_MAC08ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6588	MAC exact match address n, part 1 (ETSEC3_MAC09ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_658C	MAC exact match address n, part 2 (ETSEC3_MAC09ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6590	MAC exact match address n, part 1 (ETSEC3_MAC10ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_6594	MAC exact match address n, part 2 (ETSEC3_MAC10ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6598	MAC exact match address n, part 1 (ETSEC3_MAC11ADDR1)	32	R/W	0000_0000h	14.52.53/850

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_659C	MAC exact match address n, part 2 (ETSEC3_MAC11ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_65A0	MAC exact match address n, part 1 (ETSEC3_MAC12ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_65A4	MAC exact match address n, part 2 (ETSEC3_MAC12ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_65A8	MAC exact match address n, part 1 (ETSEC3_MAC13ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_65AC	MAC exact match address n, part 2 (ETSEC3_MAC13ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_65B0	MAC exact match address n, part 1 (ETSEC3_MAC14ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_65B4	MAC exact match address n, part 2 (ETSEC3_MAC14ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_65B8	MAC exact match address n, part 1 (ETSEC3_MAC15ADDR1)	32	R/W	0000_0000h	14.52.53/850
2_65BC	MAC exact match address n, part 2 (ETSEC3_MAC15ADDR2)	32	R/W	0000_0000h	14.52.54/851
2_6680	Transmit and receive 64-byte frame counter (ETSEC3_TR64)	32	R/W	0000_0000h	14.52.55/851
2_6684	Transmit and receive 65- to 127-byte frame counter (ETSEC3_TR127)	32	R/W	0000_0000h	14.52.56/852
2_6688	Transmit and receive 128- to 255-byte frame counter (ETSEC3_TR255)	32	R/W	0000_0000h	14.52.57/852
2_668C	Transmit and receive 256- to 511-byte frame counter (ETSEC3_TR511)	32	R/W	0000_0000h	14.52.58/853
2_6690	Transmit and receive 512- to 1023-byte frame counter (ETSEC3_TR1K)	32	R/W	0000_0000h	14.52.59/853
2_6694	Transmit and receive 1024- to 1518-byte frame counter (ETSEC3_TRMAX)	32	R/W	0000_0000h	14.52.60/853
2_6698	Transmit and receive 1519- to 1522-byte good VLAN frame count (ETSEC3_TRMGV)	32	R/W	0000_0000h	14.52.61/854
2_669C	Receive byte counter (ETSEC3_RBYT)	32	R/W	0000_0000h	14.52.62/854
2_66A0	Receive packet counter (ETSEC3_RPKT)	32	R/W	0000_0000h	14.52.63/855
2_66A4	Receive FCS error counter (ETSEC3_RFCS)	32	R/W	0000_0000h	14.52.64/855
2_66A8	Receive multicast packet counter (ETSEC3_RMCA)	32	R/W	0000_0000h	14.52.65/855
2_66AC	Receive broadcast packet counter (ETSEC3_RBCA)	32	R/W	0000_0000h	14.52.66/856
2_66B0	Receive control frame packet counter (ETSEC3_RXCF)	32	R/W	0000_0000h	14.52.67/856

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_66B4	Receive PAUSE frame packet counter (ETSEC3_RXPF)	32	R/W	0000_0000h	14.52.68/857
2_66B8	Receive unknown OP code counter (ETSEC3_RXUO)	32	R/W	0000_0000h	14.52.69/857
2_66BC	Receive alignment error counter (ETSEC3_RALN)	32	R/W	0000_0000h	14.52.70/857
2_66C0	Receive frame length error counter (ETSEC3_RFLR)	32	R/W	0000_0000h	14.52.71/858
2_66C4	Receive code error counter (ETSEC3_RCDE)	32	R/W	0000_0000h	14.52.72/858
2_66C8	Receive carrier sense error counter (ETSEC3_RCSE)	32	R/W	0000_0000h	14.52.73/859
2_66CC	Receive undersize packet counter (ETSEC3_RUND)	32	R/W	0000_0000h	14.52.74/859
2_66D0	Receive oversize packet counter (ETSEC3_ROVR)	32	R/W	0000_0000h	14.52.75/859
2_66D4	Receive fragments counter (ETSEC3_RFRG)	32	R/W	0000_0000h	14.52.76/860
2_66D8	Receive jabber counter (ETSEC3_RJBR)	32	R/W	0000_0000h	14.52.77/860
2_66DC	Receive drop counter (ETSEC3_RDRP)	32	R/W	0000_0000h	14.52.78/861
2_66E0	Transmit byte counter (ETSEC3_TBYT)	32	R/W	0000_0000h	14.52.79/861
2_66E4	Transmit packet counter (ETSEC3_TPKT)	32	R/W	0000_0000h	14.52.80/861
2_66E8	Transmit multicast packet counter (ETSEC3_TMCA)	32	R/W	0000_0000h	14.52.81/862
2_66EC	Transmit broadcast packet counter (ETSEC3_TBCA)	32	R/W	0000_0000h	14.52.82/862
2_66F0	Transmit PAUSE control frame counter (ETSEC3_TXPF)	32	R/W	0000_0000h	14.52.83/863
2_66F4	Transmit deferral packet counter (ETSEC3_TDFR)	32	R/W	0000_0000h	14.52.84/863
2_66F8	Transmit excessive deferral packet counter (ETSEC3_TEDF)	32	R/W	0000_0000h	14.52.85/863
2_66FC	Transmit single collision packet counter (ETSEC3_TSCL)	32	R/W	0000_0000h	14.52.86/864
2_6700	Transmit multiple collision packet counter (ETSEC3_TMCL)	32	R/W	0000_0000h	14.52.87/864
2_6704	Transmit late collision packet counter (ETSEC3_TLCL)	32	R/W	0000_0000h	14.52.88/865
2_6708	Transmit excessive collision packet counter (ETSEC3_TXCL)	32	R/W	0000_0000h	14.52.89/865

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_670C	Transmit total collision counter (ETSEC3_TNCL)	32	R/W	0000_0000h	14.52.90/865
2_6714	Transmit drop frame counter (ETSEC3_TDRP)	32	R/W	0000_0000h	14.52.91/866
2_6718	Transmit jabber frame counter (ETSEC3_TJBR)	32	R/W	0000_0000h	14.52.92/866
2_671C	Transmit FCS error counter (ETSEC3_TFCS)	32	R/W	0000_0000h	14.52.93/867
2_6720	Transmit control frame counter (ETSEC3_TXCF)	32	R/W	0000_0000h	14.52.94/867
2_6724	Transmit oversize frame counter (ETSEC3_TOVR)	32	R/W	0000_0000h	14.52.95/867
2_6728	Transmit undersize frame counter (ETSEC3_TUND)	32	R/W	0000_0000h	14.52.96/868
2_672C	Transmit fragments frame counter (ETSEC3_TFRG)	32	R/W	0000_0000h	14.52.97/868
2_6730	Carry register one registerCleared on read. (ETSEC3_CAR1)	32	w1c	0000_0000h	14.52.98/869
2_6734	Carry register two register 3 (ETSEC3_CAR2)	32	R/W	0000_0000h	14.52.99/871
2_6738	Carry register one mask register (ETSEC3_CAM1)	32	R/W	FE03_FFFFh	14.52.100/873
2_673C	Carry register two mask register (ETSEC3_CAM2)	32	R/W	000F_FFFDh	14.52.101/875
2_6740	Receive filer rejected packet counter (ETSEC3_RREJ)	32	R/W	0000_0000h	14.52.102/876
2_6800	Individual/group address register n (ETSEC3_IGADDR0)	32	R/W	0000_0000h	14.52.103/877
2_6804	Individual/group address register n (ETSEC3_IGADDR1)	32	R/W	0000_0000h	14.52.103/877
2_6808	Individual/group address register n (ETSEC3_IGADDR2)	32	R/W	0000_0000h	14.52.103/877
2_680C	Individual/group address register n (ETSEC3_IGADDR3)	32	R/W	0000_0000h	14.52.103/877
2_6810	Individual/group address register n (ETSEC3_IGADDR4)	32	R/W	0000_0000h	14.52.103/877
2_6814	Individual/group address register n (ETSEC3_IGADDR5)	32	R/W	0000_0000h	14.52.103/877
2_6818	Individual/group address register n (ETSEC3_IGADDR6)	32	R/W	0000_0000h	14.52.103/877
2_681C	Individual/group address register n (ETSEC3_IGADDR7)	32	R/W	0000_0000h	14.52.103/877
2_6880	Group address register n (ETSEC3_GADDR0)	32	R/W	0000_0000h	14.52.104/878

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_6884	Group address register n (ETSEC3_GADDR1)	32	R/W	0000_0000h	14.52.104/878
2_6888	Group address register n (ETSEC3_GADDR2)	32	R/W	0000_0000h	14.52.104/878
2_688C	Group address register n (ETSEC3_GADDR3)	32	R/W	0000_0000h	14.52.104/878
2_6890	Group address register n (ETSEC3_GADDR4)	32	R/W	0000_0000h	14.52.104/878
2_6894	Group address register n (ETSEC3_GADDR5)	32	R/W	0000_0000h	14.52.104/878
2_6898	Group address register n (ETSEC3_GADDR6)	32	R/W	0000_0000h	14.52.104/878
2_689C	Group address register n (ETSEC3_GADDR7)	32	R/W	0000_0000h	14.52.104/878
2_6BF8	Attribute register (ETSEC3_ATTR)	32	R/W	0000_0000h	14.52.105/879
2_6BFC	Attribute extract length and extract index register (ETSEC3_ATTRELI)	32	R/W	0000_0000h	14.52.106/881
2_6C00	Receive Queue Parameters register n (ETSEC3_RQPRM0)	32	R/W	0000_0000h	14.52.107/882
2_6C04	Receive Queue Parameters register n (ETSEC3_RQPRM1)	32	R/W	0000_0000h	14.52.107/882
2_6C08	Receive Queue Parameters register n (ETSEC3_RQPRM2)	32	R/W	0000_0000h	14.52.107/882
2_6C0C	Receive Queue Parameters register n (ETSEC3_RQPRM3)	32	R/W	0000_0000h	14.52.107/882
2_6C10	Receive Queue Parameters register n (ETSEC3_RQPRM4)	32	R/W	0000_0000h	14.52.107/882
2_6C14	Receive Queue Parameters register n (ETSEC3_RQPRM5)	32	R/W	0000_0000h	14.52.107/882
2_6C18	Receive Queue Parameters register n (ETSEC3_RQPRM6)	32	R/W	0000_0000h	14.52.107/882
2_6C1C	Receive Queue Parameters register n (ETSEC3_RQPRM7)	32	R/W	0000_0000h	14.52.107/882
2_6C44	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR0)	32	R/W	0000_0000h	14.52.108/883
2_6C4C	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR1)	32	R/W	0000_0000h	14.52.108/883
2_6C54	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR2)	32	R/W	0000_0000h	14.52.108/883
2_6C5C	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR3)	32	R/W	0000_0000h	14.52.108/883
2_6C64	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR4)	32	R/W	0000_0000h	14.52.108/883

Table continues on the next page...

ETSEC memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_6C6C	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR5)	32	R/W	0000_0000h	14.52.108/883
2_6C74	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR6)	32	R/W	0000_0000h	14.52.108/883
2_6C7C	Last Free RxBD pointer for ring n (ETSEC3_RFBPTR7)	32	R/W	0000_0000h	14.52.108/883

14.52.1 Controller ID register (ETSECx_TSEC_ID)

The controller ID register (TSEC_ID) is a read-only register. The TSEC_ID register is used to identify the eTSEC block and revision.

Address: Base address + 0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	TSEC_ID															TSEC_REV_MJ							TSEC_REV_MN										
W	[Shaded]																																
Reset	0	0	0	0	0	0	0	0	1	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TSEC_ID field descriptions

Field	Description
0–15 TSEC_ID	Value identifying the eTSEC (10/100/1000 Ethernet MAC). 0124 Unique identifier for eTSEC with 8 Rx and 8 Tx BD rings.
16–23 TSEC_REV_MJ	Value identifies the major revision of the eTSEC. 00 Initial revision
24–31 TSEC_REV_MN	Value identifies the minor revision of the eTSEC.

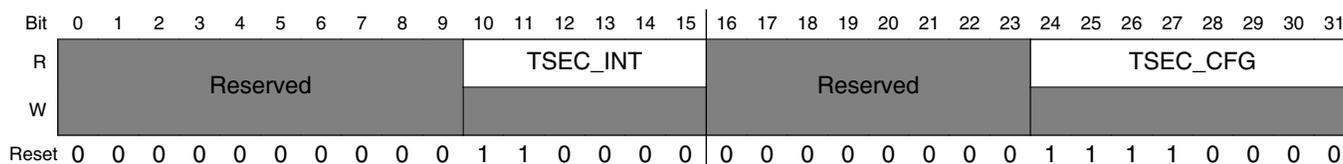
14.52.2 Controller ID register (ETSECx_TSEC_ID2)

The controller ID register (TSEC_ID2) is a read-only register. The TSEC_ID2 register is used to identify the eTSEC block configuration.

Table 14-217. TSEC_ID2 [TSEC_INT] Field Settings

Bit	Mode
10	0 - Ethernet mode not supported 1 - Ethernet supported
11-13	Reserved
14	0 - Can be configured to run in Ethernet normal/full mode 1 - Ethernet normal/full mode off
15	0 - Can be configured to run in Ethernet reduced mode 1 - Ethernet reduced mode off

Address: Base address + 4h offset



ETSECx_TSEC_ID2 field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–15 TSEC_INT	Interface mode support. See Table 14-6 for settings.
16–23 -	This field is reserved. Reserved
24–31 TSEC_CFG	Value identifies configuration options of the eTSEC. 00 eTSEC multiple ring, Rx TOE, Filer and Tx TOE supports are off F0 eTSEC multiple ring, Rx TOE, Filer and Tx TOE supports are on 30 eTSEC multiple ring support is OFF and Rx TOE, Filer and Tx TOE supports are on 50 eTSEC multiple ring and filer supports are OFF and Rx TOE and Tx TOE supports are on

14.52.3 Interrupt event register (ETSECx_IEVENT)

Interrupt events cause bits in the IEVENT register to be set. Software may poll this register at any time to check for pending interrupts. If an event occurs and its corresponding enable bit is set in the interrupt mask register (IMASK), the event also causes a hardware interrupt at the PIC. A bit in the interrupt event register is cleared by writing a 1 to that bit position. A write of 0 has no effect.

Each eTSEC can issue three kinds of hardware interrupt to the PIC:

1. Transmit data frame interrupts-Issued whenever bits TXB or TXF of IEVENT are set to 1 and either transmit interrupt coalescing is disabled or the interrupt coalescing thresholds have been met for TXF. To negate this hardware interrupt, software must clear both TXB and TXF bits.
 2. Receive data frame interrupts-Issued whenever bits RXB or RXF of IEVENT are set to 1 and either receive interrupt coalescing is disabled or the interrupt coalescing thresholds have been met for RXF. To negate this hardware interrupt, software must clear both RXB and RXF bits.
 3. Error, diagnostic, and special interrupts-Issued whenever bits MAG, GTSC, GRSC, TXC, RXC, BABR, BABT, LC, CRL , FGPI , FIR, FIQ, DPE, PERR, EBERR, TXE, XFUN, BSY, MSRO, MMRD, or MMRW of IEVENT are set to 1. Software must clear all of these bits to negate an error/diagnostic/special hardware interrupt.
- Magic Packet reception event is: MAG
 - Operational diagnostics are events on: GTSC, GRSC, TXC, and RXC
 - Interrupts resulting from errors/problems detected in the network or transceiver are: BABR, BABT, LC, and CRL
 - Interrupts resulting from internal or combination errors are: FIR, FIQ, DPE, PERR, EBERR, TXE, XFUN, and BSY
 - Special function interrupts are: FGPI, MSRO, MMRD, and MMRW

Some of the error interrupts are independently counted in the MIB block counters. Software may choose to mask off these interrupts because these errors are visible to network management through the MIB counters.

Address: Base address + 10h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	BABR	RXC	BSY	EBERR	Reserved	MSRO	GTSC	BABT	TXC	TXE	TXB	TXF	Reserved	LC	CRL	XFUN
W	w1c	w1c	w1c	w1c		w1c	w1c	w1c	w1c	w1c	w1c	w1c		w1c	w1c	w1c
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	RXB	Reserved			MAG	MMRD	MMWR	GRSC	RXF	Reserved		FGPI	FIR	FIQ	DPE	PERR
W	w1c	Reserved			w1c	w1c	w1c	w1c	w1c	Reserved		w1c	w1c	w1c	w1c	w1c
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_IEVENT field descriptions

Field	Description
0 BABR	Babbling receive error. This bit indicates that a frame was received with length in excess of the MAC's maximum frame length register while MACCFG2[Huge Frame] is set. 0 Excessive frame not received. 1 Excessive frame received.
1 RXC	Receive control interrupt. A control frame was received while MACCFG1[Rx_Flow] is set. As soon as the transmitter finishes sending the current frame, a pause operation is performed. 0 Control frame not received. 1 Control frame received.
2 BSY	Busy condition interrupt. Indicates that a frame was received and discarded due to a lack of buffers. 0 No frame received and discarded. 1 Frame received and discarded.
3 EBERR	Internal bus error. This bit indicates that a system bus error occurred while a DMA transaction was underway. As a result, transferred data is expected to be partially or completely invalid. 0 No system bus error occurred. 1 System bus error occurred.
4 -	This field is reserved. Reserved
5 MSRO	MIB counter overflow. This interrupt is asserted if the count for one of the MIB counters has exceeded the size of its register. 0 MIB count not exceeding its register size. 1 MIB count exceeds its register size.
6 GTSC	Graceful transmit stop complete. This interrupt is asserted for one of two reasons. Graceful stop means that the transmitter is put into a pause state after completion of the frame currently being transmitted.

Table continues on the next page...

ETSECx_IEVENT field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> A graceful stop, which was initiated by setting DMACTRL[GTS], is now complete. A transmission of a flow control PAUSE frame, which was initiated by setting TCTRL[TFC_PAUSE], is now complete. <p>0 No graceful stop interrupt. 1 Graceful stop requested.</p>
7 BABT	<p>Babbling transmit error. This bit indicates that the transmitted frame length has exceeded the value in the MAC's maximum frame length register and MACCFG2[Huge Frame] is cleared. Frame truncation occurs when this condition occurs.</p> <p>0 Transmitted frame length not exceeding maximum frame length. 1 Transmitted frame length exceeding maximum frame length when MACCFG2[Huge Frame] = 0.</p>
8 TXC	<p>Transmit control interrupt. This bit indicates that a control frame was transmitted.</p> <p>0 Control frame not transmitted. 1 Control frame transmitted.</p>
9 TXE	<p>Transmit error. This bit indicates that an error occurred on the transmitted channel that has caused TSTAT[THLT] to be set by the eTSEC. This bit is set whenever any transmit error occurs that causes the transmitter to halt (EBERR, LC, CRL, XFUN).</p> <p>0 No transmit channel error occurred. 1 Transmit channel error occurred.</p>
10 TXB	<p>Transmit buffer. This bit indicates that a transmit buffer descriptor was updated whose I (interrupt) bit was set in its status word and was not the last buffer descriptor of the frame.</p> <p>0 No transmit buffer descriptor updated. 1 Transmit buffer descriptor updated.</p>
11 TXF	<p>Transmit frame interrupt. This bit indicates that a frame was transmitted and that the last corresponding transmit buffer descriptor (TxBD) was updated. This only occurs if the I (interrupt) bit in the status word of the buffer descriptor is set. The specific transmit queue that was updated has its TXF bit set in TSTAT.</p> <p>0 No frame transmitted/TxBD not updated. 1 Frame transmitted/TxBD updated.</p>
12 -	<p>This field is reserved. Reserved</p>
13 LC	<p>Late collision. This bit indicates that a collision occurred beyond the collision window (slot time) in half-duplex mode. The frame is truncated with a bad CRC and the remainder of the frame is discarded.</p> <p>0 No late collision occurred. 1 Late collision occurred.</p>
14 CRL	<p>Collision retry limit. This bit indicates that the number of successive transmission collisions has exceeded the MAC's half-duplex register's retransmission maximum count (HAFDUP[Retransmission Maximum]). The frame is discarded without being transmitted and the queue halts (TSTAT[THLT n] set to 1). This only occurs while in half-duplex mode.</p> <p>0 Successive transmission collisions do not exceed maximum. 1 Successive transmission collisions exceed maximum.</p>
15 XFUN	<p>Transmit FIFO underrun. This bit indicates that the transmit FIFO became empty before the complete frame was transmitted.</p>

Table continues on the next page...

ETSECx_IEVENT field descriptions (continued)

Field	Description
	0 Transmit FIFO not underrun. 1 Transmit FIFO underrun.
16 RXB	Receive buffer. This bit indicates that a receive buffer descriptor was updated which had the I (Interrupt) bit set in its status word and was not the last buffer descriptor of the frame. 0 Receive buffer descriptor not updated. 1 Receiver buffer descriptor updated.
17–19 -	This field is reserved. Reserved
20 MAG	Magic Packet detected when the eTSEC is in Magic Packet detection mode (MACCFG2[MPEN] = 1). 0 No Magic Packet received, or Magic Packet mode was not enabled. 1 A Magic Packet was received while in Magic Packet mode. MACCFG2[MPEN] is also cleared upon receiving the Magic Packet.
21 MMRD	MII management read completion 0 MII management read not issued or in process. 1 MII management read completed that was initiated by a user through the MII Scan or Read cycle command.
22 MMWR	MII management write completion 0 MII management write not issued or in process. 1 MII management write completed that was initiated by a user write to the MIIMCON register.
23 GRSC	Graceful receive stop complete. This interrupt is asserted if a graceful receive stop is completed. It allows the user to know if the system has completed the stop and it is safe to write to receive registers (status, control or configuration registers) that are used by the system during normal operation. 0 Graceful stop not completed. 1 Graceful stop completed.
24 RXF	Receive frame interrupt. This bit indicates that a frame was received and the last receive buffer descriptor (RxBD) in that frame was updated. This occurs either if the I (interrupt) bit in the buffer descriptor status word is set, or an overrun error occurs. The specific receive queue that was updated has its RXF bit set in RSTAT. 0 Frame not received. 1 Frame received.
25–26 -	This field is reserved. Reserved
27 FGPI	Filer generated general purpose interrupt on a set of filer rule match. This bit will be set upon reception of a frame that matches a GPI rule sequence that is specified in the filer. It is synchronized with the setting of RXF. 0 No filer generated interrupt has occurred. 1 The filer has accepted a frame via a matching rule that the RQFCR[GPI] bit set.
28 FIR	The receive queue filer result is invalid, either because not enough time between frames was available to find a matching rule, or no entry in the filer table could be matched.

Table continues on the next page...

ETSECx_IEVENT field descriptions (continued)

Field	Description
	<p>0 Receive queue filer reached a definite result; however, bit FIQ may still be set if a frame was filed to a disabled RxBD ring.</p> <p>1 Receive queue filer was unable to reach a definite result. In this case, bit FIQ is also set if no entry in the filer table could provide a rule match.</p>
29 FIQ	<p>Filed frame to invalid receive queue. This bit indicates that either the receive queue filer chose to DMA a received frame to a disabled RxBD ring, or that no rule in the filer table could be matched.</p> <p>0 Received frames filed to valid queues or rejected. Note that a frame may be rejected if the filer has insufficient time to reach a conclusive result between frames, in which case bit FIR is set.</p> <p>1 Received frames filed to RxBD rings that are not enabled. The frame is discarded. If bit FIR is also set this indicates that the filer exhausted all of its table entries without a rule match.</p>
30 DPE	<p>Internal data parity error. This bit indicates that the eTSEC has detected a parity error on its stored data, which is likely to compromise the validity of recently transferred frames.</p> <p>0 No parity errors detected.</p> <p>1 Data held in the FIFO or filer arrays is expected to be corrupted due to a parity error.</p>
31 PERR	<p>Receive frame parse error for TCP/IP off-load. This bit indicates that a received frame could not be parsed unambiguously, due to encapsulated header type fields contradicting each other.</p> <p>0 Received frame parsed successfully.</p> <p>1 Received frame parse revealed header inconsistencies.</p>

14.52.4 Interrupt mask register (ETSECx_IMASK)

The interrupt mask register provides control over which possible interrupt events in the IEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset. If the corresponding bits in both the IEVENT and IMASK registers are set, the PIC receives an interrupt (for each eTSEC these are grouped into transmit, receive, and error/diagnostic interrupts). The interrupt signal remains asserted until either the IEVENT bit is cleared, by writing a 1 to it, or by writing a 0 to the corresponding IMASK bit.

Address: Base address + 14h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R					Reserved								Reserved				
W	BREN	RXCEN	BSYEN	EBERREN	Reserved	MSROEN	GTSCEN	BTEN	TXCEN	TXEEN	TXBEN	TXFEN	Reserved	LCEN	CRLN	XFUNEN	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R		Reserved								Reserved							
W	RXBEN	Reserved			MAGEN	MMRDEN	MMWREN	GRSCEN	RXFEN	Reserved			FGPIEN	FIREN	FIQEN	DPEEN	PERREN
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

ETSECx_IMASK field descriptions

Field	Description
0 BREN	Babbling receiver interrupt enable
1 RXCEN	Receive control interrupt enable
2 BSYEN	Busy interrupt enable

Table continues on the next page...

ETSECx_IMASK field descriptions (continued)

Field	Description
3 EBERREN	Ethernet controller bus error enable
4 -	This field is reserved. Reserved
5 MSROEN	MIB counter overflow interrupt enable
6 GTSCEN	Graceful transmit stop complete interrupt enable
7 BTEN	Babbling transmitter interrupt enable
8 TXCEN	Transmit control interrupt enable
9 TXEEN	Transmit error interrupt enable
10 TXBEN	Transmit buffer interrupt enable
11 TXFEN	Transmit frame interrupt enable
12 -	This field is reserved. Reserved
13 LCEN	Late collision enable
14 CRLEN	Collision retry limit enable
15 XFUNEN	Transmit FIFO underrun enable
16 RXBEN	Receive buffer interrupt enable
17–19 -	This field is reserved. Reserved
20 MAGEN	Magic packet received interrupt enable
21 MMRDEN	MII management read completion interrupt enable
22 MMWREN	MII management write completion interrupt enable
23 GRSCEN	Graceful receive stop complete interrupt enable
24 RXFEN	Receive frame interrupt enable
25–26 -	This field is reserved. Reserved
27 FGPIEN	Filer general purpose interrupt enable
28 FIREN	Filer invalid result interrupt enable

Table continues on the next page...

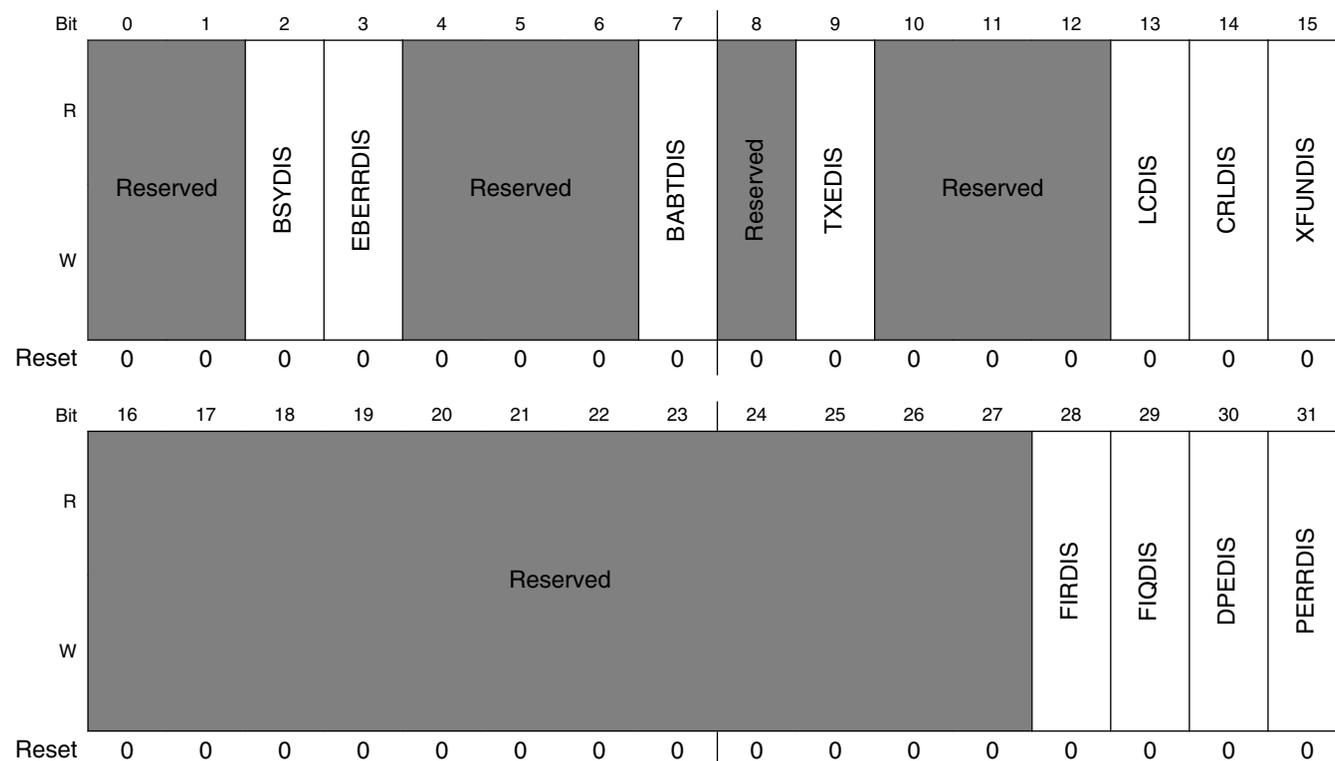
ETSECx_IMASK field descriptions (continued)

Field	Description
29 FIQEN	Filed frame to invalid queue interrupt enable
30 DPEEN	Data parity error interrupt enable
31 PERREN	Receive frame parse error enable

14.52.5 Error disabled register (ETSECx_EDIS)

The error disabled register allows the user to disable an error interruption, possibly to avoid spurious error indications external to the eTSECs.

Address: Base address + 18h offset



ETSECx_EDIS field descriptions

Field	Description
0-1 -	This field is reserved. Reserved

Table continues on the next page...

ETSECx_EDIS field descriptions (continued)

Field	Description
2 BSYDIS	<p>Busy disable.</p> <p>0 Allow eTSEC to report IEVENT[BSY] status and halt buffer descriptor queue if BSY condition occurs. 1 Do not set IEVENT[BSY] and do not halt buffer descriptor queue if BSY condition occurs.</p>
3 EBERRDIS	<p>Ethernet controller bus error disable.</p> <p>0 Allow eTSEC to report IEVENT[EBERR] status and halt buffer descriptor queue if EBERR condition occurs. 1 Do not set IEVENT[EBERR] and do not halt buffer descriptor queue if EBERR condition occurs.</p>
4–6 -	<p>This field is reserved. Reserved</p>
7 BABTDIS	<p>Babbling transmit error disable.</p> <p>0 Allow eTSEC to report IEVENT[BABT] status and set the buffer descriptor TR field. 1 Do not set IEVENT[BABT] nor the buffer descriptor TR field.</p>
8 -	<p>This field is reserved. Reserved</p>
9 TXEDIS	<p>Transmit error disable.</p> <p>0 Allow eTSEC to report IEVENT[TXE] status. 1 Do not set IEVENT[TXE] if TXE condition occurs.</p>
10–12 -	<p>This field is reserved. Reserved</p>
13 LCDIS	<p>Late collision disable.</p> <p>0 Allow eTSEC to report IEVENT[LC] status, set the buffer descriptor LC field, and halt buffer descriptor queue if LC condition occurs. 1 Do not set IEVENT[LC] nor the buffer descriptor LC field, and do not halt buffer descriptor queue if LC condition occurs.</p>
14 CRLDIS	<p>Collision retry limit disable.</p> <p>0 Allow eTSEC to report IEVENT[CRL] status, set the buffer descriptor RL field, and halt buffer descriptor queue if CRL condition occurs. 1 Do not set IEVENT[CRL] nor the buffer descriptor RL field, and do not halt buffer descriptor queue if CRL condition occurs.</p>
15 XFUNDIS	<p>Transmit FIFO underrun disable.</p> <p>0 Allow eTSEC to report IEVENT[XFUN] status, set the buffer descriptor UN field, and halt buffer descriptor queue if XFUN condition occurs. 1 Do not set IEVENT[XFUN] nor the buffer descriptor UN field, and do not halt buffer descriptor queue if XFUN condition occurs.</p>
16–27 -	<p>This field is reserved. Reserved</p>
28 FIRDIS	<p>Filer invalid result error disable.</p> <p>0 Allow eTSEC to report IEVENT[FIR] status. 1 Do not set IEVENT[FIR] if eTSEC fails to reach a definite filer result when attempting to file a received frame, but discard the frame silently.</p>

Table continues on the next page...

ETSECx_EDIS field descriptions (continued)

Field	Description
29 FIQDIS	Filed frame to invalid queue error disable. 0 Allow eTSEC to report IEVENT[FIQ] status. 1 Do not set IEVENT[FIQ] if eTSEC attempts to file a received frame to an invalid (disabled) RxBD ring, but discard the frame silently.
30 DPEDIS	Data parity error disable. 0 Allow eTSEC to report IEVENT[DPE] status. 1 Do not set IEVENT[DPE] if a parity error occurs in eTSEC's FIFO or filer arrays.
31 PERRDIS	Receive frame parse error disable. 0 Allow eTSEC to report IEVENT[PERR] status. 1 Do not set IEVENT[PERR] if a parse error occurs on a received frame.

14.52.6 Ethernet control register (ETSECx_ECNTL)

ECNTL is a register writable by the user to reset, configure, and initialize the eTSEC. Note that the GMIIM, TBIM, RPM, and RMM fields are read-only, having been set after sampling signals at power-on-reset.

The different interface configurations indicated by registers ECNTL and MACCFG2 are summarized in the following table.

Table 14-222. eTSEC Interface Configurations

Interface Mode	ECNTL Field							MACCFG2 Field
	FIFM ¹	GMIIM	TBIM	RPM	R100M	RMM	SGMIIM	I/F Mode
MII 10/100 Mbps	0	0	0	0	0	0	0	01
RMII 100 Mbps	0	0	0	0	1	1	0	01
RMII 10 Mbps	0	0	0	0	0	1	0	01
GMII 1Gbps ²	0	1	0	0	0	0	0	10
RGMI 1Gbps	0	1	0	1	0	0	0	10
RGMI 100 Mbps	0	1	0	1	1	0	0	01
RGMI 10 Mbps	0	1	0	1	0	0	0	01
TBI 1Gbps	0	0	1	0	0	0	0	10

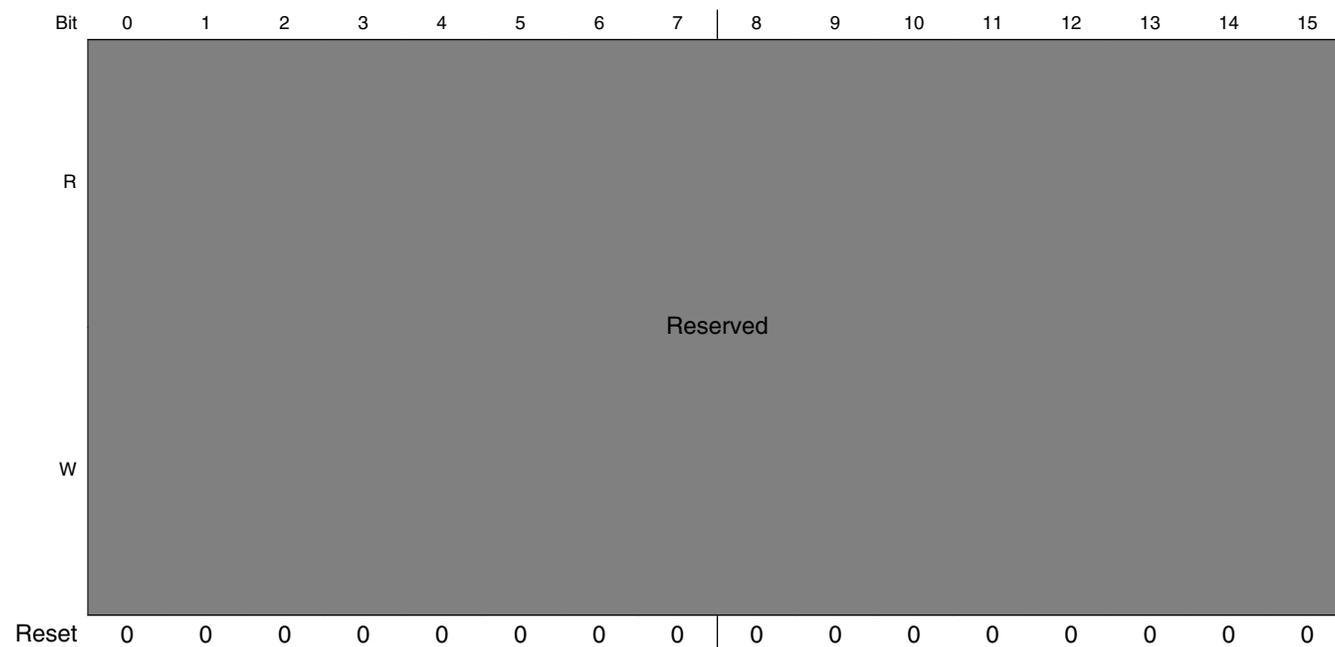
Table continues on the next page...

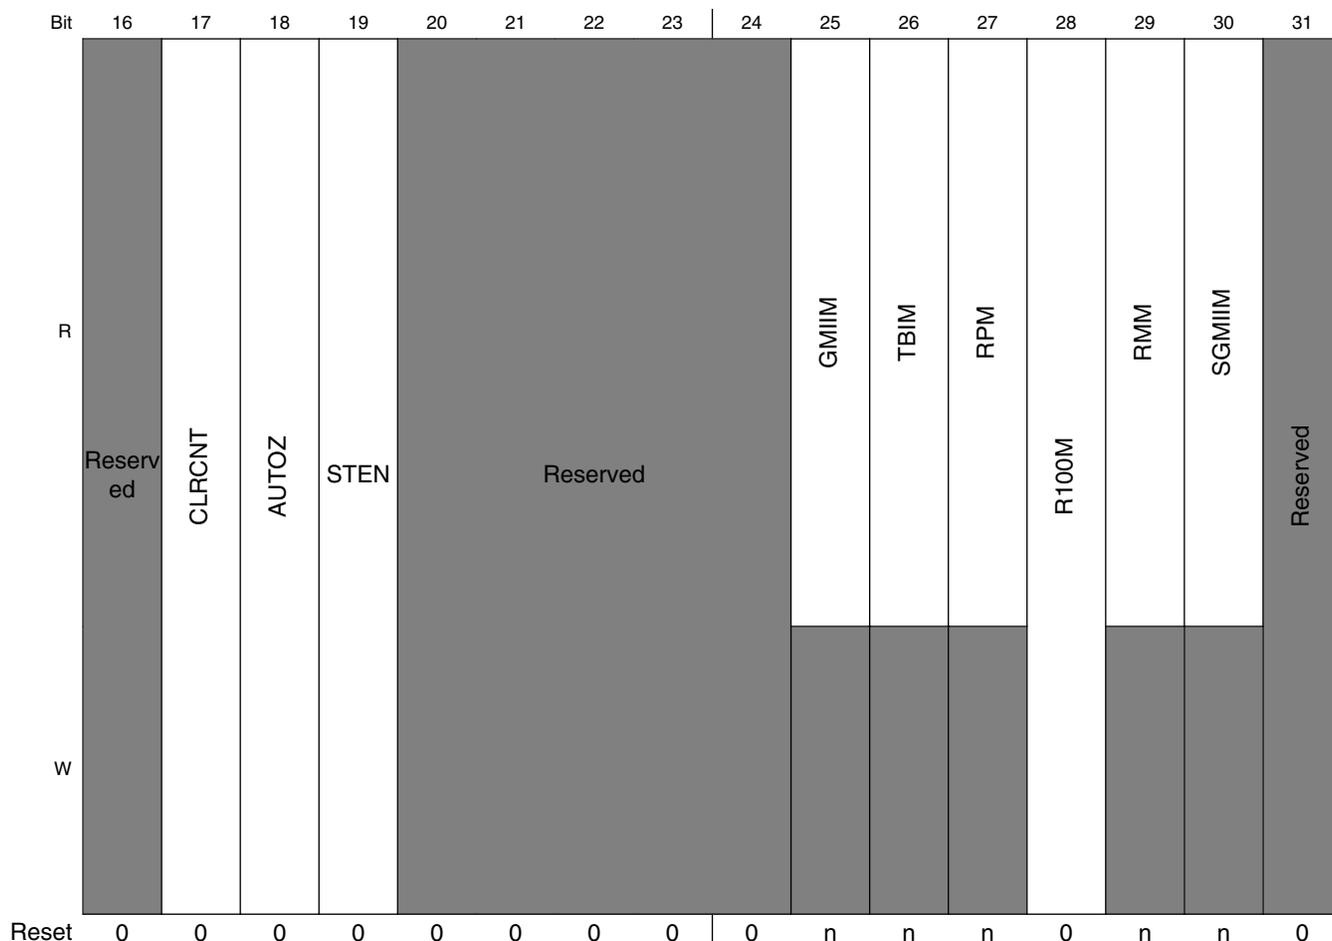
Table 14-222. eTSEC Interface Configurations (continued)

Interface Mode	ECNTRL Field							MACCFG2 Field
	FIFM ¹	GMIIM	TBIM	RPM	R100M	RMM	SGMIIM	I/F Mode
RTBI 1Gbps	0	0	1	1	0	0	0	10
SGMII 1 Gbps	0	0	1	0	0	0	1	10
SGMII 100 Mbps	0	0	1	0	1	0	1	01
SGMII 10 Mbps	0	0	1	0	0	0	1	01

- 1. FIFM bit is not supported.
- 2. See MII 10/100 Mbps mode for GMII 10/100 Mbps 'fall-back' mode.
- 2. See MII 10/100 Mbps mode for GMII 10/100 Mbps 'fall-back' mode.
- 2. See MII 10/100 Mbps mode for GMII 10/100 Mbps 'fall-back' mode.
- 2. See MII 10/100 Mbps mode for GMII 10/100 Mbps 'fall-back' mode.

Address: Base address + 20h offset





ETSECx_ECNTL field descriptions

Field	Description
0–16 -	This field is reserved. Reserved
17 CLRCNT	Clear all statistics counters and carry registers. This bit is self-resetting. 0 Allow MIB counters to continue to increment and keep any overflow indicators. 1 Reset all MIB counters and CAR1 and CAR2.
18 AUTOZ	Automatically zero MIB counter values and carry registers. This is a steady state signal and must be set prior to enabling the Ethernet controller and must not be changed without proper care. 0 The user must write the addressed counter zero after a host read. 1 The addressed counter value is automatically cleared to zero after a host read.
19 STEN	MIB counter statistics enabled. This is a steady state signal and must be set prior to enabling the Ethernet controller and must not be changed without proper care.

Table continues on the next page...

ETSECx_ECNTL field descriptions (continued)

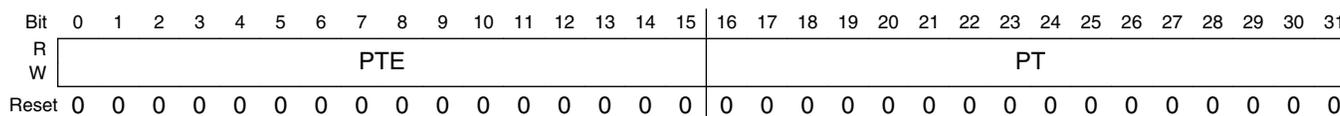
Field	Description
	<p>0 Statistics not enabled</p> <p>1 Enables internal counters to update</p>
20–24 -	This field is reserved. Reserved
25 GMIIM	<p>GMIIM interface mode. If this bit is set, a PHY with a gigabit GMII-type interface is expected to be connected. If cleared, a PHY with a non-gigabit or non-GMII interface is expected. The user should then set MACCFG2[I/F Mode] accordingly. The state of this status bit is defined during power-on reset. See Power-on reset configuration .</p> <p>0 RMII or MII mode interface expected</p> <p>1 RGMII or GMII mode interface expected</p>
26 TBIM	<p>Ten-bit interface mode. If this bit is set, ten-bit interface mode is enabled. This bit can be pin-configured at reset to set or clear. See Power-on reset configuration .</p> <p>0 Non ten-bit mode interface</p> <p>1 Ten-bit mode interface</p>
27 RPM	<p>Reduced-pin mode for Gigabit interfaces. If this bit is set, a reduced-pin interface is expected on Ethernet interfaces. RPM and RMM are never set together. This register can be pin-configured at reset to 0 or 1. See Power-on reset configuration .</p> <p>0 SGMII or GMII or MII or TBI in non-reduced-pin mode configuration</p> <p>1 RGMII or RTBI reduced-pin mode</p>
28 R100M	<p>RGMII /RMII 100 mode. This bit is ignored unless SGMII, RPM or RMM are set and MACCFG2[I/F Mode] is assigned to 10/100 (01).</p> <p>This bit must be cleared for 1-Gbps SGMII operation.</p> <p>0 RGMII is in 10 Mbps mode , RMII is in 10 Mbps mode, and every 10th RMII Reference clock is used to transfer data, SGMII is in 10 Mbps mode, and every 100th SGMII Reference clock is used to transfer data</p> <p>1 RGMII is in 100 Mbps mode , RMII is in 100 Mbps mode, and data is transferred on every Reference clock, SGMII is in 100 Mbps mode, and every 10th SGMII Reference clock is used to transfer data</p>
29 RMM	<p>Reduced-pin mode for 10/100 interfaces. If this bit is set, an RMII pin interface is expected. RMM must be 0 if RPM = 1. This register can be pin-configured at reset to 0 or 1. See Power-on reset configuration .</p> <p>0 Non-RMII interface mode</p> <p>1 RMII interface mode</p>
30 SGMIIM	<p>Serial GMII mode. If this bit is set, a SGMII pin interface is expected to be connected via an on chip SerDes.</p> <p>This register can be pin-configured at reset to 0 or 1. See Power-on reset configuration . Note that eTSEC1 does not support SGMII, so SGMIIM is always 0 for that.</p> <p>0 SGMII mode disabled. eTSEC connected via a parallel interface.</p> <p>1 SGMII mode enabled.</p>
31 -	This field is reserved. Reserved

1. FIFM bit is not supported.
2. See MII 10/100 Mbps mode for GMII 10/100 Mbps 'fall-back' mode.

14.52.7 Pause time value register (ETSECx_PTV)

PTV is a 32-bit register written by the user to store the pause duration used when the eTSEC initiates an IEEE 802.3 PAUSE control frame through TCTRL[TFC_PAUSE]. The low-order 16 bits (PT) represent the pause time and the high-order 16 bits (PTE) represent the extended pause control parameter. The pause time is measured in units of *pause_quanta*, equal to 512 bit times. The pause time can range from 0 to 65,535 *pause_quanta*, or 0 to 33,553,920 bit times. See [Flow control](#) for additional details.

Address: Base address + 28h offset



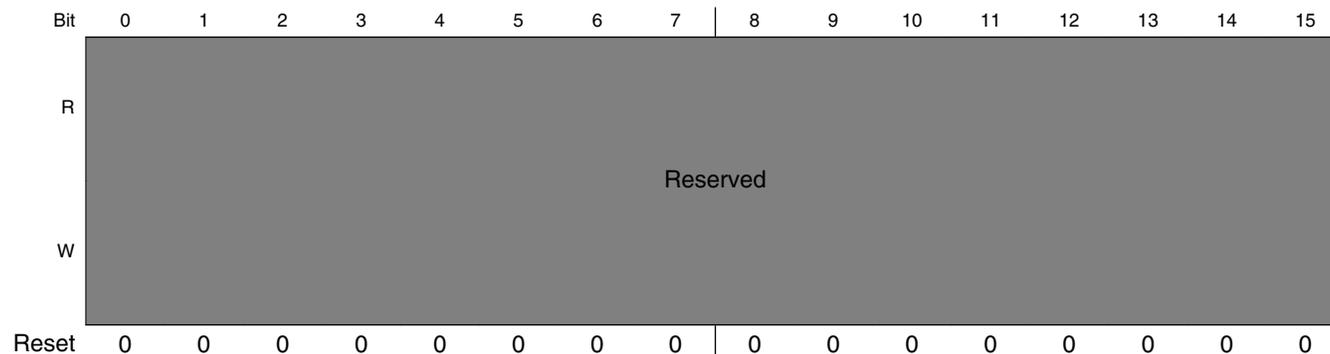
ETSECx_PTV field descriptions

Field	Description
0–15 PTE	Extended pause control. This field allows software to add a 16-bit additional control parameter into the PAUSE frame to be sent when TCTRL[TFC_PAUSE] is set. Note that current IEEE 802.3 PAUSE frame format requires this parameter to be cleared.
16–31 PT	Pause time value. Represents the 16-bit pause quanta (that is, 512 bit times). This pause value is used as part of the PAUSE frame to be sent when TCTRL[TFC_PAUSE] is set. See Flow control for more information.

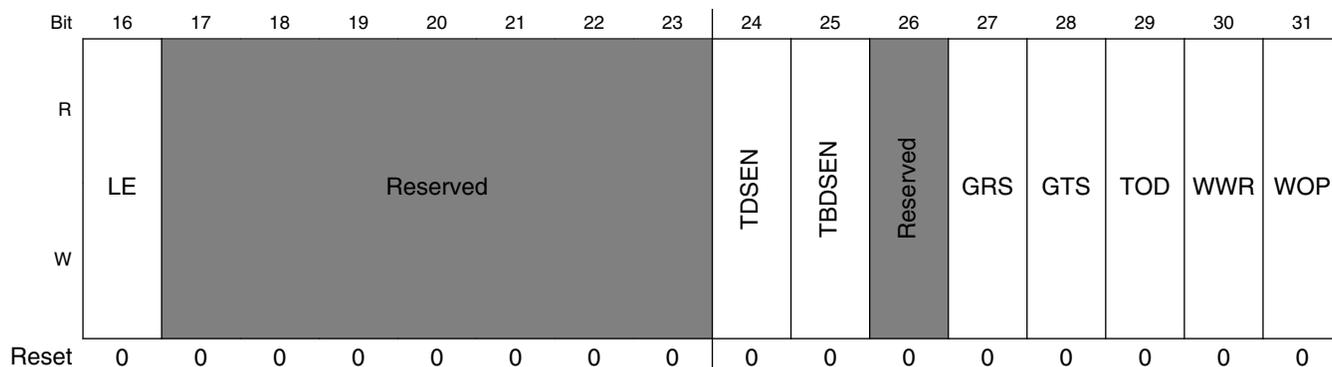
14.52.8 DMA control register (ETSECx_DMACTRL)

DMACTRL is writable by the user to configure the DMA block.

Address: Base address + 2Ch offset



eTSEC memory map/register definition



ETSECx_DMACTRL field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16 LE	Little-endian descriptor mode enable. This bit controls both the reading and writing of descriptors; data buffers are always transferred in network byte order. 0 RxBDs and TxBDs are interpreted with big-endian byte ordering, as shown in Data buffer descriptors . 1 RxBDs and TxBDs are interpreted with little-endian byte ordering. That is, the 16 bits of flags are considered a complete half-word unit, the buffer length is considered another complete half-word unit, and the buffer pointer is considered a complete word unit.
17–23 -	This field is reserved. Reserved
24 TDSEN	Tx Data snoop enable. 0 Disables snooping of all transmit frames from memory. 1 Enables snooping of all transmit frames from memory.
25 TBDSSEN	TxBD snoop enable. 0 Disables snooping of all transmit BD memory accesses. 1 Enables snooping of all transmit BD memory accesses.
26 -	This field is reserved. Reserved
27 GRS	Graceful receive stop. If this bit is set, the Ethernet controller stops receiving frames following completion of the frame currently being received. (That is, after a valid end of frame was received). The contents of the Rx FIFO are then written to memory, and the IEVENT[GRSC] is set to indicate that all current receive buffers have been closed. Because the receive enable bit of the MAC may still be set, the MAC may continue to receive but the eTSEC ignores the receive data until GRS is cleared. If this bit is cleared, the eTSEC scans the input data stream for the start of a new frame (preamble sequence and start of frame delimiter) and the first valid frame received uses the next RxBD. If GRS is set, the user must monitor the graceful receive stop complete (GRSC) bit in the IEVENT register to insure that the graceful receive stop was completed. The user can then clear IEVENT[GRSC] and can write to receive registers that are accessible to both user and the eTSEC hardware without fear of conflict. 0 eTSEC scans input data stream for valid frame. 1 eTSEC stops receiving frames following completion of current frame.
28 GTS	Graceful transmit stop. If this bit is set, the Ethernet controller stops transmission after all frames that are currently in the Tx FIFO or scheduled have been transmitted, and the GTSC interrupt in the IEVENT register is asserted. A frame that has started reading buffer descriptors or data from memory is read to

Table continues on the next page...

ETSECx_DMACTRL field descriptions (continued)

Field	Description
	<p>completion and transmitted before the GTSC interrupt occurs. However, if no frame has been scheduled for transmission and the Tx FIFO is empty, the GTSC interrupt is asserted immediately. Once transmission has completed, clearing GTS "restart" transmit.</p> <p>0 Controller continues. 1 Controller stops transmission after completion of current frame.</p>
29 TOD	<p>Transmit on demand for TxBD ring 0. This bit is applicable only to the transmitter, and requires both TCTRL[TXSCHED] = 00 and DMACTRL[WOP] = 0. If 1 is written to this bit, the eTSEC immediately begins fetching the next TxBD from ring 0, avoiding waiting the normal polling time to check the TxBD's R bit. This bit is always read as 0.</p> <p>0 eTSEC continues waiting for the TxBD ring 0 poll timer to expire. 1 eTSEC immediately fetches a new TxBD from ring 0 .</p>
30 WWR	<p>Write with response. This bit gives the user the assurance that a BD was updated in memory before it receives an interrupt concerning a transmit or receive frame.</p> <p>0 Do not wait for acknowledgement from system for BD writes before setting IEVENT bits. 1 Before setting IEVENT bits TXB, TXF, TXE, XFUN, LC, CRL, RXB, RXF, the eTSEC waits for acknowledgement from system that the transmit or receive BD being updated was stored in memory.</p>
31 WOP	<p>Wait or poll for TxBD ring 0. This bit, which is applicable only to the transmitter and when TCTRL[TXSCHED] = 00, provides the user the option for the eTSEC to periodically poll TxBDs or to wait for software to tell eTSEC to fetch a buffer descriptor. While operating in the "Wait" mode, the eTSEC allows two additional reads of a descriptor which is not ready before entering a halt state. No interrupt is driven. To resume transmission, software must clear TSTAT[THLT].</p> <p>0 Poll TxBD on ring 0 every 512 serial clocks. 1 Do not poll, but wait for TSTAT[THLT] to be cleared by the user.</p>

14.52.9 TBI PHY address register (ETSECx_TBIPA)

The TBIPA is writable by the user to assign a physical address to the TBI (or RTBI) for MII management configuration. The TBI registers are accessed at the offset of TBIPA. For detailed descriptions of the TBI registers (the MII register set for the ten-bit interface) refer to [Ten-bit interface \(TBI\)](#) .

Address: Base address + 30h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TBIPA																
W	Reserved															TBIPA																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

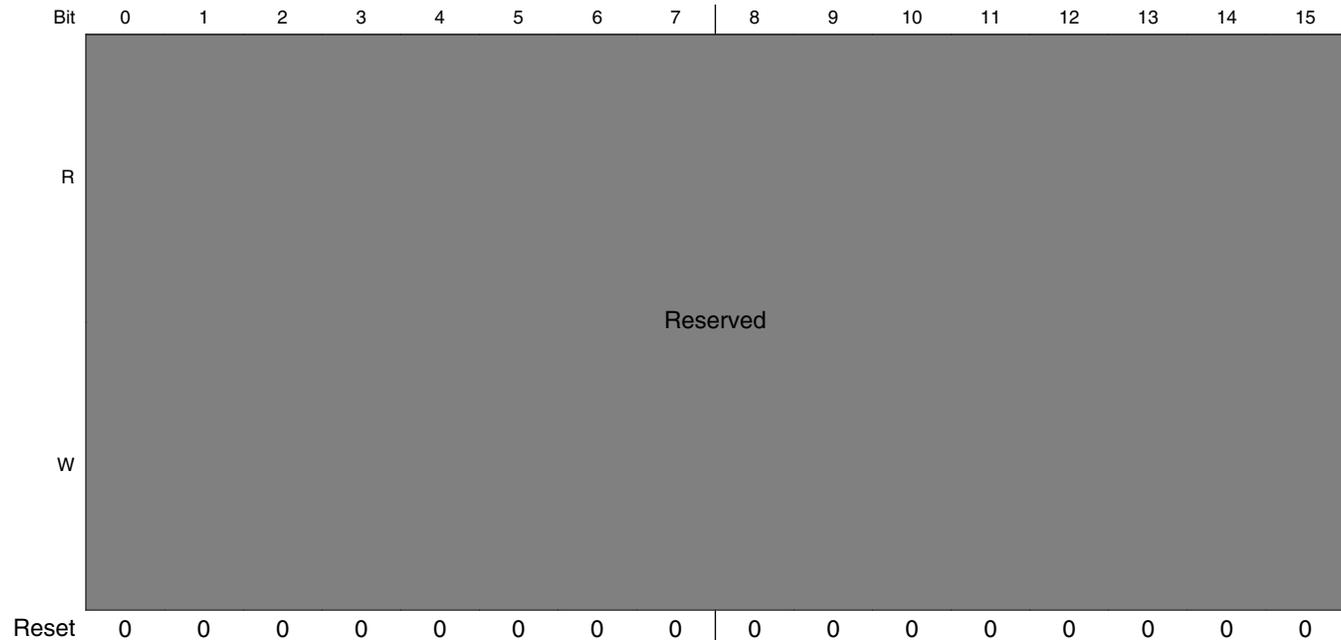
ETSECx_TBIPA field descriptions

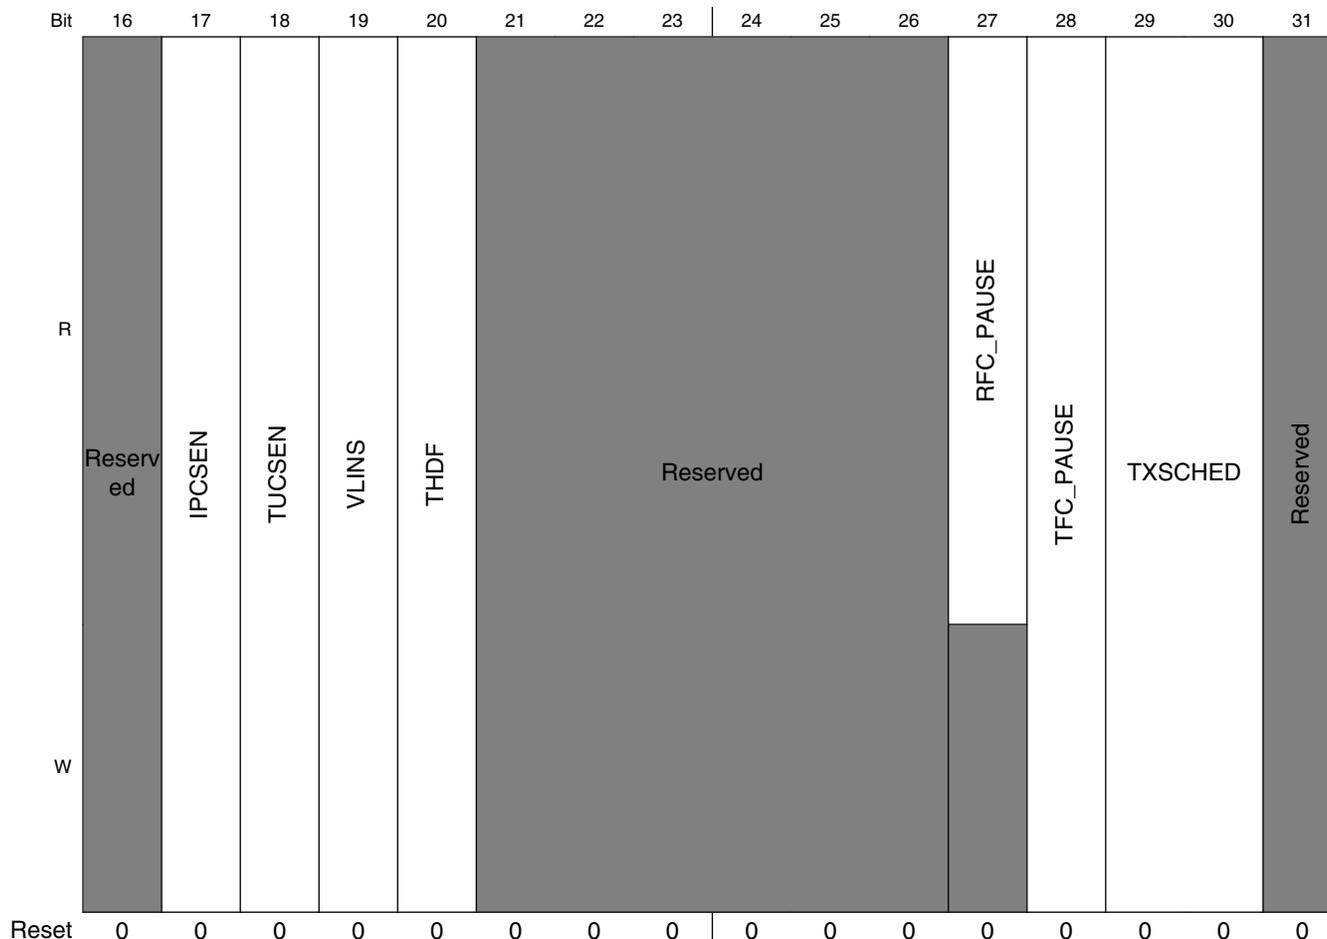
Field	Description
0–26 -	This field is reserved. Reserved
27–31 TBIPA	This field is used to program the PHY address of the ten-bit interface's MII management bus. To access the TBI register the user must write the TBIPA value to the MIIMADD [PHY Address] register located in the MAC register section. PHY Address 0 is reserved. Refer to MII management address (ETSEC_MIIMADD) ."

14.52.10 Transmit control register (ETSECx_TCTRL)

This register is writable by the user to configure the transmit block.

Address: Base address + 100h offset





ETSECx_TCTRL field descriptions

Field	Description
0–16 -	This field is reserved. Reserved
17 IPCSEN	IP header checksum generation enable. When set, the eTSEC offloads IPv4 header checksum generation. See Transmit path off-load and Tx PTP packet parsing . 0 IP header checksum generation is disabled even if enabled in a transmit frame control block. 1 IP header checksum generation is performed for IPv4 headers as determined by the settings in the current transmit frame control block.
18 TUCSEN	TCP/UDP header checksum generation enable. When set, the eTSEC offloads TCP or UDP header checksum generation. See Transmit path off-load and Tx PTP packet parsing . 0 TCP or UDP header checksum generation is disabled even if enabled in a transmit frame control block. 1 TCP or UDP header checksum generation is performed as determined by the settings in the current transmit frame control block.
19 VLINS	VLAN (IEEE Std. 802.1Q) tag insertion enable. Applicable only for transmission through the Ethernet MAC.

Table continues on the next page...

ETSECx_TCTRL field descriptions (continued)

Field	Description
	<p>0 Do not insert a VLAN tag into the frame.</p> <p>1 Insert a VLAN tag into the frame. If the frame FCB has a valid VLAN field, use the FCB to source the VLAN control word, otherwise take the default VLAN control word from register DFVLAN.</p>
20 THDF	<p>Transmit half-duplex flow control under software control for 10-/100-Mbps half-duplex media. This bit is not self-resetting.</p> <p>0 Disable back pressure</p> <p>1 Back pressure is applied to media by raising carrier</p>
21–26 -	<p>This field is reserved. Reserved</p>
27 RFC_PAUSE	<p>Receive flow control pause frame (written by the eTSEC). This read-only status bit is set if a flow control pause frame was received and the transmitter is paused for the duration defined in the received pause frame. This bit automatically clears after the pause duration is complete.</p> <p>0 Pause duration complete.</p> <p>1 Flow control pause frame received.</p>
28 TFC_PAUSE	<p>Transmit flow control pause frame. Set this bit to transmit a PAUSE frame. If this bit is set, the MAC stops transmission of data frames after the currently transmitting frame completes. Next, the MAC transmits a pause control frame with the duration value obtained from the PTV register. The TXC event occurs after sending the pause control frame. Finally, the controller clears TFC_PAUSE and resumes transmitting data frames as before. Note that pause control frames can still be transmitted if the Tx controller is stopped due to user assertion of DMACTRL[GTS] or reception of a PAUSE frame.</p> <p>0 No request for Tx PAUSE frame pending or transmission complete.</p> <p>1 Software request for Tx PAUSE frame pending.</p>
29–30 TXSCHED	<p>Transmit ring scheduling algorithm. This field determines which scheme the transmit scheduler uses to arbitrate between the enabled TxBD rings. The scheme chosen also controls how the DMACTRL and TQUEUE bits are interpreted. Ring polling is supported only by mode 00; the other modes require software to restart rings with the TSTAT register. TCP/IP offload can be enabled with any scheduling mode.</p> <p>00 Single polled ring mode. TxBD ring 0 is the only ring serviced, even if other rings are enabled and ready. In this scheduler mode, the DMACTRL[WOP] and DMACTRL[TOD] bits control polling and retry behavior. This mode supports ring polling, and allows fetching of a non-ready TxBD to be retried twice.</p> <p>01 Priority scheduling mode. Frames from enabled TxBD rings are serviced in ascending ring index order.</p> <p>10 Modified weighted round-robin scheduling mode. Each TxBD ring is polled in sequence for frames that are ready for transmission. If a non-ready TxBD is fetched from a ring, that ring is removed from the scheduling pool until software re-enables it. Ready frames are repeatedly transmitted from a chosen ring until its transmission quota is exhausted. The transmission quota for TxBD ring <i>n</i> is set to $WT_n \times 64$ bytes, where WT_n is a weight from the TR03WT/TR47WT registers. If a ring transmits more data than its quota allows, the excess is deducted from its quota on the next transmission opportunity, thereby preventing large frames from monopolizing the eTSEC bandwidth.</p> <p>11 Reserved</p>
31 -	<p>This field is reserved. Reserved</p>

14.52.11 Transmit status register (ETSECx_TSTAT)

This register is read/write-one-to-clear and is written by the eTSEC to convey DMA status information for each TxBD ring. The halt bit only has meaning for enabled rings. After processing transmit-related interrupts, software should use TSTAT to restart transmission from rings that may have been affected by the interrupt condition. In particular, an error condition that prevents eTSEC from continuing transmission halts DMA from all rings, including the ring that gave rise to the error.

Address: Base address + 104h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	THLT0	THLT1	THLT2	THLT3	THLT4	THLT5	THLT6	THLT7	Reserved							
W	w1c															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	TXF0	TXF1	TXF2	TXF3	TXF4	TXF5	TXF6	TXF7	Reserved							
W	w1c															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TSTAT field descriptions

Field	Description
0 THLT0	<p>Transmit halt of ring 0. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN0], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> Invalid BD or data address Uncorrectable error on BD or data read

Table continues on the next page...

ETSECx_TSTAT field descriptions (continued)

Field	Description
	<p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0
<p>1 THLT1</p>	<p>Transmit halt of ring 1. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN1], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0
<p>2 THLT2</p>	<p>Transmit halt of ring 2. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN2], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0
<p>3 THLT3</p>	<p>Transmit halt of ring 3. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN3], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0

Table continues on the next page...

ETSECx_TSTAT field descriptions (continued)

Field	Description
4 THLT4	<p>Transmit halt of ring 4. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN4], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0
5 THLT5	<p>Transmit halt of ring 5. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN5], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0
6 THLT6	<p>Transmit halt of ring 6. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN6], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0
7 THLT7	<p>Transmit halt of ring 7. Set by the eTSEC if is no longer processing transmit frames from this TxBD ring, and DMA from this ring is disabled. To re-start transmission from this TxBD ring, this bit must be cleared</p>

Table continues on the next page...

ETSECx_TSTAT field descriptions (continued)

Field	Description
	<p>by writing 1 to it. This bit is set only on a general error condition (as in IEVENT[TXE]), regardless of TQUEUE[EN7], or if no ready TxBDs can be fetched. DMACTRL[GTS] being set by the user does not cause this bit to be set.</p> <p>Software should examine the halted queue's buffer descriptors for repeatable error conditions before taking it out of the halt state. Failure to do so may cause an effective livelock, in which the error condition recurs and halts all queues again.</p> <p>Repeatable error conditions which cause halt include:</p> <p>Bus error:</p> <ul style="list-style-type: none"> • Invalid BD or data address • Uncorrectable error on BD or data read <p>TxBD programming errors:</p> <ul style="list-style-type: none"> • Ready=1 and length=0
8–15 -	This field is reserved. Reserved
16 TXF0	Transmit frame event occurred on ring 0. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
17 TXF1	Transmit frame event occurred on ring 1. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
18 TXF2	Transmit frame event occurred on ring 2. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
19 TXF3	Transmit frame event occurred on ring 3. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
20 TXF4	Transmit frame event occurred on ring 4. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
21 TXF5	Transmit frame event occurred on ring 5. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
22 TXF6	Transmit frame event occurred on ring 6. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
23 TXF7	Transmit frame event occurred on ring 7. Set by the eTSEC if IEVENT[TXF] was set in relation to transmitting a frame from this ring.
24–31 -	This field is reserved. Reserved

14.52.12 Default VLAN control word (ETSECx_DFVLAN)

This register defines the default value for the VLAN Ethertype and control word when VLAN tags are automatically inserted by the eTSEC, and no per-frame VLAN data is supplied by software. On receive, this register defines a customizable VLAN Ethertype for automatic deletion. Note that an Ethertype of 0x8808 (Control Word) is not permitted as a custom VLAN tag. Frames with an Ethertype of 0x8808 are dropped by the receiver. In the case of frames containing stacked VLAN tags, this register defines the tag associated with the outer or metropolitan area VLAN.

Address: Base address + 108h offset

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R																	
W	TAG																
Reset	1	0	0	0	0	0	0	1		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	PRI			CFI	VID												
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

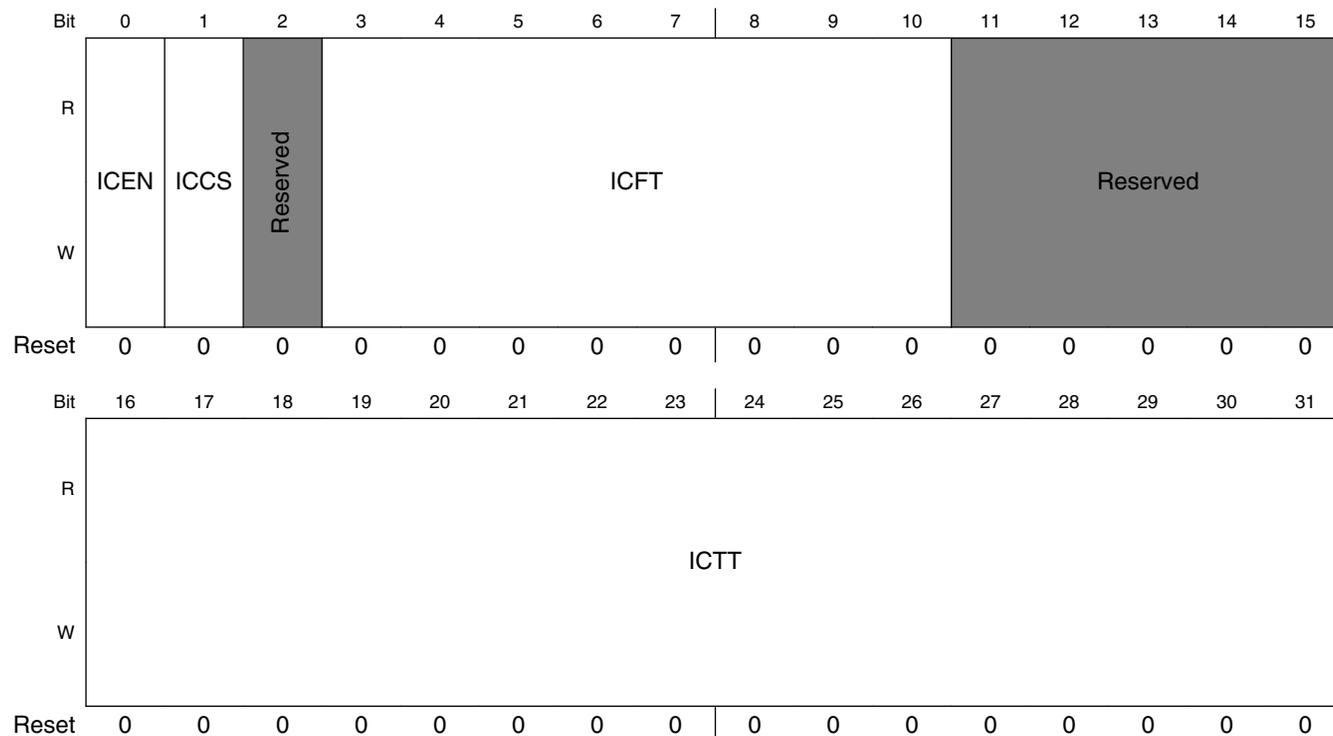
ETSECx_DFVLAN field descriptions

Field	Description
0–15 TAG	This is the default Ethertype used to tag VLAN frames. On transmit, this tag is inserted ahead of the VLAN control word; TAG should be set to 0x8100 for IEEE 802.1Q VLAN. On receive, an Ethertype matching TAG or an Ethertype of 0x8100 marks a VLAN-tagged frame. Note that if using DFVLAN to set a custom ethertype (that is, using a value other than 0x8100), packets received with a custom tag are not counted by any of the RMON counters. Affected counters include TRMGV, RMCA, RBCA, RXCF, RXPF, RXUO, RALN, RFLR, ROVR, RJBR, TMCA, TBCA, TXPF, TXCF.
16–18 PRI	This is the default value used for the IEEE Std. 802.1p frame priority.
19 CFI	This is the default value used for the IEEE Std. 802.1Q canonical format indicator.
20–31 VID	This is the default value used for the virtual-LAN identifier in VLAN-tagged frames. A value of zero is defined as the null VLAN, however field PRI may be still set independently.

14.52.13 Transmit interrupt coalescing register (ETSECx_TXIC)

The TXIC register enables and configures the operational parameters for interrupt coalescing associated with transmitted frames.

Address: Base address + 110h offset



ETSECx_TXIC field descriptions

Field	Description
0 ICEN	<p>Interrupt coalescing enable</p> <p>0 Interrupt coalescing is disabled. Interrupts are raised as they are received.</p> <p>1 Interrupt coalescing is enabled. If the eTSEC transmit frame interrupt is enabled (IMASK[TXFEN] is set), an interrupt is raised when the threshold number of frames is reached (defined by TXIC[ICFT]) or when the threshold timer expires (determined by TXIC[ICTT]).</p>
1 ICCS	<p>Interrupt coalescing timer clock source.</p> <p>0 The coalescing timer advances count every 64 eTSEC Tx interface clocks (TSEC n_GTX_CLK).</p> <p>1 The coalescing timer advances count every 64 system clocks. This mode is recommended for FIFO operation. The term 'system clock' refers to CCB clock/2.</p>
2 -	<p>This field is reserved.</p> <p>Reserved</p>
3-10 ICFT	<p>Interrupt coalescing frame count threshold. While interrupt coalescing is enabled (TXIC[ICEN] is set), this value determines how many frames are transmitted before raising an interrupt. The eTSEC threshold</p>

Table continues on the next page...

ETSECx_TXIC field descriptions (continued)

Field	Description
	counter is reset to ICFT following an interrupt. The value of ICFT must be greater than zero to avoid unpredictable behavior.
11–15 -	This field is reserved. Reserved
16–31 ICTT	Interrupt coalescing timer threshold. While interrupt coalescing is enabled (TXIC[ICEN] is set), this value determines the maximum amount of time after transmitting a frame before raising an interrupt. If frames have been transmitted but the frame count threshold has not been met, an interrupt is raised when the threshold timer reaches zero. The threshold timer is reset to the value in this field and begins counting down upon transmission of the first frame having its TxBD[I] bit set. The threshold value is represented in units of 64 clock periods as specified by the timer clock source (TXIC[ICCS]). The value of ICTT must be greater than zero to avoid unpredictable behavior.

14.52.14 Transmit queue control register (ETSECx_TQUEUE)

The TQUEUE register selectively enables each of the TxBD rings 0-7. By default, TxBD ring 0 is enabled.

Address: Base address + 114h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	EN0	EN1	EN2	EN3	EN4	EN5	EN6	EN7	Reserved							
W																
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TQUEUE field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16 EN0	Transmit queue 0 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
17 EN1	Transmit queue 1 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
18 EN2	Transmit queue 2 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.

Table continues on the next page...

ETSECx_TQUEUE field descriptions (continued)

Field	Description
19 EN3	Transmit queue 3 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
20 EN4	Transmit queue 4 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
21 EN5	Transmit queue 5 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
22 EN6	Transmit queue 6 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
23 EN7	Transmit queue 7 enable. 0 TxBD ring is not queried for transmission. In effect the transmit queue is disabled. 1 TxBD ring is queried for transmission.
24–31 -	This field is reserved. Reserved

14.52.15 TxBD Rings 0-3 round-robin weightings (ETSECx_TR03WT)

When modified weighted round-robin Tx scheduling is enabled (TCTRL[TXSCHEd] = 10), this register determines the weighting applied to each transmit queue for queues 0 to 3. For priority-based scheduling, TR03WT has no effect. A description of how queue weights affect eTSEC's round-robin algorithm appears in [Modified weighted round-robin queuing \(MWRR\)](#) .

Address: Base address + 140h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	WT0								WT1								WT2								WT3							
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TR03WT field descriptions

Field	Description
0–7 WT0	Weighting value for TxBD ring 0 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT0 x 64 bytes of data are scheduled for transmission from TxBD ring 0. Clearing this field prevents transmission.
8–15 WT1	Weighting value for TxBD ring 1 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT1 x 64 bytes of data are scheduled for transmission from TxBD ring 1. Clearing this field prevents transmission.
16–23 WT2	Weighting value for TxBD ring 2 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT2 x 64 bytes of data are scheduled for transmission from TxBD ring 2. Clearing this field prevents transmission.
24–31 WT3	Weighting value for TxBD ring 3 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT3 x 64 bytes of data are scheduled for transmission from TxBD ring 3. Clearing this field prevents transmission.

14.52.16 TxBD Rings 4-7 round-robin weightings (ETSECx_TR47WT)

When modified weighted round-robin Tx scheduling is enabled (TCTRL[TXSCHED] = 10), this register determines the weighting applied to each enabled transmit queue for queues 4 to 7. For priority-based scheduling, TR47WT has no effect. A description of how queue weights affect eTSEC's modified weighted round-robin algorithm appears in [Modified weighted round-robin queuing \(MWRR\)](#).

Address: Base address + 144h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	WT4								WT5								WT6								WT7							
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

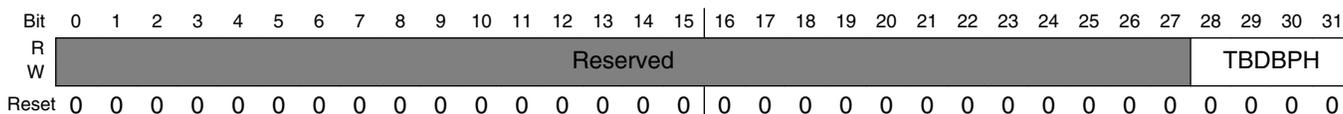
ETSECx_TR47WT field descriptions

Field	Description
0–7 WT4	Weighting value for TxBD ring 4 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT4 x 64 bytes of data are scheduled for transmission from TxBD ring 4. Clearing this field prevents transmission.
8–15 WT5	Weighting value for TxBD ring 5 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT5 x 64 bytes of data are scheduled for transmission from TxBD ring 5. Clearing this field prevents transmission.
16–23 WT6	Weighting value for TxBD ring 6 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT6 x 64 bytes of data are scheduled for transmission from TxBD ring 6. Clearing this field prevents transmission.
24–31 WT7	Weighting value for TxBD ring 7 when TCTRL[TXSCHED] = 10. On each round of the Tx scheduler, a minimum of WT7 x 64 bytes of data are scheduled for transmission from TxBD ring 7. Clearing this field prevents transmission.

14.52.17 Tx data buffer pointer high bits (ETSECx_TBDBPH)

The TBDBPH register is written by the user with the most significant address bits common to all TxBD buffer addresses, TxBD[Data Buffer Pointer]. As a consequence, all Tx buffers must be placed in a 4 gigabyte segment of memory whose base address is prefixed by the bits in TBDBPH. The TxBD ring itself can reside in a different memory region (based at TBASEH). (See [TxBD base address high bits \(ETSEC_TBASEH\)](#) .)

Address: Base address + 180h offset



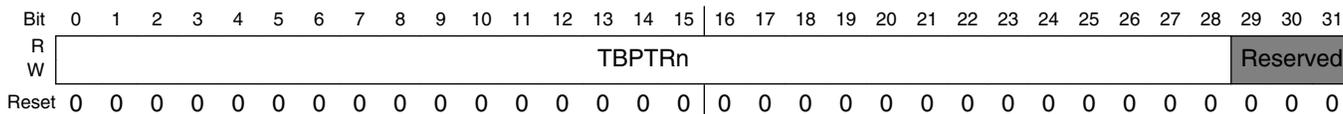
ETSECx_TBDBPH field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 TBDBPH	Most significant bits common to all data buffer addresses contained in TxBDs. The user must initialize TBDBPH before enabling the eTSEC transmit function.

14.52.18 TxBD pointer for ring n (ETSECx_TBPTRn)

TBPTR0-TBPTR7 each contains the low-order 32 bits of the next transmit buffer descriptor address for their respective TxBD ring. These registers takes on the value of their ring's associated TBASE when the TBASE register is written by software. Software must not write TBPTR0-TBPTR7 while eTSEC is actively transmitting frames. However, TBPTR0- TBPTR7 can be modified when the transmitter is disabled or when no Tx buffer is in use (after a GRACEFUL STOP TRANSMIT command is issued and the frame completes its transmission) in order to change the next TxBD eTSEC transmits.

Address: Base address + 184h offset + (8d × i), where i=0d to 7d



ETSECx_TBPTRn field descriptions

Field	Description
0–28 TBPTRn	Current TxBD pointer for TxBD ring <i>n</i> . Points to the current BD being processed or to the next BD the transmitter uses when it is idling. When the end of the TxBD ring is reached, eTSEC initializes TBPTR <i>n</i> to the value in the corresponding TBASE <i>n</i> . The TBPTR register is internally written by the eTSEC's DMA controller during transmission. The pointer increments by eight (bytes) each time a descriptor is closed successfully by the eTSEC. Note that the three least significant bits of this register are read-only and zero. After an error condition, the eTSEC returns TBPTR <i>n</i> to point to the first BD of the frame partially transmitted.
29–31 -	This field is reserved. Reserved

14.52.19 TxBD base address high bits (ETSECx_TBASEH)

The TBASEH register is written by the user with the most significant address bits common to all TxBD addresses, including TBASE0-TBASE7 and TBPTR0-TBPTR7. As a consequence, all TxBD rings must be placed in a 4 Gbyte segment of memory whose base address is prefixed by the bits in TBASEH. Data buffers are located in a potentially different region, based at TBDBPH. (See [Tx data buffer pointer high bits \(ETSEC_TBDBPH\)](#) .)

Address: Base address + 200h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TBASEH																
W	Reserved															TBASEH																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

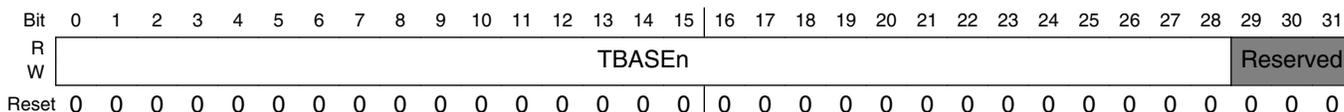
ETSECx_TBASEH field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 TBASEH	Most significant bits common to all TxBD addresses-except data buffer pointers. The user must initialize TBASEH before enabling the eTSEC transmit function.

14.52.20 TxBD base address of ring *n* (ETSECx_TBASEn)

The TBASEn registers are written by the user with the base address of each TxBD ring *n*. Each such value must be divisible by eight, since the three least significant bits always write as 000.

Address: Base address + 204h offset + (8d × i), where i=0d to 7d



ETSECx_TBASEn field descriptions

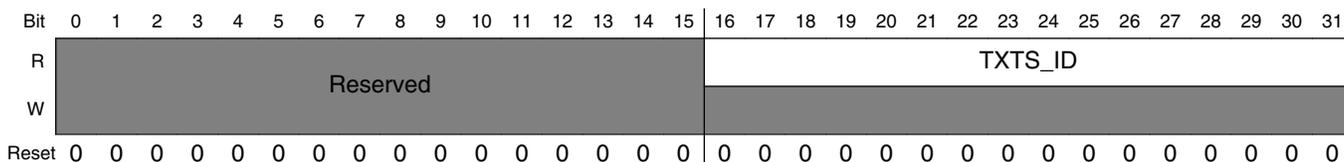
Field	Description
0–28 TBASEn	Transmit base for ring <i>n</i> . TBASEn defines the starting location in the memory map for the eTSEC TxBDs. This field must be 8-byte aligned. Together with setting the W (wrap) bit in the last BD, the user can select how many BDs to allocate for the transmit packets. The user must initialize TBASEn before enabling the eTSEC transmit function on the associated ring.
29–31 -	This field is reserved. Reserved

14.52.21 Tx time stamp identification tag (set *n*) register (ETSECx_TMR_TXTSn_ID)

Transmit time stamp identification register (TMR_TXTSn_ID). This register holds the identification number of the transmitted frame corresponding to the timestamp captured in TMR_TXTSn_H/L. Each time the eTSEC is instructed to capture the timestamp of an outgoing frame via TxFCB[PTP] the associated field in TxFCB[PTP_ID] is stored in this register, overwriting the previous value.

This register is read only in normal operation.

Address: Base address + 280h offset + (4d × i), where i=0d to 1d



ETSECx_TMR_TXTSn_ID field descriptions

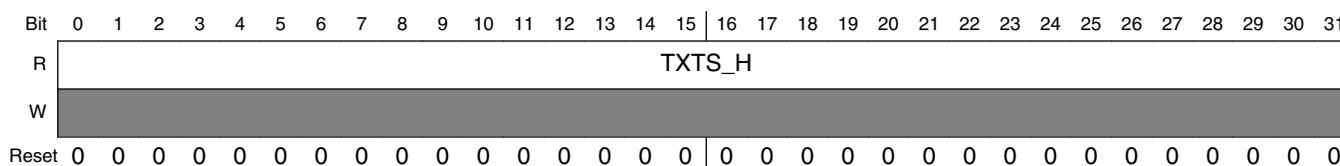
Field	Description
0–15 -	This field is reserved. Reserved
16–31 TXTS_ID	Tx time stamp identification field

14.52.22 Tx time stamp high (ETSECx_TMR_TXTSn_H)

Transmit stamp register high (TMR_TXTSn_H). This register, together with TMR_TXTSn_L, holds the value of the TMR_CNT_H/L when a frame tagged for timestamp capture (via Tx FCB[PTP]) is transmitted. Upon transmission of the start of frame symbol of such a frame, the value in TMR_CNT_H/L is copied into TMR_TXTSn_H/L.

This register is read only in normal operation.

Address: Base address + 2C0h offset + (8d × i), where i=0d to 1d



ETSECx_TMR_TXTSn_H field descriptions

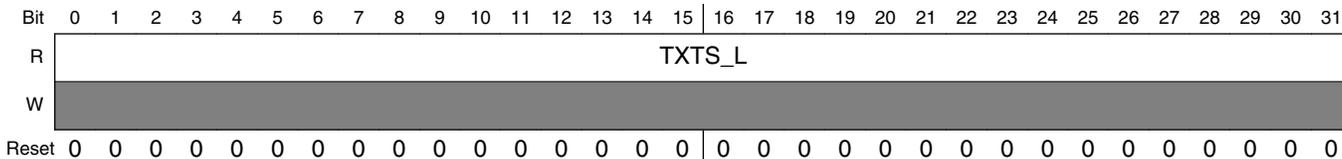
Field	Description
0–31 TXTS_H	Time stamp field of the transmitted PTP packet's start of frame detection (upper half)

14.52.23 Tx time stamp low (ETSECx_TMR_TXTSn_L)

Transmit stamp register low (TMR_TXTSn_L). This register, together with TMR_TXTSn_H, holds the value of the TMR_CNT_H/L when a frame tagged for timestamp capture (via Tx FCB[PTP]) is transmitted. Upon transmission of the start of frame symbol of such a frame, the value in TMR_CNT_H/L is copied into TMR_TXTSn_H/L.

This register is read only in normal operation.

Address: Base address + 2C4h offset + (8d × i), where i=0d to 1d



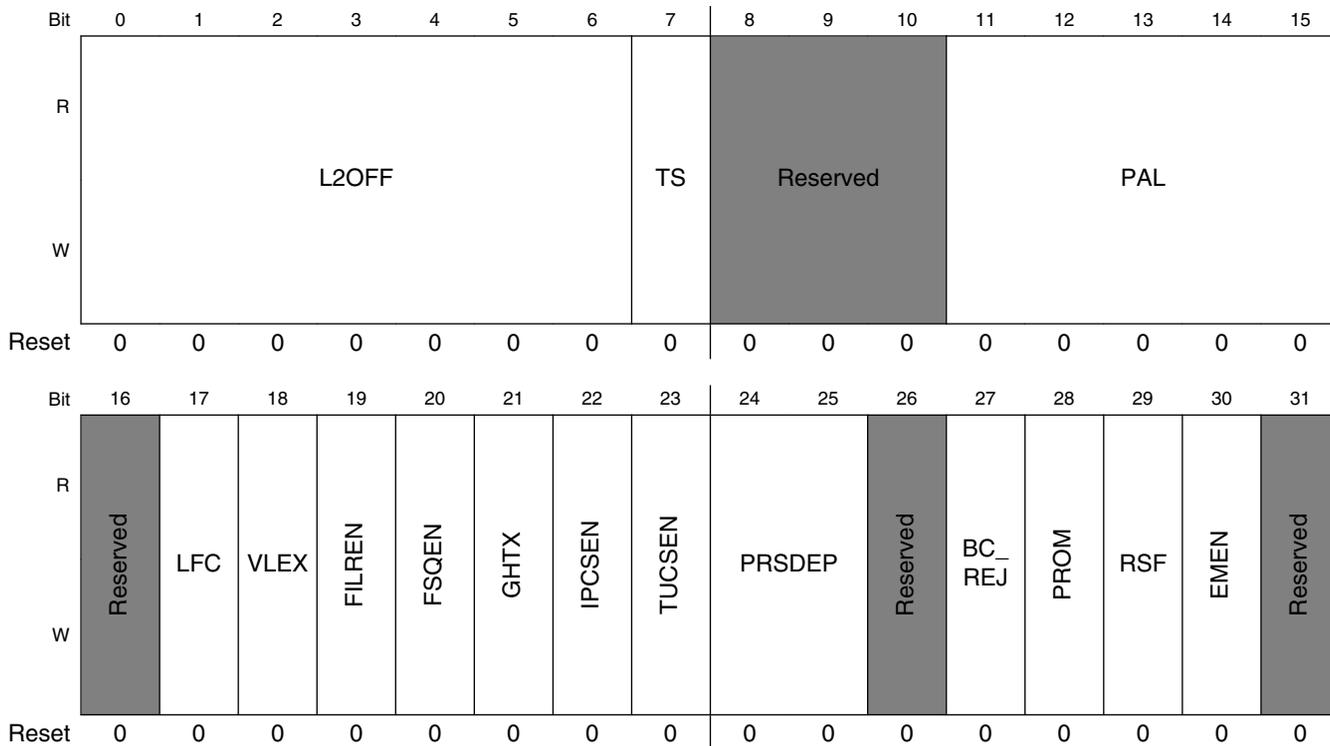
ETSECx_TMR_TXTSn_L field descriptions

Field	Description
0–31 TXTS_L	Time stamp field of the transmitted PTP packet's start of frame detection (lower half)

14.52.24 Receive control register (ETSECx_RCTRL)

The RCTRL register is programmed by the user and controls the operational mode of the receiver. It must be written only after a system reset (at initialization) or after a graceful receive stop has completed.

Address: Base address + 300h offset



ETSECx_RCTRL field descriptions

Field	Description
0–6 L2OFF	<p>Layer 2 offset. The number of octet pairs from the start of the frame that the parser should expect to see before the first byte of the Ethernet DA.</p> <p>For frames received over Ethernet, the start of frame is regarded as the SFD symbol.</p> <p>For packets received through the FIFO packet interface the start of frame is regarded as the first octet of receive data.</p> <p>The user may think of this value as representing the length - in multiples of two bytes - of a 'shim' header that is inserted between the SFD and DA. By writing to RCTRL with a mask of 0xFE00_0000, the even byte length restriction is guaranteed.</p> <p>For normal frames, this field should be left as 0.</p>
7 TS	<p>Time stamp incoming packets as padding bytes. PAL field is set to 8 if the PAL field is programmed to less than 8. Must be set to zero if TMR_CTRL[TE]=0.</p>
8–10 -	<p>This field is reserved. Reserved</p>
11–15 PAL	<p>Packet alignment padding length. If not zero, PAL (1-31) bytes of zero padding are inserted before the start of each received frame, but following the RxFCB if parsing is enabled. For Ethernet where optional preamble extraction is enabled, the padding appears before the preamble, otherwise the padding precedes the layer 2 header. The value of PAL can be set so that the start of the IP header in the receive data buffer is aligned to a 32-bit boundary. Normally, setting PAL = 2 provides minimal padding to ensure such alignment of the IP header.</p> <p>Note that the minimum zero padding value for this field should be PAL-8 if the TS field is set and 0 when PAL is < 8.</p>
16 -	<p>This field is reserved. Reserved</p>
17 LFC	<p>Lossless flow control. When set, the eTSEC determines the number of free BDs (through RQPARMn[LEN] and RBTPTRn) in each active ring. Should the free BD count in an active ring drop below its setting for RQPARMn[FBTHR], the eTSEC asserts link layer flow control.</p> <p>For full-duplex Ethernet connections, the eTSEC emits a pause frame as if TCTRL[TFC_PAUSE] was set. For FIFO packet interface connections, the RFC signal is asserted.</p> <p>0 Disabled. This is the default 1 Enabled, calculate the free BDs in each active ring and assert link layer flow control if required.</p>
18 VLEX	<p>Enable automatic VLAN tag extraction and deletion from Ethernet frames. Note that VLEX must be cleared if L2OFF is non-zero.</p> <p>Note that if PRSDEP is cleared, VLEX must be cleared as well. (VLAN tag extraction is only supported when the parser is enabled.)</p> <p>0 Do not delete VLAN tags from received Ethernet frames. 1 If a VLAN tag is seen after the Ethernet source address, and PRSDEP is non-zero, delete the VLAN tag and return the VLAN control word in the frame control block returned with this frame.</p>
19 FILREN	<p>Filer enable. When set, the receive frame filer is enabled. This file accepted frames to a particular RxBD ring according to rules defined in the filer table. In this case, PRSDEP must not be cleared.</p> <p>Note that if PRSDEP is cleared, FILREN must be cleared as well.</p> <p>0 Do not search the receive queue filer table for received frames. All received frames are sent to RxBD ring 0 by default. 1 Search the receive queue filer table for received frames, and let the filer determine the index of the RxBD ring for each frame.</p>

Table continues on the next page...

ETSECx_RCTRL field descriptions (continued)

Field	Description
20 FSQEN	<p>Enable single-queue mode for the receive frame filer. This bit is ignored unless FILREN is also set.</p> <p>0 The filer chooses the RxBD ring using the least significant bits of the virtual queue ID as a ring index. 1 The filer always attempts to file received frames to ring 0, regardless of virtual queue ID. This mode is intended for operating the filer as a packet classification engine.</p>
21 GHTX	<p>Group address hash table extend. By default, the group address hash table is 256 entries (as defined by registers GADDR0-GADDR7); registers IGADDR0-IGADDR7 are then used to define the individual address hash table. When this bit is set, the hash table is extended to a total of 512 entries (IGADDR0-IGADDR7 are then the first 256 entries of the extended 512-entry group address hash table).</p> <p>0 Both the individual and group hash functions are the 8 MSBs of the CRC-32 of the Ethernet destination address. 1 The group hash function is the 9 MSBs of the CRC-32 of the Ethernet destination address. The individual address hash function is unavailable.</p>
22 IPCSEN	<p>IP Checksum verification enable. See Receive path off-load .</p> <p>0 IPv4 header checksums are not verified by the eTSEC-even if layer 3 parsing is enabled. 1 Perform IPv4 header checksum verification if PRSDEP > 01.</p>
23 TUCSEN	<p>TCP or UDP Checksum verification enable. See Receive path off-load .</p> <p>0 TCP or UDP checksums are not verified by the eTSEC-even if layer 4 parsing is enabled. 1 Perform TCP or UDP checksum verification if PRSDEP = 11.</p>
24–25 PRSDEP	<p>Parser control. The level of parser layer recognition is determined as follows:</p> <p>If this field is non-zero, a TOE frame control block is prepended to the received frame, and the first RxBD points to the FCB.</p> <p>Note that if PRSDEP is cleared, VLEX must be cleared as well. (VLAN tag extraction is only supported when the parser is enabled.) Also, if PRSDEP is cleared, FILREN must also be cleared.</p> <p>00 Parser disabled. Receive frame filer must also be disabled by clearing RCTRL[FILREN]. 01 Only L2 (Ethernet) protocols are recognized. 10 L2 and L3 (IP) protocols are recognized . over any interface not configured as a FIFO interface. 11 L2, L3, and L4 (TCP/UDP) protocols are recognized over any interface not configured as a FIFO interface. .</p>
26 -	<p>This field is reserved. Reserved</p>
27 BC_REJ	<p>Broadcast frame reject. If this bit is set, frames with DA (destination address) = FFFF_FFFF_FFFF are rejected unless RCTRL[PROM] is set. If both BC_REJ and RCTRL[PROM] are set, then frames with broadcast DA are accepted and the M (MISS) bit is set in the receive BD.</p>
28 PROM	<p>Promiscuous mode. All Ethernet frames, regardless of destination address, are accepted.</p>
29 RSF	<p>Receive short frame mode. When set, enables the reception of frames shorter than 64 bytes.</p> <p>0 Ethernet frames less than 64B in length are silently dropped. 1 Frames more than 16B and less than 64B in length are accepted upon a DA match.</p> <p>Note that frames less than or equal to 16B in length are always silently dropped.</p>
30 EMEN	<p>Exact match MAC address enable. If this bit is set, the MAC01ADDR1-MAC15ADDR1 and MAC01ADDR2-MAC15ADDR2 registers are recognized as containing MAC addresses aliasing the MAC's station address. Setting this bit therefore allows eTSEC to receive Ethernet frames having a destination address matching one of these 15 addresses.</p>

Table continues on the next page...

ETSECx_RCTRL field descriptions (continued)

Field	Description
31 -	This field is reserved. Reserved

14.52.25 Receive status register (ETSECx_RSTAT)

The eTSEC writes to this register under the following conditions:

- A frame interrupt event occurred on one or more RxBD rings
- The receiver runs out of descriptors due to a busy condition on a RxBD ring
- The receiver was halted because an error condition was encountered while receiving a frame

Writing 1 to any bit of this register clears it. Software should clear the QHLT bit to take eTSEC's receiver function out of halt state for the associated queue.

Address: Base address + 304h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved								QHLT0	QHLT1	QHLT2	QHLT3	QHLT4	QHLT5	QHLT6	QHLT7
W	Reserved								w1c							
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								RXF0	RXF1	RXF2	RXF3	RXF4	RXF5	RXF6	RXF7
W	Reserved								w1c							
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RSTAT field descriptions

Field	Description
0-7 -	This field is reserved. Reserved
8 QHLT0	RxBD queue 0 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT0 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving.

Table continues on the next page...

ETSECx_RSTAT field descriptions (continued)

Field	Description
	0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
9 QHLT1	RxBD queue 1 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT1 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
10 QHLT2	RxBD queue 2 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT2 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
11 QHLT3	RxBD queue 3 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT3 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
12 QHLT4	RxBD queue 4 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT4 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
13 QHLT5	RxBD queue 5 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT5 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
14 QHLT6	RxBD queue 6 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT6 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
15 QHLT7	RxBD queue 7 is halted. It is a hardware-initiated stop indication. (DMACTRL[GRS] being set by the user does not cause a QHLT7 to be set.). The current frame and all other frames directed to a halted queue are discarded. A write with a value of 1 re-enables the queue for receiving. 0 This queue is enabled for reception. (That is, it is not halted) 1 All controller receive activity to this queue is halted.
16–23 -	This field is reserved. Reserved
24 RXF0	Receive frame event occurred on ring 0. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
25 RXF1	Receive frame event occurred on ring 1. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.

Table continues on the next page...

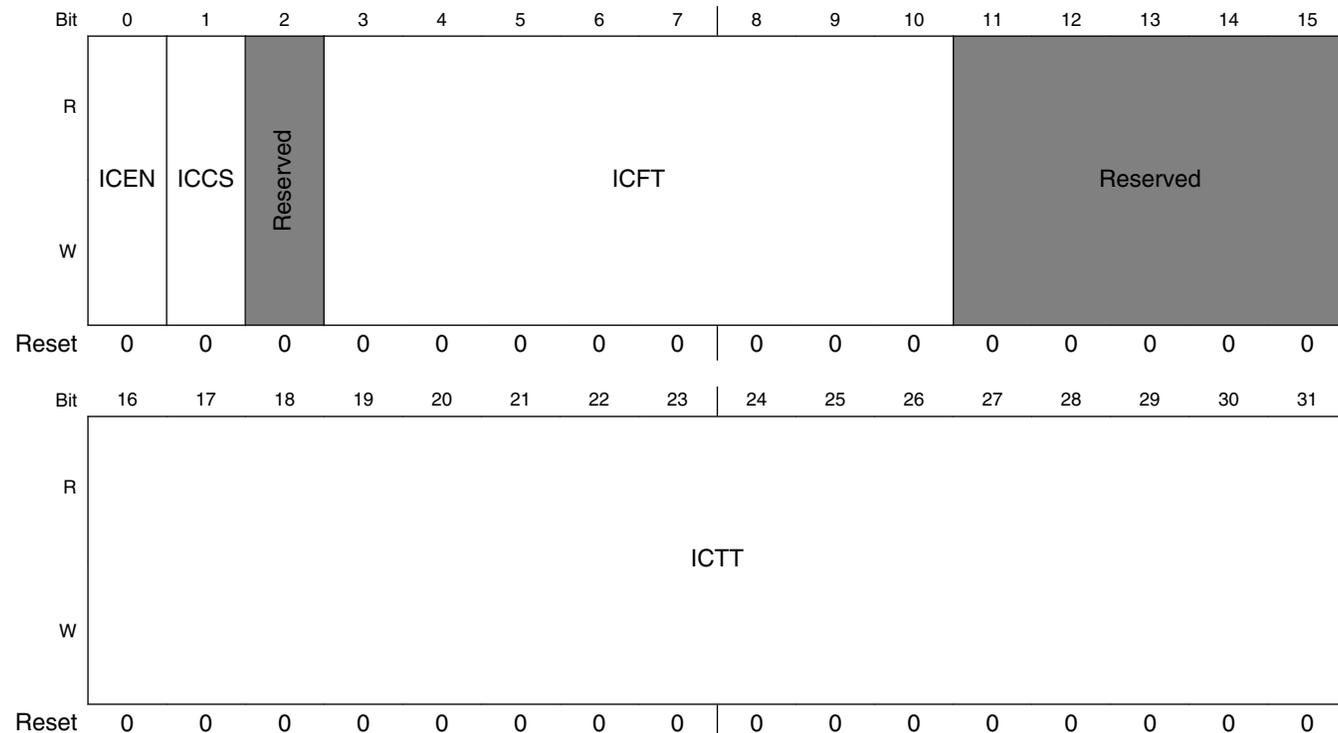
ETSECx_RSTAT field descriptions (continued)

Field	Description
26 RXF2	Receive frame event occurred on ring 2. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
27 RXF3	Receive frame event occurred on ring 3. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
28 RXF4	Receive frame event occurred on ring 4. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
29 RXF5	Receive frame event occurred on ring 5. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
30 RXF6	Receive frame event occurred on ring 6. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.
31 RXF7	Receive frame event occurred on ring 7. Set by the eTSEC if IEVENT[RXF] was set in relation to receiving a frame to this ring.

14.52.26 Receive interrupt coalescing register (ETSECx_RXIC)

The RXIC register enables and configures the operational parameters for interrupt coalescing associated with received frames.

Address: Base address + 310h offset



ETSECx_RXIC field descriptions

Field	Description
0 ICEN	Interrupt coalescing enable 0 Interrupt coalescing is disabled. Interrupts are raised as they are received. 1 Interrupt coalescing is enabled. If the eTSEC receive frame interrupt is enabled (IMASK[RXFEN] is set), an interrupt is raised when the threshold number of frames is reached (defined by RXIC[ICFT]) or when the threshold timer expires (determined by RXIC[ICTT]).
1 ICCS	Interrupt coalescing timer clock source. 0 The coalescing timer advances count every 64 eTSEC Rx interface clocks (TSECn_GTX_CLK). 1 The coalescing timer advances count every 64 system clocks. This mode is recommended for FIFO operation. The term 'system clock' refers to CCB clock/2.
2 -	This field is reserved. Reserved
3–10 ICFT	Interrupt coalescing frame count threshold. While interrupt coalescing is enabled (RXIC[ICE] is set), this value determines how many frames are received before raising an interrupt. The eTSEC threshold counter is reset to ICFT following an interrupt. The value of ICFT must be greater than zero to avoid unpredictable behavior.
11–15 -	This field is reserved. Reserved
16–31 ICTT	Interrupt coalescing timer threshold. While interrupt coalescing is enabled (RXIC[ICE] is set), this value determines the maximum amount of time after receiving a frame before raising an interrupt. If frames have been received but the frame count threshold has not been met, an interrupt is raised when the threshold timer reaches zero. The threshold timer is reset to the value in this field and begins counting down upon receiving the first frame having its RxBD[I] bit set. The threshold value is represented in units equal to 64 periods of the clock specified by RXIC[ICCS]. ICTT must be greater than zero to avoid unpredictable behavior.

14.52.27 Receive queue control register. (ETSECx_RQUEUE)

The RQUEUE register enables each of the RxBD rings 0-7. By default, RxBD ring 0 is enabled.

Address: Base address + 314h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved								EX0	EX1	EX2	EX3	EX4	EX5	EX6	EX7
W	Reserved															
Reset	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								EN0	EN1	EN2	EN3	EN4	EN5	EN6	EN7
W	Reserved															
Reset	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0

ETSECx_RQUEUE field descriptions

Field	Description
0-7 -	This field is reserved. Reserved
8 EX0	Receive queue 0 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
9 EX1	Receive queue 1 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
10 EX2	Receive queue 2 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
11 EX3	Receive queue 3 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
12 EX4	Receive queue 4 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
13 EX5	Receive queue 5 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
14 EX6	Receive queue 6 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
15 EX7	Receive queue 7 extract enable. 0 Data transferred by DMA to this RxBD ring is not extracted to cache. 1 Data transferred by DMA to this RxBD ring undergoes extraction according to ATTR register.
16-23 -	This field is reserved. Reserved
24 EN0	Receive queue 0 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
25 EN1	Receive queue 1 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
26 EN2	Receive queue 2 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.

Table continues on the next page...

ETSECx_RQUEUE field descriptions (continued)

Field	Description
27 EN3	Receive queue 3 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
28 EN4	Receive queue 4 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
29 EN5	Receive queue 5 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
30 EN6	Receive queue 6 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.
31 EN7	Receive queue 7 enable. 0 RxBD ring is not queried for reception. In effect the receive queue is disabled. 1 RxBD ring is queried for reception.

14.52.28 Receive bit field extract control register (ETSECx_RBIFX)

The RBIFX register provides a set of four 6-bit offsets for locating up to four octets in a received frame and passing them to the receive queue filer as the user-defined ARB property. Through RBIFX a custom ARB filer property can be constructed from arbitrary bytes, which allows frame filing on the basis of bitfields not ordinarily provided to the filer, such as bits from the Ethernet preamble or TCP flags. The value of property ARB is the concatenation of {B0, B1, B2, B3} to 32-bits, where B0-B3 are the bytes as defined by RBIFX.

Address: Base address + 330h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	BOCTL		B0OFFSET						B1CTL		B1OFFSET					
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	B2CTL		B2OFFSET						B3CTL		B3OFFSET					
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

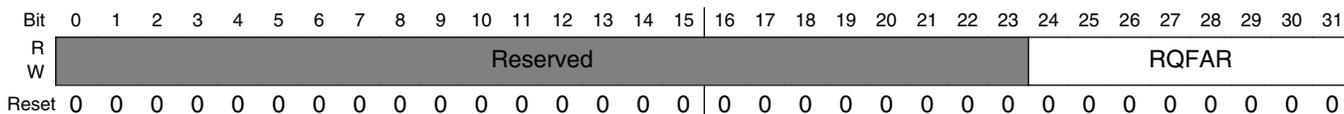
ETSECx_RBIFX field descriptions

Field	Description
0-1 B0CTL	<p>Location of byte 0 of property ARB.</p> <p>00 Byte 0 is not extracted, and appears as zero in property ARB.</p> <p>01 Byte 0 is located in the received frame at offset (B0OFFSET - 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B0OFFSET less than 8 are reserved in FIFO modes.</p> <p>10 Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 2 header.</p> <p>11 Byte 0 is located in the received frame at offset B0OFFSET bytes from the byte after the last byte of the layer 3 header.</p>
2-7 B0OFFSET	Offset relative to the header defined by B0CTL that locates byte 0 of property ARB. An effective offset of zero points to the first byte of the specified header.
8-9 B1CTL	<p>Location of byte 1 of property ARB.</p> <p>00 Byte 1 is not extracted, and appears as zero in property ARB.</p> <p>01 Byte 1 is located in the received frame at offset (B1OFFSET - 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B1OFFSET less than 8 are reserved in FIFO modes.</p> <p>10 Byte 0 is located in the received frame at offset B1OFFSET bytes from the byte after the last byte of the layer 2 header.</p> <p>11 Byte 0 is located in the received frame at offset B1OFFSET bytes from the byte after the last byte of the layer 3 header .</p>
10-15 B1OFFSET	Offset relative to the header defined by B1CTL that locates byte 1 of property ARB. An effective offset of zero points to the first byte of the specified header.
16-17 B2CTL	<p>Location of byte 2 of property ARB.</p> <p>00 Byte 2 is not extracted, and appears as zero in property ARB.</p> <p>01 Byte 2 is located in the received frame at offset (B2OFFSET - 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B2OFFSET less than 8 are reserved in FIFO modes.</p> <p>10 Byte 0 is located in the received frame at offset B2OFFSET bytes from the byte after the last byte of the layer 2 header.</p> <p>11 Byte 0 is located in the received frame at offset B2OFFSET bytes from the byte after the last byte of the layer 3 header .</p>
18-23 B2OFFSET	Offset relative to the header defined by B2CTL that locates byte 2 of property ARB. An effective offset of zero points to the first byte of the specified header.
24-25 B3CTL	<p>Location of byte 3 of property ARB.</p> <p>00 Byte 3 is not extracted, and appears as zero in property ARB.</p> <p>01 Byte 3 is located in the received frame at offset (B3OFFSET - 8) bytes from the first byte of the Ethernet DA. In non-FIFO modes, a negative effective offset points to bytes of the standard Ethernet preamble. Values of B3OFFSET less than 8 are reserved in FIFO modes.</p> <p>10 Byte 0 is located in the received frame at offset B3OFFSET bytes from the byte after the last byte of the layer 2 header.</p> <p>11 Byte 0 is located in the received frame at offset B3OFFSET bytes from the byte after the last byte of the layer 3 header .</p>
26-31 B3OFFSET	Offset relative to the header defined by B3CTL that locates byte 3 of property ARB. An effective offset of zero points to the first byte of the specified header.

14.52.29 Receive queue filing table address register (ETSECx_RQFAR)

RQFAR contains the index of the current, indirectly accessible entry of the received queue filer table. Each table entry occupies a pair of 32-bit words, denoted RQCTRL and RQPROP. To access the RQCTRL and RQPROP words of entry *n*, write *n* to RQFAR. Then read or write the indexed RQCTRL and RQPROP words by reading or writing the RQFCR and RQFPR registers, respectively.

Address: Base address + 334h offset



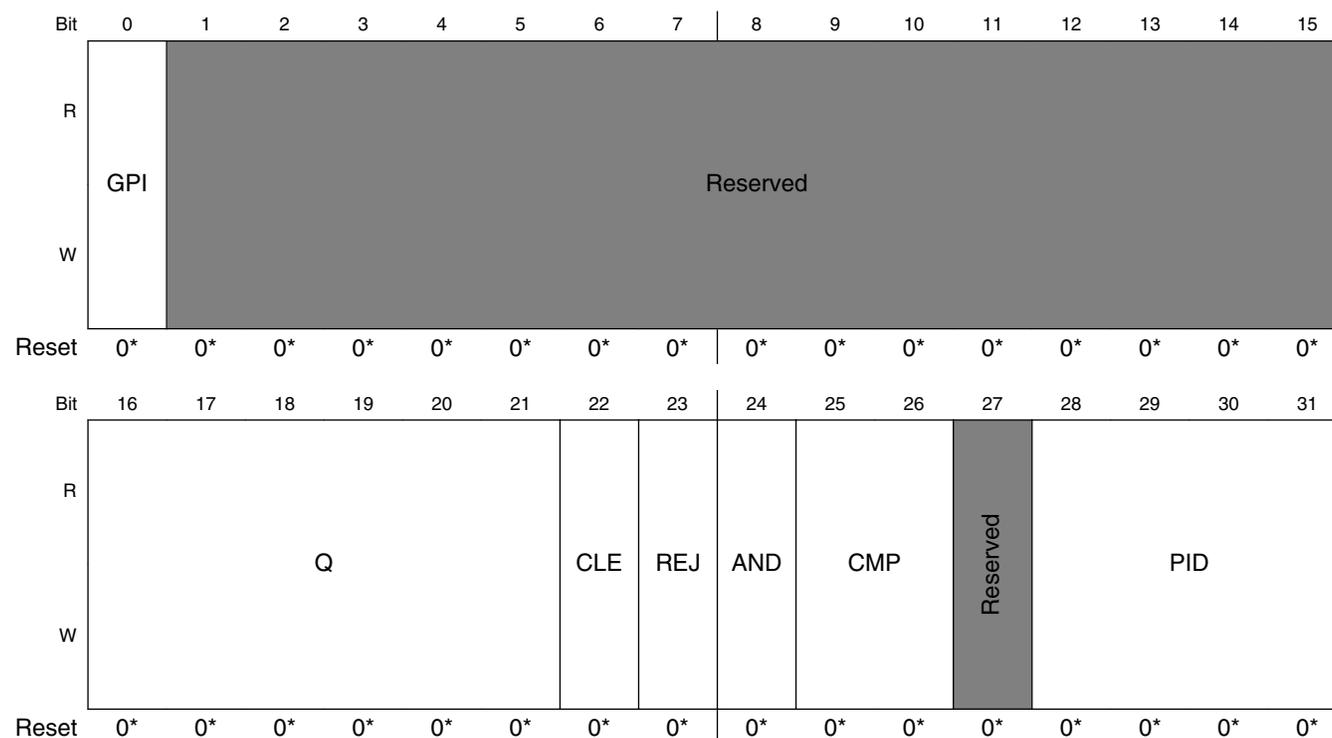
ETSECx_RQFAR field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24–31 RQFAR	Current index of receive queue filer table, which spans a total of 256 entries.

14.52.30 Receive queue filing table control register (ETSECx_RQFCR)

RQFCR is accessed to read or write the RQCTRL words in entries of the receive queue filer table. The table entries are described in greater detail in [Receive queue filer](#). The word accessed through RQFCR is defined by the current value of RQFAR.

Address: Base address + 338h offset



* Notes:

- Reset undefined

ETSECx_RQFCR field descriptions

Field	Description
0 GPI	General purpose interrupt. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the filer will instruct the Rx descriptor controller to set IEVENT[FGPI] when the corresponding receive frame is written to memory. If the timer is enabled (TMR_CTRL[TE] = 1), then TMR_PEVENT[RXP] will also be set.
1–15 -	This field is reserved. Reserved, should be written with zero.
16–21 Q	Receive queue index, from 0 to 63, inclusive, written into the Rx frame control block associated with the received frame. When a property matches the value in the RQPROP entry at this index, and REJ = 0 and AND = 0, the frame is sent to either RxBD ring 0 (if RCTRL[FSQEN] = 1) or the RxBD ring with index (Q

Table continues on the next page...

ETSECx_RQFCR field descriptions (continued)

Field	Description
	mod 8) and the filing table search is terminated. In the case where RCTRL[FSQEN] = 0, 8 virtual receive queues are overlaid on every RxBD ring, and software needs to consult the RQ field of the Rx frame control block to determine which virtual receive queue was chosen.
22 CLE	Cluster entry/exit (used in combination with AND bit). This bit brackets clusters, marking the start and end entries of a cluster. Clusters cannot be nested. 0 Regular RQCTRL entry. 1 If entry matches and AND = 1, treat subsequent entries as belonging to a nested cluster and enter the cluster; otherwise skip all entries up to and including the next cluster exit. If AND = 0, exit current cluster.
23 REJ	Reject frame. This bit and its specified action are ignored if AND = 1. 0 If entry matches, accept frame and file it to RxBD ring Q. 1 If entry matches, reject frame and discard it, ignoring Q.
24 AND	AND, in combination with CLE, REJ, and PID match, determines whether the filer will accept or reject a frame, defer evaluation until the next rule, exit a cluster, or skip a rule or set of rules. If CLE is zero: 0 Match property[PID] against RQPROP. If matched, accept or reject frame based on REJ. Otherwise skip to next rule. 1 Match property[PID] against RQPROP. If matched, defer evaluation to next rule. Otherwise, skip all rules up to and including the next rule with AND = 0. If the next rule with AND = 0 has CLE = 1, then also exit cluster. If CLE is one: 0 Match property[PID] against RQPROP. If matched, accept or reject frame based on REJ. Otherwise, exit cluster. 1 Match property[PID] against RQPROP. If matched, enter cluster. Otherwise, skip all rules up to and including the next rule while CLE = 1 and AND = 0.
25–26 CMP	Comparison operation to perform on the RQPROP entry at this index when PID > 0. The property value extracted by the frame parser is masked by the 32-bit <i>mask_register</i> prior to comparison against RQPROP. However, the property value is not permanently altered by the value in <i>mask_register</i> . By default, <i>mask_register</i> is initialized to 0xFFFF_FFFF before each frame is processed. In the case where PID = 0, CMP is interpreted as follows: 00/01 Filer <i>mask_register</i> is set to all 32 bits of RQPROP, and this entry always <i>matches</i> . 10/11 Filer <i>mask_register</i> is set to all 32 bits of RQPROP, and this entry always <i>fails to match</i> . In the case where PID > 0, CMP is interpreted as follows (& is bit-wise AND operator): 00 <i>property</i> [PID] & <i>mask_register</i> = RQPROP 01 <i>property</i> [PID] & <i>mask_register</i> >= RQPROP 10 <i>property</i> [PID] & <i>mask_register</i> != RQPROP 11 <i>property</i> [PID] & <i>mask_register</i> < RQPROP
27 -	This field is reserved. Reserved, should be written with zero.
28–31 PID	Property identifier. The value in the RQFPR[RQPROP] entry at this index is interpreted according to PID (see Receive queue filing table property register (ETSEC_RQFPR)).

14.52.31 Receive queue filing table property register (ETSECx_RQFPR)

NOTE

This section describes the layout of RQFPR with property ID 1 only; for a description of its use with property IDs 0 and 2–15, see [Table 14-59](#).

RQFPR is accessed to read or write the RQPROP words in entries of the receive queue filer table. The table entries are described in greater detail in [Receive queue filer](#). The word accessed through RQFPR is defined by the current value of RQFAR. The below figure and table describe the fields of the RQFPR register according to property ID, designated in RQFCR[PID].

Table 14-270. RQFPR Field Descriptions

PID ¹	Bit	Name	Description
0000	0-31	MASK	Mask bits to be written to Filer <i>mask_register</i> for masking of property values. The rule match/fail status for this PID is determined by RQCTRL[<i>CMP</i>]. Since <i>mask_register</i> is bit-wise ANDed with properties, every bit of MASK that is cleared also results in the corresponding property bit being cleared in comparisons. Therefore setting MASK to 0xFFFF_FFFF ensures that all property bits participate in rule matches.
0010	0-7	ARB	User-defined arbitrary bit field property: byte 0 extracted. Defaults to 0x00.
	8-15		User-defined arbitrary bit field property: byte 1 extracted. Defaults to 0x00.
	16-23		User-defined arbitrary bit field property: byte 2 extracted. Defaults to 0x00.
	24-31		User-defined arbitrary bit field property: byte 3 extracted. Defaults to 0x00.

Table continues on the next page...

Table 14-270. RQFPR Field Descriptions (continued)

PID ¹	Bit	Name	Description
0011	0-7	-	Reserved, should be written with zero.
	8-31	DAH	Destination MAC address, most significant 24 bits. Defaults to 0x000000.
0100	0-7	-	Reserved, should be written with zero.
	8-31	DAL	Destination MAC address, least significant 24 bits. Defaults to 0x000000.
0101	0-7	-	Reserved, should be written with zero.
	8-31	SAH	Source MAC address, most significant 24 bits. Defaults to 0x000000.
0110	0-7	-	Reserved, should be written with zero.
	8-31	SAL	Source MAC address, least significant 24 bits. Defaults to 0x000000.
0111	0-15	-	Reserved, should be written with zero.

Table continues on the next page...

Table 14-270. RQFPR Field Descriptions (continued)

PID ¹	Bit	Name	Description
0111	16-31	ETY	<p>Ethertype of next layer protocol, that is, last ethertype if layer 2 headers nest. Defaults to 0xFFFF.</p> <p>Using the filer to match ETY does not work in the case of PPPoE packets, because the PPPoE ethertype in the original packet, 0x8864, is always overwritten with the PPP protocol field. Thus, matches on ETY == 0x8864 always fail.</p> <p>Instead, software should use PID=1 fields IP4 (ETY = 0x0021) and IP6 (ETY = 0x0057) to distinguish PPPoE session packets carrying IPv4 and IPv6 datagrams. Other PPP protocols are encoded in the ETY field, but many of them overlap with real ethertype definitions. Consult IANA and IEEE for possible ambiguities.</p> <p>Packets with a value in the length/type field greater than 1500 and less than 1536 are treated as payload length. If the eTSEC is used in a network where there are packets carrying a type designation between 1500 and 1536 (note there are none currently publicly defined by IANA), then the S/W must confirm the parser and filer results by checking the type/length field after the packet has been written to memory to see if it falls in this range.</p> <p>Note that the eTSEC filer gets multiple packet attributes as a result of parsing the packet. The behavior of the eTSEC is that it pulls the innermost ethertype found in the packet; this means that in many supported protocols that have inner etherypes, in order to file based on the outer ethertype, arbitrary extraction should be used instead of the ETY PID. There are four cases that need to be highlighted.</p> <p>The jumbo ethertype (0x8870)-In this case, the eTSEC assumes that the</p>
826			<p>Freescale Semiconductor, Inc.</p>

Table 14-270. RQFPR Field Descriptions (continued)

PID ¹	Bit	Name	Description
1000	0-19	-	Reserved, should be written with zero.
	20-31	VID	VLAN network identifier (as per IEEE Std 802.1Q). This value defaults to 0x000 if no VLAN tag was found, or the VLAN tag contained only priority information.
1001	0-28	-	Reserved, should be written with zero.
	29-31	PRI	VLAN user priority (as per IEEE Std 802.1p). This value defaults to 000 (best effort priority) if no VLAN tag was found.
1010	0-23	-	Reserved, should be written with zero.
	24-31	TOS	IPv4 header Type Of Service field or IPv6 Traffic Class field. This value defaults to 0x00 (default RFC 2474 best-effort behavior) if no IP header appeared. Note that for IPv6 the Traffic Class field is extracted using the IP header definition in RFC 2460. IPv6 headers formed using the earlier RFC 1883 have a different format and must be handled with software. (Software should acknowledge the PIC=1 IP6 bit to distinguish proper alignment of the TOS field.)
1011	0-23	-	Reserved, should be written with zero.
	24-31	L4P	Layer 4 protocol identifier as per published IANA specification. This is the last recognized protocol type recognized in the case of IPv6 extension headers. This value defaults to 0xFF to indicate that no layer 4 header was recognized (possibly due to absence of an IP header).

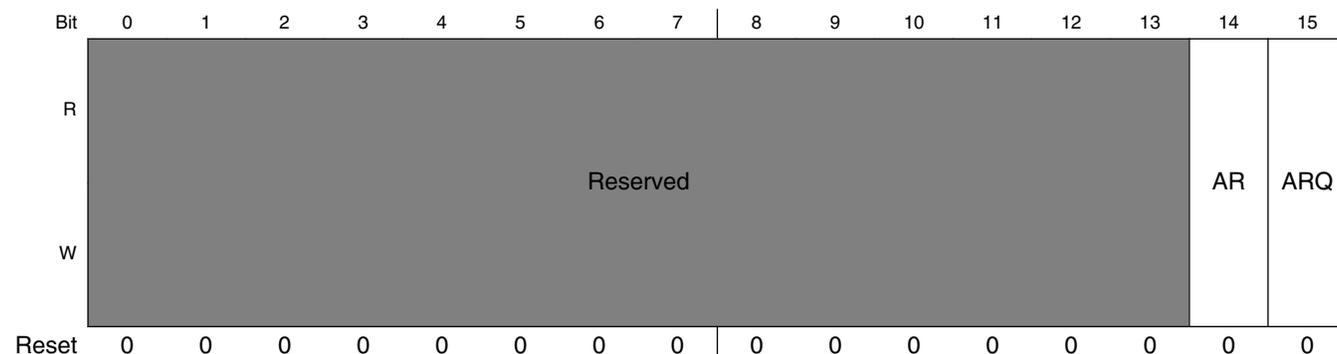
Table continues on the next page...

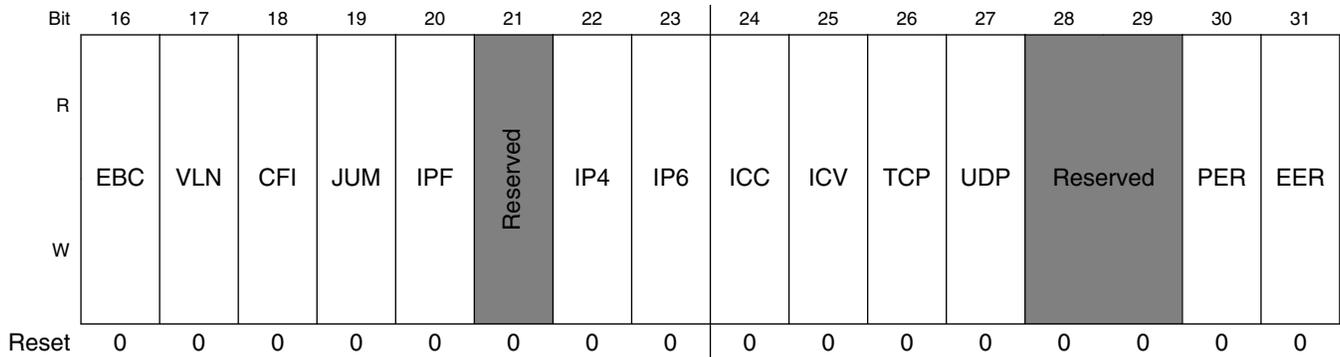
Table 14-270. RQFPR Field Descriptions (continued)

PID ¹	Bit	Name	Description
1100	0-31	DIA	Destination IP address. If an IPv4 header was found, this is the entire destination address. If an IPv6 header was found, this is the 32 most significant bits of the 128-bit destination address. This value defaults to all zeros if no IP header appeared.
1101	0-31	SIA	Source IP address. If an IPv4 header was found, this is the entire source address. If an IPv6 header was found, this is the 32 most significant bits of the 128-bit source address. This value defaults to all zeros if no IP header appeared.
1110	0-15	-	Reserved, should be written with zero.
	16-31	DPT	Destination port number for TCP or UDP headers. This value defaults to 0x0000 if no TCP or UDP headers were recognized.
1111	0-15	-	Reserved, should be written with zero.
	16-31	SPT	Source port number for TCP or UDP headers. This value defaults to 0x0000 if no TCP or UDP headers were recognized.

1. PID is the property identifier field of the filer table control entry (see RQFCR[PID]) at the same index.
1. PID is the property identifier field of the filer table control entry (see RQFCR[PID]) at the same index.
1. PID is the property identifier field of the filer table control entry (see RQFCR[PID]) at the same index.
1. PID is the property identifier field of the filer table control entry (see RQFCR[PID]) at the same index.

Address: Base address + 33Ch offset





ETSECx_RQFPR field descriptions

Field	Description
0–13 -	This field is reserved. Reserved (PID 0001)
14 AR	Set if an ARP response packet is seen. (PID 0001)
15 ARQ	Set if an ARP request packet is seen. (PID 0001)
16 EBC	Set if the destination Ethernet address is to the broadcast address. (PID 0001)
17 VLN	Set if a VLAN tag (Ethertype DFVLAN[TAG] or 0x8100) was seen in the frame. (PID 0001)
18 CFI	Set to the value of the Canonical Format Indicator in the VLAN control tag if VLAN is set, zero otherwise. Note that this field is valid only if the first encountered Ethertype in a frame is VLAN. (PID 0001)
19 JUM	Set if a jumbo Ethernet frame was parsed. (PID 0001)
20 IPF	Set if a fragmented IPv4 or IPv6 header was encountered. See the descriptions of receive FCB fields IP and PRO in Receive path off-load for more information on determining the status of received packets for which IPF is set. (PID 0001)
21 -	This field is reserved. Reserved (PID 0001)
22 IP4	Set if an IPv4 header was parsed. (PID 0001)
23 IP6	Set if an IPv6 header was parsed. (PID 0001)
24 ICC	Set if the IPv4 header checksum was checked. (PID 0001)

Table continues on the next page...

ETSECx_RQFPR field descriptions (continued)

Field	Description
25 ICV	Set if the IPv4 header checksum was verified correct. (PID 0001)
26 TCP	Set if a TCP header was parsed. (PID 0001)
27 UDP	Set if a UDP header was parsed. (PID 0001)
28–29 -	This field is reserved. Reserved. (PID 0001)
30 PER	Set on a parse error, such as header inconsistency. (PID 0001)
31 EER	Set on an Ethernet framing error that prevents parsing. (PID 0001)

1. PID is the property identifier field of the filter table control entry (see RQFCR[PID]) at the same index.

14.52.32 Maximum receive buffer length register (ETSECx_MRBLR)

The MRBLR register is written by the user. It informs the eTSEC how much space is in the receive buffer pointed to by the RxBD.

Address: Base address + 340h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															MRBL										-						
W	Reserved															MRBL										-						
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

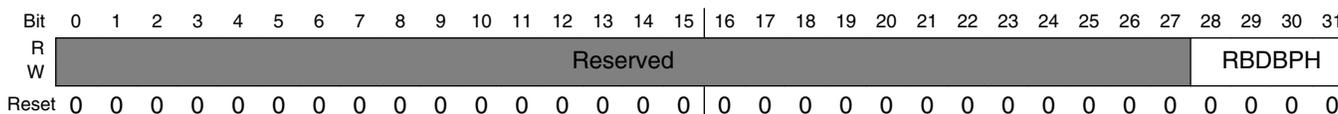
ETSECx_MRBLR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–25 MRBL	Maximum receive buffer length. MRBL is the number of bytes that the eTSEC receiver writes to the receive buffer. The MRBL register is written by the user with a multiple of 64 for all modes. The eTSEC can write fewer bytes to the buffer than the value set in MRBL if a condition such as an error or end-of-frame occurs, but it never exceeds the MRBL value; therefore, user-supplied buffers must be at least as large as the MRBL. MRBL must be set, together with the number of buffer descriptors, to ensure adequate space for received frames. See Maximum frame length (ETSEC_MAXFRM) ," for further discussion.
26–31 -	To ensure that MRBL is a multiple of 64, these bits are reserved and should be cleared.

14.52.33 Rx data buffer pointer high bits (ETSECx_RBDBPH)

The RBDBPH register is written by the user with the most significant address bits common to all RxBD buffer addresses, RxBD[Data Buffer Pointer]. As a consequence, Rx buffers must be placed in a 4 Gbyte segment of memory whose base address is prefixed by the bits in RBDBPH. The RxBD ring itself can reside in a different memory region (based at RBASEH).

Address: Base address + 380h offset



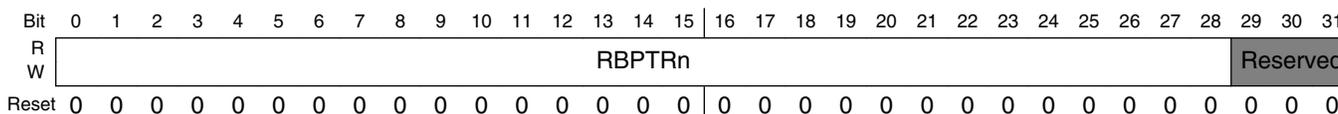
ETSECx_RBDBPH field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 RBDBPH	Most significant bits common to all data buffer addresses contained in RxBDs. The user must initialize RBDBPH before enabling the eTSEC receive function.

14.52.34 RxBD pointer for ring n (ETSECx_RBPTRn)

BPTR0-BPTR7 each contains the low-order 32 bits of the next receive buffer descriptor address for their respective RxBD ring. These registers takes on the value of their ring's associated RBASE when the RBASE register is written by software. Software must not write RBPTR *n* while eTSEC is actively receiving frames. However, RBPTR *n* can be modified when the receiver is disabled or when no Rx buffer is in use (after a GRACEFUL STOP RECEIVE command is issued and the frame completes its reception) in order to change the next RxBD eTSEC receives.

Address: Base address + 384h offset + (8d × i), where i=0d to 7d



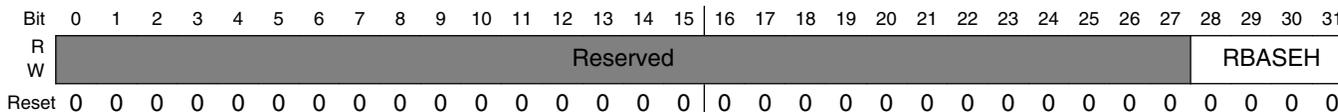
ETSECx_RBPTRn field descriptions

Field	Description
0–28 RBPTRn	Current RxBD pointer for RxBD ring <i>n</i> . Points to the current BD being processed or to the next BD the receiver uses when it is idling. After reset or when the end of the RxBD ring is reached, eTSEC initializes RBPTR <i>n</i> to the value in the corresponding RBASE <i>n</i> . The RBPTR register is internally written by the eTSEC's DMA controller during reception. The pointer increments by 8 (bytes) each time a descriptor is closed successfully by the eTSEC. Note that the 3 least-significant bits of this register are read only and zero.
29–31 -	This field is reserved. Reserved

14.52.35 RxBD base address high bits (ETSECx_RBASEH)

The RBASEH register is written by the user with the most significant address bits common to all RxBD addresses, including RBASE0-RBASE7 and RBPTR0-RBPTR7. As a consequence, RxBD rings must be placed in a 4 Gbyte segment of memory whose base address is prefixed by the bits in RBASEH. However, Rx data buffers may potentially reside in a different memory region based at RBDBPH.

Address: Base address + 400h offset



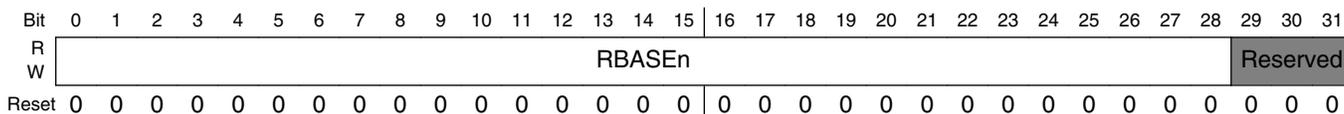
ETSECx_RBASEH field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 RBASEH	Most significant bits common to all RxBD addresses-except data buffer pointers. The user must initialize RBASEH before enabling the eTSEC receive function.

14.52.36 RxBD base address of ring n (ETSECx_RBASEn)

The RBASEn registers are written by the user with the base address of each RxBD ring n. Each such value must be divisible by eight, since the 3 least-significant bits always write as 000.

Address: Base address + 404h offset + (8d × i), where i=0d to 7d



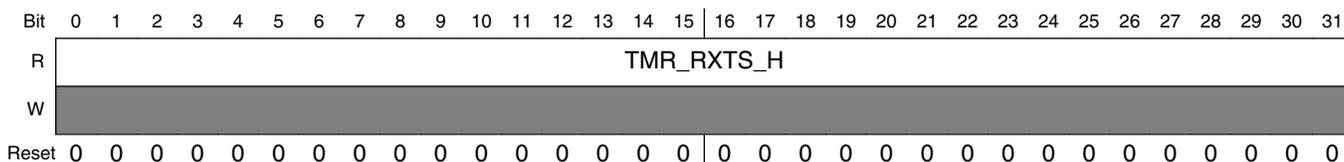
ETSECx_RBASEn field descriptions

Field	Description
0–28 RBASEn	Receive base for ring n. RBASE defines the starting location in the memory map for the eTSEC RxBDs. This field must be 8-byte aligned. Together with setting the W (wrap) bit in the last BD, the user can select how many BDs to allocate for the receive packets. The user must initialize RBASE before enabling the eTSEC receive function on the associated ring.
29–31 -	This field is reserved. Reserved

14.52.37 Rx timer time stamp register high (ETSECx_TMR_RXTS_H)

Receive time stamp register high (TMR_RXTS_H). This register, together with TMR_RXTS_L, holds the value present in TMR_CNT_H/L when the eTSEC detects a new incoming Ethernet frame. This register is only updated when the precision time stamp logic is enable via TMR_CTRL[TE]. This register is read only in normal operation.

Address: Base address + 4C0h offset



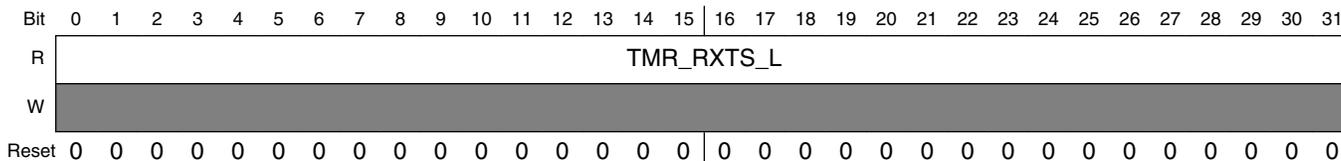
ETSECx_TMR_RXTS_H field descriptions

Field	Description
0–31 TMR_RXTS_H	Value of the eTSEC precision timer upon detection of a start of frame symbol for the received frame. (upper half)

14.52.38 Rx timer time stamp register low (ETSECx_TMR_RXTS_L)

Receive time stamp register low (TMR_RXTS_L). This register, together with TMR_RXTS_H, holds the value present in TMR_CNT_H/L when the eTSEC detects a new incoming Ethernet frame. This register is only updated when the precision time stamp logic is enable via TMR_CTRL[TE]. This register is read only in normal operation.

Address: Base address + 4C4h offset



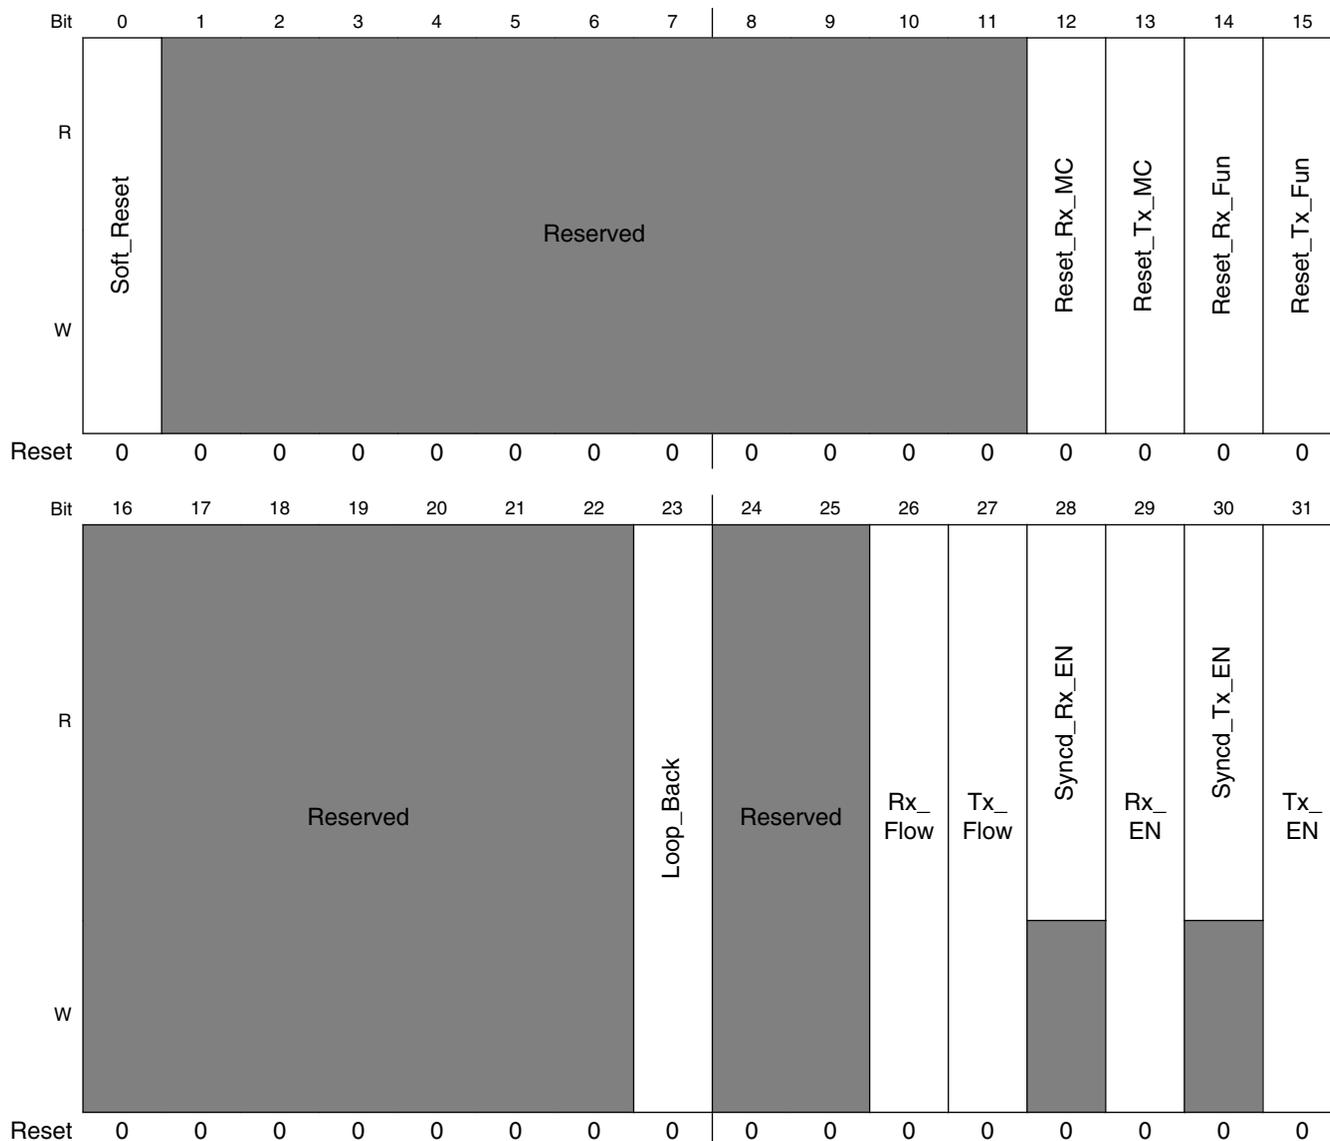
ETSECx_TMR_RXTS_L field descriptions

Field	Description
0–31 TMR_RXTS_L	Value of the eTSEC precision timer upon detection of a start of frame symbol for the received frame. (lower half)

14.52.39 MAC configuration register 1 (ETSECx_MACCFG1)

MACCFG1 is written by the user. Note that the reset bits in MACCFG1 are not self-resetting. If software sets any reset bit in MACCFG1, it must leave the bit asserted at least three Tx clocks, then clear the bit before resuming normal operation.

Address: Base address + 500h offset



ETSECx_MACCFG1 field descriptions

Field	Description
0 Soft_Reset	Soft reset. This bit is cleared by default. See Soft reset and reconfiguring procedure for more information on setting this bit. 0 Normal operation. 1 Place the entire MAC in reset except for the host interface.
1–11 -	This field is reserved. Reserved
12 Reset_Rx_MC	Reset receive MAC control block. This bit is cleared by default. 0 Normal operation. 1 Place the receive part of the MAC in reset. This block detects control frames and contains the pause timers.

Table continues on the next page...

ETSECx_MACCFG1 field descriptions (continued)

Field	Description
13 Reset_Tx_MC	Reset transmit MAC control block. This bit is cleared by default. 0 Normal operation. 1 Place the transmit part of the MAC in reset. This block multiplexes data and control frame transfers. It also responds to XOFF PAUSE control frames.
14 Reset_Rx_Fun	Reset receive function block. This bit is cleared by default. 0 Normal operation. 1 Place the receive function in reset. This block performs the receive frame protocol.
15 Reset_Tx_Fun	Reset transmit function block. This bit is cleared by default. 0 Normal operation. 1 Place the transmit function in reset. This block performs the frame transmission protocol.
16–22 -	This field is reserved. Reserved
23 Loop_Back	Loop back. This bit is cleared by default. 0 Normal operation. 1 Loop back the MAC transmit outputs to the MAC receive inputs.
24–25 -	This field is reserved. Reserved
26 Rx_Flow	Receive flow. This bit is cleared by default. Must be 0 if MACCFG2[Full Duplex] = 0. NOTE: Should not be set when operating in Half-Duplex mode 0 The receive MAC control ignores PAUSE flow control frames. 1 The receive MAC control detects and acts on PAUSE flow control frames.
27 Tx_Flow	Transmit flow. This bit is cleared by default. Must be 0 if MACCFG2[Full Duplex] = 0. NOTE: Should not be set when operating in Half-Duplex mode 0 The transmit MAC control may not send PAUSE flow control frames if requested by the system. 1 The transmit MAC control may send PAUSE flow control frames if requested by the system.
28 Syncd_Rx_EN	Receive enable synchronized to the receive stream. (Read-only) 0 Frame reception is not enabled. 1 Frame reception is enabled.
29 Rx_EN	Receive enable. This bit is cleared by default. If set, prior to clearing this bit, set DMACTRL[GRS] then confirm subsequent occurrence of the graceful receive stop interrupt (IEVENT[GRSC] is set). 0 The MAC may not receive frames from the PHY. 1 The MAC may receive frames from the PHY.
30 Syncd_Tx_EN	Transmit enable synchronized to the transmit stream. (Read-only) 0 Frame transmission is not enabled. 1 Frame transmission is enabled.

Table continues on the next page...

ETSECx_MACCFG1 field descriptions (continued)

Field	Description
31 Tx_EN	Transmit enable. This bit is cleared by default. If set, prior to clearing this bit, set DMACTRL[GTS] then confirm subsequent occurrence of the graceful transmit stop interrupt (IEVENT[GTSC] is set). 0 The MAC may not transmit frames from the system. 1 The MAC may transmit frames from the system.

14.52.40 MAC configuration register 2 (ETSECx_MACCFG2)

The MACCFG2 register is written by the user.

Address: Base address + 504h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Preamble_Length			Reserved		I_F_Mode		PreAM_RxEN	PreAM_TxEN	Huge_Frame	Length_check	MPEN	PAD_CRC	CRC_EN	Full_Duplex	
W																
Reset	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_MACCFG2 field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–19 Preamble_Length	This field determines the length in bytes of the preamble field preceding each Ethernet start-of-frame delimiter byte. Values from 0x3 to 0xF are supported by the controller . The default value of 0x7 should not be altered in order to guarantee reliable operation with IEEE 802.3 compliant hardware.
20–21 -	This field is reserved. Reserved
22–23 I_F_Mode	This field determines the type of interface to which the MAC is connected. Its default is 00. 00 Reserved bit mode 01 Nibble mode (10/100 Mbps) 10 Byte mode(1000 Mbps) 11 Reserved

Table continues on the next page...

ETSECx_MACCFG2 field descriptions (continued)

Field	Description																				
24 PreAM_RxEN	<p>User defined preamble enable for received frames. This bit is cleared by default.</p> <p>0 The MAC skips the Ethernet preamble without returning it. 1 The MAC recovers the received Ethernet preamble and passes it to the driver at the start of each received frame. If the preamble is less than 7 bytes, 0's are prepended to pad it to 7 bytes. Not applicable to mode. RMII 10/100</p>																				
25 PreAM_TxEN	<p>User defined preamble enable for transmitted frames. This bit is cleared by default.</p> <p>0 The MAC generates a standard Ethernet preamble. 1 If a user-defined preamble has been passed to the MAC it is transmitted instead of the standard preamble. Otherwise the standard Ethernet preamble is generated. The Preamble Length field should be left at its default setting if a user-defined preamble is transmitted. Not applicable to mode. RMII 10/100</p>																				
26 Huge_Frame	<p>Huge frame enable. This bit is cleared by default.</p> <p>Note that if Huge Frame is cleared, the user must ensure that adequate buffer space is allocated for received frames. See Maximum frame length (ETSEC_MAXFRM) for further information.</p> <table border="1"> <thead> <tr> <th>Frame type</th> <th>Frame length</th> <th>Packet truncation</th> <th>Buffer descriptor updated</th> </tr> </thead> <tbody> <tr> <td>Receive or transmit</td> <td>> maximum frame length</td> <td>yes</td> <td>yes</td> </tr> <tr> <td>Receive</td> <td>= maximum frame length</td> <td>no</td> <td>no</td> </tr> <tr> <td>Transmit</td> <td>= maximum frame length</td> <td>no</td> <td>yes</td> </tr> <tr> <td>Receive or transmit</td> <td>< maximum frame length</td> <td>no</td> <td>no</td> </tr> </tbody> </table> <p>0 Limit the length of frames received to less than or equal to the maximum frame length value (MAXFRM[Maximum Frame]) and limit the length of frames transmitted to less than the maximum frame length. See Buffer descriptors for further details of buffer descriptor bit updating. 1 Frames are transmitted and received regardless of their relationship to the maximum frame length.</p>	Frame type	Frame length	Packet truncation	Buffer descriptor updated	Receive or transmit	> maximum frame length	yes	yes	Receive	= maximum frame length	no	no	Transmit	= maximum frame length	no	yes	Receive or transmit	< maximum frame length	no	no
Frame type	Frame length	Packet truncation	Buffer descriptor updated																		
Receive or transmit	> maximum frame length	yes	yes																		
Receive	= maximum frame length	no	no																		
Transmit	= maximum frame length	no	yes																		
Receive or transmit	< maximum frame length	no	no																		
27 Length_check	<p>Length check. This bit is cleared by default.</p> <p>0 No length field checking is performed. 1 The MAC checks the frame's length field on receive to ensure it matches the actual data field length. Transmitted frames are not checked.</p>																				
28 MPEN	<p>Magic packet enable for Ethernet modes. This bit is cleared by default. MPEN should be enabled only after GRACEFUL RECEIVE STOP and GRACEFUL TRANSMIT STOP are completed successfully (in other words, transmission and reception have stopped).</p> <p>0 Normal receive behavior on receive, or Magic Packet mode has exited with reception of a valid Magic Packet. 1 Commence Magic Packet detection by the MAC provided that frame reception is enabled in MACCFG1. In this mode the MAC ignores all received frames until the specific Magic Packet frame is received, at which point this bit is cleared by the eTSEC, and a maskable interrupt through IEVENT[MAG] occurs.</p>																				
29 PAD_CRC	<p>Pad and append CRC. This bit is cleared by default. This bit must be set when in half-duplex mode (MACCFG2[Full Duplex] is cleared).</p>																				

Table continues on the next page...

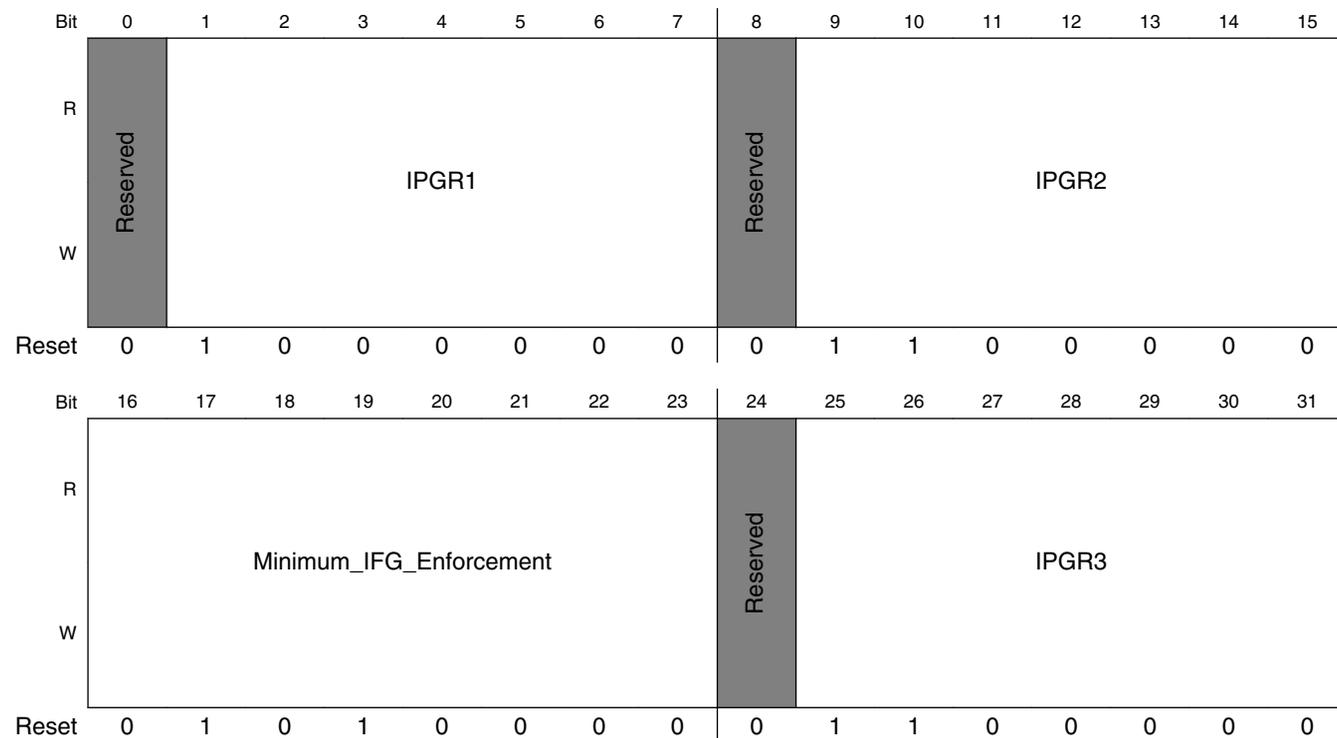
ETSECx_MACCFG2 field descriptions (continued)

Field	Description
	0 Frames presented to the MAC have a valid length and contain a CRC. 1 The MAC pads all transmitted short frames and appends a CRC to every frame regardless of padding requirement.
30 CRC_EN	CRC enable. If the configuration bit PAD/CRC ENABLE or the per-packet PAD/CRC ENABLE is set, CRC ENABLE is ignored. This bit is cleared by default. 0 Frames presented to the MAC have a valid length and contain a valid CRC. 1 The MAC appends a CRC on all frames. Clear this bit if frames presented to the MAC have a valid length and contain a valid CRC.
31 Full_Duplex	Full duplex configure. This bit is cleared by default. 0 The MAC operates in half-duplex mode only. 1 The MAC operates in full-duplex mode.

14.52.41 Inter-packet/inter-frame gap register (ETSECx_IPGIFG)

The IPGIFG register is written by the user.

Address: Base address + 508h offset



ETSECx_IPGIFG field descriptions

Field	Description
0 -	This field is reserved. Reserved
1-7 IPGR1	Non-Back-to-Back Inter-Packet-Gap, Part 1 This is a programmable field representing the optional carrier sense window referenced in IEEE 802.3/4.2.3.2.1 'carrier deference'. If carrier is detected during the timing of IPGR1, the MAC defers to carrier. If, however, carrier becomes active after IPGR1, the MAC continues timing IPGR2 and transmits, knowingly causing a collision, thus ensuring fair access to medium. Its range of values is 0x00 to IPGR2. Its default is 0x40 (64d) which follows the two-thirds/one-third guideline.
8 -	This field is reserved. Reserved
9-15 IPGR2	Non-Back-to-Back Inter-Packet-Gap, Part 2
16-23 Minimum_IFG_Enforcement	This is a programmable field representing the minimum number of bits of IFG to enforce between frames. A frame is dropped whose IFG is less than that programmed. The default setting of 0x50 (80d) represents half of the nominal minimum IFG which is 160 bits.
24 -	This field is reserved. Reserved
25-31 IPGR3	Back-to-Back Inter-Packet-Gap

14.52.42 Half-duplex control (ETSECx_HAFDUP)

The HAFDUP register is written by the user.

Address: Base address + 50Ch offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved								Alternate_BEB_Truncation				Alt_BEB	BP_No_BackOff	No_BackOff	Excess_Defer
W	Reserved															
Reset	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	1
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Retransmission_Maximum				Reserved				Collision_Window							
W					Reserved											
Reset	1	1	1	1	0	0	0	0	0	0	1	1	0	1	1	1

ETSECx_HAFDUP field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–11 Alternate_BEB_ Truncation	This field is used while ALTERNATE BINARY EXPONENTIAL BACKOFF ENABLE is set. The value programmed is substituted for the Ethernet standard value of ten. Its default is 0xA.
12 Alt_BEB	Alternate binary exponential backoff. This bit is cleared by default. 0 The Tx MAC follows the standard binary exponential back off rule. 1 The Tx MAC uses the ALTERNATE BINARY EXPONENTIAL BACKOFF TRUNCATION setting instead of the 802.3 standard tenth collision. The standard specifies that any collision after the tenth uses one less than 210 as the maximum backoff time.
13 BP_No_BackOff	Back pressure no backoff. This bit is cleared by default. 0 The Tx MAC follows the binary exponential back off rule. 1 The Tx MAC immediately re-transmits, following a collision, during back pressure operation.
14 No_BackOff	No backoff. This bit is cleared by default. 0 The Tx MAC follows the binary exponential back off rule. 1 The Tx MAC immediately re-transmits following a collision.
15 Excess_Defer	Excessively deferred. This bit is set by default. 0 The Tx MAC aborts the transmission of a packet that is excessively deferred. 1 The Tx MAC allows the transmission of a packet that is excessively deferred.
16–19 Retransmission_ Maximum	This is a programmable field specifying the number of retransmission attempts following a collision before aborting the packet due to excessive collisions. The standard specifies the attempt limit to be 0xF (15d). Its default value is 0xF.
20–25 -	This field is reserved. Reserved
26–31 Collision_Window	This is a programmable field representing the slot time or collision window during which collisions occur in properly configured networks. Because the collision window starts at the beginning of transmission, the preamble and SFD are included. Its default of 0x37 (55d) corresponds to the count of frame bytes at the end of the window.

14.52.43 Maximum frame length (ETSECx_MAXFRM)

The MAXFRM register is written by the user.

Address: Base address + 510h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															Maximum_Frame																
W	Reserved															Maximum_Frame																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0

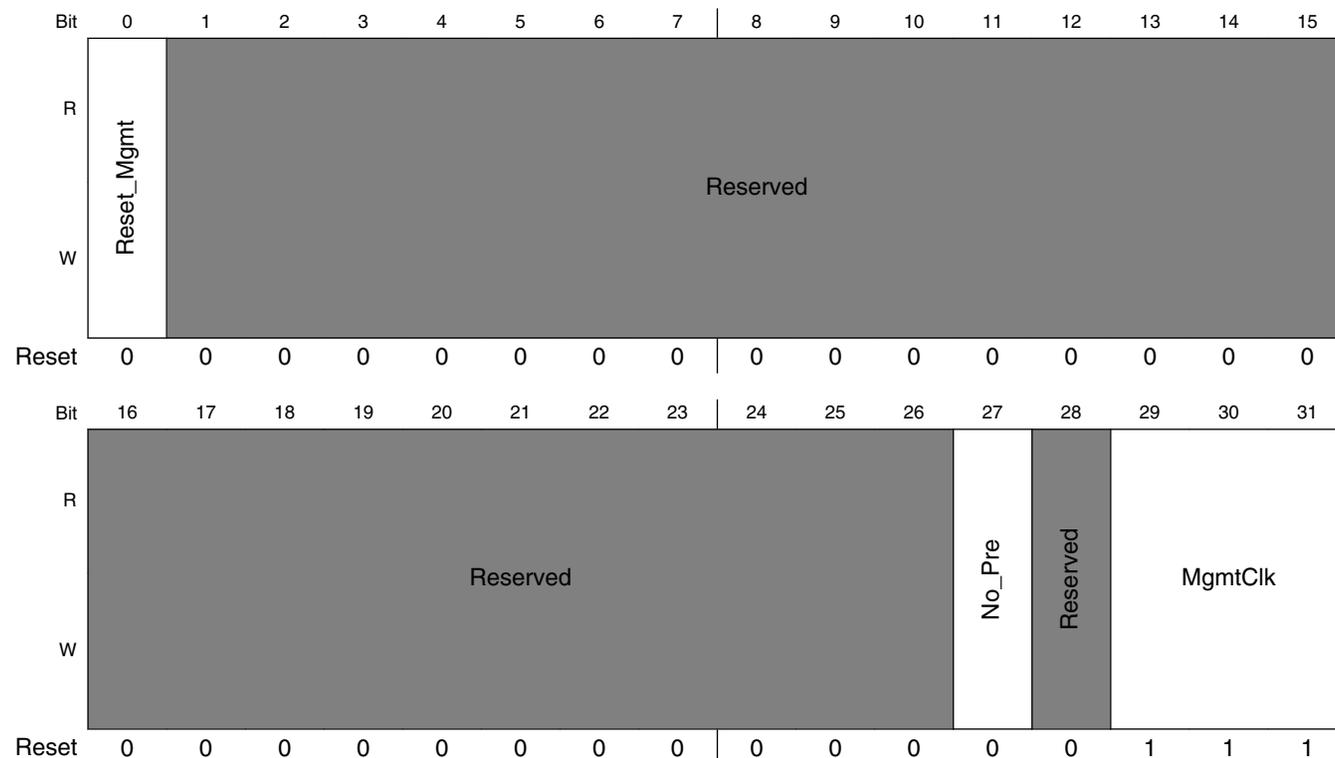
ETSECx_MAXFRM field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 Maximum_Frame	This field is set to 0x0600 (1536 bytes) by default and always must be set to a value greater than or equal to 0x0040 (64 bytes) , but not greater than 0x2580 (9600 bytes). It sets the maximum Ethernet frame size in both the transmit and receive directions. (Refer to MACCFG2[Huge Frame].) Note that if MACCFG2[Huge Frame] = 0, the value of this field must be less than or equal to MRBLR[MRBL] x (minimum number of RxBDs per ring). See MAC configuration register 2 (ETSEC_MACCFG2) , " Maximum receive buffer length register (ETSEC_MRBLR) ," and Receive buffer descriptors (RxBD) .

14.52.44 MII management configuration (ETSECx_MIIMCFG)

The MIIMCFG register is written by the user to configure all MII management operations. Note that MII management hardware is shared by all eTSECs. Thus, only through the MIIM registers of eTSEC1 can external PHYs be accessed and configured. Note: when an eTSEC is configured to use TBI/RTBI, configuration of the TBI/ RTBI (described in [Ten-bit interface \(TBI\)](#)) is done through the MIIM registers for that eTSEC. For example, if a TBI/RTBI interface is required on eTSEC2, then the MIIM registers starting at offset 0x2_5520 are used to configure it.

Address: Base address + 520h offset



ETSECx_MIIMCFG field descriptions

Field	Description
0 Reset_Mgmt	Reset management. This bit is cleared by default. 0 Allow the MII MGMT to perform mgmt read/write cycles if requested through the host interface. 1 Reset the MII MGMT.
1–26 -	This field is reserved. Reserved
27 No_Pre	Preamble suppress. This bit is cleared by default. 0 The MII MGMT performs Mgmt read/write cycles with 32 clocks of preamble. 1 The MII MGMT suppresses preamble generation and reduces the Mgmt cycle from 64 clocks to 32 clocks. This is in accordance with IEEE 802.3/22.2.4.4.2.
28 -	This field is reserved. Reserved
29–31 MgmtClk	This field determines the clock frequency of the MII management clock (EC_MDC). Its default value is 111. NOTE: The eTSEC system clock is derived from platform clock/2. 00x 1/4 of the eTSEC system clock divided by 8 010 1/6 of the eTSEC system clock divided by 8 011 1/8 of the eTSEC system clock divided by 8 100 1/10 of the eTSEC system clock divided by 8 101 1/14 of the eTSEC system clock divided by 8 110 1/20 of the eTSEC system clock divided by 8 111 1/28 of the eTSEC system clock divided by 8

14.52.45 MII management command (ETSECx_MIIMCOM)

The MIIMCOM register is written by the user.

Address: Base address + 524h offset

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R	Reserved																
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	Reserved														Scan_Cycle	Read_Cycle	
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

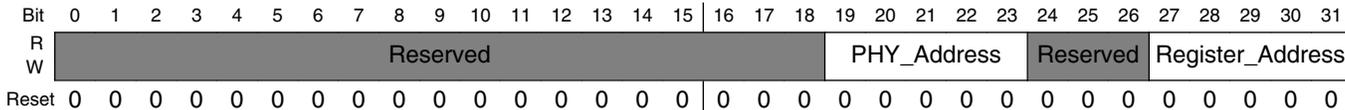
ETSECx_MIIMCOM field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 Scan_Cycle	Scan cycle. This bit is cleared by default. 0 Normal operation. 1 The MII management continuously performs read cycles. This is useful for monitoring link fail, for example.
31 Read_Cycle	Read cycle. This bit is cleared by default but is not self-clearing once set. 0 Normal operation. 1 The MII management performs a single read cycle upon the transition of this bit from 0 to 1 using the PHY address (at MIIMADD[PHY Address]) and the register address (at MIIMADD[Register Address]). The 0-to-1 transition of this bit also causes the MIIMIND[Busy] bit to be set. The read is complete when the MIIMIND[Busy] bit clears. Data is returned in register MIIMSTAT[PHY Status].

14.52.46 MII management address (ETSECx_MIIMADD)

The MIIMADD register is written by the user.

Address: Base address + 528h offset



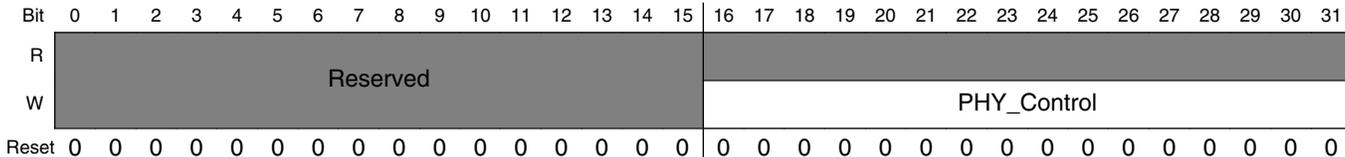
ETSECx_MIIMADD field descriptions

Field	Description
0–18 -	This field is reserved. Reserved
19–23 PHY_Address	This field represents the 5-bit PHY address field of Mgmt cycles. Up to 31 PHYs can be addressed (0 is reserved). Its default value is 0x00.
24–26 -	This field is reserved. Reserved
27–31 Register_Address	This field represents the 5-bit register address field of Mgmt cycles. Up to 32 registers can be accessed. Its default value is 0x00.

14.52.47 MII management control (ETSECx_MIIMCON)

MIIMCON is written by the user.

Address: Base address + 52Ch offset



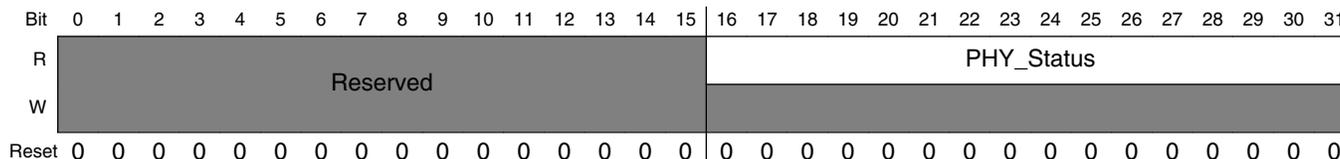
ETSECx_MIIMCON field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 PHY_Control	If written, an MII Mgmt write cycle is performed using this 16-bit data, the pre-configured PHY address (at MIIMADD[PHY Address]) and the register address (at MIIMADD[Register Address]). Its default value is 0x0000.

14.52.48 MII management status (ETSECx_MIIMSTAT)

The MIIMSTAT register is read only by the user.

Address: Base address + 530h offset



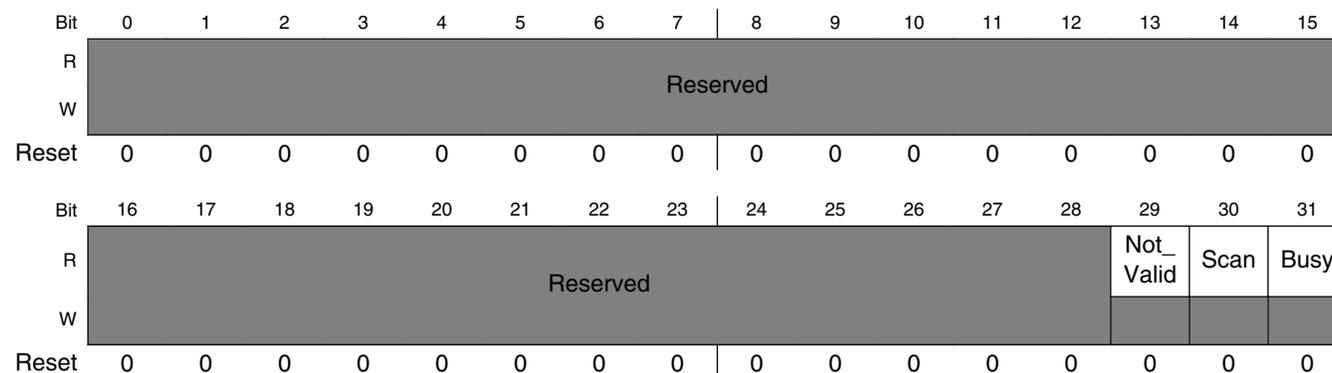
ETSECx_MIIMSTAT field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 PHY_Status	Following an MII Mgmt read cycle, the 16-bit data can be read from this location. Its default value is 0x0000.

14.52.49 MII management indicator (ETSECx_MIIMIND)

The MIIMIND register is read-only by the user.

Address: Base address + 534h offset



ETSECx_MIIMIND field descriptions

Field	Description
0–28 -	This field is reserved. Reserved

Table continues on the next page...

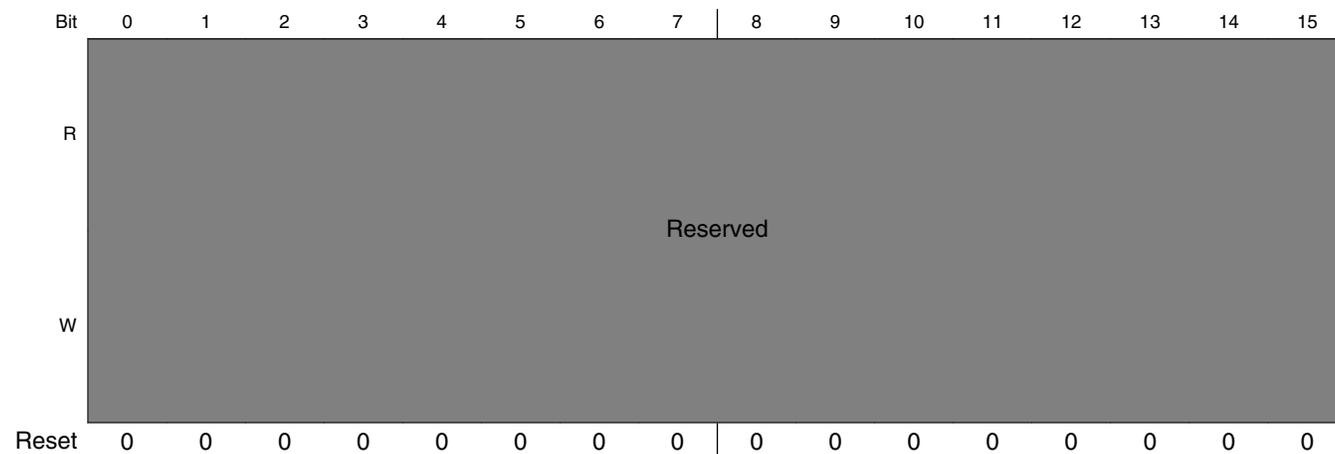
ETSECx_MIIMIND field descriptions (continued)

Field	Description
29 Not_Valid	Not valid. 0 MII Mgmt read cycle has completed and the read data is valid. 1 MII Mgmt read cycle has not completed and the read data is not yet valid.
30 Scan	Scan in progress. 0 A scan operation (continuous MII Mgmt read cycles) is not in progress. 1 A scan operation (continuous MII Mgmt read cycles) is in progress.
31 Busy	Busy. 0 MII Mgmt block is not currently performing an MII Mgmt read or write cycle. 1 MII Mgmt block is currently performing an MII Mgmt read or write cycle.

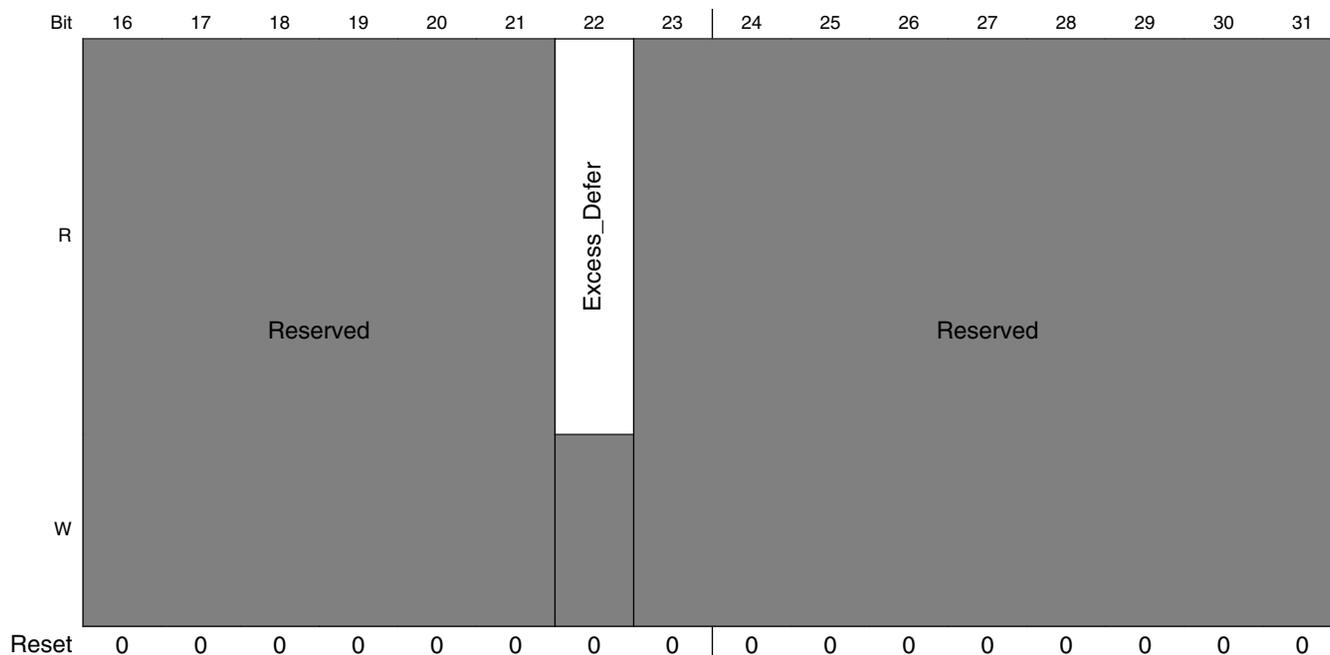
14.52.50 Interface status (ETSECx_IFSTAT)

Interface status (ETSEC_IFSTAT) shows the IFSTAT register.

Address: Base address + 53Ch offset



ETSEC memory map/register definition



ETSECx_IFSTAT field descriptions

Field	Description
0–21 -	This field is reserved. Reserved
22 Excess_Defer	Excessive transmission defer. This bit latches high and is cleared when read. This bit is cleared by default. 0 Normal operation. 1 A transmit frame has been excessively deferred (6,071 nibble-times in 100 Mbps modes or 24,287 bit-times in 10 Mbps modes)
23–31 -	This field is reserved. Reserved

14.52.51 MAC station address register 1 (ETSECx_MACSTNADDR1)

The MACSTNADDR1 register is written by the user. The value of the station address written into MACSTNADDR1 and MACSTNADDR2 is byte reversed from how it would appear in the DA field of a frame in memory. For example, for a station address of 0x12345678ABCD, MACSTNADDR1 is set to 0xCDAB7856 and MACSTNADDR2 is set to 0x34120000.

Address: Base address + 540h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_MACSTNADDR1 field descriptions

Field	Description
0–7 Station_Address_6th_Octet	This field holds the sixth octet of the station address. The sixth octet (station address bits 40-47) defaults to a value of 0x0.
8–15 Station_Address_5th_Octet	This field holds the fifth octet of the station address. The fifth octet (station address bits 32-39) defaults to a value of 0x0.
16–23 Station_Address_4th_Octet	This field holds the fourth octet of the station address. The fourth octet (station address bits 24-31) defaults to a value of 0x0.
24–31 Station_Address_3rd_Octet	This field holds the third octet of the station address. The third octet (station address bits 16-23) defaults to a value of 0x0.

14.52.52 MAC station address register 2 (ETSECx_MACSTNADDR2)

The MACSTNADDR2 register is written by the user.

Address: Base address + 544h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_MACSTNADDR2 field descriptions

Field	Description
0–7 Station_Address_2nd_Octet	This field holds the second octet of the station address. The second octet (station address bits 8-15) defaults to a value of 0x0.
8–15 Station_Address_1st_Octet	This field holds the first octet of the station address. The first octet (station address bits 0-7) defaults to a value of 0x0.
16–31 -	This field is reserved. Reserved

14.52.53 MAC exact match address n, part 1 (ETSECx_MACnADDR1)

The MAC01ADDR1-MAC15ADDR1 registers are written by the user with the unicast or multicast addresses aliasing the MAC. The value of the address written into MAC *x* ADDR1 and MAC *n* ADDR2 is byte reversed from how it would appear in the DA field of a frame in memory. For example, for a MAC address of 0x12345678ABCD, MACnADDR1 is set to 0xCDAB7856 and MACnADDR2 is set to 0x34120000. For any valid, non-zero MAC address received, exact match registers can be excluded individually by clearing them to all zero bytes.

Address: Base address + 548h offset + (8d × i), where i=0d to 14d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Exact_Match_Address_6th_Octet							Exact_Match_Address_5th_Octet							Exact_Match_Address_4th_Octet							Exact_Match_Address_3rd_Octet										
W	Exact_Match_Address_6th_Octet							Exact_Match_Address_5th_Octet							Exact_Match_Address_4th_Octet							Exact_Match_Address_3rd_Octet										
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_MACnADDR1 field descriptions

Field	Description
0–7 Exact_Match_Address_6th_Octet	Holds the sixth octet of the exact match address. The sixth octet (destination address bits 40-47) defaults to a value of 0x0.
8–15 Exact_Match_Address_5th_Octet	Holds the fifth octet of the exact match address. The fifth octet (destination address bits 32-39) defaults to a value of 0x0.
16–23 Exact_Match_Address_4th_Octet	Holds the fourth octet of the exact match address. The fourth octet (destination address bits 24-31) defaults to a value of 0x0.

Table continues on the next page...

ETSECx_MACnADDR1 field descriptions (continued)

Field	Description
24–31 Exact_Match_ Address_3rd_ Octet	Holds the third octet of the exact match address. The third octet (destination address bits 16-23) defaults to a value of 0x0.

14.52.54 MAC exact match address n, part 2 (ETSECx_MACnADDR2)

The MAC01ADDR2-MAC15ADDR2 registers are written by the user with the unicast or multicast addresses aliasing the MAC.

Address: Base address + 54Ch offset + (8d × i), where i=0d to 14d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Exact_Match_Address_2nd_Octet								Exact_Match_Address_1st_Octet								Reserved															
W	Exact_Match_Address_2nd_Octet								Exact_Match_Address_1st_Octet								Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_MACnADDR2 field descriptions

Field	Description
0–7 Exact_Match_ Address_2nd_ Octet	This field holds the second octet of the exact match address. The second octet (destination address bits 8-15) defaults to a value of 0x0.
8–15 Exact_Match_ Address_1st_ Octet	This field holds the first octet of the exact match address. The first octet (destination address bits 0-7) defaults to a value of 0x0.
16–31 -	This field is reserved. Reserved

14.52.55 Transmit and receive 64-byte frame counter (ETSECx_TR64)

Address: Base address + 680h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										TR64																					
W	Reserved										TR64																					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TR64 field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TR64	Transmit and receive 64-byte frame counter-Increment for each good or bad frame transmitted and received which is 64 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

14.52.56 Transmit and receive 65- to 127-byte frame counter (ETSECx_TR127)

Address: Base address + 684h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										TR127																					
W	Reserved										TR127																					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TR127 field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TR127	Transmit and receive 65- to 127-byte frame counter-Increments for each good or bad frame transmitted and received which is 65-127 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

14.52.57 Transmit and receive 128- to 255-byte frame counter (ETSECx_TR255)

Address: Base address + 688h offset

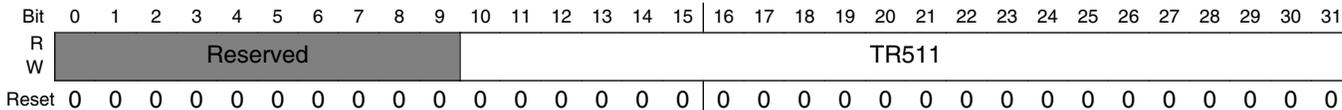
Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										TR255																					
W	Reserved										TR255																					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TR255 field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TR255	Transmit and receive 128- to 255-byte frame counter-Increments for each good or bad frame transmitted and received which is 128-255 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

14.52.58 Transmit and receive 256- to 511-byte frame counter (ETSECx_TR511)

Address: Base address + 68Ch offset

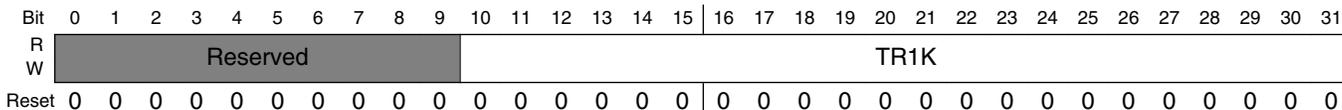


ETSECx_TR511 field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TR511	Increments for each good or bad frame transmitted and received which is 256-511 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

14.52.59 Transmit and receive 512- to 1023-byte frame counter (ETSECx_TR1K)

Address: Base address + 690h offset

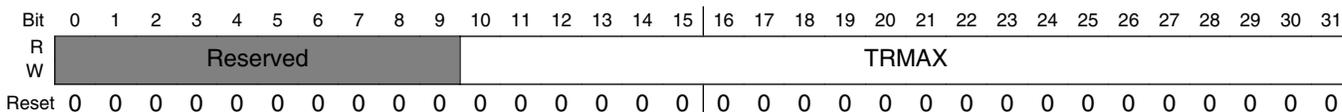


ETSECx_TR1K field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TR1K	Increments for each good or bad frame transmitted and received which is 512-1023 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

14.52.60 Transmit and receive 1024- to 1518-byte frame counter (ETSECx_TRMAX)

Address: Base address + 694h offset



ETSECx_TRMAX field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TRMAX	Increments for each good or bad frame transmitted and received which is 1024-1518 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

14.52.61 Transmit and receive 1519- to 1522-byte good VLAN frame count (ETSECx_TRMGV)

Address: Base address + 698h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TRMGV field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TRMGV	Increments for each good or bad frame transmitted and received which is 1519-1522 bytes in length, inclusive (excluding preamble and SFD but including FCS bytes).

14.52.62 Receive byte counter (ETSECx_RBYT)

Address: Base address + 69Ch offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RBYT field descriptions

Field	Description
0–31 RBYT	Receive byte counter. The statistic counter register increments by the byte count of frames received , including those in bad packets, excluding preamble and SFD but including FCS bytes.

14.52.63 Receive packet counter (ETSECx_RPKT)

Address: Base address + 6A0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																RPKT															
W	Reserved																RPKT															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RPKT field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 RPKT	Receive packet counter. Increments for each frame received packet (including bad packets, all unicast, broadcast, and multicast packets).

14.52.64 Receive FCS error counter (ETSECx_RFCS)

Address: Base address + 6A4h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																RFCS															
W	Reserved																RFCS															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RFCS field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RFCS	Receive FCS error counter. In Ethernet mode, increments for each frame received that has an integral 64-1518 length and contains a frame check sequence error.

14.52.65 Receive multicast packet counter (ETSECx_RMCA)

Address: Base address + 6A8h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																RMCA															
W	Reserved																RMCA															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RMCA field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 RMCA	Receive multicast packet counter. Increments for each multicast frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN), excluding broadcast frames. This count does not include range/length errors.

14.52.66 Receive broadcast packet counter (ETSECx_RBCA)

Address: Base address + 6ACh offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RBCA																
W	Reserved															RBCA																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RBCA field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 RBCA	Receive broadcast packet counter. Increments for each broadcast frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN), excluding multicast frames. Does not include range/length errors.

14.52.67 Receive control frame packet counter (ETSECx_RXCF)

Address: Base address + 6B0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RXCF																
W	Reserved															RXCF																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RXCF field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RXCF	Receive control frame packet counter. Increments for each MAC control frame received (PAUSE and unsupported) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

14.52.68 Receive PAUSE frame packet counter (ETSECx_RXPF)

Address: Base address + 6B4h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RXPF																
W	Reserved															RXPF																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RXPF field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RXPF	Receive PAUSE frame packet counter. Increments each time a PAUSE MAC control frame is received with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

14.52.69 Receive unknown OP code counter (ETSECx_RXUO)

Address: Base address + 6B8h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RXUO																
W	Reserved															RXUO																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RXUO field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RXUO	Receive unknown opcode counter. Increments each time a MAC control frame is received which contains an opcode other than PAUSE, but the frame has valid CRC and length 64 to 1518 (non VLAN) or 1522 (VLAN).

14.52.70 Receive alignment error counter (ETSECx_RALN)

Address: Base address + 6BCh offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RALN																
W	Reserved															RALN																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RALN field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RALN	Receive alignment error counter. Increments for each received frame from 64 to 1518 (non VLAN) or 1522 (VLAN) which contains an invalid FCS and is not an integral number of bytes.

14.52.71 Receive frame length error counter (ETSECx_RFLR)

Address: Base address + 6C0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RFLR																
W	Reserved															RFLR																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RFLR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RFLR	Receive frame length error counter. Increments for each frame received in which the 802.3 length field did not match the number of data bytes actually received (46-1500 bytes). The counter does not increment if the length field is not a valid 802.3 length, such as an Ethertype value. Frames tagged with a single VLAN tag are checked for valid length based on bytes 17-18 (rather than 13-14). Frames tagged (stacked) with multiple VLAN tags are not checked for valid length.

14.52.72 Receive code error counter (ETSECx_RCDE)

Address: Base address + 6C4h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RCDE																
W	Reserved															RCDE																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RCDE field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RCDE	Receive code error counter. Increments each time a valid carrier is present and at least one invalid data symbol is detected.

14.52.73 Receive carrier sense error counter (ETSECx_RCSE)

Address: Base address + 6C8h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RCSE																
W	Reserved															RCSE																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RCSE field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RCSE	Receive false carrier counter. Counts the number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface. The count represented by an instance of this object is incremented at most once per transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt. The event is reported along with the statistics generated on the next received frame, as defined by a 1 on TSEC <i>n</i> _RX_ER and an 0xE on TSEC <i>n</i> _RXD. Only one false carrier condition can be detected and logged between frames.

14.52.74 Receive undersize packet counter (ETSECx_RUND)

Address: Base address + 6CCh offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RUND																
W	Reserved															RUND																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RUND field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RUND	Receive undersize packet counter. Increments each time a frame is received which is less than 64 bytes in length and contains a valid FCS and were otherwise well formed. This count does not include range length errors.

14.52.75 Receive oversize packet counter (ETSECx_ROVR)

Address: Base address + 6D0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															ROVR																
W	Reserved															ROVR																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_ROVR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 ROVR	Receive oversize packet counter. Increments each time a frame is received which exceeded 1518 (non VLAN) or 1522 (VLAN) and contains a valid FCS and was otherwise well formed.

14.52.76 Receive fragments counter (ETSECx_RFRG)

Address: Base address + 6D4h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RFRG																
W	Reserved															RFRG																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RFRG field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RFRG	Receive fragments counter. Increments for each frame received which is less than 64 bytes in length and contains an invalid FCS. This includes integral and non-integral lengths.

14.52.77 Receive jabber counter (ETSECx_RJBR)

Address: Base address + 6D8h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															RJBR																
W	Reserved															RJBR																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RJBR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RJBR	Receive jabber counter. Increments for frames received which exceed 1518 (non VLAN) or 1522 (VLAN) bytes and contain an invalid FCS. This includes alignment errors.

14.52.78 Receive drop counter (ETSECx_RDRP)

Address: Base address + 6DCh offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W	Reserved															RDRP																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RDRP field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 RDRP	Receive dropped packets counter. Increments for frames received which are streamed to system but are later dropped due to lack of system resources.

14.52.79 Transmit byte counter (ETSECx_TBYT)

Address: Base address + 6E0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W	TBYT																															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TBYT field descriptions

Field	Description
0–31 TBYT	Transmit byte counter. Increments by the number of bytes that were put on the wire including fragments of frames that were involved with collisions. This count does not include preamble/SFD or jam bytes, except for half-duplex flow control (back-pressure triggered by TCTRL[THDF]=1). For THDF, the sum total of 'phantom' preamble bytes transmitted for flow control purposes is included in the TBYT increment value of the next frame to be transmitted, up to 65,535 bytes of frame and phantom preamble. Note that the value of TBYT may be greater than the actual number of bytes transmitted if the frame is truncated because it exceeds MAXFRM.

14.52.80 Transmit packet counter (ETSECx_TPKT)

Address: Base address + 6E4h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W	Reserved															TPKT																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TPKT field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TPKT	Transmit packet counter. Increments for each transmitted packet (including bad packets, excessive deferred packets, excessive collision packets, late collision packets, all unicast, broadcast, and multicast packets).

14.52.81 Transmit multicast packet counter (ETSECx_TMCA)

Address: Base address + 6E8h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										TMCA																					
W	Reserved										TMCA																					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TMCA field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TMCA	Transmit multicast packet counter. Increments for each multicast valid frame transmitted (excluding broadcast frames) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

14.52.82 Transmit broadcast packet counter (ETSECx_TBCA)

Address: Base address + 6ECh offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										TBCA																					
W	Reserved										TBCA																					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TBCA field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–31 TBCA	Transmit broadcast packet counter. Increments for each broadcast frame transmitted (excluding multicast frames) with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

14.52.83 Transmit PAUSE control frame counter (ETSECx_TXPF)

Address: Base address + 6F0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TXPF																
W	Reserved															TXPF																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TXPF field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 TXPF	Transmit PAUSE frame packet counter. Increments each time a valid PAUSE MAC control frame is transmitted with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

14.52.84 Transmit deferral packet counter (ETSECx_TDFR)

Address: Base address + 6F4h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																			TDFR												
W	Reserved																			TDFR												
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TDFR field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TDFR	Transmit deferral packet counter. Increments for each frame, which was deferred on its first transmission attempt. This count does not include frames involved in collisions.

14.52.85 Transmit excessive deferral packet counter (ETSECx_TEDF)

Address: Base address + 6F8h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																			TEDF												
W	Reserved																			TEDF												
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TEDF field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TEDF	Transmit excessive deferral packet counter. Increments for frames aborted which were deferred for an excessive period of time (3036 byte times).

14.52.86 Transmit single collision packet counter (ETSECx_TSCL)

Address: Base address + 6FCh offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TSCL																
W	Reserved															TSCL																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TSCL field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TSCL	Transmit single collision packet counter. Increments for each frame transmitted which experienced exactly one collision during transmission.

14.52.87 Transmit multiple collision packet counter (ETSECx_TMCL)

Address: Base address + 700h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TMCL																
W	Reserved															TMCL																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TMCL field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TMCL	Transmit multiple collision packet counter. Increments for each frame transmitted which experienced 2-15 collisions (including any late collisions) during transmission as defined using the Half_Duplex[RETRANSMISSION MAXIMUM] field.

14.52.88 Transmit late collision packet counter (ETSECx_TLCL)

Address: Base address + 704h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																TLCL															
W	Reserved																TLCL															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TLCL field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TLCL	Transmit late collision packet counter. Increments for each frame transmitted which experienced a late collision during a transmission attempt. Late collisions are defined using the collision window field of the half-duplex [26:31] register.

14.52.89 Transmit excessive collision packet counter (ETSECx_TXCL)

Address: Base address + 708h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																TXCL															
W	Reserved																TXCL															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TXCL field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TXCL	Transmit excessive collision packet counter. Increments for each frame that experienced 16 collisions during transmission and was aborted.

14.52.90 Transmit total collision counter (ETSECx_TNCL)

Address: Base address + 70Ch offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																TNCL															
W	Reserved																TNCL															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TNCL field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TNCL	Transmit total collision counter. Increments by the number of collisions experienced during the transmission of a frame as defined as the simultaneous presence of signals on the DO and RD circuits (That is, transmitting and receiving at the same time). Note: This count does not include collisions that result in an excessive collision condition.

14.52.91 Transmit drop frame counter (ETSECx_TDRP)

Address: Base address + 714h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TDRP																
W	Reserved															TDRP																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TDRP field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 TDRP	Transmit drop frame counter. Increments each time a memory error or an underrun has occurred.

14.52.92 Transmit jabber frame counter (ETSECx_TJBR)

Address: Base address + 718h offset

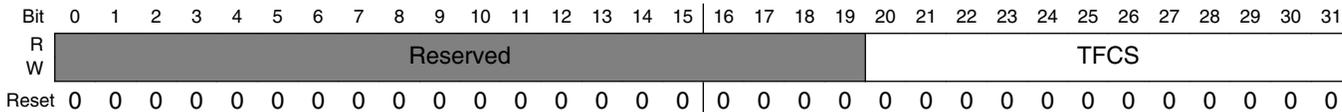
Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TJBR																
W	Reserved															TJBR																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TJBR field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TJBR	Transmit jabber frame counter. Increments for each oversized transmitted frame with an incorrect FCS value.

14.52.93 Transmit FCS error counter (ETSECx_TFCS)

Address: Base address + 71Ch offset

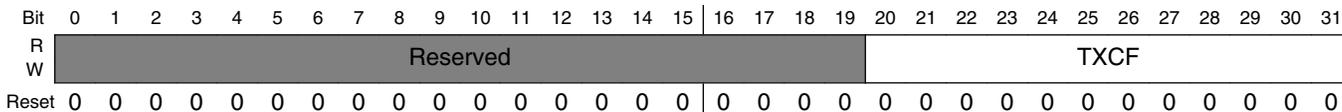


ETSECx_TFCS field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TFCS	Transmit FCS error counter. Increments for every valid sized packet with an incorrect FCS value.

14.52.94 Transmit control frame counter (ETSECx_TXCF)

Address: Base address + 720h offset

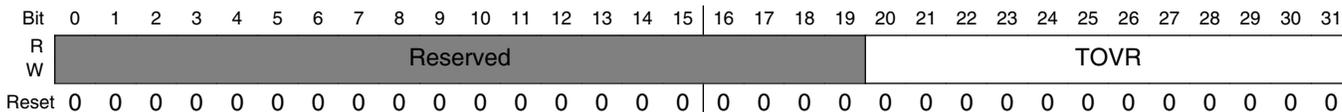


ETSECx_TXCF field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TXCF	Transmit control frame counter. Increments for every control frame with valid CRC and of lengths 64 to 1518 (non VLAN) or 1522 (VLAN).

14.52.95 Transmit oversize frame counter (ETSECx_TOVR)

Address: Base address + 724h offset



ETSECx_TOVR field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TOVR	Transmit oversize frame counter. Increments each time a frame is transmitted which exceeds 1518 bytes (non VLAN) or 1522 bytes (VLAN) with a correct FCS value.

14.52.96 Transmit undersize frame counter (ETSECx_TUND)

Address: Base address + 728h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TUND																
W	Reserved															TUND																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TUND field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TUND	Transmit undersize frame counter. Increments for every frame less then 64 bytes, with a correct FCS value.

14.52.97 Transmit fragments frame counter (ETSECx_TFRG)

Address: Base address + 72Ch offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															TFRG																
W	Reserved															TFRG																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_TFRG field descriptions

Field	Description
0–19 -	This field is reserved. Reserved
20–31 TFRG	Transmit fragment counter. Increments for every frame less then 64 bytes, with an incorrect FCS value.

14.52.98 Carry register one register Cleared on read. (ETSECx_CAR1)

Carry register bits are cleared on carry register writes when the respective bits are set.

Address: Base address + 730h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R	C164	C1127	C1255	C1511	C11K	C1MAX	C1MGV	Reserved								C1REJ	C1RBY
W	w1c	Reserved								w1c	w1c						
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	C1RPK	C1RFC	C1RMC	C1RBC	C1RXC	C1RXP	C1RXU	C1RAL	C1RFL	C1RCD	C1RCS	C1RUN	C1ROV	C1RFR	C1RJB	C1RDR	
W	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c								
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

ETSECx_CAR1 field descriptions

Field	Description
0 C164	Carry register 1 TR64 counter carry bit
1 C1127	Carry register 1 TR127 counter carry bit
2 C1255	Carry register 1 TR255 counter carry bit
3 C1511	Carry register 1 TR511 counter carry bit
4 C11K	Carry register 1 TR1K counter carry bit
5 C1MAX	Carry register 1 TRMAX counter carry bit

Table continues on the next page...

ETSECx_CAR1 field descriptions (continued)

Field	Description
6 C1MGV	Carry register 1 TRMGV counter carry bit
7–13 -	This field is reserved. Reserved
14 C1REJ	Carry register 1 RREJ counter carry bit
15 C1RBY	Carry register 1 RBYT counter carry bit
16 C1RPK	Carry register 1 RPKT counter carry bit
17 C1RFC	Carry register 1 RFCS counter carry bit
18 C1RMC	Carry register 1 RMCA counter carry bit
19 C1RBC	Carry register 1 RBCA counter carry bit
20 C1RXC	Carry register 1 RXCF counter carry bit
21 C1RXP	Carry register 1 RXPf counter carry bit
22 C1RXU	Carry register 1 RXUO counter carry bit
23 C1RAL	Carry register 1 RALN counter carry bit
24 C1RFL	Carry register 1 RFLR counter carry bit
25 C1RCD	Carry register 1 RCDE counter carry bit
26 C1RCS	Carry register 1 RCSE counter carry bit
27 C1RUN	Carry register 1 RUND counter carry bit
28 C1ROV	Carry register 1 ROVR counter carry bit
29 C1RFR	Carry register 1 RFRG counter carry bit
30 C1RJB	Carry register 1 RJBR counter carry bit
31 C1RDR	Carry register 1 RDRP counter carry bit

14.52.99 Carry register two register 3 (ETSECx_CAR2)

Carry register bits are cleared on carry register write when the respective bits are set.

Address: Base address + 734h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved											C2TJB	C2TFC	C2TCF	C2TOV	
W	Reserved											w1c	w1c	w1c	w1c	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	C2TUN	C2TFG	C2TBY	C2TPK	C2TMC	C2TBC	C2TPF	C2TDF	C2TED	C2TSC	C2TMA	C2TLC	C2TXC	C2TNC	Reserved	C2TDP
W	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c	Reserved	w1c
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_CAR2 field descriptions

Field	Description
0–11 -	This field is reserved. Reserved
12 C2TJB	Carry register 2 TJBR counter carry bit
13 C2TFC	Carry register 2 TFCS counter carry bit
14 C2TCF	Carry register 2 TXCF counter carry bit
15 C2TOV	Carry register 2 TOVR counter carry bit
16 C2TUN	Carry register 2 TUND counter carry bit
17 C2TFG	Carry register 2 TFRG counter carry bit
18 C2TBY	Carry register 2 TBYT counter carry bit
19 C2TPK	Carry register 2 TPKT counter carry bit
20 C2TMC	Carry register 2 TMCA counter carry bit
21 C2TBC	Carry register 2 TBCA counter carry bit
22 C2TPF	Carry register 2 TXPF counter carry bit
23 C2TDF	Carry register 2 TDFR counter carry bit
24 C2TED	Carry register 2 TEDF counter carry bit
25 C2TSC	Carry register 2 TSCL counter carry bit
26 C2TMA	Carry register 2 TMCL counter carry bit
27 C2TLC	Carry register 2 TLCL counter carry bit
28 C2TXC	Carry register 2 TXCL counter carry bit
29 C2TNC	Carry register 2 TNCL counter carry bit
30 -	This field is reserved. Reserved, should be cleared
31 C2TDP	Carry register 2 TDRP counter carry bit

14.52.100 Carry register one mask register (ETSECx_CAM1)

While one of the below mask bits are cleared, the corresponding carry bit in CAR1 is allowed to cause interrupt indications in register IEVENT[MSR0]. These bits all default to a set state.

Address: Base address + 738h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R									Reserved								
W	M164	M1127	M1255	M1511	M11k	M1MAX	M1MGV								M1REJ	M1RBY	
Reset	1	1	1	1	1	1	1	0	0	0	0	0	0	0	1	1	
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R																	
W	M1RPK	M1RFC	M1RMC	M1RBC	M1RXC	M1RXP	M1RXU	M1RAL	M1RFL	M1RCD	M1RCS	M1RUN	M1ROV	M1RFR	M1RJB	M1RDR	
Reset	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	

ETSECx_CAM1 field descriptions

Field	Description
0 M164	Mask register 1 TR64 counter carry bit mask
1 M1127	Mask register 1 TR127 counter carry bit mask
2 M1255	Mask register 1 TR255 counter carry bit mask
3 M1511	Mask register 1 TR511 counter carry bit mask
4 M11k	Mask register 1 TR1K counter carry bit mask
5 M1MAX	Mask register 1 TRMAX counter carry bit mask
6 M1MGV	Mask register 1 TRMGV counter carry bit mask
7–13 -	This field is reserved. Reserved
14 M1REJ	Mask register 1 RREJ counter carry bit mask
15 M1RBY	Mask register 1 RBYT counter carry bit mask

Table continues on the next page...

ETSECx_CAM1 field descriptions (continued)

Field	Description
16 M1RPK	Mask register 1 RPKT counter carry bit mask
17 M1RFC	Mask register 1 RFCS counter carry bit mask
18 M1RMC	Mask register 1 RMCA counter carry bit mask
19 M1RBC	Mask register 1 RBCA counter carry bit mask
20 M1RXC	Mask register 1 RXCF counter carry bit mask
21 M1RXP	Mask register 1 RXPF counter carry bit mask
22 M1RXU	Mask register 1 RXUO counter carry bit mask
23 M1RAL	Mask register 1 RALN counter carry bit mask
24 M1RFL	Mask register 1 RFLR counter carry bit mask
25 M1RCD	Mask register 1 RCDE counter carry bit mask
26 M1RCS	Mask register 1 RCSE counter carry bit mask
27 M1RUN	Mask register 1 RUND counter carry bit mask
28 M1ROV	Mask register 1 ROVR counter carry bit mask
29 M1RFR	Mask register 1 RFRG counter carry bit mask
30 M1RJB	Mask register 1 RJBR counter carry bit mask
31 M1RDR	Mask register 1 RDRP counter carry bit mask

14.52.101 Carry register two mask register (ETSECx_CAM2)

While one of the below mask bits are cleared, the corresponding carry bit in CAR2 is allowed to cause interrupt indications in register IEVENT[MSR0]. These bits default to a set state.

Address: Base address + 73Ch offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved											M2TJB	M2TFC	M2TCF	M2TOV	
W	Reserved											M2TJB	M2TFC	M2TCF	M2TOV	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	M2TUN	M2TFG	M2TBY	M2TPK	M2TMC	M2TBC	M2TPF	M2TDF	M2TED	M2TSC	M2TMA	M2TLC	M2TXC	M2TNC	Reserved	M2TDP
W	M2TUN	M2TFG	M2TBY	M2TPK	M2TMC	M2TBC	M2TPF	M2TDF	M2TED	M2TSC	M2TMA	M2TLC	M2TXC	M2TNC	Reserved	M2TDP
Reset	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	1

ETSECx_CAM2 field descriptions

Field	Description
0–11 -	This field is reserved. Reserved
12 M2TJB	Mask register 2 TJBR counter carry bit mask
13 M2TFC	Mask register 2 TFCS counter carry bit mask
14 M2TCF	Mask register 2 TXCF counter carry bit mask
15 M2TOV	Mask register 2 TOVR counter carry bit mask
16 M2TUN	Mask register 2 TUND counter carry bit mask

Table continues on the next page...

ETSECx_CAM2 field descriptions (continued)

Field	Description
17 M2TFG	Mask register 2 TFRG counter carry bit mask
18 M2TBY	Mask register 2 TBYT counter carry bit mask
19 M2TPK	Mask register 2 TPKT counter carry bit mask
20 M2TMC	Mask register 2 TMCA counter carry bit mask
21 M2TBC	Mask register 2 TBCA counter carry bit mask
22 M2TPF	Mask register 2 TXPF counter carry bit mask
23 M2TDF	Mask register 2 TDFR counter carry bit mask
24 M2TED	Mask register 2 TEDF counter carry bit mask
25 M2TSC	Mask register 2 TSCL counter carry bit mask
26 M2TMA	Mask register 2 TMCL counter carry bit mask
27 M2TLC	Mask register 2 TLCL counter carry bit mask
28 M2TXC	Mask register 2 TXCL counter carry bit mask
29 M2TNC	Mask register 2 TNCL counter carry bit mask
30 -	This field is reserved. Reserved
31 M2TDP	Mask register 2 TDRP counter carry bit mask

14.52.102 Receive filer rejected packet counter (ETSECx_RREJ)

Address: Base address + 740h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										RREJ																					
W	Reserved										RREJ																					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSECx_RREJ field descriptions

Field	Description
0–9 -	This field is reserved. Reserved

Table continues on the next page...

ETSECx_RREJ field descriptions (continued)

Field	Description
10–31 RREJ	Receive filter rejected packet counter. Increments for each frame with valid CRC received, but rejected by the receive queue filter—either due to a matching rule that asserted the REJ flag or due to filtering to a RxBDRing that was not enabled (see IEVENT[FIQ] error).

14.52.103 Individual/group address register n (ETSECx_IGADDRn)

The IGADDR *n* registers are written by the user. Together these registers represent, depending on RCTRL[GHTX], either the 256 entries of the individual address hash table, or the first 256 entries of the extended group address hash table used in the address recognition process. The user can enable a hash entry by setting the appropriate bit. A hash table hit occurs if the DA CRC-32 result points to an enabled hash entry.

Address: Base address + 800h offset + (4d × i), where i=0d to 7d



ETSECx_IGADDRn field descriptions

Field	Description
0–31 IGADDRn	Represents the 32-bit value associated with the corresponding register. When RCTRL[GHTX] = 0, IGADDR0 contains entries 0-31 of the 256-entry individual hash table and IGADDR7 represents entries 224-255. When RCTRL[GHTX] = 1, IGADDR0 contains entries 0-31 of the 512-entry extended group hash table and IGADDR7 represents entries 224-255.

14.52.104 Group address register n (ETSECx_GADDRn)

The GADDR *n* registers are written by the user. Together these registers represent, depending on RCTRL[GHTX], either the 256 entries of the group address hash table, or the last 256 entries of the extended group address hash table used in the address recognition process. The user can enable a hash entry by setting the appropriate bit. A hash table hit occurs if the DA CRC result points to an enabled hash entry.

Address: Base address + 880h offset + (4d × i), where i=0d to 7d



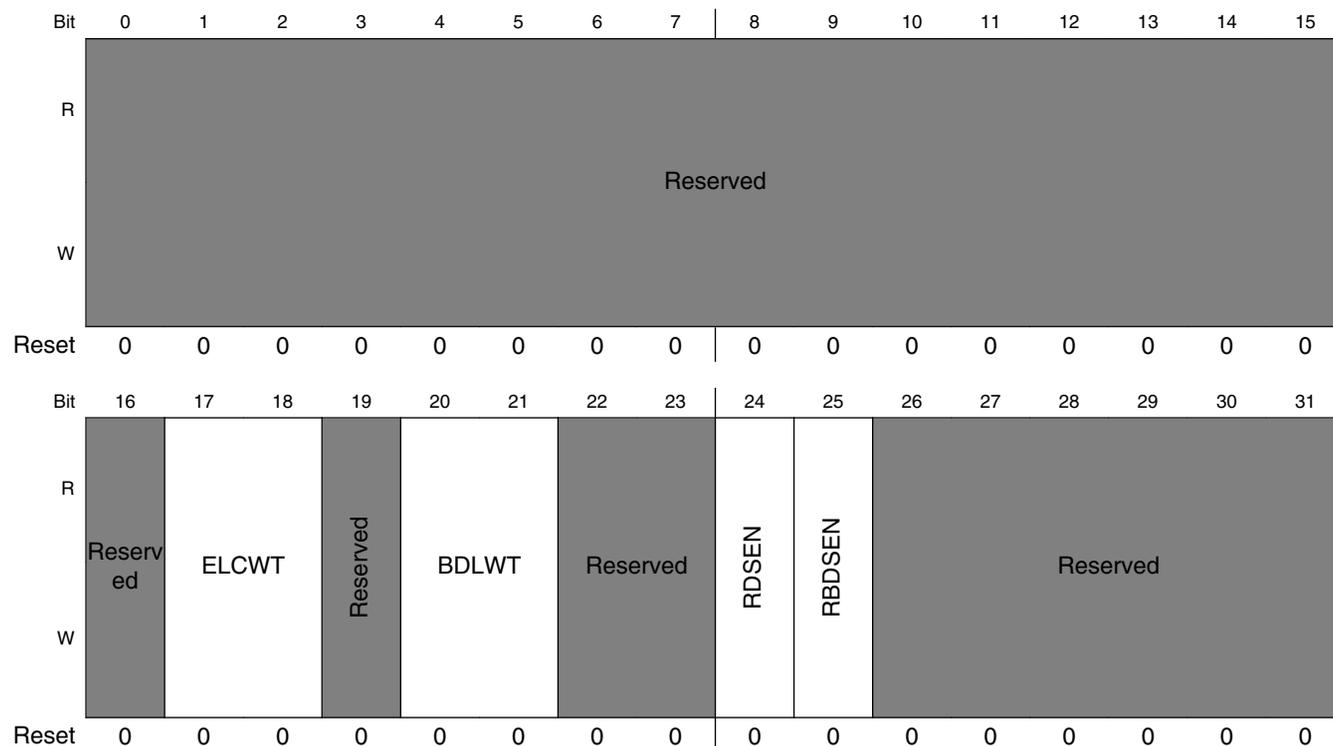
ETSECx_GADDRn field descriptions

Field	Description
0–31 GADDRn	Represents the 32-bit value associated with the corresponding register. When RCTRL[GHTX] = 0, GADDR0 contains entries 0-31 of the 256-entry group hash table and GADDR7 represents entries 224-255. When RCTRL[GHTX] = 1, GADDR0 contains entries 256-287 of the 512-entry extended group hash table and GADDR7 represents entries 480-511.

14.52.105 Attribute register (ETSECx_ATTR)

The attribute register defines memory access attributes and transaction types used to access buffer descriptors, to write receive data, and to read transmit data. Snoop enable attributes may be set for reading buffer descriptors and for reading transmit data. Buffer descriptors may be written with attributes that cause allocation into the L2 cache. Similarly, broad sections of a receive frame header may have attributes attached that cause allocation in the L2 cache. This process of specifying a region of each frame to stash into the L2 cache is referred to as *extraction*, which is specified in conjunction with register ATTRELI. ATTR[ELCWT] only has meaning if ATTRELI[EL] is non-zero. It is important to note that even though portions of received frames may be stashed to L2 cache, this is only a performance optimization as entire frames are still written to off-chip memory regardless of settings in ATTR.

Address: Base address + BF8h offset



ETSECx_ATTR field descriptions

Field	Description
0–16 -	This field is reserved. Reserved

Table continues on the next page...

ETSECx_ATTR field descriptions (continued)

Field	Description
17–18 ELCWT	<p>Extracted L2 cache write type. Specifies the write transaction type to perform for the extracted data. For maximum performance, it is recommended that if ELCWT is set to allocate, BDLWT should also be set to allocate. Writes to cache are always performed with snoop.</p> <p>00 No allocation performed. 01 Reserved 10 Allocate L2 cache line. 11 Reserved.</p>
19 -	This field is reserved. Reserved
20–21 BDLWT	<p>Buffer descriptor L2 cache write type. specifies the write transaction type to perform for the bufferdescriptor for a receive frame. Writes to cache are always performed with snoop.</p> <p>00 No allocation performed. 01 Reserved 10 Allocate L2 cache line. 11 Reserved.</p>
22–23 -	This field is reserved. Reserved
24 RDSEN	<p>Rx data snoop enable</p> <p>0 Disables snooping of all receive frames data to memory . unless ELCWT specifies L2 allocation 1 Enables snooping of all receive frames data to memory.</p>
25 RBDSEN	<p>RxBD snoop enable.</p> <p>0 Disables snooping of all receive BD memory accesses unless BDLWT specifies L2 allocation . 1 Enables snooping of all receive BD memory accesses.</p>
26–31 -	This field is reserved. Reserved

14.52.106 Attribute extract length and extract index register (ETSECx_ATTRELI)

The ATTRELI registers are written by the user to specify the extract index and extract length for extracting received frames. The extract length is typically set to the expected length of extracted packet headers.

Address: Base address + BFCh offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved		EL										-			
W	Reserved		EL										-			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved		EI										-			
W	Reserved		EI										-			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

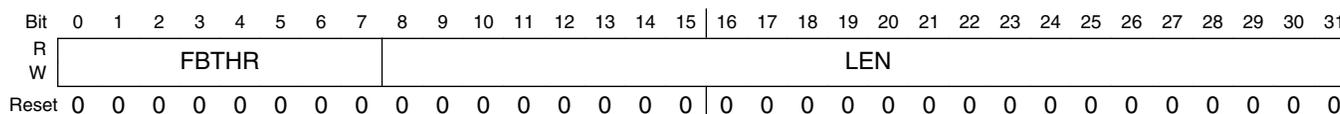
ETSECx_ATTRELI field descriptions

Field	Description
0–1 -	This field is reserved. Reserved
2–12 EL	Extracted length. Specifies the number of bytes, as a multiple of 8 bytes, to extract from the receive frame. The DMA controller uses this field to perform extraction. If cleared, no extraction is performed.
13–15 -	To ensure that EL is a multiple of 8 bytes, these bits should be written with zero.
16–17 -	This field is reserved. Reserved
18–25 EI	Extracted index. Points to the first byte, as a multiple of 64 bytes, within the receive frame as sent to memory from which to begin extracting data.
26–31 -	To ensure that EI is a multiple of 8 bytes, these bits should be written with zero.

14.52.107 Receive Queue Parameters register n (ETSECx_RQPRMn)

The RQPRMn registers specify the minimum number of BDs required to prevent flow control being asserted and the total number of Rx BDs in their respective ring. Whenever the free BD count calculated by the eTSEC for any active ring drops below the value of RQPRMn[FBTHR] for that ring, link level flow control is asserted. Software must not write to RQPRMn while LFC is enabled and the eTSEC is actively receiving frames. However, software may modify these registers after disabling LFC by clearing RCTRL[LFC]. Note that packets may be lost due to lack of RxBDs while RCTRL[LFC] is clear. Software can prevent packet loss by manually generating pause frames (through TCTRL[TFC_PAUSE]) to cover the time when RCTRL[LFC] is clear.

Address: Base address + C00h offset + (4d × i), where i=0d to 7d



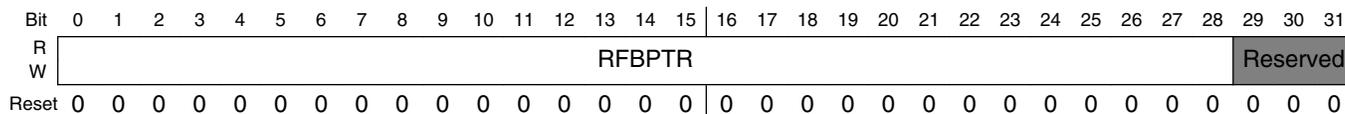
ETSECx_RQPRMn field descriptions

Field	Description
0–7 FBTHR	Free BD threshold. Minimum number of BDs required for normal operation. If the eTSEC calculated number of free BDs drops below this threshold, link layer flow control is asserted.
8–31 LEN	Ring length. Total number of Rx BDs in this ring.

14.52.108 Last Free RxBD pointer for ring n (ETSECx_RFBPTRn)

The RFBPTRn registers specify the location of the last free buffer descriptor in their respective ring. These registers live in the same 32b address space - and must share the same 4 most significant bits - as RBPTRn. That is, RFBPTRn and its associated RBPTRn must remain in the same 256MB page. Like RBPTRn, whenever RBASEn is updated, RFBPTRn is initialized to the value of RBASEn. This indicates that the ring is completely empty. As buffers are freed and their respective BDs are returned (by setting the EMPTY bit) to the ring, software is expected to update this register. The eTSEC then performs modulo arithmetic involving RBASEn, RBPTRn and RFBPTRn to determine the number of free BDs remaining in the ring. If, at any stage, the value written to RFBPTRn matches that of the respective RBPTRn the eTSEC free BD calculation assumes that the ring is now completely empty. For more information on the recommended use of these registers, see [Back pressure determination through free buffers](#) .

Address: Base address + C44h offset + (8d × i), where i=0d to 7d



ETSECx_RFBPTRn field descriptions

Field	Description
0–28 RFBPTR	Pointer to the last free BD in RxBD Ring n. When RBASEn is updated, eTSEC initializes RFBPTR n to the value in the corresponding RBASE n . Software may update this register at any time to inform the eTSEC the location of the last free BD in the ring. Note that the 3 least-significant bits of this register are read only and zero.
29–31 -	This field is reserved. Reserved.

14.5.3 ETSEC_TBI Memory Map/Register Definition

This section describes the TBI MII registers. All of the TBI registers are 16 bits wide. The TBI registers are accessed at the offset of the TBI physical address. The eTSEC's TBI physical address is stored in the TBIPA register. Writing to the TBI registers is performed in a way similar to writing to an external PHY, using the MII management interface. Using TBIPA in place of the PHY address, in the MIIMADD[PHY Address]

field, and setting the MIIMADD[Register Address] to the appropriate address offset that corresponds to the register that one wants to read or write (see the following table), the user can read (set MIIMCOM[read cycle]) or write (writing to MIIMCON[PHY control]) to the TBI block. Refer to the TBI physical address register in [Controller ID register \(ETSEC_TSEC_ID\)](#) and the TBI MII register set in the following table. Notice that jitter diagnostics and TBI control are not IEEE 802.3 required registers and are only used for test and control of the eTSEC TBI block. The TBI's TBI control register (TBI) is for configuring the eTSEC ten-bit interface block. However, because this TBI block has an MII management interface (just like any other PHY), it has an IEEE 802.3 register called the control register (CR).

ETSEC_TBI memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
0	Control (ETSEC_TBI_CR)	16	R/W	1140h	14.53.1/885
1	Status (ETSEC_TBI_SR)	16	R	0149h	14.53.2/886
4	AN advertisement (ETSEC_TBI_ANA)	16	R/W	0000h	14.53.3/887
5	AN link partner base page ability (ETSEC_TBI_ANLPBPA)	16	R	0000h	14.53.4/889
6	AN expansion (ETSEC_TBI_ANEX)	16	R	0000h	14.53.5/891
7	AN next page transmit (ETSEC_TBI_ANNPT)	16	R/W	0000h	14.53.6/891
8	AN link partner ability next page (ETSEC_TBI_ANLPANP)	16	R	0000h	14.53.7/892
F	Extended status (ETSEC_TBI_EXST)	16	R	A000h	14.53.8/893
10	Jitter diagnostics (ETSEC_TBI_JD)	16	R/W	0000h	14.53.9/894
11	TBI control (ETSEC_TBI_TBICON)	16	R/W	0010h	14.53.10/896

14.53.1 Control (ETSEC_TBI_CR)

Address: 0h base + 0h offset = 0h

Bit	0	1	2	3	4	5	6	7
Read	PHY_Reset	Reserved	Speed_0	AN_Enable	Reserved		Reset_AN	Full_Duplex
Write								
Reset	0	0	0	1	0	0	0	1
Bit	8	9	10	11	12	13	14	15
Read	Reserved	Speed_1	Reserved					
Write								
Reset	0	1	0	0	0	0	0	0

ETSEC_TBI_CR field descriptions

Field	Description												
0 PHY_Reset	PHY reset. This bit is cleared by default. This bit is self-clearing. 0 Normal operation. 1 The internal state of the TBI is reset. This in turn may change the state of the TBI link partner.												
1 -	This field is reserved. Reserved												
2 Speed_0	Speed selection. This bit defaults to a cleared state and should always be cleared, which corresponds to 1000 Mbps speed. Setting this field controls the speed at which the TBI operates. The table for Speed[1] provides the appropriate encoding. Its default is bit[2] = '0'; bit[9] = '1'.												
3 AN_Enable	Auto-negotiation enable. This bit is set by default. 0 The values programmed in bits 2, 7 and 9 determine the operating condition of the link. 1 Auto-negotiation process enabled.												
4-5 -	This field is reserved. Reserved												
6 Reset_AN	Reset auto-negotiation. This bit is cleared by default and is self-clearing. 0 Normal operation. 1 The auto-negotiation process restarts. This action is only available if auto-negotiation is enabled.												
7 Full_Duplex	Duplex mode. This bit is set by default. 0 Reserved. 1 Full-duplex operation.												
8 -	This field is reserved. Reserved, should be cleared.												
9 Speed_1	Speed selection. This bit defaults to a set state and should always be set, which corresponds to 1000 Mbps speed. Setting this field controls the speed at which the TBI operates. The following table provides the appropriate encoding. Its default is bit[2] = '0'; bit[9] = '1'.												
	<table border="1"> <thead> <tr> <th>Maximum Operating Speed</th> <th>Bit 2</th> <th>Bit 9</th> </tr> </thead> <tbody> <tr> <td>Reserved</td> <td>0</td> <td>0</td> </tr> <tr> <td>Reserved</td> <td>1</td> <td>0</td> </tr> <tr> <td>1000 Mbps</td> <td>0</td> <td>1</td> </tr> </tbody> </table>	Maximum Operating Speed	Bit 2	Bit 9	Reserved	0	0	Reserved	1	0	1000 Mbps	0	1
Maximum Operating Speed	Bit 2	Bit 9											
Reserved	0	0											
Reserved	1	0											
1000 Mbps	0	1											

Table continues on the next page...

ETSEC_TBI_CR field descriptions (continued)

Field	Description		
	Maximum Operating Speed	Bit 2	Bit 9
	Reserved	1	1
10–15 -	This field is reserved. Reserved		

14.53.2 Status (ETSEC_TBI_SR)

Address: 0h base + 1h offset = 1h

Bit	0	1	2	3	4	5	6	7
Read	Reserved							Extend_Status
Write	Reserved							Reserved
Reset	0	0	0	0	0	0	0	1
Bit	8	9	10	11	12	13	14	15
Read	Reserved	No_Pre	AN_Done	Remote_Fault	AN_Ability	Link_Status	Reserved	Extend_Ability
Write	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved
Reset	0	1	0	0	1	0	0	1

ETSEC_TBI_SR field descriptions

Field	Description
0–6 -	This field is reserved. Reserved, should be cleared.
7 Extend_Status	This bit indicates that PHY status information is also contained in the Register 15, Extended Status Register. Returns 1 on read. This bit is read-only.
8 -	This field is reserved. Reserved, should be cleared.
9 No_Pre	MF preamble suppression enable. This bit indicates whether or not the PHY is capable of handling MII management frames without the 32-bit preamble field. Returns 1, indicating support for suppressed preamble MII management frames. This bit is read-only.
10 AN_Done	Auto-negotiation complete. This bit is read-only and is cleared by default. 0 Either the auto-negotiation process is underway or the auto-negotiation function is disabled. 1 The auto-negotiation process has completed.
11 Remote_Fault	Remote fault. This bit is read-only and is cleared by default. Each read of the status register clears this bit. 0 Normal operation. 1 A remote fault condition was detected. This bit latches high in order for software to detect the condition.

Table continues on the next page...

ETSEC_TBI_SR field descriptions (continued)

Field	Description
12 AN_Ability	Auto-negotiation ability. While read as set, this bit indicates that the PHY has the ability to perform auto-negotiation. While read as cleared, this bit indicates the PHY lacks the ability to perform auto-negotiation. Returns 1 on read. This bit is read-only.
13 Link_Status	Link status. This bit is read-only and is cleared by default. 0 A valid link is not established. This bit latches low allowing for software polling to detect a failure condition. 1 A valid link is established.
14 -	This field is reserved. Reserved, should be cleared.
15 Extend_Ability	Extended capability. This bit indicates that the PHY contains the extended set of registers (those beyond control and status). Returns 1 on read. This bit is read-only.

14.53.3 AN advertisement (ETSEC_TBI_ANA)

Table 14-851. PAUSE Priority Resolution

Local Device		Link Partner		Local Resolution	Link Partner Resolution
PAUSE	ASM_DIR	PAUSE	ASM_DIR		
0	0	x	x	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
0	1	0	x	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
0	1	1	0	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
0	1	1	1	Enable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Enable PAUSE receive
1	0	0	x	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
1	0	1	x	Enable PAUSE transmit Enable PAUSE receive	Enable PAUSE transmit Enable PAUSE receive

Table continues on the next page...

Table 14-851. PAUSE Priority Resolution (continued)

Local Device		Link Partner		Local Resolution	Link Partner Resolution
PAUSE	ASM_DIR	PAUSE	ASM_DIR		
1	1	0	0	Disable PAUSE transmit Disable PAUSE receive	Disable PAUSE transmit Disable PAUSE receive
1	1	0	1	Disable PAUSE transmit Enable PAUSE receive	Enable PAUSE transmit Disable PAUSE receive
1	1	1	x	Enable PAUSE transmit Enable PAUSE receive	Enable PAUSE transmit Enable PAUSE receive

Address: 0h base + 4h offset = 4h

Bit	0	1	2	3	4	5	6	7
Read	Next_Page	Reserved	Remote_Fault		Reserved			Pause
Write								
Reset	0	0	0	0	0	0	0	0
Bit	8	9	10	11	12	13	14	15
Read	Pause	Half_Duplex	Full_Duplex	Reserved				
Write								
Reset	0	0	0	0	0	0	0	0

ETSEC_TBI_ANA field descriptions

Field	Description												
0 Next_Page	Next page configuration. The local device sets this bit to either request next page transmission or advertise next page exchange capability. 0 The local device wishes not to engage in next page exchange. 1 The local device has no next pages but wishes to allow reception of next pages. If the local device has no next pages and the link partner wishes to send next pages, the local device shall send null message codes and have the message page set to 0b000_0000_0001, as defined in annex 28C.												
1 -	This field is reserved. Reserved. (Ignore on read)												
2-3 Remote_Fault	The local device's remote fault condition is encoded in bits 2 and 3 of the base page. Values are shown in the following table. The default value is 00. Indicate a fault by setting a non-zero remote fault encoding and re-negotiating.												
	<table border="1"> <thead> <tr> <th>RF1 bit[3]</th> <th>RF2 bit[2]</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No error, link okay</td> </tr> <tr> <td>0</td> <td>1</td> <td>Offline</td> </tr> <tr> <td>1</td> <td>0</td> <td>Link_Failure</td> </tr> </tbody> </table>	RF1 bit[3]	RF2 bit[2]	Description	0	0	No error, link okay	0	1	Offline	1	0	Link_Failure
RF1 bit[3]	RF2 bit[2]	Description											
0	0	No error, link okay											
0	1	Offline											
1	0	Link_Failure											

Table continues on the next page...

ETSEC_TBI_ANA field descriptions (continued)

Field	Description		
	RF1 bit[3]	RF2 bit[2]	Description
	1	1	Auto-Negotiation_Error
4-6 -	This field is reserved. Reserved, should be cleared.		
7-8 Pause	The local device's PAUSE capability is encoded in bits 7 and 8, and the decodes are shown in the following table. For priority resolution information consult Table 14-851 .		
	PAUSE bit[8]	ASM_DIR bit[7]	Capability
	0	0	No PAUSE
	0	1	Asymmetric PAUSE toward link partner
	1	0	Symmetric PAUSE
	1	1	Both symmetric PAUSE and Asymmetric PAUSE toward local device
9 Half_Duplex	Half-duplex capability. 0 Designates local device as not capable of half-duplex operation. 1 Designates local device as capable of half-duplex operation.		
10 Full_Duplex	Full-duplex capability. 0 Designates the local device as not capable of full-duplex operation. 1 Designates the local device as capable of full-duplex operation.		
11-15 -	This field is reserved. Reserved, should be cleared.		

14.53.4 AN link partner base page ability (ETSEC_TBI_ANLPBPA)

Address: 0h base + 5h offset = 5h

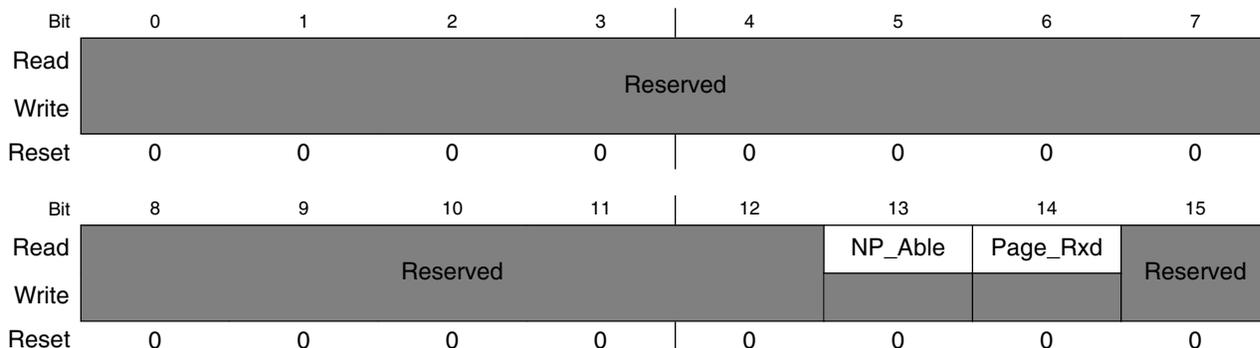
Bit	0	1	2	3	4	5	6	7
Read	Next_Page	Reserved	Remote_Fault		Reserved			Pause
Write	Reserved		Reserved		Reserved			Reserved
Reset	0	0	0	0	0	0	0	0
Bit	8	9	10	11	12	13	14	15
Read	Pause	Half_Duplex	Full_Duplex	Reserved				
Write	Reserved			Reserved				
Reset	0	0	0	0	0	0	0	0

ETSEC_TBI_ANLPBPA field descriptions

Field	Description															
0 Next_Page	<p>Next page. This bit is read-only. The link partner sets or clears this bit.</p> <p>0 Link partner has no subsequent next pages or is not capable of receiving next pages. 1 Link partner either requesting next page transmission or indicating the capability to receive next pages.</p>															
1 -	<p>This field is reserved. Reserved. (Ignore on read)</p>															
2-3 Remote_Fault	<p>The link partner's remote fault condition is encoded in bits 2 and 3 of the base page. Values are shown in the remote fault encoding field table below. This bit is read-only.</p> <table border="1"> <thead> <tr> <th>RF1 bit[3]</th> <th>RF2 bit[2]</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No error, link okay</td> </tr> <tr> <td>0</td> <td>1</td> <td>Offline</td> </tr> <tr> <td>1</td> <td>0</td> <td>Link_Failure</td> </tr> <tr> <td>1</td> <td>1</td> <td>Auto-Negotiation_Error</td> </tr> </tbody> </table>	RF1 bit[3]	RF2 bit[2]	Description	0	0	No error, link okay	0	1	Offline	1	0	Link_Failure	1	1	Auto-Negotiation_Error
RF1 bit[3]	RF2 bit[2]	Description														
0	0	No error, link okay														
0	1	Offline														
1	0	Link_Failure														
1	1	Auto-Negotiation_Error														
4-6 -	<p>This field is reserved. Reserved, should be cleared.</p>															
7-8 Pause	<p>Encoding of the link partner's PAUSE capability is shown in the PAUSE encoding table below. For priority resolution information consult. This bit is read-only.</p> <table border="1"> <thead> <tr> <th>PAUSE bit[8]</th> <th>ASM_DIR bit[7]</th> <th>Capability</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>No PAUSE</td> </tr> <tr> <td>0</td> <td>1</td> <td>Asymmetric PAUSE toward link partner</td> </tr> <tr> <td>1</td> <td>0</td> <td>Symmetric PAUSE</td> </tr> <tr> <td>1</td> <td>1</td> <td>Both symmetric PAUSE and Asymmetric PAUSE toward local device</td> </tr> </tbody> </table>	PAUSE bit[8]	ASM_DIR bit[7]	Capability	0	0	No PAUSE	0	1	Asymmetric PAUSE toward link partner	1	0	Symmetric PAUSE	1	1	Both symmetric PAUSE and Asymmetric PAUSE toward local device
PAUSE bit[8]	ASM_DIR bit[7]	Capability														
0	0	No PAUSE														
0	1	Asymmetric PAUSE toward link partner														
1	0	Symmetric PAUSE														
1	1	Both symmetric PAUSE and Asymmetric PAUSE toward local device														
9 Half_Duplex	<p>Half-duplex capability. This bit is read-only.</p> <p>0 Link partner is not capable of half-duplex mode. 1 Link partner is capable of half-duplex mode.</p>															
10 Full_Duplex	<p>Full-duplex capability. This bit is read-only.</p> <p>0 Link partner is not capable of full-duplex mode. 1 Link partner is capable of full-duplex mode.</p>															
11-15 -	<p>This field is reserved. Reserved, should be cleared.</p>															

14.53.5 AN expansion (ETSEC_TBI_ANEX)

Address: 0h base + 6h offset = 6h



ETSEC_TBI_ANEX field descriptions

Field	Description
0–12 -	This field is reserved. Reserved, should be cleared.
13 NP_Able	Next page able. This bit is read-only and returns 1 on read. While read as set, indicates local device supports next page function.
14 Page_Rxd	Page received. This bit is read-only. The bit clears on a read to the register. 0 Normal operation. 1 A new page was received and stored in the applicable AN link partner ability or AN next page register. This bit latches high in order for software to detect while polling.
15 -	This field is reserved. Reserved, should be cleared.

14.53.6 AN next page transmit (ETSEC_TBI_ANNPT)

Address: 0h base + 7h offset = 7h



ETSEC_TBI_ANNPT field descriptions

Field	Description
0 Next_Page	Next page indication. [Reference MII bit 7.15 in IEEE 802.3, 2000 Edition Clause 28.2.4] 0 Last page. 1 Additional next pages to follow.
1 -	This field is reserved. Reserved. (Ignore on read)
2 Msg_Page	Message page. [Reference MII bit 7.13] 0 Unformatted page. 1 Message page.
3 Ack2	Acknowledge 2. Used by the next page function to indicate that the device has the ability to comply with the message. [Reference MII bit 7.12] 0 The local device cannot comply with message. 1 The local device complies with message.
4 Toggle	Toggle. Used to ensure synchronization with the link partner during next page exchange. This bit always takes the opposite value of the toggle bit of the previously-exchanged link code word. The initial value in the first next page transmitted is the inverse of bit 11 in the base link code word. [Reference MII bit 7.11] This bit is read-only. 0 Toggle bit of the previously-exchanged link code word was 1. 1 Toggle bit of the previously-exchanged link code word was 0.
5–15 Message_Unformatted_Code_Field	Message pages are formatted pages that carry a pre-defined message code, which is enumerated in IEEE 802.3u/Annex 28C. Unformatted code fields take on an arbitrary value. [Reference MII field 7.10:0]

14.53.7 AN link partner ability next page (ETSEC_TBI_ANLPANP)

Address: 0h base + 8h offset = 8h

Bit	0	1	2	3	4	5	6	7
Read	Next_Page	Reserved	Msg_Page	Ack2	Toggle	Message_Unformatted_Code_Field		
Write								
Reset	0	0	0	0	0	0	0	0
Bit	8	9	10	11	12	13	14	15
Read	Message_Unformatted_Code_Field							
Write								
Reset	0	0	0	0	0	0	0	0

ETSEC_TBI_ANLPANP field descriptions

Field	Description
0 Next_Page	Next page. The link partner sets and clears this bit.

Table continues on the next page...

ETSEC_TBI_ANLPANP field descriptions (continued)

Field	Description
	0 Last page from link partner 1 Additional next pages to follow
1 -	This field is reserved. Reserved. (Ignore on read)
2 Msg_Page	Message page. 0 Unformatted page 1 Message page
3 Ack2	Acknowledge 2. Indicates the link partner's ability to comply with the message. 0 Link partner cannot comply with message. 1 Link partner complies with message.
4 Toggle	Toggle. Used to ensure synchronization with the link partner during next page exchange. This bit always takes the opposite value of the toggle bit of the previously-exchanged link code word. The initial value in the first next page transmitted is the inverse of bit 11 in the base link code word. This bit is read-only. 0 Toggle bit of the previously-exchanged link code word was 1. 1 Toggle bit of the previously-exchanged link code word was 0.
5–15 Message_Un- formatted_Code_ Field	Message pages are formatted pages that carry a pre-defined message code, which is enumerated in IEEE 802.3u/Annex 28C. Unformatted code fields take on an arbitrary value.

14.53.8 Extended status (ETSEC_TBI_EXST)

Address: 0h base + Fh offset = Fh

Bit	0	1	2	3	4	5	6	7
Read	1000X_Full	1000X_Half	1000T_Full	1000T_Half	Reserved			
Write					Reserved			
Reset	1	0	1	0	0	0	0	0
Bit	8	9	10	11	12	13	14	15
Read	Reserved							
Write	Reserved							
Reset	0	0	0	0	0	0	0	0

ETSEC_TBI_EXST field descriptions

Field	Description
0 1000X_Full	1000X full-duplex capability. Returns 1 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-X full-duplex mode. 1 PHY can operate in 1000BASE-X full-duplex mode.

Table continues on the next page...

ETSEC_TBI_EXST field descriptions (continued)

Field	Description
1 1000X_Half	1000X half-duplex capability. Returns 0 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-X half-duplex mode. 1 PHY can operate in 1000BASE-X half-duplex mode.
2 1000T_Full	1000T full-duplex capability. Returns 1 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-T full-duplex mode. 1 PHY can operate in 1000BASE-T full-duplex mode.
3 1000T_Half	1000T half-duplex capability. Returns 0 on read. This bit is read-only. 0 PHY cannot operation in 1000BASE-T half-duplex mode. 1 PHY can operate in 1000BASE-T half-duplex mode.
4–15 -	This field is reserved. Reserved

14.53.9 Jitter diagnostics (ETSEC_TBI_JD)

Annex 36A in IEEE 802.3z describes several jitter test patterns. These can be configured to be sent by writing the jitter diagnostics register. See the register description for more information. It may be wise to auto-negotiate and advertise a remote fault signaling of offline prior to beginning the test patterns.

Address: 0h base + 10h offset = 10h

Bit	0	1	2	3	4	5	6	7
Read	Jitter_Enable	Jitter_Select			Reserved		Custom_Jitter_Pattern	
Write								
Reset	0	0	0	0	0	0	0	0
Bit	8	9	10	11	12	13	14	15
Read	Custom_Jitter_Pattern							
Write								
Reset	0	0	0	0	0	0	0	0

ETSEC_TBI_JD field descriptions

Field	Description
0 Jitter_Enable	Jitter enable. This bit is cleared by default. 0 Normal transmit operation. 1 Enable the TBI to transmit the jitter test patterns defined in IEEE 802.3z 36A.
1–3 Jitter_Select	Selects the jitter pattern to be transmitted in diagnostics mode. Encoding of this field is shown in the following table. Default is 00.

Table continues on the next page...

ETSEC_TBI_JD field descriptions (continued)

Field	Description			
	Jitter Pattern Select	bit[1]	bit[2]	bit[3]
	User defined uses custom jitter pattern, bits 6-15	0	0	0
	High frequency (+/- D21.5) 101010101010101010 101010101010101010 ...	0	0	1
	Mixed frequency (+/- K28.5) 11111010110000010100 11111010110000010100 ...	0	1	0
	Low frequency 11111000001111100000 11111000001111100000 ...	0	1	1
	Complex pattern (10'h17c, 10'h0c9, 10'h0e5, 10'h2a3, 10'h17c,...)	1	0	0
	Square Wave (- K28.7) 00111110000011111000 00111110000011111000 ...	1	0	1
	Reserved	1	1	0
	Reserved	1	1	1
4-5 -	This field is reserved. Reserved			
6-15 Custom_Jitter_Pattern	Used in conjunction with jitter (pattern) select and jitter (diagnostic) enable; set this field to the desired custom pattern which is continuously transmitted. Its default is 0x000.			

14.53.10 TBI control (ETSEC_TBI_TBICON)

Address: 0h base + 11h offset = 11h

Bit	0	1	2	3	4	5	6	7
Read	Soft_Reset	Reserved	Disable_Rx_Dis	Disable_Tx_Dis	Reserved			AN_Sense
Write								
Reset	0	0	0	0	0	0	0	0
Bit	8	9	10	11	12	13	14	15
Read	Reserved		Clock_Select	MI_Mode	Reserved			
Write								
Reset	0	0	0	1	0	0	0	0

ETSEC_TBI_TBICON field descriptions

Field	Description
0 Soft_Reset	Soft reset. This bit is cleared by default. 0 Normal operation. 1 Resets the functional modules in the TBI.
1 -	This field is reserved. Reserved. (Ignore on read)
2 Disable_Rx_Dis	Disable receive disparity. This bit is cleared by default. 0 Normal operation. 1 Disables the running disparity calculation and checking in the receive direction.
3 Disable_Tx_Dis	Disable transmit disparity. This bit is cleared by default. 0 Normal operation. 1 Disables the running disparity calculation and checking in the transmit direction.
4-6 -	This field is reserved. Reserved
7 AN_Sense	Auto-negotiation sense enable. This bit is cleared by default. 0 IEEE 802.3z Clause 37 behavior is desired, which results in the link not completing. 1 Allow the auto-negotiation function to sense either a Gigabit MAC in auto-negotiation bypass mode or an older Gigabit MAC without auto-negotiation capability. If sensed, auto-negotiation complete becomes true; however, the page received is low, indicating no page was exchanged. Management can then act accordingly.
8-9 -	This field is reserved. Reserved
10 Clock_Select	Clock select. This bit selects how the on-chip TBI PHY is clocked. This bit is cleared by default.

Table continues on the next page...

ETSEC_TBI_TBICON field descriptions (continued)

Field	Description
	<p>0 The TBI PHY is clocked by dual split-phase 62.5 MHz receive clocks. These external signals must be provided via TBI receive clock 0 (TSEC <i>n</i>_RX_CLK) and TBI receive clock 1 (TSEC <i>n</i>_TX_CLK). If operating in SGMII mode, clearing this bit effectively disables the TBI PHY clock.</p> <p>1 The TBI PHY is clocked by a single 125 MHz receive clock (required for SGMII operation) . This single clock , if operating in a non-SGMII (parallel) mode, must be provided via the TBI receive clock 0 (TSEC <i>n</i>_RX_CLK) external signal. If operating in SGMII mode, this clock is provided on-chip by the SerDes block.</p>
11 MI_Mode	<p>This bit describes the configuration mode of the TBI. The user reads a 1 while the TBI is configured in GMII/MII mode (connected to a GMII/MII PHY) and a 0 while configured in TBI mode (connected to a 1000BASE-X SerDes). Its value is the inverse of ECNTRL[TBIM].</p> <p>0 TBI mode. 1 GMII mode.</p>
12–15 -	<p>This field is reserved. Reserved</p>

14.5.4 ETSEC_1588 Memory Map/Register Definition

ETSEC_1588 memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4E00	Timer control register (ETSEC_1588_TMR_CTRL)	32	R/W	0001_0000h	14.54.1/898
2_4E04	Time stamp event register (ETSEC_1588_TMR_TEVENT)	32	w1c	0000_0000h	14.54.2/901
2_4E08	Timer event mask register (ETSEC_1588_TMR_TEMASK)	32	R/W	0000_0000h	14.54.3/903
2_4E0C	Time stamp event register (ETSEC_1588_TMR_PEVENT)	32	R/W	0000_0000h	14.54.4/904
2_4E10	Timer event mask register (ETSEC_1588_TMR_PEMASK)	32	R/W	0000_0000h	14.54.5/905
2_4E14	Time stamp status register (ETSEC_1588_TMR_STAT)	32	R/W	0000_0000h	14.54.6/906
2_4E18	Timer counter high register (ETSEC_1588_TMR_CNT_H)	32	R/W	0000_0000h	14.54.7/906
2_4E1C	Timer counter low register (ETSEC_1588_TMR_CNT_L)	32	R/W	0000_0000h	14.54.8/907
2_4E20	Timer drift compensation addend register (ETSEC_1588_TMR_ADD)	32	R/W	0000_0000h	14.54.9/908
2_4E24	Timer accumulator register (ETSEC_1588_TMR_ACC)	32	R/W	0000_0000h	14.54.10/909
2_4E28	Timer prescale (ETSEC_1588_TMR_PRSC)	32	R/W	0000_0002h	14.54.11/909

Table continues on the next page...

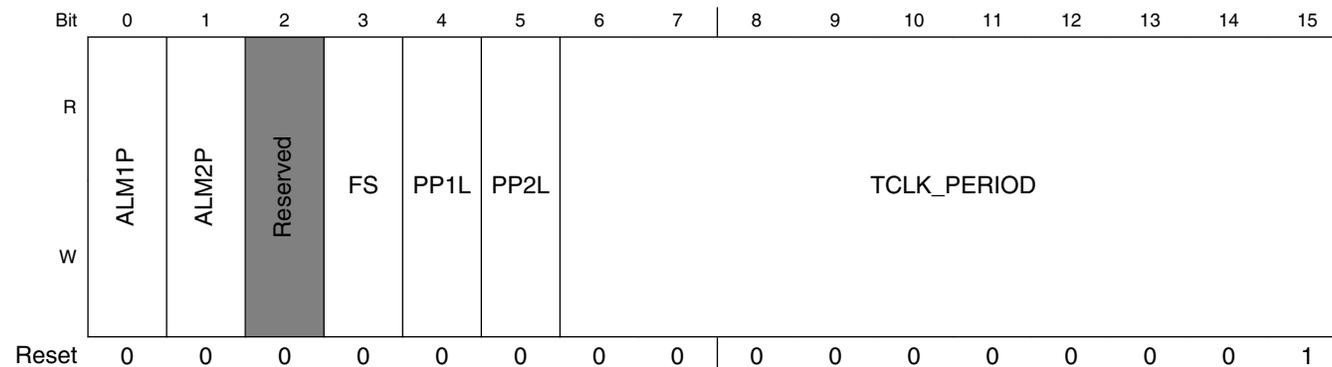
ETSEC_1588 memory map (continued)

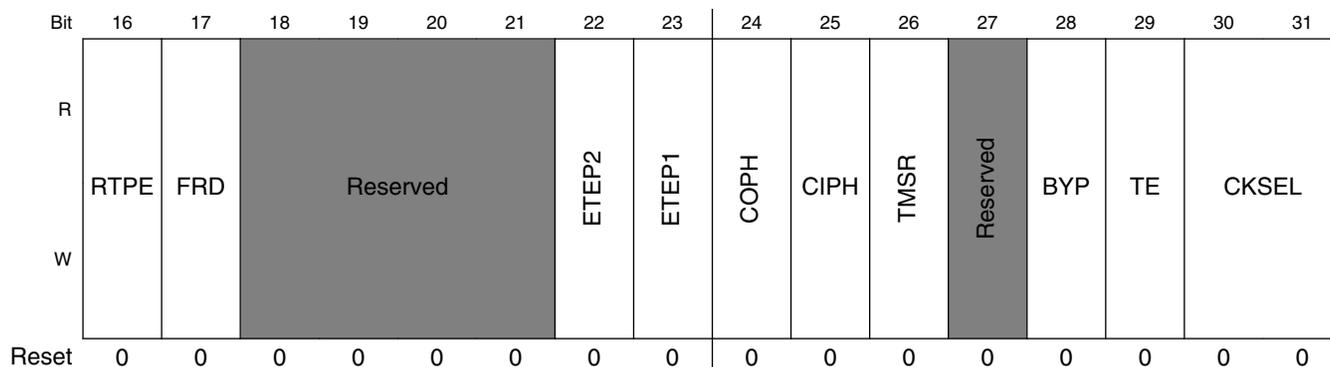
Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_4E30	Timer offset high (ETSEC_1588_TMROFF_H)	32	R/W	0000_0000h	14.54.12/910
2_4E34	Timer offset low (ETSEC_1588_TMROFF_L)	32	R/W	0000_0000h	14.54.13/910
2_4E40	Timer alarm n high register (ETSEC_1588_TMR_ALARM1_H)	32	R/W	FFFF_FFFFh	14.54.14/911
2_4E44	Timer alarm n low register (ETSEC_1588_TMR_ALARM1_L)	32	R/W	FFFF_FFFFh	14.54.15/911
2_4E48	Timer alarm n high register (ETSEC_1588_TMR_ALARM2_H)	32	R/W	FFFF_FFFFh	14.54.14/911
2_4E4C	Timer alarm n low register (ETSEC_1588_TMR_ALARM2_L)	32	R/W	FFFF_FFFFh	14.54.15/911
2_4E80	Timer fixed period interval (ETSEC_1588_TMR_FIPER1)	32	R/W	FFFF_FFFFh	14.54.16/912
2_4E84	Timer fixed period interval (ETSEC_1588_TMR_FIPER2)	32	R/W	FFFF_FFFFh	14.54.16/912
2_4EA0	Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTS1_H)	32	R/W	0000_0000h	14.54.17/913
2_4EA4	Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTS1_L)	32	R/W	0000_0000h	14.54.18/914
2_4EA8	Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTS2_H)	32	R/W	0000_0000h	14.54.17/913
2_4EAC	Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTS2_L)	32	R/W	0000_0000h	14.54.18/914

14.54.1 Timer control register (ETSEC_1588_TMR_CTRL)

This register is used to reset, configure, and initialize the eTSEC precision timer clock. The control of all timer function is performed via programming eTSEC1. The register in eTSEC1 is shared for all eTSECs.

Address: 0h base + 2_4E00h offset = 2_4E00h





ETSEC_1588_TMR_CTRL field descriptions

Field	Description
0 ALM1P	Alarm1 output polarity 0 active high output 1 active low output
1 ALM2P	Alarm2 output polarity 0 active high output 1 active low output
2 -	This field is reserved. Reserved
3 FS	FIPER start indication 0 Fiper is enabled through timer enable. 1 Fiper is enabled through timer enable and alarm indication.
4 PP1L	Fiper1 pulse loopback mode enabled. 0 Trigger1 input is based upon normal external trigger input. 1 Fiper1 pulse is looped back into Trigger1 input.
5 PP2L	Fiper2 pulse loopback mode enabled. 0 Trigger2 input is based upon normal external trigger input. 1 Fiper2 pulse is looped back into Trigger2 input.
6–15 TCLK_PERIOD	1588 timer reference clock period. The timer clock counter will increment by TCLK_PERIOD every time the accumulator register overflows. This clock period must be larger than the clock period of the timer reference clock. For applications where user does not want the clock period to be added, they can program this field to 1 to count the clock ticks. This field defaulted to 1 to count overflow ticks. For nanosecond granularity on 1588 timer counter rate, the TCLK_PERIOD should be calculated using the following equation: $TCLK_PERIOD = 10^9 / Nominal_Frequency$
16 RTPE	Record Tx Time-Stamp to PAL Enable. When set, and FCB[PTP] is set, the 8-byte time-stamp for the packet is written to the PAL located in external memory location at an offset of 16 bytes from the start of the Data Buffer Pointer of the first TxBD. For guidelines on using the RTPE bit, refer to Time-stamp insertion on transmit packets .
17 FRD	FIPER Realignment Disable

Table continues on the next page...

ETSEC_1588_TMR_CTRL field descriptions (continued)

Field	Description
	0 Fiper Realignment is enabled. 1 Fiper Realignment is disabled.
18–21 -	This field is reserved. Reserved
22 ETEP2	External trigger 2 edge polarity. The external trigger input signal is TSEC_1588_TRIG_IN2 (see " Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTSn_H) "). 0 Time stamp on the rising edge of the external trigger 1 Time stamp on the falling edge of the external trigger
23 ETEP1	External trigger 1 edge polarity. The external trigger input signal is TSEC_1588_TRIG_IN1(see " Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTSn_H) ") 0 time stamp on the rising edge of the external trigger 1 time stamp on the falling edge of the external trigger
24 COPH	Generated clock (TSEC_1588_CLK_OUT) output phase. 0 non-inverted divided clock is output 1 inverted divided clock is output
25 CIPH	Oscillator input clock phase. 0 non-inverted timer input clock 1 invertedtimer input clock (NOTE: This setting is reserved if CKSEL=01.)
26 TMSR	Timer soft reset. When enabled, it resets all the timer registers and state machines. 0 normal operation 1 place entire timer in reset except control and config registers NOTE: Prior to initiating timer reset (setting TMSR), must gracefully stop receiver (See MACCFG1[RX_EN] description). User programmable registers are not reset by the soft reset, for example, TMR_CTRL, TMR_TEMASK, TMR_PEMASK, TMR_ADD, TMR_PRSC, TMROFF_H/L, TMR_ALARMn, and TMR_FIPERn.
27 -	This field is reserved. Reserved
28 BYP	Bypass drift compensated clock 0 64-bit clock counter is incremented on the accumulator overflow 1 64-bit clock counter is directly driven from the external oscillator ignoring accumulator overflow
29 TE	1588 timer enable. If not enabled, all the timer registers and state machines are disabled. 0 timer not enabled 1 timer enabled and resume normal operation
30–31 CKSEL	1588 Timer reference clock source select. Note that the 1588 reference clock must be no slower than 1/7 the Rx_clk frequency. The default clock select is eTSEC system clock, which is always active when eTSEC is enabled. The user must ensure the corresponding clock source is active before changing the 1588 refclk selection to external reference, RTC, or TX clock. Selecting an inactive 1588 reference clock may cause boundedly undefined behavior in the Ethernet controller and on accesses to the 1588 registers. 00 External high precision timer reference clock (TSEC_1588_CLK_IN)

Table continues on the next page...

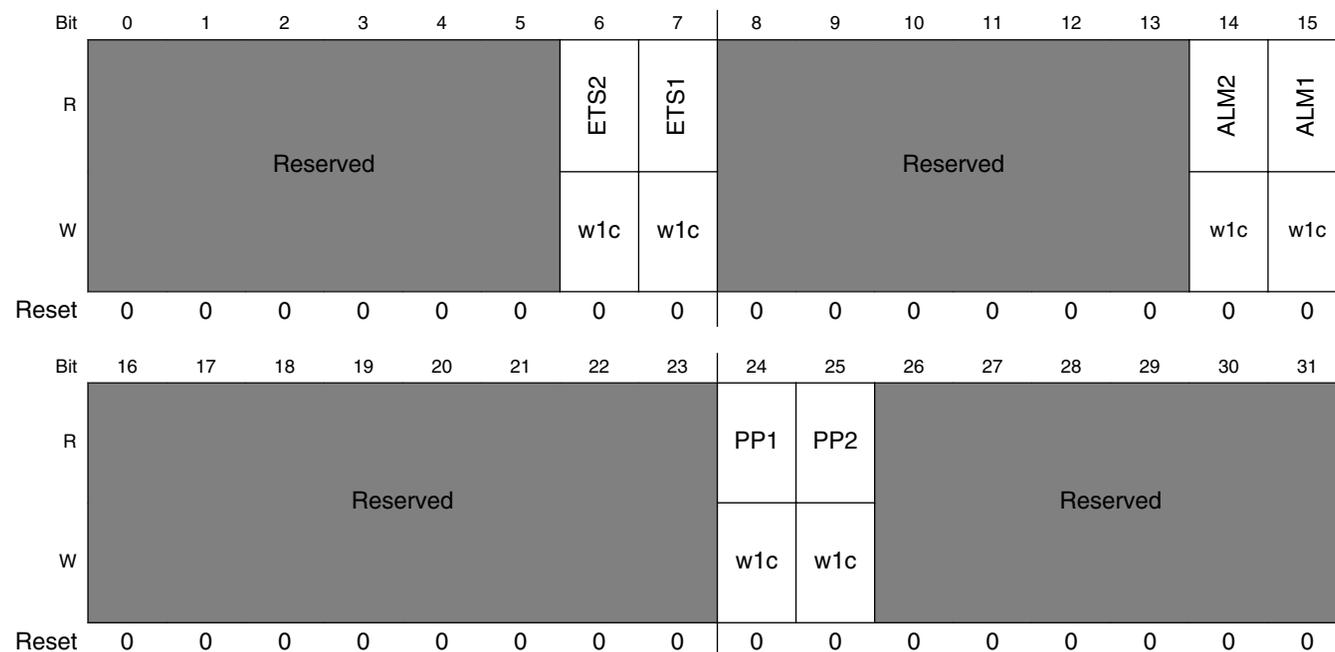
ETSEC_1588_TMR_CTRL field descriptions (continued)

Field	Description
01	eTSEC system clock
10	eTSEC1 transmit clock
11	RTC clock input

14.54.2 Time stamp event register (ETSEC_1588_TMR_TEVENT)

The eTSEC precision timer implementation can generate additional interrupts that are independent of the frame based events that controlled via IEVENT. The timer interrupts are not affected by any interrupt coalescing that may be specified in TXIC/RXIC. Software may poll this register at any time to check for pending interrupts. If an event occurs and its corresponding enable bit is set in the event mask register (TEMASK), the event also causes a hardware interrupt at the PIC. A bit in the timer event register is cleared by writing a 1 to that bit position.

Address: 0h base + 2_4E04h offset = 2_4E04h



ETSEC_1588_TMR_TEVENT field descriptions

Field	Description
0-5 -	This field is reserved. Reserved

Table continues on the next page...

ETSEC_1588_TMR_TEVENT field descriptions (continued)

Field	Description
6 ETS2	External trigger 2 timestamp sampled 0 external trigger timestamp not sampled 1 external trigger timestamp sampled
7 ETS1	External trigger 1 timestamp sampled 0 external trigger timestamp not sampled 1 external trigger timestamp sampled
8–13 -	This field is reserved. Reserved
14 ALM2	Current time equaled alarm time register 2 0 alarm time has not be reached yet 1 alarm time has been reached
15 ALM1	Current time equaled alarm time register 1 0 alarm time has not be reached yet 1 alarm time has been reached
16–23 -	This field is reserved. Reserved
24 PP1	Indicates that a periodic pulse has been generated based on FIPER1 register. 0 periodic pulse not generated 1 periodic pulse generated
25 PP2	Indicates that a periodic pulse has been generated based on FIPER2 register. 0 periodic pulse not generated 1 periodic pulse generated
26–31 -	This field is reserved. Reserved

14.54.3 Timer event mask register (ETSEC_1588_TMR_TEMASK)

Timer event mask register. The event mask register provides control over which possible interrupt events in the TMR_TEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset.

Address: 0h base + 2_4E08h offset = 2_4E08h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved						ETS2EN	ETS1EN	Reserved						ALM2EN	ALM1EN
W	Reserved						ETS2EN	ETS1EN	Reserved						ALM2EN	ALM1EN
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								PP1EN	PP2EN	Reserved					
W	Reserved								PP1EN	PP2EN	Reserved					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSEC_1588_TMR_TEMASK field descriptions

Field	Description
0–5 -	This field is reserved. Reserved
6 ETS2EN	External trigger 2 timestamp sample event enable
7 ETS1EN	External trigger 1 timestamp sample event enable
8–13 -	This field is reserved. Reserved
14 ALM2EN	Timer ALM1 event enable
15 ALM1EN	Timer ALM2 event enable
16–23 -	This field is reserved. Reserved
24 PP1EN	Periodic pulse event 1 enable
25 PP2EN	Periodic pulse event 2 enable
26–31 -	This field is reserved. Reserved

14.54.4 Time stamp event register (ETSEC_1588_TMR_PEVENT)

The eTSEC precision timer logic can generate interrupts upon the capture of a timestamp due to either transmission or reception of a frame. If an event occurs and its corresponding enable bit is set in the event mask register (PEMASK), the event also causes a hardware interrupt at the PIC. A bit in the timer event register is cleared by writing a 1 to that bit position.

Address: 0h base + 2_4E0Ch offset = 2_4E0Ch

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved																	
W	Reserved																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved						TXP2	TXP1		Reserved						RXP		
W	Reserved						TXP2	TXP1		Reserved						RXP		
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0

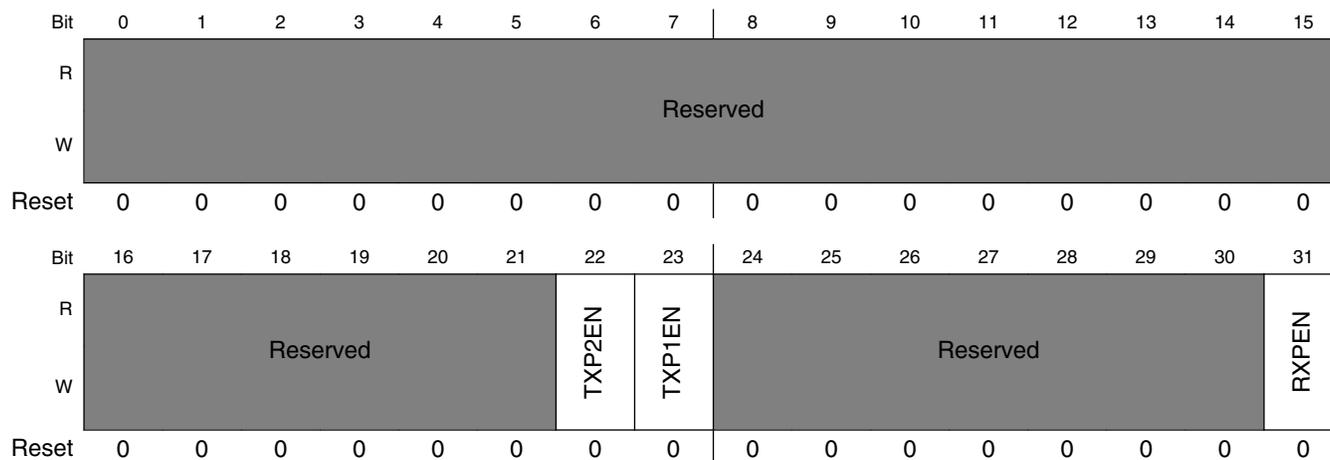
ETSEC_1588_TMR_PEVENT field descriptions

Field	Description
0–21 -	This field is reserved. Reserved
22 TXP2	Indicates that a PTP frame has been transmitted and its timestamp is stored in TXTS2 register. 0 PTP packet not transmitted 1 PTP packet has been transmitted
23 TXP1	Indicates that a PTP frame has been transmitted and its timestamp is stored in TXTS1 register. 0 PTP packet not transmitted 1 PTP packet has been transmitted
24–30 -	This field is reserved. Reserved
31 RXP	Indicates that a PTP frame has been received 0 PTP packet not received 1 PTP packet has been received

14.54.5 Timer event mask register (ETSEC_1588_TMR_PEMASK)

Timer event mask register. The event mask register provides control over which possible interrupt events in the TMR_PEVENT register are permitted to participate in generating hardware interrupts to the PIC. All implemented bits in this register are R/W and cleared upon a hardware reset.

Address: 0h base + 2_4E10h offset = 2_4E10h



ETSEC_1588_TMR_PEMASK field descriptions

Field	Description
0–21 -	This field is reserved. Reserved
22 TXP2EN	Transmit PTP packet event 2 enable
23 TXP1EN	Transmit PTP packet event 1 enable
24–30 -	This field is reserved. Reserved
31 RXPEN	Receive PTP packet event enable

14.54.6 Time stamp status register (ETSEC_1588_TMR_STAT)

This register requires the eTSEC filer to be enabled (via RCTRL[FILREN]). When eTSEC generates an interrupt based on the timestamp event for a received packet, the queue ID which the incoming packet will be sent to is captured in this register. This register update is synchronized with the RXF interrupt of the corresponding received packet. Writing 1 to any bit of this register clears it.

Address: 0h base + 2_4E14h offset = 2_4E14h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																STAT_VEC															
W	Reserved																STAT_VEC															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSEC_1588_TMR_STAT field descriptions

Field	Description
0–25 -	This field is reserved. Reserved
26–31 STAT_VEC	Timer general purpose status vector. It will store the 6-bit queue number generated by the filer. User to decode this status vector. For example, user can encode received PTP packet message types (Sync, Delay_req, Follow_up, Delay_resp, Management) in the filer virtual queue field.

14.54.7 Timer counter high register (ETSEC_1588_TMR_CNT_H)

The timer counter registers (TMR_CNT_H/L) represent accurate time in terms clock ticks or in nano-seconds. Writes to these registers will override the previous time. The register in eTSEC1 is shared for all eTSECs. These are read/write registers.

Address: 0h base + 2_4E18h offset = 2_4E18h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	TMR_CNT_H																															
W	TMR_CNT_H																															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ETSEC_1588_TMR_CNT_H field descriptions

Field	Description
0–31 TMR_CNT_H	Value of the current time counter, upper half. Current time is calculated by adding TMROFF_H/L with the TMR_CNT_H/L counter. This register can be written through the register writes. Writes to the TMR_CNT_L register copies the written value into the shadow TMR_CNT_L register. Writes to the TMR_CNT_H register copies the values written into the

ETSEC_1588_TMR_CNT_H field descriptions (continued)

Field	Description
	<p>shadow TMR_CNT_H register. Contents of the shadow registers are copied into the TMR_CNT_L and TMR_CNT_H registers following a write into the TMR_CNT_H register. Writes to these registers have precedence over the timer increment. The user must write to TMR_CNT_L register first.</p> <p>Reads from the TMR_CNT_L register copy the entire 64-bit clock time of the read enable into the TMR_CNT_H/L shadow registers. Read instruction from the TMR_CNT_H register reads the value stored in the TMR_CNT_H shadow register. The user must read the TMR_CNT_L register first to get correct 64-bit TMR_CNT_H/L counter values.</p>

14.54.8 Timer counter low register (ETSEC_1588_TMR_CNT_L)

The timer counter registers (TMR_CNT_H/L) represent accurate time in terms clock ticks or in nano-seconds. Writes to these registers will override the previous time. The register in eTSEC1 is shared for all eTSECs. These are read/write registers.

Address: 0h base + 2_4E1Ch offset = 2_4E1Ch



ETSEC_1588_TMR_CNT_L field descriptions

Field	Description
0-31 TMR_CNT_L	<p>Value of the current time counter, lower half.</p> <p>Current time is calculated by adding TMROFF_H/L with the TMR_CNT_H/L counter. This register can be written through the register writes. Writes to the TMR_CNT_L register copies the written value into the shadow TMR_CNT_L register. Writes to the TMR_CNT_H register copies the values written into the shadow TMR_CNT_H register. Contents of the shadow registers are copied into the TMR_CNT_L and TMR_CNT_H registers following a write into the TMR_CNT_H register. Writes to these registers have precedence over the timer increment. The user must write to TMR_CNT_L register first.</p> <p>Reads from the TMR_CNT_L register copy the entire 64-bit clock time of the read enable into the TMR_CNT_H/L shadow registers. Read instruction from the TMR_CNT_H register reads the value stored in the TMR_CNT_H shadow register. The user must read the TMR_CNT_L register first to get correct 64-bit TMR_CNT_H/L counter values.</p>

14.54.9 Timer drift compensation addend register (ETSEC_1588_TMR_ADD)

Timer drift compensation addend register (TMR_ADD) is used to hold timer frequency compensation value (FreqCompensationValue). The nominal frequency of the clock counter is determined by the FreqDivRatio and the clock frequency (FreqClock). This register is programmed with $2^{32} / \text{FreqDivRatio}$. Frequency division ratio (FreqDivRatio) is the ratio between the frequency of the oscillator (TimerOsc) and the desired clock frequency (NominalFreq). FreqDivRatio is a design constant chosen to be greater than 1.0001. The ADDEND value is added to the 32-bit accumulator register at every rising edge of the oscillator clock (TimerOsc). The clock counter is incremented at every carry pulse of the accumulator. Only one of this register is required for the entire group of eTSECs.

Address: 0h base + 2_4E20h offset = 2_4E20h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

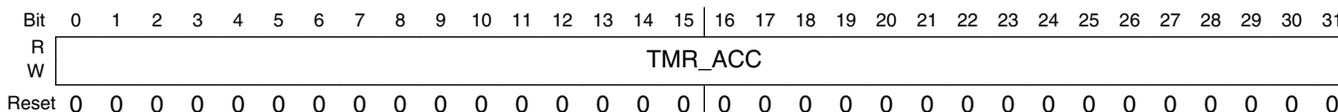
ETSEC_1588_TMR_ADD field descriptions

Field	Description
0–31 ADDEND	Timer drift compensation addend register value. It is programmed with a value of $2^{32} / \text{FreqDivRatio}$. For example, TimerOsc = 50 MHz NominalFreq = 40 MHz FreqDivRatio = 1.25 ADDEND = $\text{ceil}(2^{32} / 1.25) = 0xCCCC_CCCD$

14.54.10 Timer accumulator register (ETSEC_1588_TMR_ACC)

Timer accumulator register accumulates the value of the addend register into it. An overflow pulse of the accumulator is used to increment the timer clock by TMR_CTRL[TCLK_PERIOD]. This register is read only in normal operation. The register in eTSEC1 is shared for all eTSECs.

Address: 0h base + 2_4E24h offset = 2_4E24h



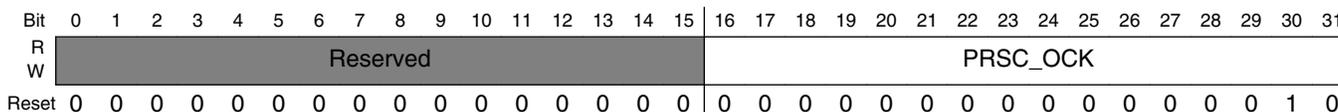
ETSEC_1588_TMR_ACC field descriptions

Field	Description
0–31 TMR_ACC	32-bit timer accumulator register

14.54.11 Timer prescale (ETSEC_1588_TMR_PRSC)

Timer generated output clock prescale register. It is used to adjust output clock frequency that is put onto the 1588 clock output signal . The register in eTSEC1 is shared for all eTSECs.

Address: 0h base + 2_4E28h offset = 2_4E28h



ETSEC_1588_TMR_PRSC field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 PRSC_OCK	Output clock division/prescale factor. Output clock is generated by dividing the timer input clock by this number. Programmed value in this field must be greater than 1. Any value less than 1 is treated as 2.

14.54.12 Timer offset high (ETSEC_1588_TMROFF_H)

The timer offset registers are used to provide current time by adding their value to the clock counter.

Address: 0h base + 2_4E30h offset = 2_4E30h



ETSEC_1588_TMROFF_H field descriptions

Field	Description
0–31 TMROFF_H	Offset value of the clock counter, upper half. Current time in is calculated by adding TMROFF_H/L with the timer's counter TMR_CNT_H/L register.

14.54.13 Timer offset low (ETSEC_1588_TMROFF_L)

The timer offset registers are used to provide current time by adding their value to the clock counter.

Address: 0h base + 2_4E34h offset = 2_4E34h



ETSEC_1588_TMROFF_L field descriptions

Field	Description
0–31 TMROFF_L	Offset value of the clock counter, lower half. Current time in is calculated by adding TMROFF_H/L with the timer's counter TMR_CNT_H/L register.

14.54.14 Timer alarm n high register (ETSEC_1588_TMR_ALARMn_H)

Alarm time comparator registers (TMR_ALARM n _H/L). These registers hold alarm time for comparison with the current counter value plus the offset registers. There are two sets of these registers for eTSEC1 which are shared among all eTSECs.

Address: 0h base + 2_4E40h offset + (8d × i), where i=0d to 1d



ETSEC_1588_TMR_ALARMn_H field descriptions

Field	Description
0–31 ALARM_H	<p>Alarm time comparator register, upper half.</p> <p>The corresponding alarm event in TMR_TEVENT is set when the current time counter becomes equal to or greater than the alarm time compare value in TMR_ALARM n _L/H. Writing the TMR_ALARM n _L register deactivates the alarm event after it has fired. Writing the TMR_ALARM n _L followed by the TMR_ALARM n _H register rearms the alarm function with the new compare value.</p> <p>The value programmed in this register must be an integer multiple of TMR_CTRL[TCLK_PERIOD] in order to get correct result. This register is reset to all ones to avoid false alarm after reset.</p> <p>In FS (Fiper start) mode the alarm trigger is used as an indication to the fiper start down counting. Only alarm 1 supports this mode. In FS mode, alarm polarity bit should be configured to 0 (rising edge).</p>

14.54.15 Timer alarm n low register (ETSEC_1588_TMR_ALARMn_L)

Alarm time comparator registers (TMR_ALARM n _H/L). These registers hold alarm time for comparison with the current counter value plus the offset registers. There are two sets of these registers for eTSEC1 which are shared among all eTSECs.

Address: 0h base + 2_4E44h offset + (8d × i), where i=0d to 1d



ETSEC_1588_TMR_ALARMn_L field descriptions

Field	Description
0–31 ALARM_L	<p>Alarm time comparator register, lower half.</p> <p>The corresponding alarm event in TMR_TEVENT is set when the current time counter becomes equal to or greater than the alarm time compare value in TMR_ALARM <i>n</i> _L/H. Writing the TMR_ALARM <i>n</i> _L register deactivates the alarm event after it has fired. Writing the TMR_ALARM <i>n</i> _L followed by the TMR_ALARM <i>n</i> _H register rearms the alarm function with the new compare value.</p> <p>The value programmed in this register must be an integer multiple of TMR_CTRL[TCLK_PERIOD] in order to get correct result. This register is reset to all ones to avoid false alarm after reset.</p> <p>In FS (Fiper start) mode the alarm trigger is used as an indication to the fiper start down counting. Only alarm 1 supports this mode. In FS mode, alarm polarity bit should be configured to 0 (rising edge).</p>

14.54.16 Timer fixed period interval (ETSEC_1588_TMR_FIPERn)

Timer fixed interval period pulse generator register. It is used to generate periodic pulses. This register is reset with 0xFFFF_FFFF to prevent any false pulse upon initialization. The down count register loads the value programmed in the fixed period interval (FIPER). FIPER register must be programmed before the timer is enabled. At every tick of the timer accumulator overflow, the counter decrements by the value of TMR_CTRL[TCLK_PERIOD]. It generates a pulse when the down counter value reaches zero. It reloads the down counter in the cycle following a pulse.

Should a user wish to use the TMR_FIPER1 register to generate a 1 PPS event, the following setup should be used:

- Program TMR_FIPER1 to a value that will generate a pulse every second,
- Program TMR_ALARM1 to the correct time for the first PPS event
- Enable the timer

The eTSEC will then wait for TMR_ALARM1 to expire before enabling the count down of TMR_FIPER1. The end result will be that TMR_FIPER1 will pulse every second after the original timer ALARM1 expired.

In the case where the PPS signals are required to be phased aligned to the prescale output clock, the alarm value should be configured to 1 clock period less than the wanted value.

In order to keep tracking the prescale output clock, each time before enabling the FIPER, the user must reset the FIPER by writing a new value to the register. The ratio between the prescale register value and the FIPER value should be devisable by the clk period.

$$\text{FIPER_VALUE} = (\text{prescale_value} \times \text{tclk_per} \times N) - \text{tclk_per}$$

For example:

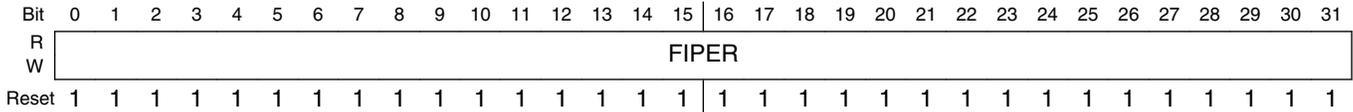
$$\text{prescale} = 9$$

clock period = 10

The FIPER can get the following values: 80, 170, 260

The two registers in eTSEC1 are shared for all eTSECs.

Address: 0h base + 2_4E80h offset + (4d × i), where i=0d to 1d



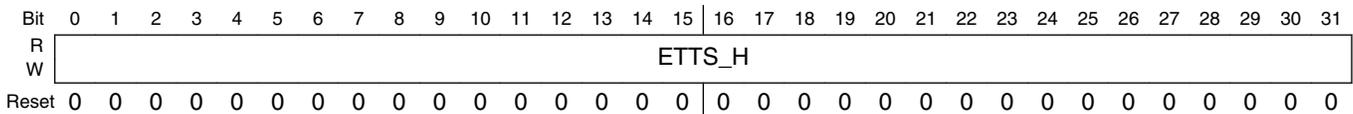
ETSEC_1588_TMR_FIPERn field descriptions

Field	Description
0–31 FIPER	Fixed interval pulse period register. This field must be programmed to an integer multiple of TMR_CTRL[TCLK_PERIOD] value to ensure a period pulse being generated correctly.

14.54.17 Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTSn_H)

General purpose external trigger -stamp register (TMR_ETTSn_H/L). These registers hold time at the programmable edge of the external trigger. The registers in eTSEC1 are shared for all eTSECs. These registers are read only in normal operation.

Address: 0h base + 2_4EA0h offset + (8d × i), where i=0d to 1d



ETSEC_1588_TMR_ETTSn_H field descriptions

Field	Description
0–31 ETTS_H	Time stamp field at the programmable edge of the external trigger. (upper half)

14.54.18 Time stamp of general purpose external trigger (ETSEC_1588_TMR_ETTSn_L)

General purpose external trigger -stamp register (TMR_ETTSn_H/L). These registers hold time at the programmable edge of the external trigger. The registers in eTSEC1 are shared for all eTSECs. These registers are read only in normal operation.

Address: 0h base + 2_4EA4h offset + (8d × i), where i=0d to 1d



ETSEC_1588_TMR_ETTSn_L field descriptions

Field	Description
0-31 ETTS_L	Time stamp field at the programmable edge of the external trigger. (lower half)

14.6 Functional description

This section describes the following:

- [Programming Model Considerations](#)
- [Connecting to physical interfaces on Ethernet](#)
- [Gigabit Ethernet controller channel operation](#)
- [TCP/IP off-load](#)
- [Quality of service \(QoS\) provision](#)
- [Lossless Flow Control Configuration Registers](#)
- [Hardware assist for IEEE 1588-compatible timestamping](#)
- [Buffer descriptors](#)

14.6.1 Programming Model Considerations

The eTSECs use a software model that is a superset of the PowerQUICC TSEC functionality and is similar to that employed by the Fast Ethernet function supported on the Freescale MPC8260 CPM FCC and in the FEC of the MPC860T.

The eTSEC device is programmed by a combination of control/status registers (CSRs) and buffer descriptors. The CSRs are used for mode control, interrupts, and to extract status information. The descriptors are used to pass data buffers and related buffer status or frame information between the hardware and software.

All accesses to and from the registers must be made as 32-bit accesses. There is no support for accesses of sizes other than 32 bits. Reads from unmapped register addresses return zero. Unless otherwise specified, the read value of reserved bits in mapped registers is not defined, and must not be assumed to be 0.

This section of the document defines the memory map and describes the registers in detail. The buffer descriptor is described in [Buffer descriptors](#).

The TBI and reduced ten-bit interface (RTBI) module MII registers are also described in this section. The TBI/RTBI registers are defined like PHY registers and, as such, are accessed through the MII management interface in the same way the PHYs are accessed. For detailed descriptions of the TBI/RTBI registers (the MII register set for the ten-bit interface) refer to [ETSEC_TBI Memory Map/Register Definition](#).

14.6.1.1 MAC Functionality

This section describes the MAC registers and provides a brief overview of the functionality that can be exercised through the use of these registers, particularly those that provide functionality not explicitly required by the IEEE 802.3 standard. All of the MAC registers are 32 bits wide.

14.6.1.1.1 Configuring the MAC

The MAC configuration registers 1 and 2 provide for configuring the MAC in multiple ways:

- Adjusting the preamble length-The length of the preamble can be adjusted from the nominal seven bytes to some other (≥ 3 byte) value. Should custom preamble insertion/extraction be configured, then this register must be left at its default value of 7.
- Varying pad/CRC combinations-Three different pad/CRC combinations are provided to handle a variety of system requirements. Simplest are frames that already have a valid frame check sequence (FCS) field. The other two options include appending a valid CRC or padding and then appending a valid CRC, resulting in a minimum frame of 64 octets. In addition to the programmable register set, the pad/CRC behavior can be dynamically adjusted on a per-packet basis.

14.6.1.1.2 Controlling CSMA/CD

The half-duplex register (HAFDUP) allows control over the carrier-sense multiple access/collision detection (CSMA/CD) logic of the eTSEC. Half-duplex mode is only supported for 10- and 100-Mbps operation. Following the completion of the packet transmission the part begins timing the inter packet gap (IPG) as programmed in the back-to-back IPG configuration register. The system is now free to begin another frame transfer.

In full-duplex mode both the carrier sense (CRS) and collision (COL) indications from the PHY are ignored, but in half-duplex mode the eTSEC defers to CRS, and following a carrier event, times the IPG using the non-back-to-back IPG configuration values that include support for the optional two-thirds/one-third CRS deferral process. This optional IPG mechanism enhances system robustness and ensures fair access to the medium. During the first two-thirds of the IPG, the IPG timer is cleared if CRS is sensed. During the final one-third of the IPG, CRS is ignored and the transmission begins once IPG is timed. The two-thirds/one-third ratio is the recommended value.

14.6.1.1.3 Handling Packet Collisions

While transmitting a packet in half-duplex mode, the eTSEC is sensitive to COL. If a collision occurs, it aborts the packet and outputs the 32-bit jam sequence. The jam sequence is comprised of several bits of the CRC, inverted to guarantee an invalid CRC upon reception. A signal is sent to the system indicating that a collision occurred and that the start of the frame is needed for retransmission. The eTSEC then backs off of the medium for a time determined by the truncated binary exponential back off (BEB) algorithm. Following this back-off time, the packet is retried. The back-off time can be skipped if configured through the half-duplex register. However, this is non-standard behavior and its use must be carefully applied. Should any one packet experience excessive collisions, the packet is aborted. The system should flush the frame and move to the next one in line. If the system requests to send a packet while the eTSEC is deferring to a carrier, the eTSEC simply waits until the end of the carrier event and the timing of IPG before it honors the request.

If packet transmission attempts experience collisions, the eTSEC outputs the jam sequence and waits some amount of time before retrying the packet. This amount of time is determined by a controlled randomization process called truncated binary exponential back-off. The amount of time is an integer number of slot times. The number of slot times to delay before the n th retransmission attempt is chosen as a uniformly-distributed random integer r in the range:

$$0 \leq r \leq 2k, \text{ where } k = \min(n, 10).$$

So after the first collision, the eTSEC backs off either 0 or 1 slot times. After the fifth collision, the eTSEC backs off between 0 and 32 slot times. After the tenth collision, the maximum number of slot times to back off is 1024. This can be adjusted through the half-duplex register. An alternate truncation point, such as 7 for instance, can be programmed. On average, the MAC is more aggressive after seven collisions than other stations on the network.

14.6.1.1.4 Controlling Packet Flow

Packet flow can be dealt with in a number of ways within eTSEC. A default retransmit attempt limit of 15 can be reduced using the half-duplex register. The slot time or collision window can be used to gate the retry window and possibly reduce the amount of transmit buffering within the system. The slot time for 10/100 Mbps is 512 bit times. Because the slot time begins at the beginning of the packet (including preamble), the end occurs around the 56th byte of the frame data. Slot time in 1000-Mbps mode is not supported.

Full-duplex flow control is provided for in IEEE 802.3x. Currently the standard does not address flow control in half-duplex environments. Common in the industry, however, is the concept of back pressure. The eTSEC implements the optional back pressure mechanism using the raise carrier method. If the system receive logic wishes to stop the reception of packets in a network-friendly way, transmit half-duplex flow control (THDF) is set (TCTRL[THDF]). If the medium is idle, the eTSEC raises carrier by transmitting preamble. Other stations on the half-duplex network then defer to the carrier.

In the event the preamble transmission happens to cause a collision, the eTSEC ensures the minimum 96-bit presence on the wire, then drops preamble and waits a back-off time depending on the value of the back-pressure-no-back-off configuration bit HAFDUP[BP No BackOff]. These transmitting-preamble-for-back pressure collisions are not counted. If HAFDUP[BP No BackOff] is set, the eTSEC waits an inter-packet gap before resuming the transmission of preamble following the collision and does not defer. If HAFDUP[BP No BackOff] is cleared, the eTSEC adheres to the truncated BEB algorithm that allows the possibility of packets being received. This also can be detrimental in that packets can now experience excessive collisions, causing them to be dropped in the stations from which they originate. To reduce the likelihood of lost packets and packets leaking through the back pressure mechanism, HAFDUP[BP No BackOff] must be set.

The eTSEC drops carrier (cease transmitting preamble) periodically to avoid excessive defer conditions in other stations on the shared network. If, while applying back pressure, the eTSEC is requested to send a packet, it stops sending preamble, and waits one IPG before sending the packet. HAFDUP[BP No BackOff] applies for any collision that occurs during the sending of this packet. Collisions for packets while half duplex back

pressure is asserted are counted. The eTSEC does not defer while attempting to send packets while in back pressure. Again, back pressure is non-standard, yet it can be effective in reducing the flow of receive packets.

14.6.1.1.5 Controlling PHY Links

Control and status to and from the PHY is provided through the two-wire MII management interface described in IEEE 802.3u. The MII management registers (MII management configuration, command, address, control, status, and indicator registers) are used to exercise this interface between a host processor and one or more PHY devices (including the TBI).

The eTSEC MII's registers provide the ability to perform continuous read cycles (called a scan cycle); although, scan cycles are not explicitly defined in the standard. If requested (by setting MIIMCOM[Scan Cycle]), the part performs repetitive read cycles of the PHY status register, for example. In this way, link characteristics may be monitored more efficiently. The different fields in the MII management indicator register (scan, not valid and busy) are used to indicate availability of each read of the scan cycle to the host from MIIMSTAT[PHY scan].

Yet another parameter that can be modified through the MII registers is the length of the MII management interface preamble. After establishing that a PHY supports preamble suppression, the host may so configure the eTSEC. While enabled, the length of MII management frames are reduced from 64 clocks to 32 clocks. This effectively doubles the efficiency of the interface.

14.6.1.2 MIB Registers

This section describes the MIB registers. The eTSEC RMON module has 37 separate statistics counters, which simply count or accumulate statistical events that occur as packets transmitted and received. These counters support RMON MIB group 1, RMON MIB group 2 if table counters, RMON MIB group 3, RMON MIB group 9, RMON MIB 2, and the IEEE 802.3 Ethernet MIB.

An interrupt can be generated upon any one counter's rollover condition through a carry interrupt output from the RMON. Each counter's rollover condition can be discretely masked from causing an interrupt by internal masking registers. In addition, each individual counter value may be reset on read access, or all counters may be simultaneously reset by setting ECNTRL[CLRCNT].

The majority of MIB counters are Ethernet-specific.

NOTE

RMON counters do not comprehend custom VLAN tagged frames. Affected counters include TRMGV, RMCA, RBCA, RXCF, RXPf, RXUO, RALN, RFLR, ROVR, RJBR, TMCA, TBCA, TXPF, TXCF. Specifically, custom VLAN tagged frames are not afforded the ability to be greater than 1518, as compared to the IEEE standard tagged frames.

NOTE

The transmit and receive frame counters (TR64, TR127, TR255, TR511, TR1K, TRMAX, and TRMGV) do not increment for aborted frames (collision retry limit exceeded, late collision, underrun, EBERR, TxFIFO data error, frame truncated due to exceeding MAXFRM, or excessive deferral).

14.6.1.3 Hash Function Registers

This section provides detailed descriptions of the registers used for hash functions. All of the registers are 32 bits wide. The DA field of every received frame is processed through a 32-bit CRC generator (CRC-32 polynomial), and the 8 or 9 most significant bits of the CRC are mapped to a hash table entry. The user can enable a hash entry by setting its bit. A hash entry usually represents a set of addresses. A hash table hit occurs if the DA CRC result points to an enabled hash entry. Software may need to further filter the address in order to eliminate false-positive hits in the hash table.

If RCTRL[GHTX] = 0, the 8 most significant bits of the CRC are used as the hash table index. In this case, registers IGADDR0-IGADDR7 comprise a 256-entry hash table exclusively for individual (unicast) address matching, while registers GADDR0-GADDR7 comprise a 256-entry hash table for group (multicast) address matching. If RCTRL[GHTX] = 1, the group hash table is extended to all 512 entries, and the 9 most significant bits of the CRC are used as the hash table index. In this case, registers IGADDR0-IGADDR7 hold hash table entries 0-255 for group addresses, while registers GADDR0-GADDR7 hold entries 256-511 of the extended group hash table.

See [Hash table algorithm](#) for more information on the hash algorithm.

14.6.1.4 Lossless Flow Control Configuration Registers

When enabled through RCTRL[LFC], the eTSEC tracks location of the last free BD in each Rx BD ring through the value of RFBPTRn. Using this pointer and the ring length stored in RQPRMn[LEN], the eTSEC continuously calculates the number of free BDs in

the ring. Whenever the calculated number of free BDs in the ring drops below the pause threshold specified in RQPRMn[FBTHR], the eTSEC issues link layer flow control. It continues to assert flow control until the free BD count for each active ring reaches or exceeds RQPRMn[FBTHR]. See [Back pressure determination through free buffers](#) for the theory of operation of these registers.

14.6.1.5 IEEE 1588-Compatible Timestamping Registers

IEEE 1588 compliant timestamping on this device is accomplished using the per-port transmit timestamping registers within each Ethernet controller memory space (See [Tx time stamp identification tag \(set n\) register \(ETSEC_TMR_TXTSn_ID\)](#), [Tx time stamp high \(ETSEC_TMR_TXTSn_H\)](#), and [Tx time stamp low \(ETSEC_TMR_TXTSn_L\)](#)) in conjunction with the following common registers, which are located within the memory space for eTSEC1. Because the common 1588 timestamping registers exist within the eTSEC1 memory space, the eTSEC1 controller must remain enabled in order to use 1588 timestamping for any Ethernet port.

14.6.1.6 Ten-Bit Interface (TBI)

This section describes the ten-bit interface (TBI), reduced ten-bit interface (RTBI), and the TBI/RTBI MII set of registers. TBI and RTBI operate in the same manner (the only difference is that RTBI has reduced I/O signalling).

14.6.1.6.1 TBI Transmit Process

The eTSEC's TBI implements the transmit portion of the physical coding sublayer as found in Clause 36 of IEEE 802.3z. In SerDes mode, packets conveyed across the GMII are encapsulated and encoded into 10-bit symbols and output to the SerDes. In GMII mode, the GMII signals are passed through to the attached GMII PHY.

14.6.1.6.1.1 Packet Encapsulation

If TX_EN is de-asserted the eTSEC outputs an idle stream. If TX_EN is asserted, a Start_of_Packet symbol is output. This symbol replaces the first byte of the preamble field. All other bytes of the packet pass through an 8B10B encoding module. After the last byte of the FCS field is signaled through the GMII, the MAC de-asserts TX_EN. The eTSEC then outputs an End_of_Packet symbol. Then, depending on the position of the End_of_Packet symbols (being in either an odd or even position) the eTSEC outputs one or two Carrier_Extend symbols. Following the last Carrier_Extend symbol, the eTSEC

resumes sending idle codes. If, during a packet, the eTSEC wishes to mark a byte invalid, TX_ER is asserted. The eTSEC, upon detection of TX_ER, substitutes the data symbol for an Error_Propagation symbol.

14.6.1.6.1.2 8B10B Encoding

Every eight-bit data octet has two (not necessarily different) ten-bit symbols associated with it. Depending on the running disparity (the cumulative difference of ones and zeros) the eTSEC module chooses the appropriate symbol.

Special encapsulation symbols are called ordered_sets. Ordered_sets are comprised of one to four ten-bit symbols. Ordered_sets can be found in Clause 36 of the IEEE 802.3z specification.

14.6.1.6.1.3 Preamble Shortening

Because the idle ordered_set comprises two symbols and begins on an even symbols boundary, packets can only begin on an even boundary. However, the GMII has no such restriction and may signal TX_EN on an odd boundary. If this happens, the eTSEC delays the Start_of_Packet symbol, effectively ignoring the first byte of preamble; thus, a seven octet preamble becomes six octets on the Ten-Bit Interface.

14.6.1.6.2 TBI Receive Process

The eTSEC's TBI Implements the receive portion of the physical coding sublayer as found in Clause 36 of IEEE 802.3z. The Receive portion includes the Synchronization state machine. In SerDes mode, the eTSEC first attempts to acquire synchronization on the link by examining received symbols. Once synchronization is acquired, received packets are decoded and sent across the Receive GMII interface. In GMII mode, the GMII signals are passed through to the MAC.

14.6.1.6.2.1 Synchronization

The eTSEC examines received symbols looking for the seven bit 'comma' string embedded in some special symbols. Both the idle ordered_set and the Configuration ordered_set contain a symbol which has the comma. Once a certain number of codes with comma are detected, the eTSEC is considered to have acquired synchronization.

14.6.1.6.2.2 Auto-Negotiation for 1000BASE-X

Once synchronization is acquired, ordered_sets are decoded. If Configuration ordered_sets are received, the eTSEC decodes the two octet data field and the sixteen-bit Configuration data is stored and used to Auto-Negotiate with the link partner. in the

Receive Configuration Register (RXCR[15:0]) an internal register used to receive all the link partners informations and used to compare to local ability during negotiation. Not visible to user. If, during Auto-Negotiation an invalid symbol is detected, Auto-Negotiation re-starts. After Auto-Negotiation is completed the TBI MII Status Register SR[AN done] in set. In this mode, packets may be received from the link partner.

14.6.2 Connecting to physical interfaces on Ethernet

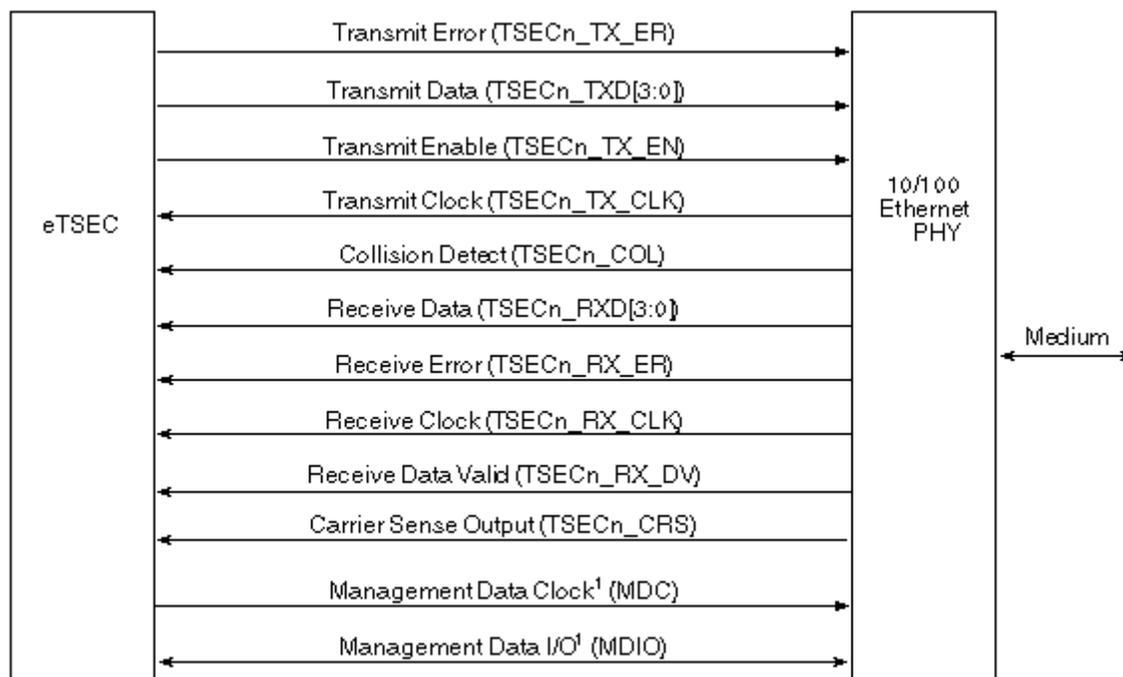
This section describes how to connect the eTSEC to various interfaces: MII, GMII, RMII, RGMII, TBI, and RTBI.

To avoid confusion, all of the buses follow the bus conventions used in the IEEE 802.3 specification because the PHYs follow the same conventions. (For instance, in the bus TSEC_n_TXD[7 :0], bit 7 is the msb and bit 0 is the lsb). If a mode does not use all input signals available to a particular eTSEC, those inputs that are not used must be pulled low on the board.

14.6.2.1 Media-independent interface (MII)

This section describes the media-independent interface (MII) intended to be used between the PHYs and the eTSEC.

The figure below depicts the basic components of the MII including the signals required to establish eTSEC module connection with a PHY.



¹ The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

Figure 14-872. eTSEC-MII connection

An MII interface has 18 signals (including the MDC and MDIO signals), as defined by the IEEE 802.3u standard, for connecting to an Ethernet PHY.

14.6.2.2 Reduced media-independent interface (RMII)

This section describes the reduced media-independent interface (RMII) intended to be used between the PHYs and the GMII MAC.

The RMII is a reduced-pin alternative to the IEEE 802.3u MII. The RMII reduces the number of signals required to interconnect the MAC and the PHY from a maximum of 18 signals (MII) to 10 signals. To accomplish this objective, the data paths are halved in width and clocked at twice the MII clock frequency, while clocks, carrier sense and error signals have been partly combined. For 100 Mbps operation, the reference clock operates at 50 MHz, whereas for 10 Mbps operation, the clock remains at 50MHz, but only every 10th cycle is used. The figure below depicts the basic components of the reduced media-independent interface and the signals required to establish an eTSEC's connection with a PHY. The RMII is implemented as defined by the RMII Specification of the RMII Consortium, as of March 20, 1998.

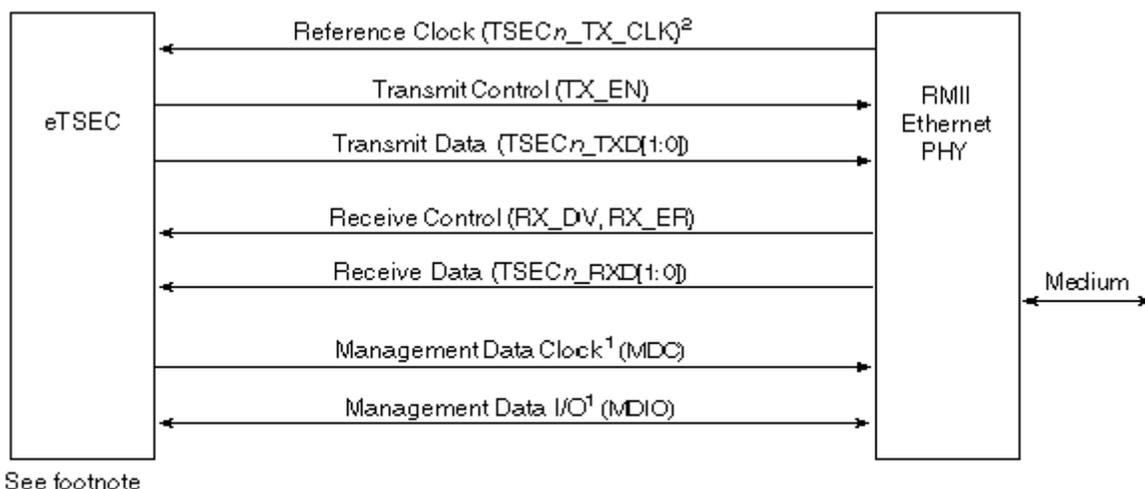


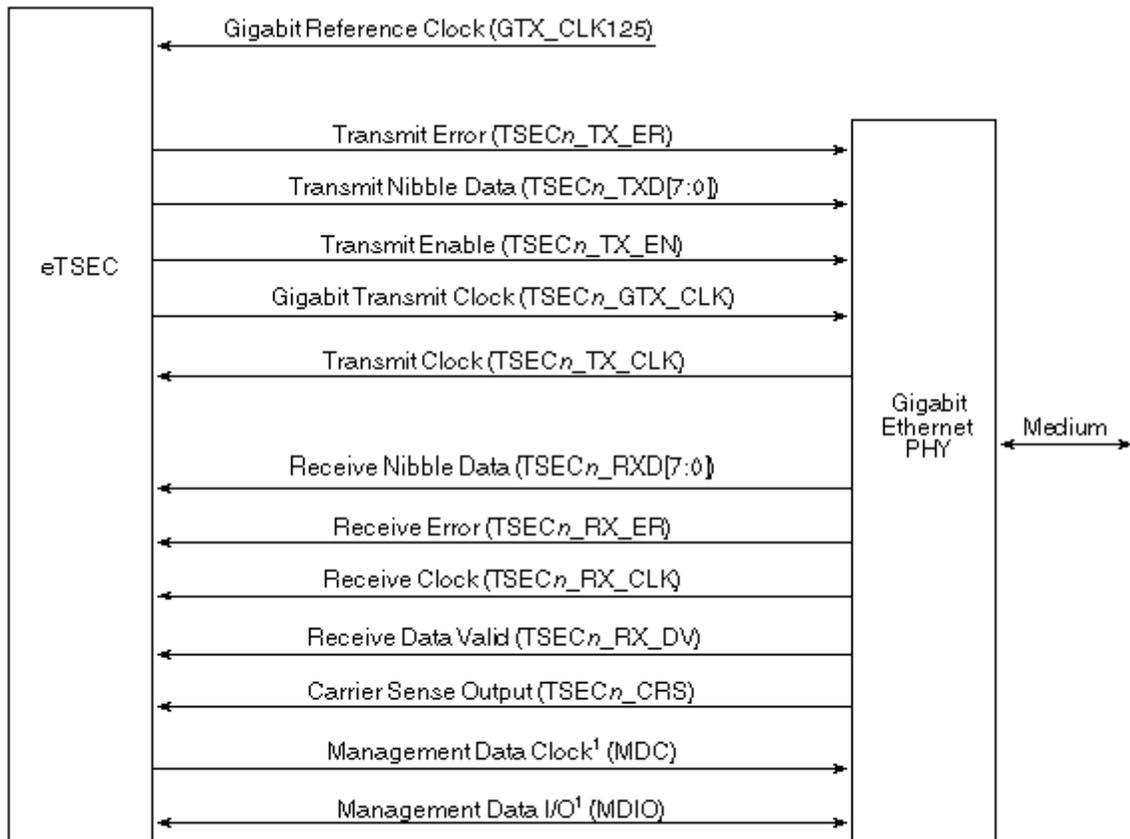
Figure 14-873. eTSEC-RMII connection

1) The management signals (MDC and MDIO) are common to all of the Ethernet controllers module connections in the system, assuming that each PHY has a different management address.

14.6.2.3 Gigabit media-independent interface (GMII)

This section describes the gigabit media-independent interface (GMII) intended to be used between the PHYs and the eTSEC.

The figure below depicts the basic components of the GMII including the signals required to establish the eTSEC module connection with a PHY.



¹ The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

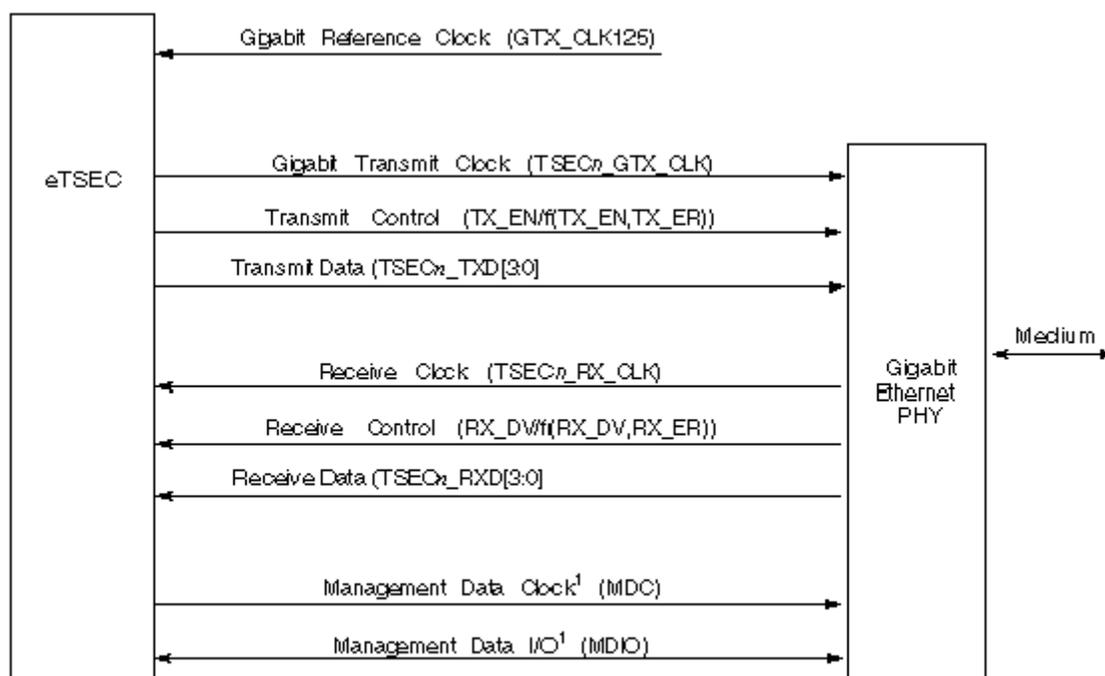
Figure 14-874. eTSEC-GMII Connection

A GMII interface has 28 signals (TSEC_n_GTX_CLK + _GTX_CLK125 included), as defined by the IEEE 802.3u standard, for connecting to an Ethernet PHY.

14.6.2.4 Reduced gigabit media-independent interface (RGMII)

This section describes the reduced gigabit media-independent interface (RGMII) intended to be used between the PHYs and the GMII MAC.

The RGMII is an alternative to the IEEE802.3u MII, the IEEE802.3z GMII and the TBI. The RGMII reduces the number of signals required to interconnect the MAC and the PHY from a maximum of 28 signals (GMII) to 15 signals (GTX_CLK125 included) in a cost effective and technology independent manner. To accomplish this objective, the data paths and all associated control signals are multiplexed using both edges of the clock. For gigabit operation, the clocks operate at 125MHz, and for 10/100 operation, the clocks operate at 2.5 MHz or 25 MHz, respectively. Note that the GTX_CLK125 input must be provided at 125 MHz for an RGMII interface, regardless of operation speed (1 Gbps, 100 Mbps, or 10 Mbps). The figure below depicts the basic components of the gigabit reduced media-independent interface and the signals required to establish the gigabit Ethernet controllers' module connection with a PHY. The RGMII is implemented as defined by the RGMII specification Version 1.2a 9/22/00.



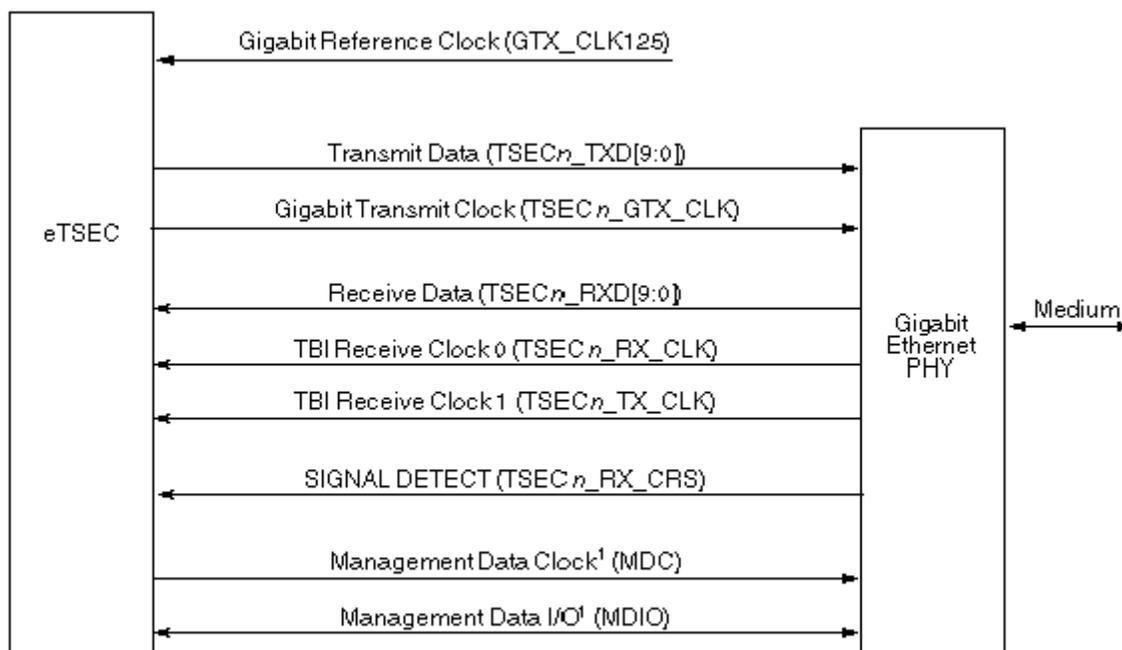
¹ The management signals (MDC and MDIO) are common to all of the gigabit Ethernet controllers module connections in the system, assuming that each PHY has a different management address.

Figure 14-875. eTSEC-RGMII connection

14.6.2.5 Ten-bit interface (TBI)

This section describes the ten-bit interface (TBI) intended to be used between the PHYs and the eTSEC to implement a standard SerDes interface for optical-fiber devices in 1000BASE-SX/LX applications.

The figure below depicts the basic components of the TBI including the signals required to establish eTSEC module connection with a PHY. RBC0 and RBC1 are differential 62.5 MHz receive clocks. If not connected to the TBI PHY, the Signal Detect (SDET) input must be tied high. This causes the eTSEC to begin auto negotiation with the SerDes immediately upon the TBI module being enabled.



¹ The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

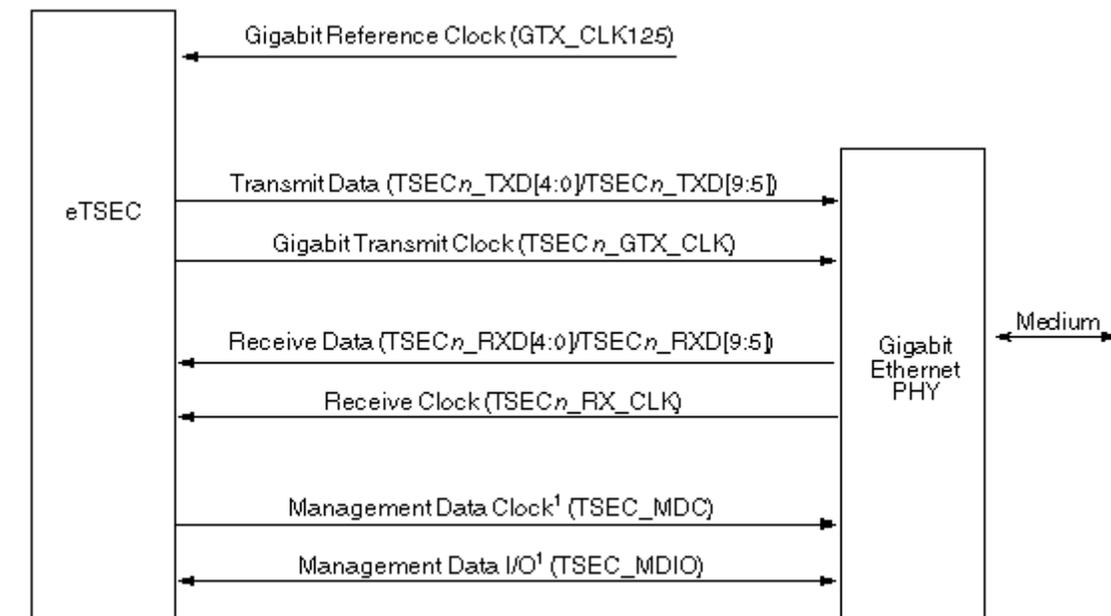
Figure 14-876. eTSEC-TBI connection

A TBI interface has 26 signals (GE_GTX_CLK125 included) for connecting to an Ethernet PHY, as defined by IEEE 802.3z GMII and TBI standards.

14.6.2.6 Reduced ten-bit interface (RTBI)

This section describes the reduced ten-bit interface (RTBI) intended to be used between the PHYs and the eTSEC to implement a reduced-pin count version of a SerDes interface for optical-fiber devices in 1000BASE-SX/LX applications.

The figure below depicts the basic components of the RTBI including the signals required to establish eTSEC module connection with a PHY. Note that in RTBI the eTSEC immediately begins auto-negotiation with the SerDes.



¹ The management signals (MDC and MDIO) are common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

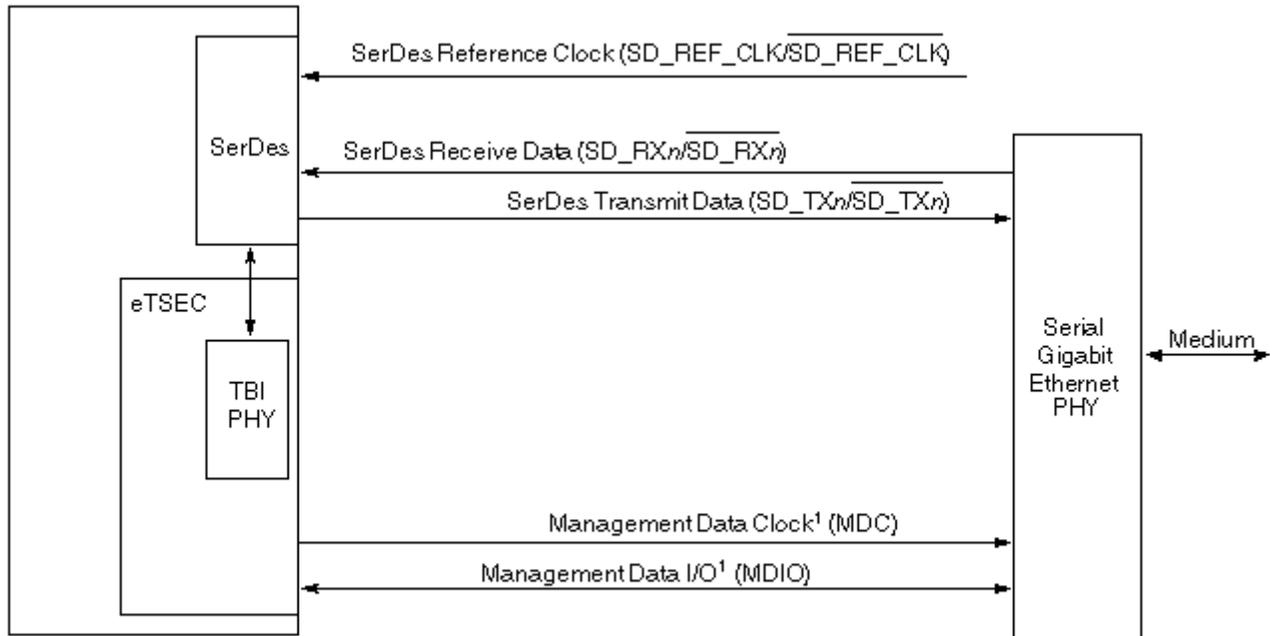
Figure 14-877. eTSEC-RTBI connection

A RTBI interface has 15 signals (GE_GTX_CLK125 included), as defined by the RGMII specification Version 1.2a 9/22/00, and is intended to be an alternative to the IEEE 802.3u MII, the IEEE 802.3z GMII and the TBI standard for connecting to an Ethernet PHY.

14.6.2.7 Serial gigabit media-independent interface (SGMII)

This section describes the serial gigabit media-independent interface (SGMII) intended to be used between a SerDes PHY and the eTSEC to implement a serial gigabit version of a media-independent interface.

The figure below depicts the basic components of the SGMII including the signals required to establish eTSEC module connection with a PHY. Note that in SGMII the eTSEC utilizes the on-chip TBI PHY in addition to the SerDes interface.



¹ The management signals (MDC and MDIO) may be common to all of the Ethernet controllers' connections in the system, assuming that each PHY has a different management address.

Figure 14-878. eTSEC-SGMII connection

14.6.2.8 Ethernet physical interfaces signal summary

The table below describes the signal multiplexing for the GMII, MII, and RMII interfaces.

Table 14-887. GMII, MII, and RMII signals multiplexing

eTSEC Signals			GMII Interface			MII Interface			RMII Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 25			Frequency [MHz] 50		
Voltage[V] 3.3/2.5			Voltage[V] 3.3			Voltage[V] 3.3			Voltage[V] 3.3		
Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals

Table continues on the next page...

Table 14-887. GMII, MII, and RMII signals multiplexing (continued)

eTSEC Signals			GMII Interface			MII Interface			RMII Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 25			Frequency [MHz] 50		
Voltage[V] 3.3/2.5			Voltage[V] 3.3			Voltage[V] 3.3			Voltage[V] 3.3		
GTX_CLK	O	1	GTX_CLK	O	1						
TX_CLK	I	1	TX_CLK	I	1	TX_CLK	I	1	REF_CLK	I	1
TxD[0]	O	1	TxD[0]	O	1	TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1	TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	TxD[2]	O	1	TxD[2]	O	1			
TxD[3]	O	1	TxD[3]	O	1	TxD[3]	O	1			
TX_EN	O	1	TX_EN	O	1	TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	TX_ER	O	1	TX_ER	O	1			
RX_CLK	I	1	RX_CLK	I	1	RX_CLK	I	1			
RxD[0]	I	1	RxD[0]	I	1	RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1	RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	RxD[2]	I	1	RxD[2]	I	1			
RxD[3]	I	1	RxD[3]	I	1	RxD[3]	I	1			
RX_DV	I	1	RX_DV	I	1	RX_DV	I	1	CRS_DV	I	1
RX_ER	I	1	RX_ER	I	1	RX_ER	I	1	RX_ER	I	1
COL	I	1				COL	I	1			
CRS	I	1				CRS	I	1			
Sum		25	Sum		23	Sum		16	Sum		8

The table below describes the signal multiplexing for RGMII, TBI, and RTBI interfaces.

Table 14-888. RGMII, TBI, and RTBI Signals Multiplexing

eTSEC Signals			RGMII Interface			TBI Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 3.3			Voltage[V] 2.5		
Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals
GTX_CLK	O	1									
TX_CLK	I	1				RX_CLK1	I	1			
TxD[0]	O	1	TxD[0]	O	1	TCG[0]	O	1	TCG[0]/TCG[5]	O	1
TxD[1]	O	1	TxD[1]	O	1	TCG[1]	O	1	TCG[1]/TCG[6]	O	1
TxD[2]	O	1	TxD[2]	O	1	TCG[2]	O	1	TCG[2]/TCG[7]	O	1

Table continues on the next page...

Table 14-888. RGMII, TBI, and RTBI Signals Multiplexing (continued)

eTSEC Signals			RGMII Interface			TBI Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 3.3			Voltage[V] 2.5		
Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals
TxD[3]	O	1	TxD[3]/	O	1	TCG[3]	O	1	TCG[3]/ TCG[8]	O	1
TX_EN	O	1	TX_CTL (TX_EN/ TX_ERR)	O	1	TCG[8]	O	1	TCG[4]/ TCG[9]	O	1
TX_ER	O	1				TCG[9]	O	1			
RX_CLK	I	1	RX_CLK	I	1	RX_CLK 0	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]/ RxD[4]	I	1	RCG[0]	I	1	RCG[0]/ RCG[5]	I	1
RxD[1]	I	1	RxD[1]/ RxD[5]	I	1	RCG[1]	I	1	RCG[1]/ RCG[6]	I	1
RxD[2]	I	1	RxD[2]/ RxD[6]	I	1	RCG[2]	I	1	RCG[2]/ RCG[7]	I	1
RxD[3]	I	1	RxD[3]/ RxD[7]	I	1	RCG[3]	I	1	RCG[3]/ RCG[8]	I	1
RX_DV	I	1	RX_CTL (RX_DV/ RX_ERR)	I	1	RCG[8]	I	1	RCG[4]/ RCG[9]	I	1
RX_ER	I	1				RCG[9]	I	1			
COL	I	1									
CRS	I	1				SDET	I	1		I	
Sum		25	Sum		12	Sum		24	Sum		12

The table below describes the signal multiplexing for RGMII and RTBI interfaces.

Table 14-889. RGMII and RTBI Signals Multiplexing

eTSEC Signals			RGMII Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 2.5		
Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals	Signals (TSECn ₋)	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1						
TxD[0]	O	1	TxD[0]	O	1	TCG[0]/TCG[5]	O	1

Table continues on the next page...

Table 14-889. RGMII and RTBI Signals Multiplexing (continued)

eTSEC Signals			RGMII Interface			RTBI Interface		
Frequency [MHz] 125			Frequency [MHz] 125			Frequency [MHz] 62.5		
Voltage[V] 3.3/2.5			Voltage[V] 2.5			Voltage[V] 2.5		
Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals	Signals (TSECn_)	I/O	No. of Signals
TxD[1]	O	1	TxD[1]	O	1	TCG[1]/TCG[6]	O	1
TxD[2]	O	1	TxD[2]	O	1	TCG[2]/TCG[7]	O	1
TxD[3]	O	1	TxD[3]	O	1	TCG[3]/TCG[8]	O	1
TX_EN	O	1	TX_CTL (TX_EN/ TX_ERR)	O	1	TCG[4]/TCG[9]	O	1
TX_ER	O	1						
RX_CLK	I	1	RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]/RxD[4]	I	1	RCG[0]/RCG[5]	I	1
RxD[1]	I	1	RxD[1]/RxD[5]	I	1	RCG[1]/RCG[6]	I	1
RxD[2]	I	1	RxD[2]/RxD[6]	I	1	RCG[2]/RCG[7]	I	1
RxD[3]	I	1	RxD[3]/RxD[7]	I	1	RCG[3]/RCG[8]	I	1
RX_DV	I	1	RX_CTL (RX_DV/ RX_ERR)	I	1	RCG[4]/RCG[9]	I	1
RX_ER	I	1						
COL	I	1						
CRS	I	1					I	
Sum		25	Sum		12	Sum		12

The table below describes the signalling for SGMII interfaces. SGMII communication using the eTSEC is accomplished through the SerDes interface.

Table 14-890. SGMII Signalling

Signals	I/O	No. of Signals	Function
SDn_RX[n]	I	2	SGMII receive data (differential)
SDn_TX[n]	O	2	SGMII transmit data (differential)
SDn_REF_CLK	I	2	Reference clock (differential)
Sum		6	-

The table below describes the signals shared by all interfaces.

Table 14-891. Shared Signals

Signals	I/O	No. of Signals	Function
MDIO	I/O	1	Management interface I/O
MDC	O	1	Management interface clock

Table continues on the next page...

Table 14-891. Shared Signals (continued)

Signals	I/O	No. of Signals	Function
GTX_CLK125	I	1	Reference clock
Sum		3	-

14.6.3 Gigabit Ethernet controller channel operation

This section describes the operation of the eTSEC.

First, the software initialization sequence is described. Next, the software (Ethernet driver) interface for transmitting and receiving frames is reviewed. Frame filtering and receive filing algorithm features are also discussed. The section concludes with interrupt handling, inter-packet gap time, and loop back descriptions.

14.6.3.1 Initialization sequence

This section describes which registers are reset due to a hard or software reset and what registers the user must initialize prior to enabling the eTSEC.

14.6.3.1.1 Hardware controlled initialization

A hard reset occurs when the system powers up.

All eTSEC's registers and control logic are reset to their default states after a hard reset has occurred. In this state, each eTSEC behaves like a PowerQUICC device, except for the absence of out-of-sequence TxBD features. That is, initially TCP/IP off-load is disabled and only single RxBD and TxBD rings are accessible.

14.6.3.1.2 User initialization

After the system has undergone a hard reset, software must initialize certain basic eTSEC registers.

Other registers can also be initialized during this time, but they are optional and must be determined based on the requirements of the system. See [Table 14-3](#) for the register list. The table below describes the minimum steps for register initialization.

Table 14-892. Steps for minimum register initialization

Description
1. Set and clear MACCFG1 [Soft_Reset]

Table continues on the next page...

Table 14-892. Steps for minimum register initialization (continued)

Description
2. Initialize MACCFG2
3. Initialize MAC station address
4. Set up the PHY using the MII Mgmt Interface
5. Configure the TBI control to TBI or GMII
6. Clear IEVENT
7. Initialize IMASK
8. Initialize RCTRL
9. Initialize DMACTRL

After the initialization of registers is performed, the user must execute the following steps in the order described below to bring the eTSEC into a functional state (out of reset):

1. Write to the MACCFG1 register and set the appropriate bits. These need to include RX_EN and TX_EN. To enable flow control, Rx_Flow and Tx_Flow should also be set.
2. For the transmission of Ethernet frames, TxBDs must first be built in memory, linked together as a ring, and pointed to by the TBASE_n registers. A minimum of nine buffer descriptors per ring is required, unless the ring is disabled. If TCP/IP offload is to be enabled, the TxBD[TOE] bit must be set for each frame.
3. Likewise, for the reception of Ethernet frames, a minimum of nine buffer descriptors per ring is required, unless the ring is disabled, and the receive queue (or queues) must be ready, with its RxBD pointed to by the RBASE_n registers. If TCP/IP offload is to be enabled, RCTRL[PRSDP] must be set to the required offload level. Both transmit and receive can be gracefully stopped after transmission and reception begins.
4. Clearing DMACTRL[GTS] triggers the transmission of frame data if the transmitter had been previously stopped. The DMACTRL[GRS] must be cleared if the receiver had been previously stopped. Refer to the DMACTRL register section, and [Data buffer descriptors](#), for more information.

14.6.3.2 Soft reset and reconfiguring procedure

Before issuing a soft-reset to and/or reconfiguring the MAC with new parameters, the user must properly shutdown the DMA and make sure it is in an idle state for the entire duration.

User must gracefully stop the DMA by setting both GRS and GTS bits in the DMACTRL register, then wait for both GRSC and GTSC bits to be set in the IEVENT register before resetting the MAC or changing parameters. Both GRS and GTS bits must be cleared before re-enabling the MAC to resume the DMA.

During the MAC configuration, if a new set of Tx buffer descriptors are used, the user must load the pointers into the TBASE registers. Likewise if a new set of Rx buffer descriptors are used, the RBASE registers must be written with new pointers.

Following is a procedure to gracefully reset and reconfigure the MAC:

1. Set GRS/GTS bits in DMACTRL register
2. Poll GRSC/GTSC bits in IEVENT register until both are set
3. Set SOFT_RESET bit in MACCFG1 register (Note that SOFT_RESET must remain set for at least 3-10 TX clocks before proceeding.)
4. Clear SOFT_RESET bit in MACCFG1 register
5. Load TBASE0-TBASE7 with new Tx BD pointers
6. Load RBASE0-RBASE7 with new Rx BD pointers
7. Setup other MAC registers (MACCFG2, MAXFRM, and so on)
8. Setup group address hash table (GADDR0-GADDR15) if address filtering is required
9. Setup receive frame filter table (through RQFAR, RQFCR, and RQFPR) if filtering to multiple RxBD rings is required
10. Setup WWR, WOP, TOD bits in DMACTRL register
11. Enable transmit queues in TQUEUE, and ensure that the transmit scheduling mode is correctly set in TCTRL.
12. Enable receive queues in RQUEUE, and optionally set TOE functionality in RCTRL.
13. Clear THLT and TXF bits in TSTAT register by writing 1 to them
14. Clear QHLT and RXF bits in RSTAT register by writing 1 to them.
15. Clear GRS/GTS bits in DMACTRL (do not change other bits)
16. Enable Tx_EN/Rx_EN in MACCFG1 register

14.6.3.3 Gigabit Ethernet frame transmission

The Ethernet transmitter requires little core intervention.

After the software driver initializes the system, the eTSEC begins to poll the first transmit buffer descriptor (TxBD) in TxBD ring 0 every 512 transmit clocks. If TxBD[R] is set, and the TxBD ring is scheduled for transmission, the eTSEC begins copying the associated transmit buffer from memory to its Tx FIFO. The transmitter takes data from the Tx FIFO and transmits data to the MAC. The MAC transmits the data through the

GMII interface to the physical media. The transmitter, once initialized, runs until the end-of-frame (EOF) condition is detected unless a collision within the collision window occurs (half-duplex mode) or an abort condition is encountered.

If the user has a frame ready to transmit, setting the DMACTRL[TOD] eliminates waiting for the next poll and a DMA transfer of the transmit data buffers can begin immediately. The transmission begins once all data for the frame is loaded into the Tx FIFO or sufficient transmit data (determined by the Tx FIFO threshold register) is in the Tx FIFO. If the line is not busy, the MAC transmit logic asserts TX_EN and sends the 7-octet preamble sequence, 1-octet start of frame delimiter, and frame information in that order. If the line is busy, the controller waits for the carrier sense signal, CRS, to remain inactive for 60 bit times (60 clocks) and transmission begins after an additional 36 bit times (96 bit times after CRS became active). In full-duplex mode, because collisions are ignored, frame transmission maintains only the interframe gap (96 bit times) regardless of CRS.

In half-duplex mode (MACCFG2[Full Duplex] is cleared) the MAC defers transmission if the line is busy (CRS asserted). Before transmitting, the MAC waits for carrier sense to become inactive, at which point it then determines if CRS remains negated for 60 clocks. If so, transmission begins after an additional 36 bit times (96 bit times after CRS originally became negated). If CRS continues to be asserted, the MAC follows a specified back-off procedure and tries to retransmit the frame until the retry limit is reached. Data stored in the Tx FIFO is re-transmitted in case of a collision. This avoids unnecessary memory traffic.

The transmitter also monitors for an abort condition and terminates the current frame if an abort condition is encountered. In full-duplex mode the protocol is independent of network activity, and only the transmit inter-frame gap must be enforced.

The transmitter implements full-duplex flow control. If a flow control frame is received, the MAC does not service the transmitter's request to send data until the pause duration is over. If the MAC is currently sending data after a pause frame has been received and processed, the MAC finishes sending the current frame, then suspends subsequent frames (except a pause frame) until the pause duration is over. In addition, the transmitter supports transmission of flow control frames through TCTRL[TFC_PAUSE]. The transmit pause frame is generated internally based on the PAUSE register that defines the pause value to be sent. Note that it is possible to send a pause frame while the pause timer has not expired.

The MAC automatically appends FCS (32-bit CRC) bytes to the frame if any of the following values are set:

- TxBD[PAD/CRC] is set in first TxBD
- TxBD[TC] is set in first TxBD

- MACCFG2[PAD/CRC] is set
- MACCFG2[CRC] is set

The Tx_EN is negated after the FCS is sent. This notifies the PHY of the need to generate the illegal Manchester encoding that signifies the end of an Ethernet frame. Following the transmission of the FCS, the Ethernet controller writes the frame status bits into the BD and clears TxBD[R]. If the end of the current buffer is reached and TxBD[L] is cleared (a frame is comprised of multiple buffer descriptors), only TxBD[R] is cleared.

For both half- and full-duplex modes, an interrupt can be issued depending on TxBD[I]. The Ethernet controller then proceeds to the next TxBD in the table. In this way, the core can be interrupted after each frame, after each buffer, or after a specific buffer is sent. If TxBD[PAD/CRC] is set, the Ethernet controller pads any frame shorter than 64 bytes with zero bytes to make up the minimum length.

To pause transmission, or rearrange the transmit queue, set DMACTRL[GTS]. This can be useful for transmitting expedited data ahead of previously-linked buffers or for error situations. If this bit is set, the eTSEC transmitter performs a graceful transmit stop. The Ethernet controller stops immediately if no transmission is in progress or continues transmission until all queued frames in the Tx FIFO have been disposed of. The IEVENT[GTSC] interrupt occurs once the graceful transmit stop operation is completed. After the DMACTRL[GTS] is cleared, the eTSEC resumes transmission with the next frame.

While the eTSEC is in 10/100Mbps mode it sends bytes least-significant nibble first and each nibble is sent lsb first. While it is in 1000Mbps mode it sends bytes LSB first.

14.6.3.4 Gigabit Ethernet frame reception

The eTSEC Ethernet receiver is designed to work with little core intervention and can perform data extraction, address recognition, CRC checking, short frame checking, and maximum frame-length checking.

After a hardware reset, the software driver clears the RSTAT register and sets MACCFG1[RX_EN]. The Ethernet receiver is enabled and immediately starts processing receive frames. The MAC checks for when TSEC_n_RX_DV is asserted and as long as TSEC_n_COL remains negated (full-duplex mode ignores TSEC_n_COL), the MAC looks for the start of a frame by searching for a valid preamble/SFD (start of frame delimiter) header, which is stripped (unless MACCFG2[PreAM RxEN] is set) and the frame begins to be processed. If a valid header is not found, the frame is ignored.

If the receiver detects the first bytes of a frame, the eTSEC controller begins to perform the frame recognition function through destination address (DA) recognition (see [Frame recognition](#)). Based on this match the frame can be accepted or rejected. The receiver can filter frames based on individual (unicast), group (multicast), and broadcast addresses. Because Ethernet receive frame data is not written to memory until the internal frame recognition algorithm is complete, system bus usage is not wasted on frames unwanted by this station.

If a frame is accepted, the Ethernet controller fetches the receive buffer descriptor (RxBD) from either queue 0 or the queue determined by the filter. If the RxBD is not being used by software (RxBD[E] is set), the eTSEC starts transferring the incoming frame. RxBD[F] is set for the first RxBD used for any particular receive frame. If the current RxBD is not available for the received frame, a receive busy error condition is raised in IEVENT[BSY].

After the buffer is filled, the eTSEC clears RxBD[E] and, if RxBD[I] is set, generates an interrupt. If the incoming frame is larger than the buffer, the Ethernet controller fetches the next RxBD in the table. If it is empty, the controller continues receiving the rest of the frame. In half-duplex mode, if a collision is detected during the frame, no RxBDs are used; thus, no collision frames are presented to the user except late collisions, which indicate LAN problems.

The RxBD length is determined by the MRBL field in the maximum receive buffer length register (MRBLR). The smallest valid value is 64 bytes, with larger values being some integral multiple of 64 bytes. During reception, the Ethernet controller checks for frames that are too short or too long. After the frame ends (CRS is negated), the receive CRC field is checked and written to the data buffer. The data length written to the last RxBD in the Ethernet frame is the length of the entire frame, which enables the software to recognize an oversized frame condition.

Receive frames are not truncated when they exceed maximum frame bytes in the MAC's maximum frame register if MACCFG2[Huge Frame] is set, yet the babbling receiver error interrupt occurs (IEVENT[BABR] is set) and RxBD[LG] is set.

After the receive frame is complete, the Ethernet controller sets RxBD[L], updates the frame status bits in the RxBD, and clears RxBD[E]. If RxBD[I] is set, the Ethernet controller next generates an interrupt (that can be masked) indicating that a frame was received and is in memory. The Ethernet controller then waits for a new frame.

To interrupt reception or rearrange the receive queue, DMACTRL[GRS] must be set. If this bit is set, the eTSEC receiver performs a graceful receive stop. The Ethernet controller stops immediately if no frames are being received or continues receiving until the current frame either finishes or an error condition occurs. The IEVENT[GRSC] interrupt event is signaled after the graceful receive stop operation is completed. While in

this mode the user can write to registers that are accessible to both the user and the eTSEC hardware without fear of conflict, and finally clear IEVENT[GRSC]. After DMACTRL[GRS] is cleared, the eTSEC scans the input data stream for the start of a new frame (preamble sequence and start of frame delimiter), it resumes receiving, and the first valid frame received is placed in the next available RxBD.

14.6.3.5 Ethernet preamble customization

By default eTSEC generates a standard Ethernet preamble sequence prior to transmitting frames.

However, the user can substitute a custom preamble sequence for the purpose of controlling switching equipment at the receiver, particularly at 100/1000 Mbps speeds. In any RMII mode only the standard preamble can be transmitted.

eTSEC normally searches for and discards the standard Ethernet preamble sequence upon receiving frames. Part of the received preamble sequence can be optionally recovered and returned as part of the frame data, making it visible to user software. Note however, that preamble cannot be recovered in any RMII mode.

Note that it is also possible for the first two bytes of custom preamble (PreOct0 and PreOct1) to be lost in during conversion to ten-bit code groups in the PCS sub-layer. Thus is it recommended that any custom preamble start at PreOct2.

14.6.3.5.1 User-defined preamble transmission

To substitute a custom preamble, the user must ensure that:

- MACCFG2[PreAm TxEN] bit is set
- The first TxBD of every frame containing a custom preamble has its PRE bit set
- An 8-byte custom preamble sequence appears before the Ethernet DA field in the first transmit data buffer

The definition of the 8-byte custom preamble sequence is shown below.

Table 14-893. Definition of custom preamble sequence

Byte Offsets	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0-1	PreOct0							PreOct1								
2-3	PreOct2							PreOct3								
4-5	PreOct4							PreOct5								
6-7	PreOct6															

functional description

The fields of the custom preamble sequence are described in the table below. It should be noted that use of preamble octets matching the standard start of frame delimiter (0xD5) can be expected to trigger premature frame reception by the receiving station.

Table 14-894. Custom preamble field descriptions

Bytes	Bits	Name	Description
0-1	0-7	PreOct0	Octet #0 of custom transmit preamble. This is the first octet of preamble sent.
	8-15	PreOct1	Octet #1 of custom transmit preamble. This is the second octet of preamble sent.
2-3	0-7	PreOct2	Octet #2 of custom transmit preamble. This is the third octet of preamble sent.
	8-15	PreOct3	Octet #3 of custom transmit preamble. This is the fourth octet of preamble sent.
4-5	0-7	PreOct4	Octet #4 of custom transmit preamble. This is the fifth octet of preamble sent.
	8-15	PreOct5	Octet #5 of custom transmit preamble. This is the sixth octet of preamble sent.
6-7	0-7	PreOct6	Octet #6 of custom transmit preamble. This is the seventh octet of preamble sent. The last octet (the start of frame delimiter) is generated by the MAC automatically.
	8-15	-	Reserved; should be cleared.

14.6.3.5.2 User-visible preamble reception

To return the received preamble, the user must ensure that:

- MACCFG2[PreAm RxEN] bit is set
- Space for an 8-byte preamble sequence is allowed before the Ethernet DA field in the first receive data buffer of each frame

The definition of the 8-byte received preamble sequence is shown below.

Table 14-895. Definition of received preamble sequence

Byte Offsets	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0-1	PreOct0							PreOct1								
2-3	PreOct2							PreOct3								
4-5	PreOct4							PreOct5								
6-7	PreOct6															

The fields of the received preamble sequence are described in the table below. Should the received preamble be shorter than the 7-octet sequence defined by IEEE Std. 802.3, initial bytes of the received preamble sequence hold undefined values. The standard start of frame delimiter (0xD5) is always omitted. Note that preamble extraction is not possible in RMII mode.

Table 14-896. Received preamble field descriptions

Bytes	Bits	Name	Description
0-1	0-7	PreOct0	Octet #0 of received preamble. This is the first octet of preamble received.
	8-15	PreOct1	Octet #1 of received preamble. This is the second octet of preamble received.
2-3	0-7	PreOct2	Octet #2 of received preamble. This is the third octet of preamble received.
	8-15	PreOct3	Octet #3 of received preamble. This is the fourth octet of preamble received.
4-5	0-7	PreOct4	Octet #4 of received preamble. This is the fifth octet of preamble received.
	8-15	PreOct5	Octet #5 of received preamble. This is the sixth octet of preamble received.
6-7	0-7	PreOct6	Octet #6 of received preamble. This is the seventh octet of preamble received. The last octet (the start of frame delimiter) is discarded.
	8-15	-	Reserved

14.6.3.6 RMON support

Using promiscuous mode, the eTSEC can automatically gather network statistics required for remote network interface monitoring.

The RMON MIB group 1, RMON MIB group 2, RMON MIB group 3, RMON MIB group 9, RMON MIB2, and the IEEE 802.3 Ethernet MIB are supported. For RMON statistics and their corresponding counters, see the memory map.

14.6.3.7 Frame recognition

The Ethernet controller performs frame recognition using destination address (DA) recognition.

A frame can be rejected or accepted based on the outcome.

14.6.3.7.1 Destination address recognition and frame filtering

The eTSEC can perform layer 2 frame filtering on the basis of destination Ethernet address (DA), as illustrated by the flowchart in the figure below.

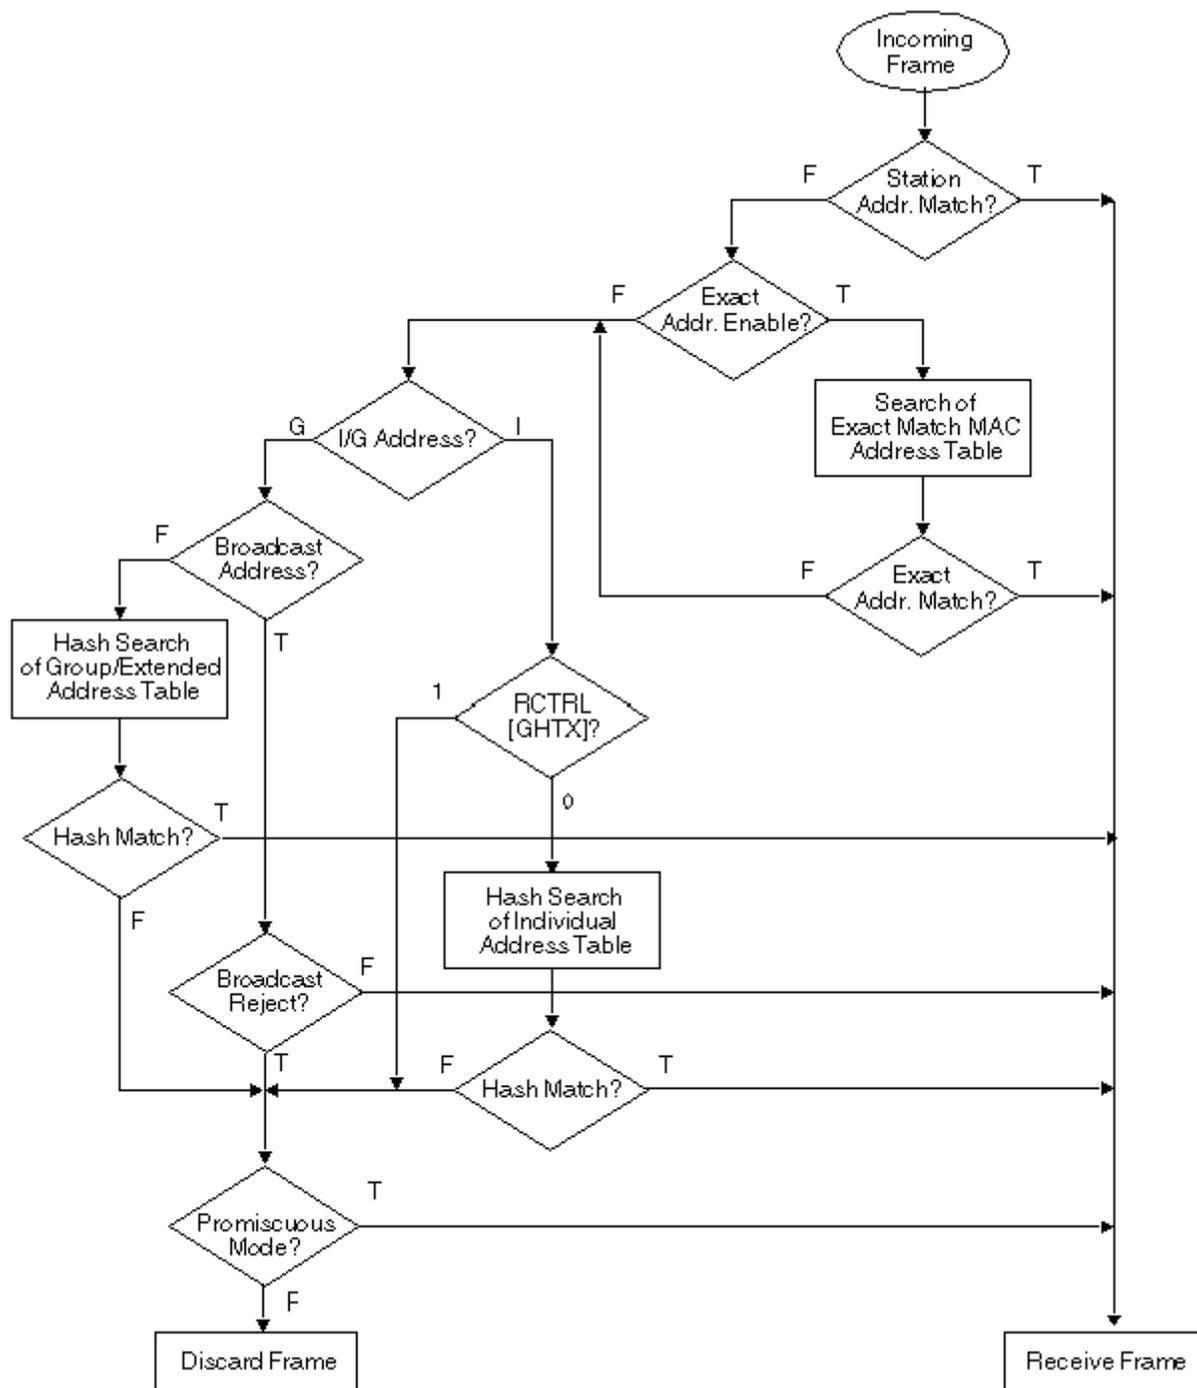


Figure 14-879. Ethernet address recognition flowchart

In promiscuous mode, the eTSEC accepts all received frames regardless of DA. Note, however, that Ethernet frame filtering simply restricts the traffic seen by the receive queue filer. Therefore even in promiscuous mode it remains possible to program the filer to reject frames based on their higher-layer header contents.

In the case of an individual address, the DA field of the received frame is compared with the physical address that the user programs in the station address registers (MACSTNADDR1 and MACSTNADDR2). If the DA does not match the station address, and exact MAC address matching is enabled through RCTRL[EMEN], the controller performs address recognition on the multiple MAC addresses written to the MACxADDR1 and MACxADDR2 registers. These virtual addresses give a particular eTSEC the ability to mirror other MACs on the network, which caters for router redundancy protocols, such as HSRP and VRRP.

If exact MAC address matching is not enabled, the eTSEC determines whether DA is a group or individual address. If DA is the standard broadcast address, and broadcast addresses are not rejected, the frame is accepted. If any other group address is received, the eTSEC looks-up the DA by means of the group hash table. The group hash table may be extended to 512 entries if RCTRL[GHTX] = 1. Otherwise, an individual address is hashed into the 256-entry individual hash table when RCTRL[GHTX] = 0.

14.6.3.7.2 Hash table algorithm

The hash table process used in the group hash filtering operates as follows.

By default, the Ethernet controller maps any 48-bit destination address into one of 256 bins, represented by the 256 bits in IGADDR0-IGADDR7 for individual addresses, and the 256 bits in GADDR0-GADDR7 for group addresses. But in the case where RCTRL[GHTX] is set, both sets of registers are combined into an extended group-only hash table of 512 bits, where IGADDR0-IGADDR7 contain the first 256 bits and GADDR0-GADDR7 contain the last 256 bits. No individual-address table exists in extended mode.

The 48-bit destination address received by the MAC is passed through the Ethernet CRC-32 algorithm to produce a hash value. The CRC polynomial used is:

$$x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

The MAC initializes its CRC register to 0xFFFFFFFF before computing a CRC on the 6 bit-reversed octets of the DA. A non-optimized sample of C code for computing the DA hash is listed below. The 9 most significant bits of the raw, uninverted CRC are used as the hash table index, H[8:0]. If RCTRL[GHTX] = 0, bits H[8:6] select one of the 8 IGADDR or GADDR registers, while bits H[5:1] select a bit within the 32-bit register. If RCTRL[GHTX] = 1, bits H[8:5] select one of the 16 registers in the {IGADDR, GADDR} set, while bits H[4:0] select a bit within the 32-bit register. For example, if H[8:5] = 7, IGADDR7 is selected, whereas H[8:5] = 9 selects GADDR1.

Sample C Code for Computing eTSEC Hash Table Indices:

functional description

```

/* Wrapper macros for 256-bucket and 512-bucket hash tables:
   Pass 6-byte Ethernet MAC address as parameter. */
#define TSEC_HASH256(macaddr) ((crc32(macaddr) >> 24) & 0xff)
#define TSEC_HASH512(macaddr) ((crc32(macaddr) >> 23) & 0x1ff)
/* CRC constants. Note: CRC-32 polynomial is bit-reversed. */
#define CRC_POLYNOMIAL 0xedb88320
#define CRC_INITIAL    0xffffffff
#define MAC_ADDRLEN    6
#define BITS_PER_BYTE  8
/* crc32() Takes the array of bytes, macaddr[], representing an
   Ethernet MAC address and returns the CRC-32 result over these bytes,
   where each byte is used in bit-reversed form (Ethernet bit order).
   Index 0 of macaddr[] is the first byte of the address on the wire.
   Test case: the result of crc32 on {0x00, 0x01, 0x02, 0x03, 0x04, 0x05}
   should be 0xad0c28f3.
*/
unsigned long crc32(unsigned char macaddr[MAC_ADDRLEN])
{
    unsigned long crc, result;
    int byte, i;
    /* CRC-32 algorithm starts by inverting first 4 bytes */
    crc = CRC_INITIAL;
    /* add each byte to running CRC accumulator */
    for (byte = 0; byte < MAC_ADDRLEN; ++byte) {
        crc ^= macaddr[byte];
        /* shift CRC right to perform but reversal on byte of address */
        for (i = 0; i < BITS_PER_BYTE; ++i)
            if (crc & 1)
                crc = (crc >> 1) ^ CRC_POLYNOMIAL;
            else
                crc >>= 1;
    }
    /* finally, reverse bits of result to get CRC in normal bit order */
    for (result = 0, i = 4*BITS_PER_BYTE-1; i >= 0; crc >>= 1, --i)
        result |= (crc & 1) << i;
    return result;
}

```

If the CRC hash table index selects a bit that is set in the hash table, the frame is accepted. If 32 group addresses are stored in the hash table and random group addresses are received, the extended hash table prevents roughly 480/512 (93.8%) of the group address frames from reaching memory. Software must further filter those that reach memory to determine if they contain the correct addresses. Alternatively, small multicast groups can be held in the exact match MAC address registers, which guarantees that only correct frames are admitted.

The effectiveness of the hash table declines as the number of addresses increases. For instance, as the number of addresses stored in the 512-bin hash table increases, the vast majority of the hash table bits are set, preventing only a small fraction of frames from reaching memory.

NOTE

The hash table cannot be used to reject frames that match a set of selected addresses because unintended addresses can map to the same bit in the hash table. The receive queue filer may be used to reject frames with unintended address hits in the hash table.

14.6.3.8 Magic Packet mode

eTSEC implements the AMD Magic Packet specification for LAN-initiated power management.

This mode is normally entered with the rest of the system in a low-power mode. Software must enable normal receive function in the Ethernet MAC, and then finally set the MACCFG2[MPEN] bit to enable Magic Packet detection before the system enters a reduced mode. While the rest of the system is operating in low-power mode, the enabled eTSEC continues to receive Ethernet frames, but discards them immediately. Upon receipt of any frame whose contents contain the valid Magic Packet sequence, the eTSEC exits out of Magic Packet mode, thus clearing MACCFG2[MPEN], and raises an error/diagnostic interrupt through IEVENT[MAG], which causes the surrounding system to wake-up. Frames received after Magic Packet mode has exited are received into software buffers as usual. Software can abort Magic Packet mode by writing 0 to MACCFG2[MPEN] at any time.

AMD specify a Magic Packet to be any Ethernet frame containing a valid Ethernet header (Destination and Source Addresses) and valid FCS (CRC-32), and whose payload includes the specific Magic Packet byte sequence at any offset from the start of frame. The specific byte sequence comprises an unbroken stream of 102 bytes, the first 6 bytes of which are 0xFFFFFFFF_FFFFFFFF, followed by 16 copies of the MAC's unique IEEE station address in the normal byte order for Ethernet addresses. For example, if the station address were 0x112233_445566, then the MAC would have to receive 0xFFFFFFFF_FFFFFFFF, 0x112233_445566, ..., 0x112233_445566 in any payload to detect a Magic Packet. Only frames addressed specifically to the MAC's station address or a valid multicast or broadcast address can be examined for the Magic Packet sequence.

14.6.3.9 Flow control

Because collisions cannot occur in full-duplex mode, gigabit Ethernet can operate at the maximum rate.

If the rate becomes too fast for a station's receiver, the station's transmitter can send flow-control frames to reduce the rate. Flow-control instructions are transferred by special frames of minimum frame size. The length/type fields of these frames have a special value.

The table below lists the flow-control frame structure.

Table 14-897. Flow control frame structure

Size [Octets]	Description	Value	Comment
7	Preamble		-
1	SFD		Start frame delimiter
6	Destination address	01-80-C2-00-00-01	Multicast address reserved for use in MAC frames (or MAC station address)
6	Source address		-
2	Length/type	88-08	Control frame type
2	MAC opcode	00-01	Pause command
2	MAC parameter		Pause time as defined by the PTV[PT] field. The pause period is measured in pause_quanta, a speed independent constant of 512 bit-times (unlike slot time). The most-significant octet is transmitted first.
2	Extended MAC parameter		Pause time extended as defined by the PTV[PTE] field. The most significant octet is transmitted first.
40	Reserved	-	-
4	FCS		Frame check sequence (CRC)

If flow-control mode is enabled (MACCFG1[Rx_Flow] is set) and the receiver identifies a pause-flow control frame, transmission stops for the time specified in the control frame. The controller completes any frame in progress before stopping transmission and does not commence counting the pause time until transmit is idle. During a pause, only a control frame can be sent (TCTRL[TFC_PAUSE] is set). Normal transmission resumes after the pause timer stops counting, or resumes immediately if a pause frame with a zero time-out is received. If another pause-control frame is received during the pause, the period changes to the new value received.

14.6.3.10 Interrupt handling

The following describes what usually occurs within a eTSEC interrupt handler:

- If an interrupt occurs, read IEVENT to determine interrupt sources. IEVENT bits to be handled in this interrupt handler are normally cleared at this time. There are three kinds of interrupts:
 - Receive data frame interrupts, when bits RXB or RXF in IEVENT are set
 - Transmit data frame interrupts, when bits TXB or TXF in IEVENT are set
 - Error, diagnostic, and special interrupts (all bits in IEVENT other than RXB, RXF, TXB, or TXF)
- Process the TxBDs to reuse them if the IEVENT[TXB, TXF or TXE] were set. Consult register bits TSTAT[TXF0-TXF7] to determine which TxBD rings gave rise to the transmit interrupt in the case of TXF. If the transmit speed is fast or the

interrupt delay is long, more than one transmit buffer may have been sent by the eTSEC; thus, it is important to check more than just one TxBD during the interrupt handler. One common practice is to process all TxBDs in the interrupt handler until one is found with R set.

- Obtain data from RxBD rings if IEVENT[RXC, RXB or RXF] is set. Consult register bits RSTAT[RXF0-RXF7] to determine which RxBD rings gave rise to the receive interrupt in the case of RXF. If the receive speed is fast or the interrupt delay is long, the eTSEC may have received more than one RxBD; thus, it is important to check more than just one RxBD during interrupt handling. Typically, all RxBDs in the interrupt handler are processed until one is found with E set. Because the eTSEC pre-fetches BDs, the BD table must be big enough so that there is always another empty BD to pre-fetch, otherwise a BSY error occurs.
- Clear any set halt or frame interrupt bits in TSTAT and RSTAT registers, or DMACTRL[GTS] and DMACTRL[GRS] by writing 1s to these bits.
- Continue normal execution.

Table 14-898. Non-error transmit interrupts

Interrupt	Description	Action Taken by the eTSEC
GTSC	Graceful transmit stop complete: transmitter is put into a pause state after completion of the frame currently being transmitted.	None
TXC	Transmit control: Instead of the next transmit frame, a control frame was sent.	None
TXB	Transmit buffer: A transmit buffer descriptor, that is not the last one in the frame, was updated in one of the enabled TxBD rings.	Programmable 'write with response' TxBD to memory before setting IEVENT[TXB].
TXF	Transmit frame: A frame from an enabled TxBD ring was transmitted and the last transmit buffer descriptor (TxBD) of that frame was updated.	Programmable 'write with response' to memory on the last TxBD before setting IEVENT[TXF].

Table 14-899. Non-error receive interrupts

Interrupt	Description	Action Taken by the eTSEC
GRSC	Graceful receive stop complete: Receiver is put into a pause state after completion of the frame currently being received.	None
RXC	Receive control: A control frame was received. As soon as the transmitter finishes sending the current frame, a pause operation is performed.	None
RXB	Receive buffer: A receive buffer descriptor, that is not the last one of the frame, was updated in one of the enabled RxBD rings.	Programmable 'write with response' RxBD to memory before setting IEVENT[RXB].
RXF	Receive frame: A frame was received to an enabled RxBD ring and the last receive buffer descriptor (RxBD) of that frame was updated.	Programmable 'write with response' to memory on the last RxBD before setting IEVENT[RXF].

14.6.3.10.1 Interrupt coalescing

Interrupt coalescing offers the user the ability to contour the behavior of the eTSEC with regard to frame interrupts. Separate but identical mechanisms exist for both transmitted frames and received frames. In either case, frame interrupts require that software set the I-bit in RxBDs or TxBDs, and disable buffer interrupts (IEVENT[RXB] or IEVENT[TXB]). Particular rings can remain free of interrupts by ensuring that the I-bit is consistently cleared in all BDs. While interrupt coalescing is enabled, a transmit or receive frame interrupt is raised either when a counter threshold-defined number of frames is received/transmitted or the timer threshold-defined period of time has elapsed, whichever occurs first. Disabling and then re-enabling interrupt coalescing forces reset of the coalescing timers and counters to reflect changes made to the threshold registers.

14.6.3.10.2 Interrupt coalescing by frame count threshold

To avoid interrupt bandwidth congestion due to frequent, consecutive interrupts, the user may enable and configure interrupt coalescing to deliberately group frame interrupts, reducing the total number of interrupts raised. The number of frames received or transmitted prior to an interrupt being raised is determined by the frame threshold field (ICFT) in the appropriate interrupt coalescing configuration register (RXIC or TXIC). The frame threshold field may be assigned a value between 1 and 255. The internal transmit or receive frame counter decrements from this initial value each time a frame is transmitted or received. Upon reaching zero, an interrupt is raised, the appropriate threshold counter is reset to the value in the ICFT field, and then eTSEC continues counting frames while the interrupt is active. The appropriate threshold counter is also reset to the value in the ICFT field if an interrupt is raised subject to the corresponding threshold timer.

14.6.3.10.3 Interrupt coalescing by timer threshold

To avoid stale frame interrupts, the user may also assign a timer threshold, beyond which any frame interrupts not yet raised are forced.

The timer threshold fields of the receive and transmit interrupt coalescing configuration registers (RXIC[ICTT] and TXIC[ICTT]) are defined in units equivalent to 64 interface clocks or system clocks, depending on the setting of the ICCS field in RXIC and TXIC.

After transmitting a frame, the transmit interrupt coalescing threshold time begins counting down from the value in TXIC[ICTT]. An interrupt is raised when the counter reaches zero. In the event of graceful transmit stop completion before the coalescing timer expires, the eTSEC issues two interrupts, the first for GTS, the second for TXF (due to timer expiration of a pending event). To prevent the second interrupt from affecting servicing of the GTS event, it is recommended that the user mask out the TXF

event during execution of the service routine. After receiving a frame, the receive interrupt coalescing threshold time begins counting down from the value in RXIC[ICTT]. An interrupt is raised when the counter reaches zero. In the event of graceful receive stop completion before the coalescing timer expires, the eTSEC issues two interrupts, the first for GRS, the second for RXF (due to timer expiration of a pending event). To prevent the second interrupt from affecting servicing of the GRS event, it is recommended that the user mask out the RXF event during execution of the service routine.

The interrupt coalescing timer thresholds (transmit and receive, operating independently) may be values ranging from 0x0001 to 0xFFFF. The table below specifies the range of possible timing thresholds subject to timer clock source, the interface or system frequency, and the value of the RXIC[ICTT] or TXIC[ICTT] field.

Table 14-900. Interrupt coalescing timing threshold ranges

ICCS (Clock Source)	eTSEC Interface Format and Frequency or eTSEC System Frequency	Interrupt Coalescing Threshold Time	
		Minimum (ICTT = 0x0001)	Maximum (ICTT = 0xFFFF)
0 (I/F clock)	10Base-T at 2.5 MHz	25.6 μ s	1.68 s
0 (I/F clock)	100Base-T at 25 MHz	2.56 μ s	168 ms
0 (I/F clock)	1000Base-T at 125 MHz	0.51 μ s	33.6 ms
1 (sys. clock)	eTSEC operating at 266 MHz	0.24 μ s	15.7 ms
1 (sys. clock)	eTSEC operating at 333 MHz	0.19 μ s	12.6 ms

The transmit timer threshold counter is reset to the value in TXIC[ICTT] and begins counting down on transmission of the frame following an interrupt.

The receive timer threshold counter is reset to the value in RXIC[ICTT] and begins counting down on receiving the frame following an interrupt.

14.6.3.11 Inter-frame gap time

If a station must transmit, it waits until the LAN becomes silent for a specified period (inter-frame gap, or IFG).

The minimum inter-packet gap (IPG) time for back-to-back transmission is set by IPGIFG[Back-to-Back Inter-Packet-Gap]. The receiver receives back-to-back frames with the minimum interframe gap (IFG) as set in IPGIFG[Minimum IFG Enforcement]. If multiple frames are ready to transmit, the Ethernet controller follows the minimum IPG as long as the following restrictions are met:

- The first TxBD pointer, TBPTR_n, of any given frame is located at a 16-byte aligned address.
- All BDs for any multiple-BD frame reside in the same cache line.

- TCP/UDP and IP Checksum generation are disabled in each frame's TxFCB, or in TCTRL, or frames are limited to 1200 bytes in length.
- Each TxBD[Data Length] is greater-than or equal to 64 bytes.

If the first TxBD alignment restriction is not met, the back-to-back IPG may be as many as 32 cycles. If the TxBD size restriction is not met, the back-to-back IPG may be significantly longer.

In half-duplex mode, after a station begins sending, it continually checks for collisions on the LAN. If a collision is detected, the station forces a jam signal (all ones) on its frame and stops transmitting. Collisions usually occur close to the beginning of a packet. The station then waits a random time period (back-off) before attempting to send again. After the back-off completes, the station waits for silence on the LAN (carrier sense negated) and then begins retransmission (retry) on the LAN. Retransmission begins 36 bit times after carrier sense is negated for at least 60 bit times. If the frame is not successfully sent within a specified number of retries, an error is indicated (collision retry limit exceeded).

14.6.3.12 Internal and external loop back

Setting MACCFG1[Loop Back] causes the MAC transmit outputs to be looped back to the MAC receive inputs.

Clearing this bit results in normal operation. This bit is cleared by default. Clearing this bit results in normal operation.

14.6.3.13 Error-handling procedure

The eTSEC reports frame reception and transmission error conditions using the channel BDs, the error counters, and the IEVENT register.

Transmission errors are described in the table below.

Table 14-901. Transmission errors

Error	Response
Transmitter underrun	Transmitter underrun can occur either after frame transmission has commenced, or in response to an incomplete sequence of TxBDs. In the former case, the controller sends 32 bits that ensure a CRC error, and terminates buffer transmission. In the latter case, the relevant transmit queue is halted. In all cases, the eTSEC closes the buffer, sets TxBD[UN], IEVENT[XFUN], and IEVENT[TXE]. The controller resumes transmission after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared).

Table continues on the next page...

Table 14-901. Transmission errors (continued)

Error	Response
Retransmission attempts limit expired	The controller terminates buffer transmission, sets TxBD[RL], closes the buffer, IEVENT[CRL], and IEVENT[TXE]. Transmission resumes after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared).
Late collision	The controller terminates buffer transmission, sets TxBD[LC], closes the buffer, IEVENT[LC], and IEVENT[TXE]. The controller resumes transmission after TSTAT[THLT] is cleared (and DMACTRL[GTS] is cleared).
Memory read error	A system bus error occurred during a DMA transaction. The controller sets IEVENT[EBERR], DMA stops sending data to the FIFO which causes an underrun error, and therefore TxBD[UN] is set, but IEVENT[XFUN] is not set. The TSTAT[THLT] is set. Transmits are continued once TSTAT[THLT] is cleared.
Data parity error	Data in the transmit FIFO was potentially corrupted. The controller sets IEVENT[DPE], but otherwise continues transmission until halted explicitly.
Babbling transmit error	A frame is transmitted which exceeds the MAC's Maximum Frame Length and MACCFG2[Huge Frame] is a 0. The controller sets IEVENT[BABT] and continues without interruption.

Reception errors are described in the table below.

Table 14-902. Reception errors

Error	Description
Overrun error	The Ethernet controller maintains an internal FIFO buffer for receiving data. If a receiver FIFO buffer overrun occurs, the controller sets RxB[OV], sets RxB[L], closes the buffer, increments the discarded frame counter (RDRP), and sets IEVENT[RXF]. The receiver then enters hunt mode (seeking start of a new frame).
Busy error	A frame is received and discarded due to a lack of buffers. The controller sets IEVENT[BSY] and increments the discarded frame counter (RDRP). In addition, the RSTAT[QHLT _n] bit is set. RDRP increments for each frame that is received while the receiver is halted due to a busy condition. The halted queue resumes reception once the RSTAT[QHLT _n] bit is cleared.
Filed frame to invalid queue error	A frame is received and discarded as a result of the filer directing it to an RxB ring that is currently not enabled. The controller sets IEVENT[FIQ] and increments the discarded frame counter (RDRP).
Parser error	<p>If the receive frame parser is enabled, a parse error can be flagged as a result of inconsistencies discovered between fields of the embedded packet headers. For example, the L2 header may indicate an IPv4 header, but the IP version number fails to match. In the event of a parse error, parsing is terminated at the inconsistent header, and the RxFCB[PERR] field indicates at which layer of the protocol stack the error was discovered. Receiver function continues regardless of parse errors, but IEVENT[PERR] is set. The receive queue filer may operate with reduced or default information in some cases; therefore, filer rule sets should be constructed so as to be tolerant of misformed frames.</p> <p>NOTE: Any values in the length/type field between 1500 and 1536 is treated as a length, however, only illegal packets exist with this length/type since these are not valid lengths and not valid types. These are treated by the MAC logic as out of range.</p> <p>Software must confirm the parser and filer results by checking the type/length field after the packet has been written to memory to see if it falls in this range.</p>
Non-octet error (dribbling bits)	The Ethernet controller handles a nibble of dribbling bits if the receive frame terminates as non-octet aligned and it checks the CRC of the frame on the last octet boundary. If there is a CRC error, the frame non-octet aligned (RxB[NO]) error is reported, IEVENT[RXF] is set, and the alignment error counter increments. The eTSEC relies on the statistics collector block to increment the receive alignment error counter (RALN). If there is no CRC error, no error is reported.

Table continues on the next page...

Table 14-902. Reception errors (continued)

Error	Description
CRC error	If a CRC error occurs, the controller sets RxB[CR], closes the buffer, and sets IEVENT[RXF]. This eTSEC relies on the statistics collector block to record the event. After receiving a frame with a CRC error, the receiver then enters hunt mode.
Memory read error	A system bus error occurred during a DMA transaction. The controller sets IEVENT[EBERR] and discards the frame and increments the discarded frame counter (RDRP). In addition the RSTAT[QHLT n] bit is set. The halted queue resumes reception once the RSTAT[QHLT n] bit is cleared.
Data parity error	Data in the receive FIFO or filter table was potentially corrupted. The controller sets IEVENT[DPE], but otherwise continues reception until halted explicitly.
Babbling receive error	A frame is received that exceeds the MAC's maximum frame length. The controller sets IEVENT[BABR] and continues.

14.6.4 TCP/IP off-load

Each eTSEC provides hardware support for accelerating the basic functions of TCP/IP packet transmission and reception.

By default, these features are disabled and must be explicitly enabled through RCTRL and TCTRL. In this configuration, the eTSEC processes frames as vanilla Ethernet frames and none of the multi-ring QoS/CoS receive services or per-frame VLAN insertion and deletion are available. Operate eTSEC in this default configuration when using existing TCP/IP stack software that has not been modified to take advantage of TOE.

TOE can be enabled independently for Rx and Tx and at various levels. Receive TOE functions are controlled by RCTRL and transmit functions through a combination of TCTRL[TUCSEN] and the Tx frame control block.

On receive, according to RCTRL[PRSDEP], eTSEC can parse frames at layer 2 of the stack only (Ethernet headers and switching headers), layers 2 to 3 (including IPv4 or IPv6), or layers 2 to 4 (including TCP and UDP). TOE provides protocol header recognition, header verification (IPv4 header checksum verification), and TCP/UDP payload checksum verification including verification of associated pseudo-header checksums. For large frames off-load of checksum verification saves a significant fraction of the CPU cycles that would otherwise be spent by the TCP/IP stack. IP packet fragmentation and re-assembly, and TCP stream establishment and tear-down are not performed in hardware. The frame parser sets RQFPR[IPF] status flag encountering a fragmented frame. The frame parser in eTSEC searches a maximum of 512 bytes from the start of a received frame when attempting to locate headers; headers deeper than 512 bytes are assumed not to exist, and any associated receive status flags in the frame control block remain cleared.

On transmit, TOE provides IPv4 and TCP/UDP header checksum generation. Like receive TOE, checksum generation reduces CPU load significantly for TCP/IP stacks modified to exploit eTSEC TOE functions. The eTSEC does not checksum transmitted packets with IPv6 routing headers or calculate TCP/UDP checksums from IP fragments. If a transmitted TCP segment requires checksum generation but IPv6 extension headers would prevent eTSEC from calculating the pseudo-header checksum, software can calculate just the pseudo-header checksum in advance and supply it to the eTSEC as part of per-frame TOE configuration.

14.6.4.1 Frame control blocks

Frame control blocks (FCBs) are 8-byte blocks of TOE control and/or status data that are passed between software (driver and TCP/IP stack) and each eTSEC.

A FCB always precedes the frame it applies to, and is present only when TOE functions are being used. As the figure below shows, the first BD of each frame points to the initial data buffer and the FCB. The initial data buffer must be at least 8 bytes long to contain the FCB without breaking it. Custom or received Ethernet preamble sequences also follow the FCB if preambles are visible.

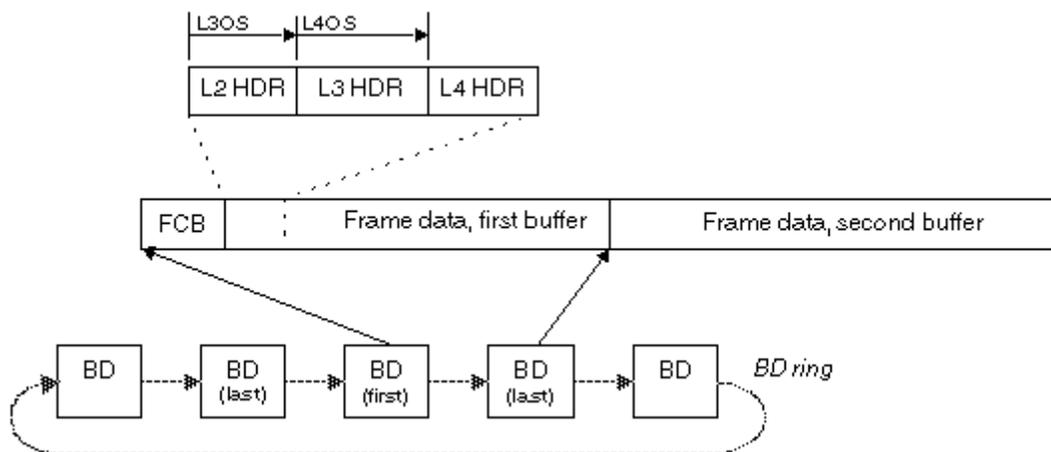


Figure 14-880. Location of frame control blocks for TOE parameters

For TxBD rings, FCBs are assumed present when the TxBD[TOE/UN] bit is set by user software. The eTSEC ignores the TxBD[TOE/UN] bit in all BDs other than those pointing to initial data buffers, therefore FCBs must not be inserted in second and subsequent data buffers. Since TxBD[TOE/UN] can be set under software discretion, TOE acceleration for transmit may be applied on a frame-by-frame basis.

functional description

In the case of RxBD rings, FCBs are inserted by the eTSEC whenever RCTRL[PRSDPEP] is set to a non-zero value. Only one FCB is inserted per frame, in the buffer pointed to by the RxBD with bit F set. TOE acceleration for receive is enabled for all frames in this case.

14.6.4.2 Transmit path off-load and Tx PTP packet parsing

TOE functions for transmit are defined by the contents of the Tx FCB. The definition for the Tx FCB is described below.

Table 14-903. Transmit frame control block

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
Offset + 0	VLN	IP	IP6	TUP	UDP	CIP	CTU	NPH									PTP
Offset + 2	L4OS								L3OS								
Offset + 4	PHCS																
Offset + 6	VLCTL																

The user instructs the Tx packet to be timestamped via setting bit 15 in the TxFCB to mark a PTP packet. TxFCB[VLCTL] can be translated as the Tx PTP packet identification number. BD[TOE] has to be set to enable transmit PTP packet time stamping. TxFCB[PTP] bit takes precedence over TxFCB[VLN] bit. It disables per packet VLAN tag insertion. On a PTP packet, VLAN tag can be inserted from the DFVLAN register. A proposed TxFCB update for the PTP packet is shown in [Table 14-916](#).

The contents of the Tx FCB are defined below.

Table 14-904. Tx frame control block description

Bytes	Bits	Name	Description
0-1	0	VLN	VLAN control word valid. This bit is ignored when the PTP bit is set. VLAN tag is read from the DFVLAN register if PTP=1. 0 Ignore VLCTL field. 1 If VLAN tag insertion is enabled for eTSEC, use the VLCTL field as the VLAN control word.
	1	IP	Layer 3 header is an IP header. 0 Ignore layer 3 and higher headers. 1 Assume that the layer 3 header is an IPv4 or IPv6 header, and take L3OS field as valid.
	2	IP6	IP header is IP version 6. Valid only if IP = 1. 0 IP header version is 4. 1 IP header version is 6.
	3	TUP	Layer 4 header is a TCP or UDP header. 0 Do not process any layer 4 header. 1 Assume that the layer 4 header is either TCP or UDP (see UDP bit), and offload checksumming on the basis that the IP header has no extension headers.
	4	UDP	UDP protocol at layer 4. 0 Layer 4 protocol is either TCP (if TUP = 1) or undefined. 1 Layer 4 protocol is UDP if TUP = 1.
	5	CIP	Checksum IP header enable. 0 Do not generate an IP header checksum. 1 Generate an IPv4 header checksum.
	6	CTU	Checksum TCP or UDP header enable. 0 Do not generate a TCP or UDP header checksum. RFC 768 advises that UDP packets not requiring checksum validation should have their checksum field set to zero. 1 Generate a TCP header checksum if IP = 1 and TUP = 1 and UDP = 0.
	7	NPH	Disable calculation of TCP or UDP pseudo-header checksum. This bit should be set if IP options need to be consulted in forming the pseudo-header checksum, as eTSEC does not examine IP options or extension headers for TCP/IP offload on transmit. 0 Calculate TCP or UDP pseudo-header checksum as normal, assuming that the IP header has no options. 1 Do not calculate a TCP or UDP pseudo-header checksum, but instead use the value in field PHCS when determining the overall TCP or UDP checksum.
	8-14	-	Reserved
15	PTP	Indication to the transmitter that this is a PTP packet. Enabling PTP disables per packet VLAN tag insertion. Instead, VLAN tag will be read from the DFVLAN when the PTP field is true. 0 Do not attempt to capture transmission event time 1 Valid PTP_ID field. When this packet is transmitted, capture the time of transmission. Must be clear if TMR_CTRL[TE] is clear.	

Table continues on the next page...

Table 14-904. Tx frame control block description (continued)

Bytes	Bits	Name	Description
2-3	0-7	L4OS	Layer 4 header offset from start of layer 3 header. The layer 4 header starts L4OS octets after the layer 3 header if it is present. The maximum layer 3 header length supported is thus 255 bytes, which may prevent TCP/IP offload on particularly large IPv6 headers.
	8-15	L3OS	Layer 3 header offset from start of frame not including the 8 bytes for this FCB. The layer 3 header starts L3OS octets from the start of the frame including any custom preamble header that may be present. The maximum layer 2 header length supported is thus 255 bytes.
4-5	0-15	PHCS	Pseudo-header checksum (16-bit one's complement sum with carry wraparound, but without result inversion) for TCP or UDP packets, calculated by software. Valid only if NPH = 1.
6-7	0-15	VLCTL/ PTP_ID	VLAN control word for insertion in the transmitted VLAN tag. Valid only if VLN = 1. Tx PTP packet identification number. This number will be copied into the Tx PTP packet time stamp identification field. PTP field takes precedence over VLN field.

14.6.4.3 Receive path off-load

Upon receive, the Rx FCB returns the status of frame parse and TOE functions applied to the accompanying frame.

The definition for the Rx FCB is described below.

Table 14-905. Receive frame control block

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	VLN	IP	IP6	TUP	CIP	CTU	EIP	ETU	-				PERR		-	GPFP
Offset + 2			RQ					PRO								
Offset + 4																
Offset + 6	VLCTL															

The contents of the Rx FCB are defined below.

Table 14-906. Rx frame control block descriptions

Bytes	Bits	Name	Description
0-1	0	VLN	VLAN tag recognized. This bit is set only if RCTRL[VLEX] is set. 0 No VLAN tag recognized. 1 IEEE Std. 802.1Q VLAN tag found; VLAN control word in VLCTL is valid.
	1	IP	IP header found at layer 3. RCTRL[PRSDPEP] must be set to 10 or 11 in order to enable IP discovery. See also IP6 bit of FCB. 0 No layer 3 header recognized. 1 An IP header was recognized at layer 3; the IANA protocol identifier for the next header can be found in PRO; see PRO for more information. If S/W is relying on the RxFCB for the parse results, any RxFCB[IP] bits set with the corresponding RxFCB[PRO] = 0xFF indicates a fragmented packet (or that this packet had a back-to-back IPv6 routing extension header). Additionally, RQFPR[IPF] (see Receive queue filing table property register (ETSEC_RQFPR)) indicates that the packet was fragmented.
	2	IP6	IP version 6 header found at layer 3. 0 No IPv6 header was found. 1 The layer 3 header was an IPv6 header provided IP = 1.
	3	TUP	TCP or UDP header found at layer 4. RCTRL[PRSDPEP] must be set to 10 or 11 in order to enable TCP/UDP discovery. 0 No layer 4 header recognized. 1 The layer 4 header was recognized as either TCP (PRO = 0x06) or UDP (PRO = 0x11).
	4	CIP	IPv4 header checksum checked. RCTRL[PRSDPEP] must be set to 10 or 11 in order to enable IPv4 checksum verification. 0 IPv4 header checksum not verified, either because verification was disabled or a valid IPv4 header could not be located. 1 IPv4 header checksum was verified by the eTSEC, and bit EIP indicates result.
	5	CTU	TCP or UDP header checksum checked. RCTRL[PRSDPEP] must be set to 11 in order to enable layer 4 checksum verification. 0 TCP or UDP header checksum not verified, either because verification was disabled or a valid TCP or UDP header could not be located. If a UDP header with zero checksum was located, this bit is cleared in accordance with RFC 768. 1 TCP or UDP header checksum was verified by the eTSEC, and ETU indicates result.
	6	EIP	IPv4 header checksum verification error. Not valid unless CIP = 1. 0 No checksum error in IPv4 header. 1 Error in header checksum only if IP = 1 and IP6 = 0.
	7	ETU	TCP or UDP header checksum verification error. Not valid unless CTU = 1. 0 No checksum error in TCP or UDP header. 1 Error in header checksum only if PRO = 0x06 or PRO = 0x11.
8-11	-		Reserved
12-13		PER	Parse error. 00 No error in L2 to L4 parse 01 Reserved 10 Inconsistent or unsupported L3 header sequence 11 Reserved
14	-		Reserved

Table continues on the next page.
P2020 QorIQ Integrated Processor Reference Manual, Rev. 2, 12/2012

Table 14-906. Rx frame control block descriptions (continued)

Bytes	Bits	Name	Description
	15	GPFP	General-purpose filer event packet. This packet was filed based on matching a GPI rule sequence.
2-3	0-1	-	Reserved
	2-7	RQ	Receive queue index. This index was selected by the eTSEC Rx Filer (from a matching Filer rule's RQCTRL[Q] field) when it accepted the associated frame. If filing is not enabled, RQ is zero. Note that the 3 least significant bits of RQ correspond with the RxBD ring index whenever RCTRL[FSQEN] = 0.
	8-15	PRO	<p>If IP = 1, PRO is set as follows:</p> <ul style="list-style-type: none"> • PRO=0xFF for a fragment header or a back to back route header • PRO=0xnn for an unrecognized header, where nn is the next protocol field • PRO=(TCP/UDP header), as defined in the IANA specification, if TCP or UDP header is found <p>If IP = 0, PRO is undefined.</p> <p>Note that the eTSEC parser logic stops further parsing when encountering an IP datagram that has indicated that it has fragmented the upper layer protocol. This in general means that there is likely no layer 4 header following the IP header and extension headers. eTSEC leaves the RxFCB[PRO] and RQFPR[L4P] fields 0xFF in this case, which usually means that there was no IP header seen. In this case RxFCB[IP] and optionally RxFCB[IP6] is set. IP header checksumming operates and performs as intended. Most of the time, the eTSEC updates the RxFCB[PRO] field and RQFPR[L4P] fields with whatever value was found in the protocol field of the IP header. See Receive queue filing table property register (ETSEC_RQFPR), for a description of RQFPR.</p>
4-5	0-15	-	Reserved
6-7	0-15	VLCTL	VLAN control word as per IEEE Std. 802.1Q. The lower 12 bits comprise the VLAN identifier. Valid only if VLN = 1.

14.6.5 Quality of service (QoS) provision

This section describes the quality of service support features of this device.

It includes a parser which extracts vital packet properties and passes them to the filer which essentially acts as a frame classifier.

14.6.5.1 Receive parser

The receive parser parses the incoming frame data and generates filer properties and frame control block (FCB).

The receive parser composes of the Ethernet header parser and L3/L4 parser.

The Ethernet header parser parses only L2 (ethertype) headers. It is enabled by RCTRL[PRSDEP] != 0. It has the following key features:

- Extraction of 48-bit MAC destination and source addresses
- Extraction and recognition of the first 2-byte ethertype field
- Extraction and recognition of the final 2-byte ethertype field
- Extraction of 2-byte VLAN control field
- Walk through MPLS stack and find layer 3 protocol
- Walk through VLAN stack and find layer 3 protocol
- Recognition of the following ethertypes for inner layer parsing
 - LLC and SNAP header
 - JUMBO and SNAP header
 - IPV4
 - IPV6
 - VLAN
 - MPLSU/MPLSM
 - PPOES
 - ARP

For stack L2 (that is, more than one ethertypes) header, the Ethernet parser traverses through the header until it finds the last valid ethertype or the ethertype is unsupported. The table below describes what the Ethernet header parser recognizes for stack L2 header.

Table 14-907. Supported stack L2 Ethernet headers

Column-Current L2 Ethertype Row-Next Supported L2 Ethertype	LLC/SNAP	JUMBO/SNAP	IPV4	IPV6	VLAN	MPLSU	MPLSM	PPOES	ARP
LLC/SNAP	N	N	Y	Y	Y	Y	Y	Y	Y
JUMBO/SNAP	N	N	Y	Y	Y	Y	Y	Y	Y
IPV4	N	N	N	N	N	N	N	N	N
IPV6	N	N	N	N	N	N	N	N	N
VLAN	Y	Y	Y	Y	Y	Y	Y	Y	Y
MPLSU	N	N	Y*	Y*	N	y	Y	N	N
MPLSM	N	N	Y*	Y*	N	Y	Y	N	N
PPOES	N	N	Y	Y	N	Y	Y	N	N
ARP	N	N	N	N	N	N	N	N	N

NOTE: * means that it is the next protocol

The L3 parser is enabled by RCTRL[PRSDEP] = 10 or 11. It begins when the Ethernet parser ends and a valid IPv4/v6 ethertype is found. The L4 header is enabled by RCTRL[PRSDEP] = 11. It begins when the L3 parser ends and a valid TCP/UDP next protocol is found and no fragment frame is found. The primary functionalities of L3(IPv4/6) and L4(TCP/UDP) parsers are as follows:

- IP recognition (v4/v6, ARP, encapsulated protocol)
- IP header checksum verification
- IPv4/6 over IPv4/6 (tunneling)-parse headers and find layer 4 protocol
- IP layer 4 protocol/next header extraction
- Stop parsing on unrecognized next header/protocol
- IPv4 support
 - IPv4 source and destination addresses
 - 8-bit IPv4 type of service
 - IP layer 4 protocol / next header support
 - IPV4
 - IPV4 Fragment. Parser stops after a fragment is found
 - TCP/UDP
- IPv6 support
 - The first 4 bytes of the IPv6 source address extraction
 - The first 4 bytes of the IPv6 destination address extraction
 - IPv6 source address hash for pseudo header calculation
 - IPv6 destination address hash for pseudo header calculation
 - 8-bit IPv6 traffic class field extraction
 - Payload length field extraction
 - IP layer 4 protocol/next header support
 - IPV6
 - IPV6 fragment. Parser stops after a fragment is found
 - IPV6 route
 - IPV6 hop/destination
 - TCP/UDP
- L4 (TCP/UDP) support
 - Extraction of 16-bit source port number extraction
 - Extraction of 16-bit destination port number extraction
 - TCP checksum calculation (including pseudo header)
 - UDP checksum calculation if the checksum field is not zero (including pseudo header)

14.6.5.2 Receive queue filer

The receive queue filer receives protocol header properties extracted from the incoming frame by the eTSEC frame parse engine.

A property is defined to be a field extracted from a packet header, such as a TCP port number or VLAN identifier. As soon as the last identifiable header has been recognized, the filer commences searching the receive queue filer table, comparing properties in the table against properties extracted from the frame. This table is illustrated in the figure below. Software populates the table with property values, stored to the RQPROP field, and indicates how to match and interpret the properties by setting flags in the RQCTRL field. The eTSEC memory map provides access to these fields by way of an address register (RQFAR) and two porthole registers (RQFCR and RQFPR).

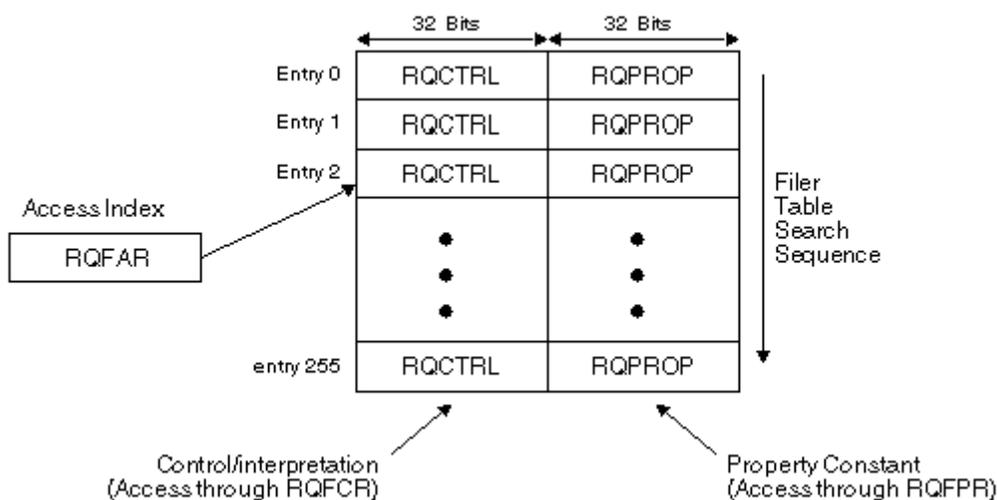


Figure 14-881. Structure of the receive queue filer table

14.6.5.2.1 Filing rules

Unless the filer is disabled, every received frame from the Ethernet MAC initiates a search of the receive queue filer table, starting at entry 0.

The table search is terminated as soon as an entry is found whose contents match a property of the frame. Accordingly, software must guarantee that at least one entry results in a match—even if only to set a default receive queue index.

Since eTSEC searches the table at a rate of two entries every system clock cycle, all 256 entries can be searched in the time taken to receive a 64-byte Ethernet frame.

The equation to calculate the maximum rules is $MaxRules = \left(\left[\frac{CCB \text{ clock}/2}{phyclk} \right] \times 2 \times (TF + IFG) \right)$ where:

- CCB = also called the platform clock

- TF = number of PHY clocks between start of one frame and when the parser is finished with parsing the frame.
- IFG = interframe gap

For example, $\text{sysclk} = 333 \text{ MHz}$, $\text{PHY clk} = 155 \text{ MHz}$, a 40-byte packet in 16-bit FIFO mode results in $(333 \div 155) \times ((40 \text{ bytes} \div (16 \text{ bits} \times (1 \text{ byte} \div 8 \text{ bits})) + 3) \times 2)$, or 98 rules. Note that in this 40-byte example, the number of rules calculated (e.g. 98) is valid only if the parser does not stop before 40 bytes of frame data is received. It is recommended for the filer to have additional rules to filter out frames that the parser stops parsing before 40 bytes of frame data is received. Note that in non-FIFO modes such as GMII running at slow platform frequency, MaxRules equation also applies. For example, using 20 for IPG, a minimum of 190.5 MHz eTSEC sysclk (381 MHz platform) is needed to process all 256 entries in GMII for a 64-byte receive frame.

Each entry of the receive queue filer table specifies a simple match rule for determining how to process the received frame. The elements of a filing rule, expressed in the RQCTRL and RQPROP fields, are summarized as follows:

- The PID field in RQCTRL identifies what property is being matched against RQPROP. The eTSEC supports 16 properties, some of which are different portions of the same header field. Reserved or unused bits in RQPROP are read as zero. See [Receive queue filing table property register \(ETSEC_RQFPR\)](#) for a list of all properties and their associated PID values.
- The Q field in RQCTRL identifies which one of 64 virtual receive queues the frame should be filed to (sent through DMA) in the event of a filing rule match that accepts the frame. The physical RxBD ring this queue maps to is controlled by the RCTRL[FSQEN] bit. If RCTRL[FSQEN] = 0, the three least significant bits of the Q field indicate which physical RxBD ring hosts the queue. If RCTRL[FSQEN] = 1, RxBD ring 0 hosts all receive queues, but the RxFCB[RQ] field allows software to distinguish queues by ID. In all cases if Q maps to a RxBD ring that is not currently enabled, the frame is discarded with an IEVENT[FIQ] error.
- The REJ field in RQCTRL controls whether the frame is to be rejected (REJ = 1) or filed (REJ = 0) upon a filing rule match. Rejected frames occupy Rx FIFO space, but do not consume memory bus cycles.
- The CMP field in RQCTRL determines how property PID is compared against RQPROP. Equality, inequality, greater-or-equal, and less-than compares are available.
- The AND field in RQCTRL allows more than one comparison in a sequence to be chained together as a Boolean AND condition. Setting AND = 1 defers evaluation of the rule until the next entry has been matched, which may, in turn, have AND set. If any comparison involving AND = 1 fails, the entire chained sequence fails. A typical use for AND is to combine a pair of comparisons in a range match; the first such entry has AND = 1, the second has AND = 0 and its values of Q and REJ take effect.

- The CLE field in RQCTRL offers a way to bracket a set of consecutive-perhaps related-rules into a rule cluster. A cluster must be preceded by a guard rule, which simply determines whether the cluster rules can be evaluated. If the guard rule succeeds and its last entry has both CLE = 1 and AND = 1, the cluster rules that follow are enabled. The cluster ends at the first entry where CLE = 1 and AND = 0, which may also belong to a rule that files or rejects a frame. If the guard rule fails, all rules in the cluster are skipped, including mask_register assignments. Clusters must not be nested.
- The GPI field offers the user the ability to interrupt the core upon matching a rule that causes a frame to be filed to memory. Once the last RxBD corresponding to that frame is written to memory, the IEVENT[FGPI] event will be asserted. This bit will be set regardless of any interrupt coalescing that may be set.

14.6.5.2.2 Comparing properties with bit masks

By default, extracted properties are compared arithmetically according to the CMP field in each RQCTRL word.

This permits point value matches in each table entry, and range checks across a pair of table entries combined with the AND attribute in RQCTRL. However, inspection of the parse flags, Ethernet preamble, and IP addresses typically requires "don't care" bit fields in the properties to be cleared as part of the comparison. The eTSEC provides a dedicated 32-bit register, known as the mask_register, for performing such masking operations. At the start of each table search by the filer, mask_register is reset to 0xFFFF_FFFF, which ensures that no masking occurs.

Filer rules may be configured to assign specific bit patterns to mask_register. Such rules can be configured to either match always (useful for implementing a default rule and specifying an associated receive queue), or fail always (which prevents termination of the filer table search). Once mask_register has been assigned, it retains its value until it is reassigned or the table search terminates. All properties are non-destructively bit-wise ANDed with mask_register prior to comparison in subsequent rules, which allows an entire cluster of rules to make use of a common mask. Individual masks for specific rules can also be created simply by combining a mask_register assignment (match always form) with a regular rule using the AND attribute.

To create a mask_register assignment rule, it is necessary to select PID = 0 in RQCTRL, and choose CMP such that the rule either matches (CMP = 01) or fails (CMP = 11). In this entry, RQPROP is then considered to be the assigned bit vector.

14.6.5.2.3 Special-case rules

It is frequently useful to create rules that are guaranteed to succeed or fail, specifically to enforce a default filing decision or act as null entries. Suggested constructions for such rules are shown in the table below.

Table 14-908. Special filer rules

Rule Description	RQCTRL Fields							RQPROP Word	RQCTRL Word ¹
	GPI	CL E	REJ	AN D	Q	CMP	PID		
Default file-Always file frame to ring Q	0	0	0	0	Q	01	0000	All zeros	0x0000_ <i>qq</i> 20
Default reject-Always discard frame	0	0	1	0	000_000	01	0000	All zeros	0x0000_0120
Empty rule in AND-Always matches	0	0/1 ²	0	1	000_000	01	0000	0xFFFF_FFFF	0x0000_00A0
Empty rule in rule set-Always fails	0	0/1 ³	0	0	000_000	11	0000	0xFFFF_FFFF	0x0000_0060

1. Hexadecimal digits *qq* denotes field Q shifted left 2 bits.
2. Set CLE = 1 if the empty rule guards a cluster.
3. Set CLE = 1 if the empty rule occurs at the end of a cluster.

14.6.5.2.4 Filer interrupt events

The filer can produce three interrupt events in IEVENT.

Event FIR indicates an error condition where the filer was unable to provide a definite result, either because no rule in the table succeeded, or because frames arrived too rapidly to complete searching of the table. Event FIQ indicates that the filer accepted a frame to a RxBD ring that was not enabled in RQCTRL (this can also occur if the filer is disabled, but RxBD ring 0-default queue or FSQEN mode queue-is not enabled). FIQ is also asserted in the case where no rule in the entire table succeeded. The various combinations of these interrupt events and their interpretation appear in the table below.

Table 14-909. Receive queue filer interrupt events

IEVENT[FIR] 1	IEVENT[FIQ]	Description
0	0	No error. The filer successfully rejected or filed a frame.
0	1	Illegal queue error. The filer accepted a frame to a RxBD ring that is disabled (including ring 0 if filing is disabled).
1	0	Partial search error. The filer did not have sufficient time to complete its search of the filer table.
1	1	No matching rule error. The filer searched all 256 entries of the filer table without finding a rule that succeeds.

A functional interrupt is provided via use of the general purpose interrupt (GPI) bit in the filer table. When a property matches the value in the RQPROP entry at this index, and $REJ = 0$ and $AND = 0$, the filer will set $IEVENT[FGPI]$ when the corresponding receive frame is written to memory. This allows the user to set up a filer rule where the core will be interrupted upon the reception of 'special' frames.

If the timer is enabled ($TMR_CTRL[TE] = 1$), then the interrupt dedicated for timer events (in addition to the usual receive, transmit and error interrupts) will be asserted.

14.6.5.2.5 Setting up the receive queue filer table

The eTSEC frame parser always provides values for all properties, even where the relevant headers are not available.

In the latter case, the filer is given default properties that can be used to avoid conflict with normal, defined property values. Accordingly, the rules in the filer table can be partitioned into rule sets such that if all rules in a given set fail (due to headers being unavailable), lower priority rule sets can be subsequently searched until either a rule set provides a match or a single default-catch-all-rule specifies a definite receive queue. For example, an IEEE 802.1p priority rule set may be followed by an IP TOS rule set, followed by a default rule; thus, if no VLAN tag appears in the received frame, the TOS rules are checked, or the default is activated should no IP header be present.

The rule cluster feature is used to conditionalize evaluation of rule sets. Typically, this avoids evaluating rules based on properties that may not be valid or relevant to the filing or filtering decision. For example, TCP-related rules might be clustered behind a guard rule that checks that a TCP header has appeared and the IP address matches our home address. Property 1-the parse flags property-is provided specifically to check the characteristics of the received frame and the parser error status. The `mask_register` is typically assigned beforehand to extract specific flags, in which case care should be taken that `mask_register` be reassigned an appropriate mask vector for following comparisons.

In many cases it is possible to write the entire filer table before using eTSEC, as the rule set is static. However, dynamic rule updates can be supported by pre-allocating partially instantiated rule sets, which software rewrites as necessary. Rules that are not instantiated should be composed of empty entries, as indicated in [Table 14-908](#). In many cases empty entries can be overwritten by software without stopping eTSEC's receive function.

14.6.5.2.6 Filer example-802.1p priority filing

This example illustrates how to file frames according to layer 2 802.1p priority.

This matches against property 1001, comparing each specific priority level in order to associate them with a RxB D ring index. Note that if a VLAN tag does not appear in the frame, the parser passes priority 0 to the filer, which always matches the rule at entry 7 and terminate the table search.

Table 14-910. Filer table example-802.1p priority filing

Table Entry	RQCTRL Fields							RQPROP	Comment	RQCTRL Word
	GPI	CL E	REJ	AN D	Q	CMP	PID			
0	0	0	0	0	000_000	00	1001	0x0000_0007	File priority 7 to ring 0	0x0000_0009
1	0	0	0	0	000_001	00	1001	0x0000_0006	File priority 6 to ring 1	0x0000_0409
2	0	0	0	0	000_010	00	1001	0x0000_0005	File priority 5 to ring 2	0x0000_0809
3	0	0	0	0	000_011	00	1001	0x0000_0004	File priority 4 to ring 3	0x0000_0C09
4	0	0	0	0	000_100	00	1001	0x0000_0003	File priority 3 to ring 4	0x0000_1009
5	0	0	0	0	000_101	00	1001	0x0000_0002	File priority 2 to ring 5	0x0000_1409
6	0	0	0	0	000_110	00	1001	0x0000_0001	File priority 1 to ring 6	0x0000_1809
7	0	0	0	0	000_111	00	1001	All zeros	File undefined 802.1p or priority 0 to ring 7-Default always matches	0x0000_1C09

14.6.5.2.7 Filer example-IP diff-serv code points filing

This example demonstrates use of rule priority for determining class selector codepoints (RFC 2474) from the IP TOS property.

An example filer table is shown below. The example relies on the fact that the first rule matched terminates the search, hence successively lower Diff-Serv codepoint ranges can be compared in each step until the default (zero or greater) range is reached. By default, property 1010 (IP TOS) takes the value 0x00 if no IP headers were recognized, therefore the table search always terminates.

Table 14-911. Filer table example-IP diff-serv code points filing

Table Entry	RQCTRL Fields							RQPROP	Comment	RQCTRL Word
	GPI	CLE	REJ	AND	Q	CMP	PID			
0	0	0	0	0	001_000	01	1010	0x0000_00E0	File class 7 to queue 8 (TOS >= 0xE0)	0x0000_202A
1	0	0	0	0	001_001	01	1010	0x0000_00C0	File class 6 to queue 9 (TOS >= 0xC0)	0x0000_242A

Table continues on the next page...

Table 14-911. Filer table example-IP diff-serv code points filing (continued)

Table Entry	RQCTRL Fields							RQPROP	Comment	RQCTRL Word
	GPI	CLE	REJ	AND	Q	CMP	PID			
2	0	0	0	0	001_010	01	1010	0x0000_00A0	File class 5 to queue 10 (TOS >= 0xA0)	0x0000_282A
3	0	0	0	0	001_011	01	1010	0x0000_0080	File class 4 to queue 11 (TOS >= 0x80)	0x0000_2C2A
4	0	0	0	0	000_100	01	1010	0x0000_0060	File class 3 to queue 4 (TOS >= 0x60)	0x0000_102A
5	0	0	0	0	001_100	01	1010	0x0000_0040	File class 2 to queue 12 (TOS >= 0x40)	0x0000_302A
6	0	0	0	0	010_100	01	1010	0x0000_0020	File class 1 to queue 20 (TOS >= 0x20)	0x0000_502A
7	0	0	0	0	011_100	01	1010	All zeros	File class 0 to queue 28 (TOS >= 0x00) or file to ring 4 by default	0x0000_702A

14.6.5.2.8 Filer example-TCP and UDP port filing

This example demonstrates rule clusters and AND-combined entries for filing packets based on transport protocol and well-known port numbers in a termination application.

An example filer table is shown below. The example contains two clusters; the first is entered only for TCP packets, the second is entered only for UDP packets. A default filing rule catches the case where neither TCP nor UDP headers are found. Each cluster compares source port number (property 1111) against a list of server ports, and files the

packets accordingly. Note that entries 1 and 2 form an AND rule for checking that the port number ≥ 20 and port number < 22 . Entries 4 and 5 are initially set up to always fail (zero port number), and thus comprise empty entries that can be used at a later time.

Table 14-912. Filer table example-TCP and UDP port filing

Table Entry	RQCTRL Fields							RQPROP	Comment	RQCTRL Word
	GPI	CL E	REJ	AN D	Q	CMP	PID			
0	0	1	0	1	000_000	00	1011	0x0000_0006	Enter cluster if layer 4 is TCP	0x0000_028B
1	0	0	0	1	000_000	01	1111	0x0000_0014	AND rule-FTP from TCP ports 20 and 21: file to ring 2	0x0000_00AF
2	0	0	0	0	000_010	11	1111	0x0000_0016		0x0000_086F
3	0	0	0	0	000_011	00	1111	0x0000_0017	telnet from TCP port 23: file to ring 3	0x0000_0C0F
4	0	0	0	0	000_000	00	1111	All zeros	<i>empty entry reserved for future use</i>	0x0000_000F
5	0	0	0	0	000_000	00	1111	All zeros	<i>empty entry reserved for future use</i>	0x0000_000F
6	0	1	0	0	000_001	01	0000	All zeros	end cluster; default TCP: file to ring 1	0x0000_0620
7	0	1	0	1	000_000	00	1011	0x0000_0011	Enter cluster if layer 4 is UDP	0x0000_028B
8	0	0	0	0	000_101	00	1111	0x0000_0801	NFS from UDP port 2049	0x0000_140F
9	0	0	0	0	000_111	00	1111	0x0000_0208	Route from UDP port 520	0x0000_000F
10	0	0	0	0	000_110	00	1111	0x0000_0045	TFTP from UDP port 69	0x0000_180F
11	0	1	0	0	000_100	01	0000	All zeros	End cluster; default UDP: file to ring 4	0x0000_1220
12	0	0	0	0	000_000	01	0000	All zeros	By default, file to ring 0	0x0000_0020

14.6.5.2.9 Filer example-interrupt from deep sleep

The example below shows how the filer can facilitate exit from deep sleep if any of the following packets arrive:

- ARP packet with Target IP address matching either of two IP addresses (either static or link local address)
- IPv4/UDP multicast DNS query
- IPv4/UDP SNMP broadcast query

These packets are also be stored in memory. All other packets are dropped.

Table 14-913. Filer example-interrupt from deep sleep

Table Entry	RQCTRL Fields							RQPROP	Comment
	GP I	CL E	RE J	AN D	Q	CMP	PID		
0	0	0	0	0	000_000	11	0000	0x0001_0000	Check for ARP request; set mask register to mask off everything but the ARP request flag.

Table continues on the next page...

Table 14-913. Filer example-interrupt from deep sleep (continued)

Table Entry	RQCTRL Fields							RQPROP	Comment
	GP I	CL E	RE J	AND	Q	CMP	PID		
1	0	1	0	1	000_000	00	0001	0x0001_0000	Check to see if ARP request flag is set by doing a =1 comparison. Enter the "ARP Request Cluster" if true.
2	0	0	0	0	000_000	11	0000	0xFFFF_FFFF	ARP Cluster: Set Mask to unmask everything (Reset mask to all F's)
3	1	0	0	0	000_001	00	1100	0XXXXX_XXXX	Compare the ARP Target IP address against "MY_IP_1," which is indicated by the user-defined value of 0XXXXX_XXXX; if they match, accept the frame and generate an interrupt.
4	1	0	0	0	000_001	00	1100	0YYYYY_YYYY	Compare the ARP Target IP address against "MY_IP_2," which is indicated by the user-defined value of 0YYYYY_YYYY; if they match, accept the frame.
5	0	1	1	0	000_000	01	0000	0x0000_0000	Default rule that will always discard the packet; inserted here because an ARP request was received, but the Target IP address did not match either local IP addresses; therefore drop the packet and exit the cluster.
6	0	0	0	0	000_000	11	0000	0x0000_02D0	Set Mask for IP4 Packet (2D0), with IPv4 checksum checked and verified, and UDP header located.
7	0	0	0	1	000_000	00	0001	0x0000_02D0	Check to see if IP4 Packet, with IPv4 checksum checked and verified, and UDP header.
8	0	0	0	1	000_000	00	0000	0xFFFF_FFFF	Set Mask to unmask everything (Reset mask to all F's).
9	0	1	0	1	000_000	00	1011	0x0000_0011	Check against UDP protocol; if this passes, enter the cluster - all packets in the cluster are IPv4 packets with UDP protocol identified as the L4 protocol type.
10	0	0	0	1	000_000	00	0011	0x00XX_XXXX	Compare upper L2 DA bits to XX_XXXX (for multicast DNS query)
11	0	0	0	1	000_000	00	0100	0x00YY_YYYY	Compare lower L2 DA bits to YY_YYYY
12	0	0	0	1	000_000	00	1100	0xZZZZ_ZZZZ	Compare L3 Destination IP address to ZZZZ_ZZZZ
13	0	0	0	1	000_000	00	1110	0x0000_XXXX	Compare L4 destination port to XXXX
14	1	0	0	0	000_001	00	1111	0x0000_YYYY	If all of the previously consecutive ANDed conditions pass, multicast DNS Query has matched.
15	0	0	0	1	000_000	00	0011	0x00XX_XXXX	Compare upper L2 DA bits to XX_XXXX (for SNMP broadcast query).
16	0	0	0	1	000_000	00	0100	0x00YY_YYYY	Compare lower L2 DA bits to YY_YYYY.
17	1	0	0	0	000_001	00	1110	0x0000_ZZZZ	If all of the previously consecutive ANDed conditions pass, SNMP broadcast Query has matched.

Table continues on the next page...

Table 14-913. Filer example-interrupt from deep sleep (continued)

Table Entry	RQCTRL Fields							RQPROP	Comment
	GP I	CL E	RE J	AND	Q	CMP	PID		
18	0	1	1	0	000_000	01	0000	0x0000_0000	Cluster End: IPv4, UDP Comparison Default rule that will always discard the packet; inserted here because an IPv4 packet with L4=UDP request was received, but the profiles didn't match anything "interesting"; therefore drop the packet and exit the cluster.
19	0	0	1	0	000_000	01	0000	0x0000_0000	Default rule that will always discard the packet; inserted here no matches for any "interesting" packets were received that are used to wake up the CPU. All packets that reach this rule are discarded.

14.6.5.3 Transmission scheduling

Each eTSEC can maintain multiple TxBD rings (or transmission queues) to satisfy QoS requirements.

The ability to choose from a number of transmission streams dynamically is especially important during periods of network congestion. Certain application such as voice and video streaming are delay sensitive, but loss insensitive. For instance, VoIP applications require little bandwidth, but are highly sensitive to latency. Conversely, FTP or SMTP protocols are delay insensitive, but loss sensitive.

eTSEC has a transmission scheduler that implements a programmable QoS regime. The scheduler is responsible for choosing which of the prefetched TxBDs shall be processed next, and accordingly issuing DMA requests to service the data stream described by the chosen BD(s). The scheduler cycle is one of:

1. decide on a TxBD queue,
2. transmit exactly one frame from that queue, and
3. return to deciding on another queue, in step 1.

If TCTRL[TXSCHED] is set to 00, no transmission scheduling occurs, and only TxBD ring 0 is polled for new data to transmit, with DMACTRL controlling waiting or polling. TCTRL[TXSCHED], if not zero, can be programmed to invoke one of two scheduling algorithms, namely priority-based queuing (PBQ), and modified weighted round-robin queuing (MWRR). In all cases where TCTRL[TXSCHED] is not zero, the scheduler can choose from among 1 to 8 TxBD rings per eTSEC, with individual rings being enabled by the setting of TQUEUE[EN0-EN7] bits. For example, TxBD rings 3, 4, and 7 may be enabled for scheduling by setting EN3, EN4, and EN7, and clearing all other EN bits.

14.6.5.3.1 Priority-based queuing (PBQ)

PBQ is the simplest scheduler decision policy. The enabled TxBD rings are assigned a priority value based on their index. Rings with a lower index have precedence over rings with higher indices, with priority assessed on a frame-by-frame basis. For example, frames in TxBD ring 0 have higher priority than frames in TxBD ring 1, and frames in TxBD ring 1 have higher priority than frames in TxBD ring 2, and so on.

The scheduling decision is then achieved as follows:

```
loop
  # start or S/W clear of TSATn
  ring = 0;
  while ring <= 7 loop
    if enabled(ring) and not ring_empty(ring) then
      transmit_frame(ring);
      ring = 0;
    else
      ring = ring + 1;
    endif
  endwhile
endloop
```

14.6.5.3.2 Modified weighted round-robin queuing (MWRR)

eTSEC implements a modified weighted round-robin scheduling algorithm across all enabled TxBD rings when TCTRL[TXSCHED] = 10.

In MWRR, the weights in the TR03WT and TR47WT registers determine the ideal size of each transmit slot, as measured in multiples of 64 bytes. Thus, to set a transmit slot of 512 bytes, a weight of 512/64 or 8 needs to be set for the ring. In this mode TxBD rings 1-7 are selected in round-robin fashion, whereas TxBD ring 0, if enabled with ready data for transmission, is always selected in between other rings so as to expedite transmission from ring 0.

The scheduling decision is then achieved as follows:

```
for ring = 1..7 and enabled(ring) loop
  credit[ring] = 0;
endloop
for ring = 1..7 and enabled(ring) loop
  if not ring_empty(0) then
    credit[0] = credit[0] + weight[0];
    while credit[0] > 0 loop
      transmit_frame(0);
      credit[0] = credit[0] - frame_size;
      if ring_empty(0) then
        credit[0] = 0;
      endif
    endwhile
  endif
  if not ring_empty(ring) then
    credit[ring] = credit[ring] + weight[ring];
  endif
endloop
```

functional description

```

while credit[ring] > 0 loop
    transmit_frame(ring);
    credit[ring] = credit[ring] - frame_size;
    if ring_empty(ring) then
        credit[ring] = 0;
    endif
endloop
endloop

```

The algorithm checks registers TQUEUE[EN0-EN7] for `enabled()`, TSTAT[THLT0-THLT7] for `ring_empty()`, and TRxWT for `weight()`. For TxB ring k , having a weight WT_k , the long term average throughput for that ring is:

rate of queue[k] ($K = 1$ to 7) = (available bandwidth) * $WT_k / (\text{sum}(WT_i) + 6WT_0)$

rate of queue(0) = (available bandwidth) * $7 * WT_0 / (\text{sum}(WT_i) + 6WT_0)$

where $i = 0$ to 7

14.6.6 Lossless flow control

The eTSEC DMA subsystem is designed to be able to support simultaneous receive and transmit traffic at gigabit line rates.

If the host memory has sufficient bandwidth to support such line rates, then the principle cause of overflow on receive traffic is due to a lack of Rx BDs. Thus, the long term receive throughput is determined by the rate at which software can process receive traffic. If a user desires to prevent dropped packets, they can inform the far-end link to stop transmission while the software processing catches up with the backlog.

To avoid overflow in the latter case, back pressure must be applied to the far-end transmitter before the Rx descriptor controller encounters a non-empty BD and halts with a BSY error. As there is lag between application of back-pressure and response of the far-end, the pause request must be issued while there are still BDs free in the ring. In the traditional eTSEC descriptor ring programming model, there is no way for hardware to know how many free BDs are available, so software must initiate any pause requests required during operation. If software is backlogged, the request may not be issued in time to prevent BSY errors. To allow the eTSEC to generate the pause request automatically, additional information (a pointer to the last free BD and ring length) is required.

14.6.6.1 Back pressure determination through free buffers

Ultimately, the rate of data reception is determined by how quickly software can release buffers back into the receive ring(s).

Each time a buffer is freed, the associated BD has its empty bit set and hardware is free to consume both. Thus the number of free BDs in a given Rx ring indicates how close hardware is to the end of that ring. To prevent data loss, back pressure should be applied when the number of free BDs drops below some critical level. The number of BDs that can be consumed by an incoming packet stream while back-pressure takes effect is determined by several factors, such as: receive traffic profile, transmit traffic profile, Rx buffer size, physical transmission time between eTSEC and far-end device and intra-device latency. Theoretically, the worst case is as follows:

$$\text{FreeBDsRequired} = \frac{\text{MaxFrameSize}}{\text{MinFrameSize} + \text{IFG}} + \frac{\text{MaxFrameSize}}{\text{RxBufferSize}} + \text{LinkDelay}$$

This case comes about when:

- The eTSEC has just started transmitting a large frame and thus cannot send out a pause frame
- Upon reception of the pause request the far-end has just started transmission of a large frame
- The eTSEC receives a burst of short frames with minimum inter-frame-gap (96bit times for Ethernet)

Once the user has determined the worst case scenario for their application, they program the required free BD threshold into the eTSEC (through RQPRM[PBTHR]). Since different BD rings may have different sizes and expected packet arrival rates, a separate threshold is provided for each active ring. It is recommended that a threshold of at least fourBDs is the practical minimum for gigabit Ethernet links.

For the Rx descriptor controller to determine the number of free BDs remaining in the ring, it needs to know the following:

1. The location of the current BD being used by hardware
2. The location of the last BD that was released (freed) by software
3. The length of the Rx BD ring.

For each active ring, the current BD pointer (RBPTR_n) is maintained by the eTSEC. Software knows both the size of the Rx ring and the location of the last freed BD. By providing the eTSEC with those values (through RQPRM[LEN] and RFBPTR respectively) the eTSEC always know how many receive buffers are available to be consumed by incoming data.

The number of guaranteed free BDs in the ring is then determined by:

When RFBPTR_n < RBPTR_n

$$\text{FreeBDs} = \text{RQPRM}_n[\text{LEN}] - \text{RBPTR}_n + \text{RFBPTR}_n$$

When $\text{RFBPTR}_n > \text{RBPTR}_n$

$$\text{FreeBDs} = \text{RFBPTR}_n - \text{RBPTR}_n$$

When $\text{RBPTR}_n = \text{RFBPTR}_n$ the number of free BDs in the ring is either one (since RFBPTR_n points to a free BD) or equal to the ring length. Since the BD pointed to by RBPTR_n may be either in use or about to be used, it is not considered in the free BD count. To resolve the case where the two pointers collide, the following logic applies:

If RBASE_n was updated and thus initializes both RBPTR_n and RFBPTR_n , the ring is deemed empty.

If RFBPTR_n is updated by a software write and matches RBPTR_n , the ring is deemed empty.

If HW updates RBPTR_n and the result matches RFBPTR_n , the ring is deemed to have one BD remaining. Upon writing this BD back to memory (indicating the buffer is occupied) the ring is deemed to be full.

Important. There is a possibility that if software is severely backlogged in updating RFBPTR_n , the hardware could wrap around the ring entirely, consume exactly the remaining number of BDs and not halt with a BSY error. If software then increments RFBPTR_n to the next address (thereby equalling RBPTR_n), the hardware assumes the ring is now empty (when in fact there is only a single BD freed up). This results in the hardware failing to maintain back pressure on the far end. Upon software incrementing RFBPTR_n a subsequent time, the wrap condition is successfully detected and hardware recognizes a nearly full ring (rather than a nearly empty one). Since software can increment RFBPTR_n by any amount, it is not possible for hardware to determine in this case whether the user has cleared the entire ring or just one BD. Users can eliminate the possibility of this condition occurring by ensuring that RFBPTR_n is incremented by at least two BDs each time (that is, clear at least two buffers whenever the RxBD unload routine is called).

Once the eTSEC determines that this threshold has been reached, back pressure is applied accordingly. The type of back pressure that is applied varies according to the physical interface that is used.

- **Half duplex Ethernet:** No support in this mode.
- **Full duplex Ethernet:** An IEEE 802.3 PAUSE frame (see [Flow control](#)) is issued as if the $\text{TCTRL}[\text{TFC_PAUSE}]$ bit was set. An internal counter tracks the time the far end controller is expected to remain in pause (based on the setting of $\text{PTV}[\text{PT}]$).

When that counter reaches half the value of PTV[PT], the eTSEC reissues a pause frame if the free BD calculation for any ring is below the threshold for that ring. For example, if PTV[PT] is set to 10 quanta, a pause frame is re-issued when five quanta have elapsed if the free BD threshold is still not met. A practical minimum for PTV[PT] of 4 quanta is recommended.

- **FIFO packet interface:** Link layer flow control is asserted through use of the RFC signal (CRS pin). Flow control is asserted for the entire time that free BD threshold is not met. The same mechanism is used for both GMII-style and encoded packet modes.

14.6.6.2 Software use of hardware-initiated back pressure

14.6.6.2.1 Initialization

Software configures RBASEn and RQPRMn[LEN] according to the parameters for that ring.

Then the number of free BDs that are required to prevent the eTSEC from automatically asserting flow control are programmed in RQPRM[FBTHR]. The receiver is then enabled.

The act of programming RBASEn initializes RFBPTRn to the start of the of the ring. When the ring is in this initial empty state, there is no concept of a last freed BD. In this case, the calculated number of free BDs is the size of the ring. Since the BD that the hardware is currently pointing to is to be considered in-use, the free BD count is actually one higher than the total available. As soon as the hardware consumes a BD (by writing it back to memory), RBPTRn advances and the free BD count reflects the correct number of available free BDs.

14.6.6.2.2 Operation

As software frees BDs from the ring, it writes the physical address of the BD just freed to RFBPTRn.

The eTSEC asserts flow control if the distance (using modulo arithmetic) between RBPTRn and RFBPTRn is $< RQPRMn[FBTHR]$. In multi-ring operation, if the free BD count of **any** active ring drops below the threshold for that ring, flow control is asserted. Once enough BDs are freed for **all** active rings to meet their respective free BD thresholds, application of back pressure cases.

The eTSEC does not issue an exit pause frame (that is, pause frame with PTV of 0x0000) once all active rings have sufficient BDs. Instead, it waits for the far-end pause timer to expire and start re-transmission.

14.6.7 Hardware assist for IEEE 1588-compatible timestamping

There is a push in industrial control applications to use Ethernet as the principal link layer for communications. This requires Ethernet to be used for both data transfer and real-time control. For real-time systems, each node is required to be synchronized to a master clock. The precision of this clock is dictated by the application, but generally needs to be of the order of <math><1\mu\text{Sec}</math> for high-speed machinery (for example, printing presses).

IEEE 1588 [1588] specifies a mechanism for synchronizing multiple nodes to a master clock. Support for 1588 can be done entirely in software running on a host CPU, but applications that require sub 10 μSec accuracy will need hardware support for accurate timestamping of incoming packets.

The eTSEC includes a new timer clock module to support the IEEE Std. 1588 timer standard. The following sections describe the features, programming model, and implementation information.

NOTE

IEEE 1588 timestamping is not supported in conjunction with the SGMII 10/100 interface mode.

14.6.7.1 Features

- 64-bit free running timer running from an external oscillator or internal clock
- Programmable timer oscillator clock selection
- Self-correcting precision timer with nano-second resolution
- Time stamp all incoming packets inline
 - Maskable interrupts on received PTP packet's filter rule match
- Time stamp transmit packets when instructed in the TxFCB
 - Maskable interrupts on transmit timestamp capture
- Two Tx time stamp registers per eTSEC with 16-bit tag for each of them to support burst mode.
- Time stamp capture on two general-purpose external triggers
 - Maskable interrupts on GPIO timestamp trigger
 - Programmable polarity of external trigger (GPIO) edge
- Two 64-bit alarm (future time) registers for future time comparison
 - Maskable interrupts on alarm

- Three programmable timer output pulse period phase aligned with 1588 timer clock
 - Maskable interrupts associated with each pulse
- Separate maskable timer interrupt event register
- Recognition of incoming PTP packet through filter rule match
- Phase aligned adjustable (divide by N) clock output
- Supports all Ethernet modes supported by the eTSEC, including full- and half-duplex modes
- Supports both master and slave modes
- Supports timestamp of nano-second resolution

14.6.7.2 Timer logic overview

The 1588 timer module can be partitioned into four different sub-modules as shown below.

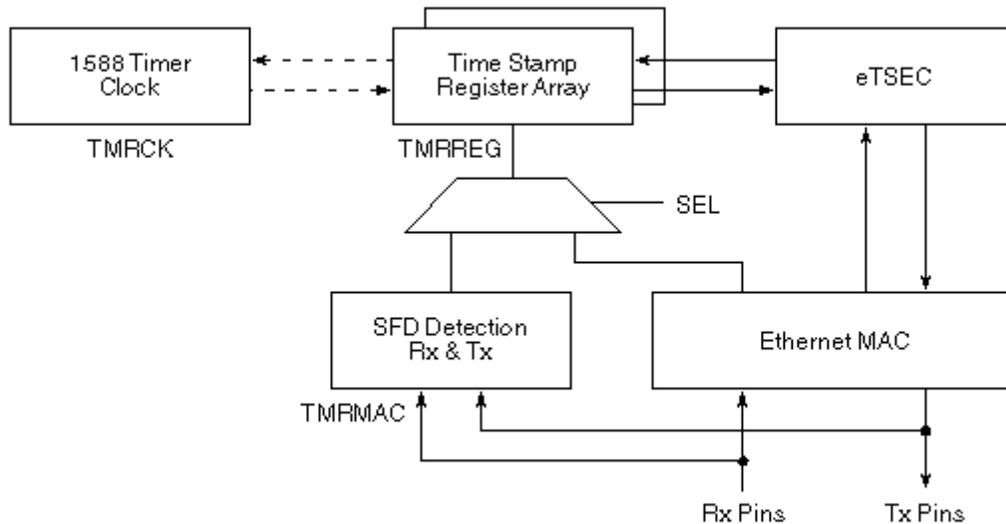
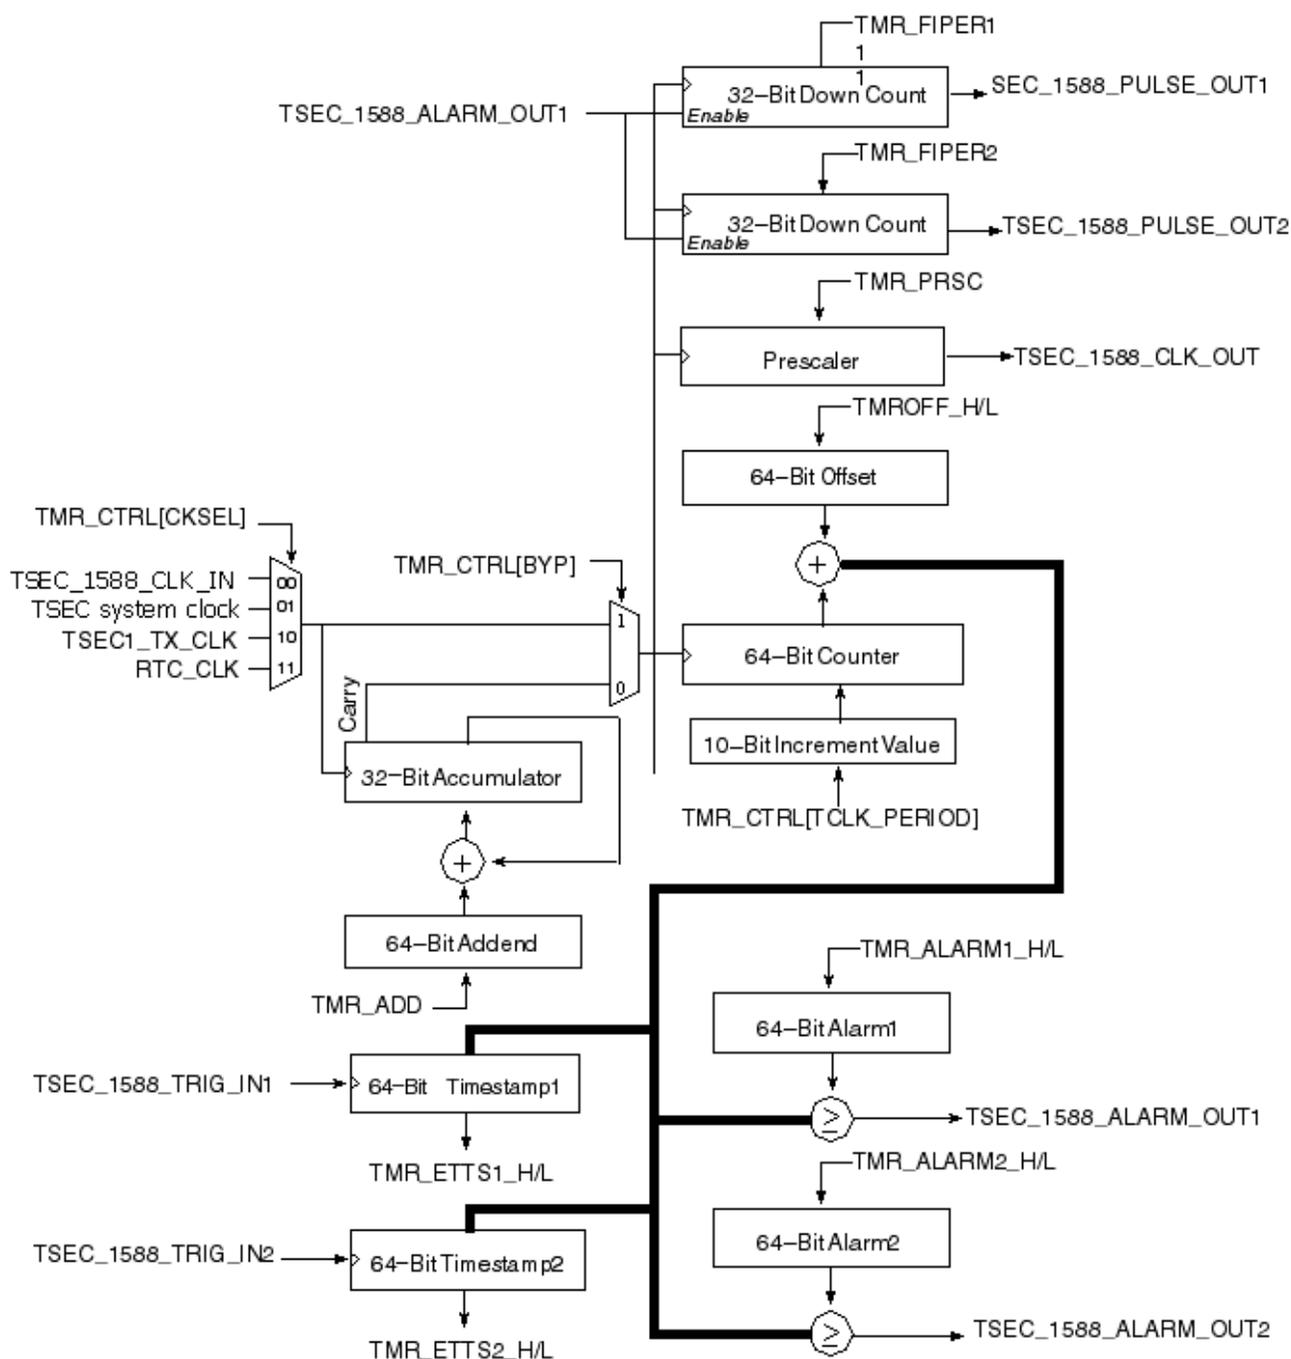


Figure 14-882. 1588 timer design partition



Note that TSEC system clock is CCB clock/2.

Figure 14-883. 1588 Current time control

The 64-bit current time is controlled by timer enable, TMR_CTRL[TE], selected clock, TMR_CTRL[CKSEL], bypass mode, TMR_CTRL[BYP], the timer clock period, TMR_CTRL[TCLK_PERIOD], and the ADDEND value, TMR_ADD registers, as shown above.

14.6.7.3 Time-stamp insertion on the received packets

Every incoming packet's 8-byte time stamp is inserted into the packet data buffer as padding alignment bytes.

Time-stamp insertion into the data buffer requires RCTRL[PAL] to be set to a value greater than or equal to 8 and the control bit RCTRL[TS] bit to be set.

14.6.7.3.1 Timestamp point

The required timestamp point, as specified in the IEEE 1588 Specification Sep-2004 (IEC 61588 First Edition), is shown below.

From this, it is clear that the end of the SFD is the critical point in the MII data stream.

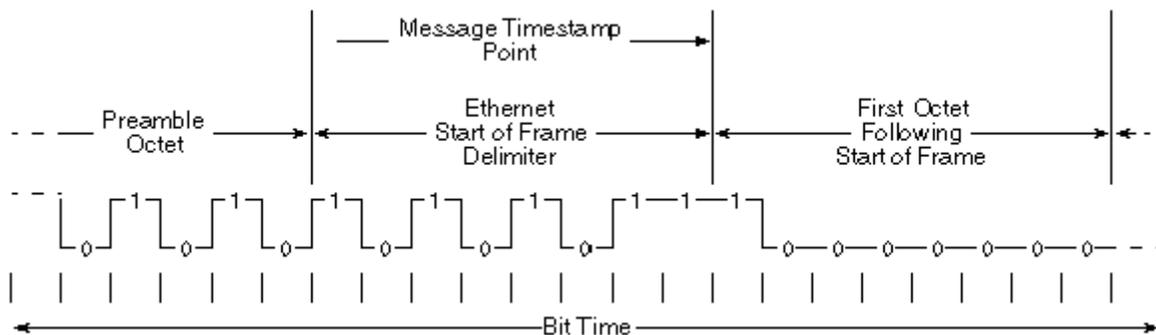


Figure 14-884. Ethernet sampling points for 1588

The sample point coincides with the cycle after the SFD (Start of Frame Delimiter) detection by the MAC. For received frames, this will be at least 4 bit times (MII) or 8 bit times (GMII) after the message timestamp point specified in [1588]. For transmission, the eTSEC sample point precedes the sample point specified in [1588] by at least 4-bit times (MII) or 8-bit times (GMII). For a particular mode, the eTSEC sample point is a consistent number of bit times relative to the SFD detection. Thus, the offset from the [1588] specified sample point can be accounted for in the PTP software implementation.

14.6.7.4 PTP packet parsing

PTP packets are typically embedded within a UDP payload with special IP source and destination address and special source and destination ports numbers.

Special fields of interest of a PTP packet are listed below.

Table 14-914. PTP payload special fields

Layer	Octet (Offset from the SFD)	Field	Value	eTSEC filer PID	Comments
Ethernet	12-13	Length/Packet	0x0800	ETY-RQPFR[PID=0111]	IPv4
IP header	22	Time to live	0x00	RBIFX- choose an arbitrary extraction byte	Must be 0
IP header	23	IP Protocol	0x11	L4P-RQPFR[PID=1011]	UDP
IP header	26-29	Source IP Address IANA defines 4 multicast address for the PTP packet		SIA-RQPFR[PID=1101]	
IP header	30-33	Destination IP Address IANA defines 4 multicast address for the PTP packet	224.0.1.129 224.0.1.130 224.0.1.131 224.0.1.132	DIA-RQPFR[PID=1100]	DefaultPTPdomain AlternatePTPdomain1 AlternatePTPdomain2 AlternatePTPdomain3
UDP header	34-35	Source port number		SPT-RQPFR[PID=1011]	
UDP header	36-37	Destination port number	319 320	DPT-RQPFR[PID=1011]	EventPort GeneralPort
UDP data	74	Control	0x0 0x1 0x2 0x3 0x4	RBIFX- choose an arbitrary extraction byte	Sync Delay_req Follow_up Delay_resp Management

A representation of the PTP packet is shown below

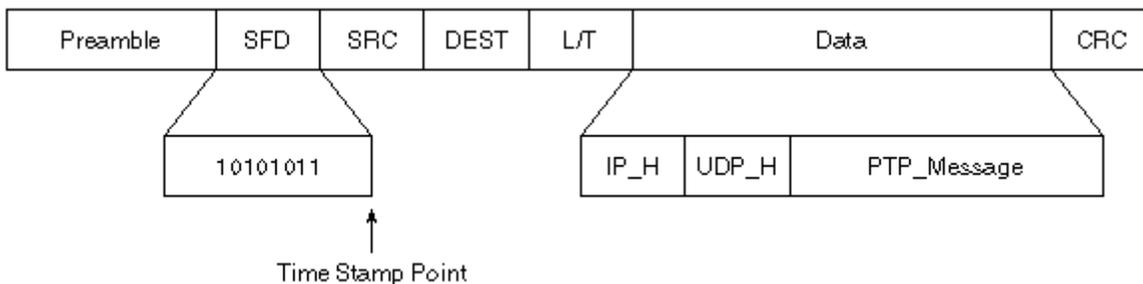


Figure 14-885. PTP packet format

14.6.7.4.1 General purpose filer rule

The eTSEC receive filer has been enhanced with the addition of a general-purpose event bit. This event bit can be used in conjunction with filing table rules to identify 1588 packets and indicate these packets by setting special timer status register bits (TMR_STAT). Additionally, 1588 packets can be easily identified by upper-layer software by using the filer to queue all PTP packets to one or more predefined virtual queues. See [Filing rules](#) for further information.

14.6.7.5 Time-stamp insertion on transmit packets

Software has the option to write the time stamp of the transmitted frame to memory in the padding alignment bytes (PAL) located between the TxFCB and the frame data.

It is required that a minimum of two TxBDs are used. The first points to the start of the 8 byte TxFCB. The second points to the start of frame data. In memory, the TxFCB, and at least the first 16 bytes of the TxPAL must be adjacent, that is, located in contiguous memory locations, as depicted in [Figure 14-886](#).

The first TxBD[TOE] bit is set. When the TMR_CTRL[Record Time-stamp In PAL Enable] and TxFCB[PTP] bits are set, the timestamp is written to memory location TxBD[Data Buffer Pointer]+16.

The second TxBD's Data Length must either contain the full frame length, or a value greater than the TxThreshold setting. Refer to [Table 14-915](#). When time-stamps are inserted into the TxPAL, the TMR_TXTS_n_H/L and TMR_TXTS_n_ID registers still function normally.

14.6.7.5.1 Interrupts

The TxPAL is updated with a time-stamp before closing the second TxBD.

The TxBD[I] bit can be set for the second TxBD frame to cause an interrupt (via IEVENT[TXF]) after the time-stamp has been written to the TxPAL.

When time-stamps are inserted into the TxPAL, the TMR_TXTS_n_H/L and TMR_TXTS_n_ID registers still function normally. Therefore, the 1588 interrupt can be triggered by using the TMR_PEVENT register bits TXP1, and TXP2.

Table 14-915. Time-stamp insertion programming requirements

Requirement	Behavior if requirement is not met
TMR_CTRL[RTPE]=1	If TMR_CTRL[RTPE]=0, then no time-stamp is written to a TxPAL.
TxBD[TOE]=1	If TxBD[TOE]=0, then no time-stamp is written to a TxPAL.

Table continues on the next page...

Table 14-915. Time-stamp insertion programming requirements (continued)

Requirement	Behavior if requirement is not met
First TxBD[Data Buffer Pointer] is 8-byte aligned	The time-stamp will be written to address First TxBD[Data Buffer Pointer] + 0x10 rounded down to the nearest 8-byte aligned address, except at 4K page boundaries, in which case the time-stamp may be invalid, and the Second TxBD close status will be lost.
First TxBD[Data Length]=8, 8 bytes for TxFCB	If L2 or frame data is included in the Length, the buffer immediately following the FCB is transmitted on the line and the frame data stored in memory will be overwritten with a time-stamp value after the frame is transmitted.
TxFCB[PTP]=1	If TxBD[PTP]=0, then no time-stamp is written to a TxPAL.
The TxFCB is followed immediately by a minimum of 16 bytes for the TxPAL	The time-stamp will be written to address First TxBD[Data Buffer Pointer] + 0x10.
Second TxBD[Data Buffer Pointers] points to start of L2 or frame data	If there is only one TxBD used to transfer a PTP frame, then no time-stamp is written to a TxPAL.
The second TxBD's Data Length must either contain the full frame length, or a value greater than 1024 bytes	If this condition is not true, the time-stamp in TxPAL is invalid.

The figure below depicts the buffer format requirements for time-stamp insertion on transmit packets.

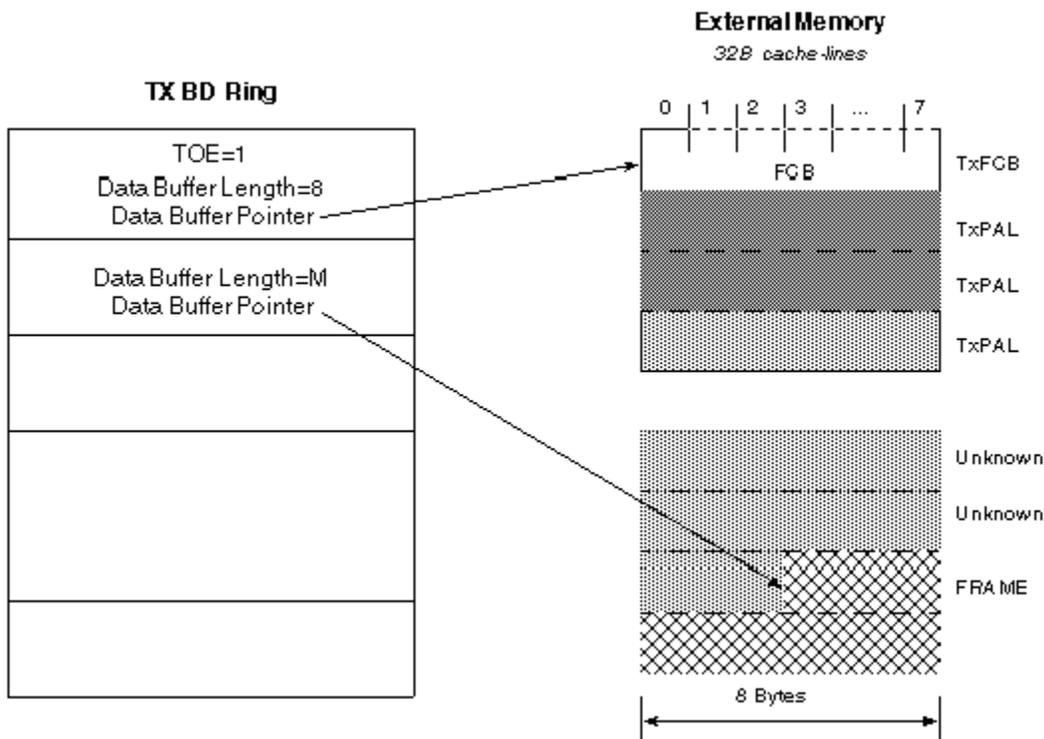


Figure 14-886. Buffer format for transmit time-stamp insertion

14.6.7.5.2 Error condition

When an error is encountered after a PTP packet has begun to be processed, the time-stamp written to the TxPAL is zero.

Subsequent frames may be flushed by eTSEC. There will be no time-stamp update to TxPAL for the subsequent flushed frames.

14.6.7.6 Tx PTP packet parsing

Software instructs the Tx packet to be timestamped via setting bit 15 in the TxFCB to mark a PTP packet. TxFCB[VLCTL] can be translated as the Tx PTP packet identification number. BD[TOE] must be set to enable transmit PTP packet time stamping. TxFCB[PTP] bit takes precedence over TxFCB[VLN] bit. It disables per packet VLAN tag insertion. On a PTP packet, a VLAN tag can be inserted from the DFVLAN register. The TxFCB for the PTP packet is shown below.

Table 14-916. Transmit frame control block

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	VLN	IP	IP6	TUP	UDP	CIP	CTU	NPH								PTP
Offset + 2	L4OS							L3OS								
Offset + 4	PHCS															
Offset + 6	VLCTL/PTP_ID															

The contents of the Tx FCB are defined below.

Table 14-917. Tx frame control block description

Bytes	Bits	Name	Description
0-1	0	VLN	VLAN control word valid. This bit is ignored when the PTP bit is set. VLAN tag is read from the DFVLAN register if PTP=1. 0 Ignore VLCTL field. 1 If VLAN tag insertion is enabled for eTSEC, use the VLCTL field as the VLAN control word.
	1	IP	Layer 3 header is an IP header. 0 Ignore layer 3 and higher headers. 1 Assume that the layer 3 header is an IPv4 or IPv6 header, and take L3OS field as valid.
	2	IP6	IP header is IP version 6. Valid only if IP = 1. 0 IP header version is 4. 1 IP header version is 6.
	3	TUP	Layer 4 header is a TCP or UDP header. 0 Do not process any layer 4 header. 1 Assume that the layer 4 header is either TCP or UDP (see UDP bit), and offload checksumming on the basis that the IP header has no extension headers.
	4	UDP	UDP protocol at layer 4. 0 Layer 4 protocol is either TCP (if TUP = 1) or undefined. 1 Layer 4 protocol is UDP if TUP = 1.
	5	CIP	Checksum IP header enable. 0 Do not generate an IP header checksum. 1 Generate an IPv4 header checksum.
	6	CTU	Checksum TCP or UDP header enable. 0 Do not generate a TCP or UDP header checksum. RFC 768 advises that UDP packets not requiring checksum validation should have their checksum field set to zero. 1 Generate a TCP header checksum if IP = 1 and TUP = 1 and UDP = 0.
	7	NPH	Disable calculation of TCP or UDP pseudo-header checksum. This bit should be set if IP options need to be consulted in forming the pseudo-header checksum, as eTSEC does not examine IP options or extension headers for TCP/IP offload on transmit. 0 Calculate TCP or UDP pseudo-header checksum as normal, assuming that the IP header has no options. 1 Do not calculate a TCP or UDP pseudo-header checksum, but instead use the value in field PHCS when determining the overall TCP or UDP checksum.
	8-14	-	Reserved
15	PTP	Indication to the transmitter that this is a PTP packet. Enabling PTP disables per packet VLAN tag insertion. Instead, VLAN tag will be read from the DFVLAN when the PTP field is true. 0 Do not attempt to capture transmission event time 1 Valid PTP_ID field. When this packet is transmitted, capture the time of transmission. Must be clear if TMR_CTRL[TE] is clear.	

Table continues on the next page...

Table 14-917. Tx frame control block description (continued)

Bytes	Bits	Name	Description
2-3	0-7	L4OS	Layer 4 header offset from start of layer 3 header. The layer 4 header starts L4OS octets after the layer 3 header if it is present. The maximum layer 3 header length supported is thus 255 bytes, which may prevent TCP/IP offload on particularly large IPv6 headers.
	8-15	L3OS	Layer 3 header offset from start of frame not including the 8 bytes for this FCB. The layer 3 header starts L3OS octets from the start of the frame including any custom preamble header that may be present. The maximum layer 2 header length supported is thus 255 bytes.
4-5	0-15	PHCS	Pseudo-header checksum (16-bit one's complement sum with carry wraparound, but without result inversion) for TCP or UDP packets, calculated by software. Valid only if NPH = 1.
6-7	0-15	VLCTL/ PTP_ID	VLAN control word for insertion in the transmitted VLAN tag. Valid only if VLN = 1. Tx PTP packet identification number. This number will be copied into the Tx PTP packet time stamp identification field. PTP field takes precedence over VLN field.

14.6.8 Buffer descriptors

The eTSEC buffer descriptor (BD) is modeled after the MPC8260 Fast Ethernet controller BD for ease of reuse across the PowerQUICC network processor family.

Drawing from the MPC8260 FEC BD programming model, the eTSEC descriptor base registers point to the beginning of BD rings. The eTSEC BD also expands upon the MPC8260 BD model to accommodate the eTSEC's unique features. However, the 8-byte data BD format is designed to be compatible with the existing MPC8260 BD model.

The eTSEC is capable of duplicating-or extracting-data directly into the L2 cache memory. This allows the processor to quickly access critical frame information as soon as the processor is ready without having to first fetch the data from main memory, which holds the master copy. This results in substantial improvement in throughput and hence reduction in latency.

14.6.8.1 Data buffer descriptors

Data buffers are used in the transmission and reception of Ethernet frames (see [Figure 14-887](#)).

Data BDs encapsulate all information necessary for the eTSEC to transmit or receive an Ethernet frame.

Within each data BD there is a status field, a data length field, and a data pointer. The BD completely describes an Ethernet packet by centralizing status information for the data packet in the status field of the BD and by containing a data BD pointer to the location of

the data buffer. Software is responsible for setting up the BDs in memory. Because of pre-fetching, a minimum of four buffer descriptors per ring are required. This applies to both the transmit and the receive descriptor rings. Transmit rings are limited to a maximum size of 65536 BDs due to BD and frame data prefetching. Software also must have the data pointer pointing to a legal memory location. Within the status field, there exists an ownership bit which defines the current state of the buffer (pointed to by the data pointer). Other bits in the status field of the buffer descriptor are used to communicate status/control information between the eTSEC and the software driver.

Because there is no next BD pointer in the transmit/receive BD (see [Figure 14-888](#)), all BDs must reside sequentially in memory. The eTSEC increments the current BD location appropriately to the next BD location to be processed. There is a wrap bit in the last BD that informs the eTSEC to loop back to the beginning of the BD chain. Software must initialize the TBASE and RBASE registers that point to the beginning transmit and receive BDs for eTSEC.

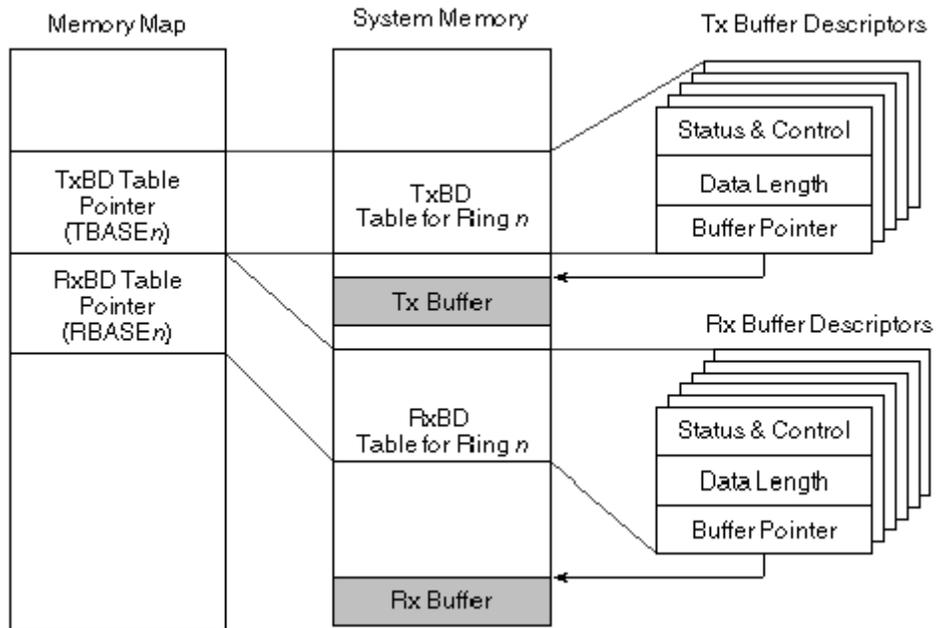


Figure 14-887. Example of eTSEC Memory Structure for BDs

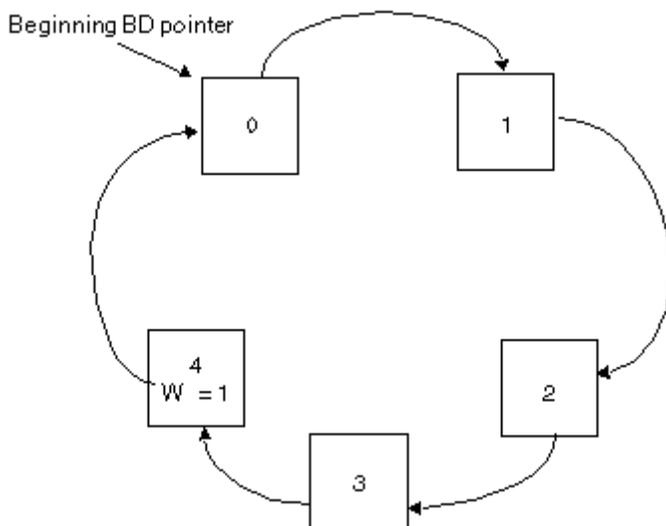


Figure 14-888. Buffer descriptor ring

14.6.8.2 Transmit data buffer descriptors (TxBD)

Data is presented to the eTSEC for transmission by arranging it in memory buffers referenced by the TxBDs.

In the TxBD the user initializes the R, PAD, W, I, L, TC, PRE, HFE, CF, and TOE bits and the length (in bytes) in the first word, and the buffer pointer in the second word. Unused fields or fields written by the eTSEC must be initialized to zero.

The eTSEC clears the R bit in the first word of the BD after it finishes using the data buffer. The transfer status bits are then updated. Additional transmit frame status can be found in statistic counters in the MIB block.

Software must expect eTSEC to prefetch multiple TxBDs, and for TCP/IP checksumming an entire frame must be read from memory before a checksum can be computed. Accordingly, the R bit of the first TxBD in a frame must not be set until at least one entire frame can be fetched from this TxBD onwards. If eTSEC prefetches TxBDs and fails to reach a last TxBD (with bit L set), it halts further transmission from the current TxBD ring and report an underrun error as IEVENT[XFUN]; this indicates that an incomplete frame was fetched, but remained unprocessed. The relevant TBPTR register points to the next unread TxBD following the error.

Table 14-918. Transmit buffer descriptor

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	R	PAD/CRC	W	I	L	TC	PRE/DEF	-	HFE/LC	CF/RL	RC				TOE/UN	TR
Offset + 2	DATA LENGTH															

Table continues on the next page...

Table 14-918. Transmit buffer descriptor (continued)

Offset + 4	TX DATA BUFFER POINTER
Offset + 6	

The TxBD definition is interpreted by eTSEC hardware as if TxBDs mapped to C data structures in the manner illustrated below.

```
typedef unsigned short uint_16; /* choose 16-bit native type */
typedef unsigned int uint_32; /* choose 32-bit native type */
typedef struct txbd_struct {
    uint_16 flags;
    uint_16 length;
    uint_32 bufptr;
} txbd;
```

Table 14-919. Transmit data buffer descriptor (TxBD) field descriptions

Offset	Bits	Name	Description
0-1	0	R	Ready, written by eTSEC and user. 0 The data buffer associated with this BD is not ready for transmission. The user is free to manipulate this BD or its associated data buffer. The eTSEC clears this bit after the buffer is transmitted or after an error condition is encountered. 1 The data buffer, which is prepared for transmission by the user, was not transmitted or is currently being transmitted. No fields of this BD may be written by the user once this bit is set.
	1	PAD/CRC	Padding for frames. (Valid only while it is set in the first BD and MACCFG2[PAD enable] is cleared). If MACCFG2[PAD enable] is set, this bit is ignored. 0 Do not add padding to short frames. 1 Add PAD to frames. PAD bytes are inserted until the length of the transmitted frame equals 64 bytes. Unlike the MPC8260 which PADs up to MINFLR value, the eTSEC PADs always up to the IEEE minimum frame length of 64 bytes. CRC is always appended to frames.
	2	W	Wrap. Written by user. 0 The next buffer descriptor is found in the consecutive location. 1 The next buffer descriptor is found at the location defined in TBASE.
	3	I	Interrupt. Written by user. 0 No interrupt is generated after this buffer is serviced. 1 IEVENT[TXB] or IEVENT[TXF] are set after this buffer is serviced. These bits can cause an interrupt if they are enabled (That is, IEVENT[TXBEN] or IEVENT[TXFEN] are set).
	4	L	Last in frame. Written by user. 0 The buffer is not the last in the transmit frame. 1 The buffer is the last in the transmit frame.
	5	TC	Tx CRC. Written by user. (Valid only while it is set in first BD and TxBD[PAD/CRC] is cleared and MACCFG2[PAD/CRC enable] is cleared and MACCFG2[CRC enable] is cleared.) If MACCFG2[PAD/CRC enable] is set or MACCFG2[CRC enable] is set, this bit is ignored in Ethernet modes. 0 End transmission immediately after the last data byte with no hardware generated CRC appended, unless TxBD[PAD/CRC] is set. 1 Transmit the CRC sequence after the last data byte.
	6	PRE	Transmit user-defined Ethernet preamble. Written by user. Valid only if set in the first BD of a frame, and MACCFG2[PreAm TxEN] is set. 0 This frame does not contain Ethernet preamble bytes for transmission. 1 This frame includes a user-defined Ethernet preamble sequence prior to the destination address in the data buffer.
		DEF	Defer indication. The eTSEC updates this bit after transmitting a frame (TxBD[L] is set) 0 This frame was not deferred. 1 This frame did not have a collision before it was sent but it was sent late because of deferring
	7	-	Reserved
	8	HFE	Huge frame enable. Written by user. Valid only if set in the first BD of a frame and MACCFG2[Huge Frame] is cleared. If MACCFG2[Huge Frame] is set, this bit is ignored. 0 Truncate transmit frame if its length is greater than the MAC's maximum frame length. 1 Allow large frames to be transmitted without truncation.
		LC	Late collision. Written by the eTSEC. 0 No late collision. 1 A collision occurred after 64 bytes are sent. The eTSEC terminates the transmission and updates LC.
		CF	Control Frame. Written by user. Valid only if set in the first BD of a frame. 0 Regular frame: transmission is deferred when eTSEC is in PAUSE

Table 14-919. Transmit data buffer descriptor (TxBD) field descriptions (continued)

Offset	Bits	Name	Description
2-3	0-15	Data Length	Data length is the number of octets the eTSEC should transmit from this BD's data buffer. It is never modified by the eTSEC. This field must be greater than zero, as zero indicates a BD not ready.
4-7	0-31	TX Data Buffer Pointer	The transmit buffer pointer contains the address of the associated data buffer. The data buffer pointer for the first BD of a TxPAL-enabled frame must be aligned on an 8-byte boundary. There are no alignment restrictions for the data buffer pointers of the second or subsequent BDs of a TxPAL-enabled frame, or for non-TxPAL frames.

14.6.8.3 Receive buffer descriptors (RxBD)

In the RxBD the user initializes the E, I, and W bits in the first word and the pointer in second word.

If the data buffer is used, the eTSEC modifies the E, L, F, M, BC, MC, LG, NO, CR, OV, and TR bits and writes the length of the used portion of the buffer in the first word. The M, BC, MC, LG, NO, CR, OV, and TR bits in the first word of the buffer descriptor are only modified by the eTSEC if the L (last BD in frame) bit is set. The first word of the RxBD contains control and status bits. Its formats are detailed below.

The number of buffer descriptors in a ring is set using the W bit to indicate that the next buffer wraps back to the beginning of the ring. See [Maximum frame length \(ETSEC_MAXFRM\)](#) for information on setting the size of the buffer ring.

Table 14-920. Receive buffer descriptor

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Offset + 0	E	RO1	W	I	L	F	0	M	BC	MC	LG	NO	SH	CR	OV	TR
Offset + 2	DATA LENGTH															
Offset + 4	RX DATA BUFFER POINTER															
Offset + 6																

The RxBD definition is interpreted by eTSEC hardware as if RxBDs mapped to C data structures in the manner illustrated below.

```
typedef unsigned short uint_16; /* choose 16-bit native type */
typedef unsigned int uint_32; /* choose 32-bit native type */
typedef struct rxbd_struct {
    uint_16 flags;
    uint_16 length;
    uint_32 bufptr;
} rxbd;
```

Table 14-921. Receive buffer descriptor field descriptions

Offset	Bits	Name	Description
0-1	0	E	Empty, written by the eTSEC (when cleared) and by the user (when set). 0 The data buffer associated with this BD is filled with received data, or data reception is aborted due to an error condition. The status and length fields have been updated as required. 1 The data buffer associated with this BD is empty, or reception is currently in progress.
	1	RO1	Receive software ownership bit. This field is reserved for use by software. This read/write bit is not modified by hardware, nor does its value affect hardware.
	2	W	Wrap, written by user. 0 The next buffer descriptor is found in the consecutive location. 1 The next buffer descriptor is found at the location defined in RBASE.
	3	I	Interrupt, written by user. 0 No interrupt is generated after this buffer is serviced. 1 IEVENT[RXB] or IEVENT[RXF] are set after this buffer is serviced. This bit can cause an interrupt if enabled (IMASK[RXBEN] or IMASK[RXFEN]). If the user wants to be interrupted only if RXF occurs, then the user must disable RXB (IMASK[RXBEN] is cleared) and enable RXF (IMASK[RXFEN] is set).
	4	L	Last in frame, written by the eTSEC. 0 The buffer is not the last in a frame. 1 The buffer is the last in a frame.
	5	F	First in frame, written by the eTSEC. 0 The buffer is not the first in a frame. 1 The buffer is the first in a frame.
	6	-	Reserved
	7	M	Miss, written by the eTSEC. (This bit is valid only if the L-bit is set and eTSEC is in promiscuous mode.) This bit is set by the eTSEC for frames that were accepted in promiscuous mode, but were flagged as a "miss" by the internal address recognition; thus, while in promiscuous mode, the user can use the M-bit to quickly determine whether the frame was destined to this station. 0 The frame was received because of an address recognition hit. 1 The frame was received because of promiscuous mode.
	8	BC	Broadcast. Written by the eTSEC. (Only valid if L is set.) Is set if the DA is broadcast (FF-FF-FF-FF-FF-FF).
	9	MC	Multicast. Written by the eTSEC. (Only valid if L is set.) Is set if the DA is multicast and not BC.
	10	LG	Rx frame length violation, written by the eTSEC (only valid if L is set). A frame length greater than or equal to the maximum frame length was recognized; in this case LG is set regardless of the setting of MACCFG2[Huge Frame]. If MACCFG2[Huge Frame] is cleared, the frame is truncated to the value programmed in the maximum frame length register. This bit is valid only if the L bit is set.
	11	NO	Rx non-octet aligned frame, written by the eTSEC (only valid if L is set). A frame that contained a number of bits not divisible by eight was received.

Table continues on the next page...

Table 14-921. Receive buffer descriptor field descriptions (continued)

Offset	Bits	Name	Description
	12	SH	Short frame, written by the eTSEC (only valid if L is set). A frame length less than the minimum 64B that is defined for Ethernet. was recognized, provided RCTRL[RSF] is set.
	13	CR	Rx CRC error, written by the eTSEC (only valid if L is set). This frame contains a CRC error and is an integral number of octets in length. This bit is also set if a receive code group error is detected.
	14	OV	Overflow, written by the eTSEC (only valid if L is set). A receive FIFO overflow occurred during frame reception. If this bit is set, the other status bits, M, LG, NO, CR and TR lose their normal meaning and are zero.
	15	TR	Truncation, written by the eTSEC (only valid if L is set). Set if the receive frame is truncated. This can happen if a frame length greater than the maximum frame length is received and MACCFG2[Huge Frame] is cleared. If this bit is set, the frame must be discarded and the other error bits must be ignored as they may be incorrect.
2-3	0-15	Data Length	Data length, written by the eTSEC. Data length is the number of octets written by the eTSEC into this BD's data buffer if L is cleared (the value is equal to MRBLR), or, if L is set, the length of the frame including CRC, FCB (if RCTRL[PRSDEP > 00]), preamble (if MACCFG2[PreAmRxEn]=1), timestamp (if RCTRL[TS]=1) and any padding (RCTRL[PAL]).
4-7	0-15	RX Data Buffer Pointer	Receive buffer pointer, written by the user. The receive buffer pointer, which always points to the first location of the associated data buffer, must be 8-byte aligned. For best performance, use 64-byte aligned receive buffer pointer addresses. The buffer must reside in memory external to the eTSEC.

14.7 Initialization/application information

14.7.1 Interface mode configuration

This section describes how to configure the eTSEC in different supported interface modes.

These include the following:

- MII
- RMII
- GMII
- RGMII
- SGMII
- TBI
- RTBI

The pinout, the data registers that must be initialized, as well as speed selection options are described.

14.7.1.1 MII interface mode

The table below describes the signal configurations required for MII interface mode.

Table 14-922. MII interface mode signal configuration

eTSEC Signals			MII Interface		
			Frequency [MHz] 25		
			Voltage [V] 3.3		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	leave unconnected		
TX_CLK	I	1	TX_CLK	I	1
TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	TxD[2]	O	1
TxD[3]	O	1	TxD[3]	O	1
TxD[4]	O	1	leave unconnected		
TxD[5]	O	1	leave unconnected		
TxD[6]	O	1	leave unconnected		
TxD[7]	O	1	leave unconnected		
TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	TX_ER	O	1
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	RxD[2]	I	1
RxD[3]	I	1	RxD[3]	I	1
RxD[4]	I	1	not used		
RxD[5]	I	1	not used		
RxD[6]	I	1	not used		
RxD[7]	I	1	not used		
RX_DV	I	1	RX_DV	I	1
RX_ER	I	1	RX_ER	I	1
COL	I	1	COL	I	1
CRS	I	1	CRS	I	1
Sum		25	Sum		16

Table 14-923. Shared MII signals

eTSEC signals	I/O	No. of signals	MII signals	I/O	No. of signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
EC GTX_CLK125	I	1	not used	I	0	Reference clock
Sum		3	Sum		2	

The table below describes the register initializations required to configure the eTSEC in MII mode.

Table 14-924. MII mode register initialization steps

Set soft_reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, for MII, half duplex operation. Set I/F Mode bit, MACCFG2[0000_0000_0000_0000_0111_0001_0000_0100] (This example has Full Duplex = 0, Preamble count = 7, PAD/CRC append = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1)
Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] Set station address to 02_60_8C_87_65_43, for example.
Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] Set station address to 02_60_8C_87_65_43, for example.
Assign a Physical address to the TBI so as to not conflict with the external PHY Physical address, TBIPA[0000_0000_0000_0000_0000_0000_0000_0101] Set to 05, for example.
Reset the management interface. MIIMCFG[1000_0000_0000_0000_0000_0000_0000_0111]
Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not less than 2.5 MHz

Table continues on the next page...

Table 14-924. MII mode register initialization steps (continued)

<p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the eTSEC MII Mgmt bus is idle.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Auxiliary Control and Status Register to configure the PHY through the Management interface (overrides configuration signals of the PHY). MIIMADD[0000_0000_0000_0000_0000_0000_0001_1100]</p>
<p>Perform an MII Mgmt write cycle to the external PHY Writing to MII Mgmt Control with 16-bit data intended for the external PHY register, MIIMCON[0000_0000_0000_0000_0000_0000_0000_0100]</p>
<p>Check to see if MII Mgmt write is complete Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Extended PHY control register #1 to set up the interface mode selection. MIIMADD[0000_0000_0000_0000_0000_0000_0001_0111]</p>
<p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY register, MIIMCON[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Mode control register to set up the interface mode selection. MIIMADD[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY register, MIIMCON[0000_0000_0000_0000_00uu_00uu_0u00_0000]</p> <p>where u is user defined based on desired configuration.</p>
<p>Check to see if MII Mgmt write is complete Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>If auto-negotiation was enabled in the PHY, check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0000_0000_0000_0001]</p> <p>The PHY Status register is at address 0x1 and in this case the PHY Address is 0x00.</p>

Table continues on the next page...

Table 14-924. MII mode register initialization steps (continued)

<p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle]. Set MIIMCOM[Read Cycle]. (Uses the PHY address (0) and Register address (1) placed in MIIMADD register), When MIIMIND[BUSY]=0, read the MIIMSTAT register and check bit 10 (AN Done and Link is up) MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0100] Other information about the link is also returned.(Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Check auto-negotiation attributes in the PHY as necessary.</p>
<p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDRn (Optional) GADDRn[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0-TBASE7, TBASE0-TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0-RBASE7, RBASE0-RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Enable Transmit Queues Initialize TQUEUE</p>
<p>Enable Receive Queues Initialize RQUEUE</p>
<p>Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]</p>

14.7.1.2 GMII interface mode

The table below describes the signal configurations required for GMII interface mode.

Table 14-925. GMII interface mode signal configuration

eTSEC Signal s			GMII Interface		
			Frequency [MHz] 125		
			Voltage [V] 3.3		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1	TX_CLK	I	1
TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	TxD[2]	O	1
TxD[3]	O	1	TxD[3]	O	1
TxD[4]	O	1	TxD[4]	O	1
TxD[5]	O	1	TxD[5]	O	1
TxD[6]	O	1	TxD[6]	O	1
TxD[7]	O	1	TxD[7]	O	1
TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	TX_ER	O	1
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	RxD[2]	I	1
RxD[3]	I	1	RxD[3]	I	1
RxD[4]	I	1	RxD[4]	I	1
RxD[5]	I	1	RxD[5]	I	1
RxD[6]	I	1	RxD[6]	I	1
RxD[7]	I	1	RxD[7]	I	1
RX_DV	I	1	RX_DV	I	1
RX_ER	I	1	RX_ER	I	1
COL	I	1	not used		
CRS	I	1	not used		
Sum		25	Sum		23

The table below describes the shared signals of the GMII interface.

Table 14-926. Shared GMII signals

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O

Table continues on the next page...

Table 14-926. Shared GMII signals (continued)

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDC	O	1	MDC	O	1	Management interface clock
EC GTX_CLK125	I	1	GTX_CLK125	I	1	Reference clock
Sum		3	Sum		3	

The table below describes the register initializations required to configure the eTSEC in GMII mode.

Table 14-927. GMII Mode Register Initialization Steps

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, for GMII, Full duplex operation. Set I/F Mode bit. MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (This example has Full Duplex = 1, Preamble count = 7, PAD/CRC append = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1)
Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] Set station address to 02_60_8C_87_65_43, for example.
Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] Set station address to 02_60_8C_87_65_43, for example.
Assign a Physical address to the TBI so as to not conflict with the external PHY Physical address, TBIPA[0000_0000_0000_0000_0000_0000_0000_0101] Set to 05, for example.
Reset the management interface, MIIMCFG[1000_0000_0000_0000_0000_0000_0000_0111]
Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not less than 2.5 MHz
Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.

Table continues on the next page...

Table 14-927. GMII Mode Register Initialization Steps (continued)

<p>Set up the MII Mgmt for a write cycle to the external PHY Auxiliary Control and Status Register to configure the PHY through the Management interface (overrides configuration signals of the PHY),</p> <p>MIIMADD[0000_0000_0000_0000_0000_0000_0001_1100]</p>
<p>Perform an MII Mgmt write cycle to the external PHY.</p> <p>Write to MII Mgmt Control with 16-bit data intended for the external PHY register,</p> <p>MIIMCON[0000_0000_0000_0000_0000_0000_0000_0100]</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Extended PHY control register #1 to set up the interface mode selection</p> <p>MIIMADD[0000_0000_0000_0000_0000_0000_0001_0111]</p>
<p>Perform an MII Mgmt write cycle to the external PHY.</p> <p>Write to MII Mgmt Control with 16-bit data intended for the external PHY register,</p> <p>MIIMCON[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Mode control register to set up the interface mode selection,</p> <p>MIIMADD[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Perform an MII Mgmt write cycle to the external PHY.</p> <p>Write to MII Mgmt Control with 16-bit data intended for the external PHY register,</p> <p>MIIMCON[0000_0000_0000_0000_0000_000u_00u1_0100_0000]</p> <p>where u is user defined based on desired configuration.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>If auto-negotiation was enabled in the PHY, check to see if PHY has completed Auto-Negotiation.</p> <p>Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0000_0000_0000_0001]</p> <p>The PHY Status register is at address 0x1 and in this case the PHY Address is 0x00</p>

Table continues on the next page...

Table 14-927. GMII Mode Register Initialization Steps (continued)

<p>Perform an MII Mgmt read cycle of Status Register. Clear MIIMCOM[Read Cycle]. Set MIIMCOM[Read Cycle]. (Uses the PHY address (0) and Register address (1) placed in MIIMADD register), When MIIMIND[BUSY]=0, Read the MIIMSTAT register and check bit 10 (AN Done and Link is up), MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0100] Other information about the link is also returned.(Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Check auto-negotiation attributes in the PHY as necessary.</p>
<p>Clear IEVENT register, IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDRn (Optional) GADDRn[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0-TBASE7, TBASE0-TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0-RBASE7, RBASE0-RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>
<p>Enable Transmit Queues Initialize TQUEUE</p>
<p>Enable Receive Queues Initialize RQUEUE</p>
<p>Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]</p>

14.7.1.3 TBI interface mode

The table below describes the signal configurations required for TBI interface mode.

Table 14-928. TBI interface mode signal configuration

eTSEC Signals			TBI Interface		
			Frequency [MHz] 62.5		
			Voltage [V] 3.3		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1	RX_CLK1	I	1
TxD[0]	O	1	TCG[0]	O	1
TxD[1]	O	1	TCG[1]	O	1
TxD[2]	O	1	TCG[2]	O	1
TxD[3]	O	1	TCG[3]	O	1
TxD[4]	O	1	TCG[4]	O	1
TxD[5]	O	1	TCG[5]	O	1
TxD[6]	O	1	TCG[6]	O	1
TxD[7]	O	1	TCG[7]	O	1
TX_EN	O	1	TCG[8]	O	1
TX_ER	O	1	TCG[9]	O	1
RX_CLK	I	1	RX_CLK0	I	1
RxD[0]	I	1	RCG[0]	I	1
RxD[1]	I	1	RCG[1]	I	1
RxD[2]	I	1	RCG[2]	I	1
RxD[3]	I	1	RCG[3]	I	1
RxD[4]	I	1	RCG[4]	I	1
RxD[5]	I	1	RCG[5]	I	1
RxD[6]	I	1	RCG[6]	I	1
RxD[7]	I	1	RCG[7]	I	1
RX_DV	I	1	RCG[8]	I	1
RX_ER	I	1	RCG[9]	I	1
COL	I	1	not used		
CRS	I	1	SDET	I	1
Sum		25	Sum		24

The table below describes the shared signals for the TBI interface.

Table 14-929. Shared TBI signals

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O

Table continues on the next page...

Table 14-929. Shared TBI signals (continued)

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDC	O	1	MDC	O	1	Management interface clock
ECGTX_CLK125	I	1	GTX_CLK125	I	1	Reference clock
Sum		3	Sum		3	

The table below describes the register initializations required to configure the eTSEC in TBI mode.

Table 14-930. TBI mode register initialization steps

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1)
Initialize MAC Station Address MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example.
Initialize MAC Station Address MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example.
Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000] set to 16, for example.
Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not less than 2.5 MHz
Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.
Set up the MII Mgmt for a read cycle to TBI's Control register (write the TBI address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] The control register (CR) is at offset address 0x0 from TBIPA.

Table continues on the next page...

Table 14-930. TBI mode register initialization steps (continued)

<p>Perform an MII Mgmt read cycle to verify state of TBI Control Register(Optional)</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the TBI address and Register address placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MIIMSTAT and look for AN Enable and other bit information.</p>
<p>Set up the MII Mgmt for a write cycle to TBI's AN Advertisement register (write the PHY address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0001_0000_0000_0100]</p> <p>The AN Advertisement register is at offset address 0x04 from the TBI's address. (in this case 0x10)</p>
<p>Perform an MII Mgmt write cycle to TBI.</p> <p>Writing to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register,</p> <p>MIIMCON[0000_0000_0000_0000_0000_0001_1010_0000]</p> <p>This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to TBI's Control register (write the PHY address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000]</p> <p>the control register (CR) is at offset address 0x00 from the TBI's address. (in this case 0x10)</p>
<p>Perform an MII Mgmt write cycle to TBI.</p> <p>Writing to MII Mgmt Control with 16-bit data intended for TBI's Control register,</p> <p>MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000]</p> <p>This enables the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation.</p> <p>Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0001_0000_0000_0001]</p> <p>The PHY Status control register is at address 0x1 and in this case the PHY Address is 0x10.</p>

Table continues on the next page...

Table 14-930. TBI mode register initialization steps (continued)

<p>Perform an MII Mgmt read cycle of Status Register.</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (2) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MIIMSTAT register and check bit 10 (AN Done)</p> <p>MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000]</p> <p>Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0110]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x10) and Register address (6) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Expansion register and check bits 13 and 14 (NP Able and Page Rx'd)</p> <p>MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0101]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x10) and Register address (5) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)</p> <p>MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_0000_000x_x110_0000]</p>
<p>Clear IEVENT register,</p> <p>IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional)</p> <p>IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional)</p> <p>MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDRn (Optional)</p> <p>GADDRn[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional)</p> <p>RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional)</p> <p>DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data</p> <p>Initialize TBASE0-TBASE7,</p> <p>TBASE0-TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>

Table continues on the next page...

Table 14-930. TBI mode register initialization steps (continued)

Initialize (Empty) Receive Descriptor ring and fill with empty buffers
Initialize RBASE0-RBASE7, RBASE0-RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Enable Transmit Queues
Initialize TQUEUE
Enable Receive Queues
Initialize RQUEUE
Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]

14.7.1.4 RGMII interface mode

The table below shows the signals configurations required for RGMII interface mode.

Table 14-931. RGMII interface mode signal configuration

eTSEC signals			RGMII interface		
			Frequency [MHz] 125		
			Voltage [V] 2.5		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1	not used		
TxD[0]	O	1	TxD[0]/TxD[4]	O	1
TxD[1]	O	1	TxD[1]/TxD[5]	O	1
TxD[2]	O	1	TxD[2]/TxD[6]	O	1
TxD[3]	O	1	TxD[3]/TxD[7]	O	1
TxD[4]	O	1	leave unconnected		
TxD[5]	O	1	leave unconnected		
TxD[6]	O	1	leave unconnected		
TxD[7]	O	1	leave unconnected		
TX_EN	O	1	TX_CTL (TX_EN/TX_ERR)	O	1
TX_ER	O	1	leave unconnected		
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RxD[0]/RxD[4]	I	1
RxD[1]	I	1	RxD[1]/RxD[5]	I	1
RxD[2]	I	1	RxD[2]/RxD[6]	I	1
RxD[3]	I	1	RxD[3]/RxD[7]	I	1
RxD[4]	I	1	not used		
RxD[5]	I	1	not used		

Table continues on the next page...

Table 14-931. RGMII interface mode signal configuration (continued)

eTSEC signals			RGMII interface		
			Frequency [MHz] 125		
			Voltage [V] 2.5		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
RxD[6]	I	1	not used		
RxD[7]	I	1	not used		
RX_DV	I	1	RX_CTL (RX_DV/RX_ERR)	I	1
RX_ER	I	1	not used		
COL	I	1	not used		
CRS	I	1	not used		
Sum		25	Sum		12

The table below describes the shared signals for the RGMII interface.

Table 14-932. Shared RGMII signals

eTSEC Signals	I/O	No. of Signals	GMII Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
GTX_CLK125	I	1	GTX_CLK125	I	1	Reference clock
Sum		3	Sum		3	

The table below describes the register initializations required to configure the eTSEC in RGMII mode.

Table 14-933. RGMII mode register initialization steps

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has RGMII 10Mbps mode, Statistics Enable = 1)
Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example.

Table continues on the next page...

Table 14-933. RGMII mode register initialization steps (continued)

<p>Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example.</p>
<p>Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000] set to 16, for example.</p>
<p>Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101] Set source clock divide by 14, for example, to insure that TSEC_MDC clock speed is not less than 2.5 MHz.</p>
<p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.</p>
<p>Set up the MII Mgmt for a write cycle to external the PHY AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0100] The AN Advertisement register is at offset address 0x04 from the external PHY address. (in this case 0x11)</p>
<p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY AN Advertisement register, MIIMCON[0000_0000_0000_0000_u0uu_uuuu_uuuu_uuuu] Where u must be selected by the user for proper system configuration.</p>
<p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0001_0000_0000] The control register (CR) is at offset address 0x00 from the external PHY address. (in this case 0x11)</p>
<p>Perform an MII Mgmt write cycle to the external PHY. Write to MII Mgmt Control with 16-bit data intended for the external PHY Control register, MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000] This enables the external PHY to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation. Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0000_0010_0000_0001] The PHY Status register is at address 0x1 and in this case the PHY Address is 0x2.</p>

Table continues on the next page...

Table 14-933. RGMII mode register initialization steps (continued)

<p>Perform an MII Mgmt read cycle of Status Register.</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (2) placed in MIIMADD register)</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MIIMSTAT register and check bit 10. (AN Done)</p> <p>MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000]</p> <p>Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0110]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x11) and Register address (6) placed in MIIMADD register)</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd)</p> <p>MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0101]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x11) and Register address (5) placed in MIIMADD register)</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)</p> <p>MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_000x_1x10_0000]</p>
<p>Clear IEVENT register,</p> <p>IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional)</p> <p>IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize MACnADDR1/2 (Optional)</p> <p>MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDRn (Optional)</p> <p>GADDRn[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional)</p> <p>RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional)</p> <p>DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize (Empty) Transmit Descriptor ring and fill buffers with Data</p> <p>Initialize TBASE0-TBASE7,</p> <p>TBASE0-TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]</p>

Table continues on the next page...

Table 14-933. RGMII mode register initialization steps (continued)

Initialize (Empty) Receive Descriptor ring and fill with empty buffers
Initialize RBASE0-RBASE7, RBASE0-RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Enable Transmit Queues
Initialize TQUEUE
Enable Receive Queues
Initialize RQUEUE
Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]

14.7.1.5 RMI interface mode

The table below shows the signals configurations required for RMI interface mode.

Table 14-934. RMI interface mode signal configuration

eTSEC Signals			RMI Interface		
			Frequency [MHz] 50		
			Voltage [V] 3.3		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	leave unconnected		
TX_CLK	I	1	REF_CLK	I	1
TxD[0]	O	1	TxD[0]	O	1
TxD[1]	O	1	TxD[1]	O	1
TxD[2]	O	1	leave unconnected		
TxD[3]	O	1	leave unconnected		
TxD[4]	O	1	leave unconnected		
TxD[5]	O	1	leave unconnected		
TxD[6]	O	1	leave unconnected		
TxD[7]	O	1	leave unconnected		
TX_EN	O	1	TX_EN	O	1
TX_ER	O	1	leave unconnected		
RX_CLK	I	1	leave unconnected		
RxD[0]	I	1	RxD[0]	I	1
RxD[1]	I	1	RxD[1]	I	1
RxD[2]	I	1	not used		
RxD[3]	I	1	not used		
RxD[4]	I	1	not used		
RxD[5]	I	1	not used		

Table continues on the next page...

Table 14-934. RMIi interface mode signal configuration (continued)

eTSEC Signals			RMIi Interface		
			Frequency [MHz] 50		
			Voltage [V] 3.3		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
RxD[6]	I	1	not used		
RxD[7]	I	1	not used		
RX_DV	I	1	CRS_DV	I	1
RX_ER	I	1	RX_ER	I	1
COL	I	1	not used		
CRS	I	1	not used		
Sum		25	Sum		8

The table below describes the shared signals for the RMIi interface.

Table 14-935. Shared RMIi signals

eTSEC Signals	I/O	No. of Signals	RMIi Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
TX_CLK	I	1	REF_CLK	I	1	Reference clock
Sum		3	Sum		3	

The table below describes the register initializations required to configure the eTSEC in RMIi mode.

Table 14-936. RMIi mode register initialization steps

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0001_0000_0101] (I/F Mode = 1, Full Duplex = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0001_0000] (Used to setup Reduced-Pin mode = 1, and TBIM = 0, statistics enable = 1)
Initialize MAC Station Address MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543 for example

Table continues on the next page...

Table 14-936. RMI mode register initialization steps (continued)

<p>Initialize MAC Station Address</p> <p>MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100]</p> <p>to 02608C:876543 for example</p>
<p>Setup the MII Mgmt clock speed,</p> <p>MIIMCFG[0000_0000_0000_0000_0000_0000_0000_1101]</p> <p>set system clock divide by 14 for example to insure that MDC clock speed = < 2.5 MHz</p>
<p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the eTSEC MII Mgmt bus is idle.</p>
<p>Set up the MII Mgmt for a write cycle to external the PHY AN Advertisement register (write the PHY address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0001_0001_0000_0100]</p> <p>The AN Advertisement register is at offset address 0x04 from the external PHY address. (in this case 0x11)</p>
<p>Perform an MII Mgmt write cycle to the external PHY.</p> <p>Write to MII Mgmt Control with 16-bit data intended for the external PHY AN Advertisement register,</p> <p>MIIMCON[0000_0000_0000_0000_u0uu_uuuu_uuuu_uuuu]</p> <p>Where u must be selected by the user for proper system configuration.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to the external PHY Control register (write the PHY address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0001_0001_0000_0000]</p> <p>The control register is at offset address 0x00 from the external PHY address. (in this case 0x11)</p>
<p>Perform an MII Mgmt write cycle to the external PHY.</p> <p>Write to MII Mgmt Control with 16-bit data intended for the external PHY Control register,</p> <p>MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000]</p> <p>This enables the external PHY to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation.</p> <p>Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0000_0010_0000_0001]</p> <p>The PHY Status register is at address 0x1 and in this case the PHY Address is 0x2.</p>

Table continues on the next page...

Table 14-936. RMII mode register initialization steps (continued)

<p>Perform an MII Mgmt read cycle of Status Register.</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (1) placed in MIIMADD register)</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MIIMSTAT register and check bit 10. (AN Done)</p> <p>MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000]</p> <p>Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0110]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x11) and Register address (6) placed in MIIMADD register)</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd)</p> <p>MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0001_0000_0101]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x11) and Register address (5) placed in MIIMADD register)</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)</p> <p>MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_0000_000x_x110_0000]</p>
<p>Setting up the MII Mgmt for a write cycle to TBI MII Mgmt register (write the TBI's address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0001_0000_0000_1011]</p> <p>the TBI control register is at offset address 0x11 from TBIPA</p>
<p>Perform an MII Mgmt write cycle</p> <p>Writing to MII Mgmt Control with 16-bit data intended for TBI's MII Mgmt control register (TBI control),</p> <p>MIIMCON[0000_0000_0000_0000_0000_0010_0001_0000]</p> <p>This configures the TBI control to GMII mode and AN sense</p>
<p>Check to see if MII Mgmt write is complete</p> <p>Read MII Mgmt Indicator register and check for Busy = 0,</p> <p>MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicate that the write cycle was completed</p>

Table continues on the next page...

Table 14-936. RMI mode register initialization steps (continued)

<p>Perform an MII Mgmt read cycle (Optional)</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the TBI address and Register address placed in MIIMADD register),</p> <p>read the MIIMSTAT register and verify that</p> <p>MIIMSTAT ---> [0000_0000_0000_0000_0000_0010_0001_0000]</p>
<p>Check to see if PHY has completed Auto-Negotiation</p> <p>Setting up the MII Mgmt for a read cycle to PHY's MII Mgmt register (write the PHY's address and Register address),</p> <p>MIIMADD[0000_0000_0000_0000_0000_0010_0000_0010]</p> <p>the PHY Status control register is at address 0x2 and lets say the PHY Address is 0x2</p>
<p>Perform an MII Mgmt read cycle of Status Register</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (2) placed in MIIMADD register),</p> <p>read the MIIMSTAT register and check bit 10 (AN Done)</p> <p>MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000]</p> <p>other information about the link is also returned (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register</p> <p>MIIMADD[0000_0000_0000_0000_0000_0010_0000_0110]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (6) placed in MIIMADD register),</p> <p>read the MII Mgmt AN Expansion register and check bits 13 and 14 (NP Able and Page Rx'd)</p> <p>MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register (Optional)</p> <p>MIIMADD[0000_0000_0000_0000_0000_0010_0000_0101]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (5) placed in MIIMADD register),</p> <p>read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10 (Half and Full Duplex)</p> <p>MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_0000_00X_1110_0000]</p>
<p>Clear IEVENT register,</p> <p>IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize IMASK (Optional)</p> <p>IMASK[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize GADDRn (Optional)</p> <p>GADDRn[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize RCTRL (Optional)</p> <p>RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>
<p>Initialize DMACTRL (Optional)</p> <p>DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]</p>

Table continues on the next page...

Table 14-936. RMI mode register initialization steps (continued)

Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0-TBASE7, TBASE0-TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0-RBASE7, RBASE0-RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Enable transmit queues Initialize TQUEUE
Enable receive queues Initialize RQUEUE
Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]

14.7.1.6 RTBI interface mode

The table below describes the signal configurations required for RTBI interface mode.

Table 14-937. RTBI interface mode signal configuration

eTSEC Signal s			RTBI Interface		
			Frequency [MHz] 125		
			Voltage [V] 2.5		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
GTX_CLK	O	1	GTX_CLK	O	1
TX_CLK	I	1	not used		
TxD[0]	O	1	TCG[0]/TCG[5]	O	1
TxD[1]	O	1	TCG[1]/TCG[6]	O	1
TxD[2]	O	1	TCG[2]/TCG[7]	O	1
TxD[3]	O	1	TCG[3]/TCG[8]	O	1
TxD[4]	O	1	leave unconnected		
TxD[5]	O	1	leave unconnected		
TxD[6]	O	1	leave unconnected		
TxD[7]	O	1	leave unconnected		
TX_EN	O	1	TCG[4]/TCG[9]	O	1
TX_ER	O	1	leave unconnected		
RX_CLK	I	1	RX_CLK	I	1
RxD[0]	I	1	RCG[0]/RCG[5]	I	1
RxD[1]	I	1	RCG[1]/RCG[6]	I	1
RxD[2]	I	1	RCG[2]/RCG[7]	I	1

Table continues on the next page...

Table 14-937. RTBI interface mode signal configuration (continued)

eTSEC Signal s			RTBI Interface		
			Frequency [MHz] 125		
			Voltage [V] 2.5		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
RxD[3]	I	1	RCG[3]/RCG[8]	I	1
RxD[4]	I	1	not used		
RxD[5]	I	1	not used		
RxD[6]	I	1	not used		
RxD[7]	I	1	not used		
RX_DV	I	1	RCG[4]/RCG[9]	I	1
RX_ER	I	1	not used		
COL	I	1	not used		
CRS	I	1	not used		
Sum		25	sum		12

The table below describes the shared signals for the RTBI interface.

Table 14-938. Shared RTBI signals

eTSEC Signals	I/O	No. of Signals	GMI I Signals	I/O	No. of Signals	Function
MDIO	I/O	1	MDIO	I/O	1	Management interface I/O
MDC	O	1	MDC	O	1	Management interface clock
ECGTX_CLK125	I	1	GTX_CLK125	I	1	Reference clock
Sum		3	Sum		3	

The table below describes the register initializations required to configure the eTSEC in RTBI mode.

Table 14-939. RTBI mode register initialization steps

Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0000_0000] (This example has Statistics Enable = 1)

Table continues on the next page...

Table 14-939. RTBI mode register initialization steps (continued)

<p>Initialize MAC Station Address, MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000]</p> <p>to 02608C:876543, for example.</p>
<p>Initialize MAC Station Address, MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100]</p> <p>to 02608C:876543, for example.</p>
<p>Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000]</p> <p>set to 16, for example.</p>
<p>Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_0000_0000_0000_0101]</p> <p>Set source clock divide by 14, for example, to insure that TSEC_MDC clock speed is not less than 2.5 MHz.</p>
<p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the eTSEC MII Mgmt bus is idle.</p>
<p>Set up the MII Mgmt for a read cycle to TBI's Control register (write the TBI's address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000]</p> <p>The control register (CR) is at offset address 0x0 from TBIPA.</p>
<p>Perform an MII Mgmt read cycle to verify state of TBI Control Register(Optional)</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the TBI address and Register address placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MIIMSTAT and look for AN Enable and other bit information.</p>
<p>Set up the MII Mgmt for a write cycle to TBI's AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0100]</p> <p>The AN Advertisement register is at offset address 0x04 from the TBI's address. (in this case 0x10)</p>
<p>Perform an MII Mgmt write cycle to TBI.</p> <p>Write to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register, MIIMCON[0000_0000_0000_0000_0000_0001_1010_0000]</p> <p>This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000]</p> <p>The control register (CR) is at offset address 0x00 from the TBI's address. (in this case 0x10)</p>

Table continues on the next page...

Table 14-939. RTBI mode register initialization steps (continued)

<p>Perform an MII Mgmt write cycle to TBI.</p> <p>Writing to MII Mgmt Control with 16-bit data intended for TBI's Control register, MIIMCON[0000_0000_0000_0000_0001_0010_0000_0000]</p> <p>This enables the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation.</p> <p>Set up the MII Mgmt for a read cycle to the PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0001]</p> <p>The PHY Status control register is at address 0x1 and in this case the PHY Address is 0x10.</p>
<p>Perform an MII Mgmt read cycle of Status Register.</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (2) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MIIMSTAT register and check bit 10 (AN Done)</p> <p>MIIMSTAT ---> [0000_0000_0000_0000_0000_0000_0010_0000]</p> <p>Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0110]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x10) and Register address (6) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Expansion register and check bits 13 and 14. (NP Able and Page Rx'd)</p> <p>MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0101]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x10) and Register address (5) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY]=0,</p> <p>read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)</p> <p>MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_0000_00x_x110_0000]</p>
<p>Clear IEVENT register,</p> <p>IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>

Table continues on the next page...

Table 14-939. RTBI mode register initialization steps (continued)

Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize GADDRn (Optional) GADDRn[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0-TBASE7, TBASE0-TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0-RBASE7, RBASE0-RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Enable Transmit Queues Initialize TQUEUE
Enable Receive Queues Initialize RQUEUE
Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]

14.7.1.7 SGMII interface support

Table 14-940. SGMII interface signal configuration (4-Wire)

SerDes Signals			SGMII Interface		
Frequency [MHz] 1250			Frequency [MHz] 1250		
Voltage [V] LVDS			Voltage [V] LVDS		
Signals	I/O	No. of Signals	Signals	I/O	No. of Signals
TXn/ TXn̄	O	2	TXD	O	2
RXn/ RXn̄	I	2	RXD	I	2
Sum		4	Sum		4

SGMII mode initialization sequence requires additional initialization for the SerDes. An example of SGMII mode initialization sequence is shown in the table below.

NOTE

SGMII mode utilizes the internal TBI PHY. The internal TBI PHY only auto-negotiates at 1 Gbps. However, 10 Mbps and 100 Mbps speeds are supported in SGMII mode. It is recommended that the external PHY inform the MAC if the desired link speed is not 1 Gbps. Software can perform MII management cycles to determine the external PHY link speed and program ECNTRL and MACCFG2 accordingly.

Table 14-941. SGMII Mode Register Initialization Steps

Initialize SerDes to select SGMII. The initialization sequence should be prepended with SerDes initialization.
Set Soft_Reset, MACCFG1[1000_0000_0000_0000_0000_0000_0000_0000]
Clear Soft_Reset, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACCFG2, MACCFG2[0000_0000_0000_0000_0111_0010_0000_0101] (I/F Mode = 2, Full Duplex = 1) (Set I/F mode = 1 in SGMII 10/100 Mbps speed)
Initialize ECNTRL, ECNTRL[0000_0000_0000_0000_0001_0000_0010_0010] (This example has Statistics Enable = 1, TBIM = 1, SGMIIIM = 1) (Set R100M = 1 in SGMII 100 Mbps speed)
Initialize MAC Station Address MACSTNADDR2[0110_0000_0000_0010_0000_0000_0000_0000] to 02608C:876543, for example.
Initialize MAC Station Address MACSTNADDR1[0100_0011_0110_0101_1000_0111_1000_1100] to 02608C:876543, for example.
Assign a Physical address to the TBI, TBIPA[0000_0000_0000_0000_0000_0000_0001_0000] set to 16, for example.
Setup the MII Mgmt clock speed, MIIMCFG[0000_0000_0000_0000_000_0000_0000_0101] set source clock divide by 14 for example to insure that MDC clock speed is not less than 2.5 MHz
Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the eTSEC MII Mgmt bus is idle.

Table continues on the next page...

Table 14-941. SGMII Mode Register Initialization Steps (continued)

<p>Set up the MII Mgmt for a read cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] the control register (CR) is at offset address 0x00 from the TBI's address.</p>
<p>Perform an MII Mgmt read cycle to verify state of TBI Control Register (optional) Clear MIIMCOM[Read Cycle] Set MIIMCOM[Read Cycle] (Uses the TBI address and Register address placed in MIIMADD register), When MIIMIND[BUSY] = 0, read the MIIMSTAT and look for AN Enable and other bit information.</p>
<p>Set up the MII Mgmt for a write cycle to TBICON register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0001_0001] The TBICON register is at offset address 0x11 from the TBI's address.</p>
<p>Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBICON register, MIIMCON[0000_0000_0000_0000_0000_0000_0010_0000] This sets TBI in single clock mode and MII Mode off to enable communication with SerDes.</p>
<p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p>
<p>Set up the MII Mgmt for a write cycle to TBI's AN Advertisement register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0100] The AN Advertisement register is at offset address 0x04 from the TBI's address.</p>
<p>Perform an MII Mgmt write cycle to TBI. Writing to MII Mgmt Control with 16-bit data intended for TBI's AN Advertisement register, MIIMCON[0000_0000_0000_0000_0000_0001_1010_0000] This advertises to the Link Partner that the TBI supports PAUSE and Full Duplex mode and does not support Half Duplex mode.</p>
<p>Check to see if MII Mgmt write is complete. Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000] This indicates that the write cycle was completed.</p>
<p>Additional SerDes setup as required</p>
<p>Set up the MII Mgmt for a write cycle to TBI's Control register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0000] the control register (CR) is at offset address 0x00 from the TBI's address.</p>

Table continues on the next page...

Table 14-941. SGMII Mode Register Initialization Steps (continued)

<p>Perform an MII Mgmt write cycle to TBI.</p> <p>Writing to MII Mgmt Control with 16-bit data intended for TBI's Control register, MIIMCON[0000_0000_0000_0000_0001_0011_0100_0000]</p> <p>This enables the TBI to restart Auto-Negotiations using the configuration set in the AN Advertisement register.</p>
<p>Check to see if MII Mgmt write is complete.</p> <p>Read MII Mgmt Indicator register and check for Busy = 0, MIIMIND ---> [0000_0000_0000_0000_0000_0000_0000_0000]</p> <p>This indicates that the write cycle was completed.</p>
<p>Check to see if PHY has completed Auto-Negotiation.</p> <p>Set up the MII Mgmt for a read cycle to PHY MII Mgmt register (write the PHY address and Register address), MIIMADD[0000_0000_0000_0000_0001_0000_0000_0001]</p> <p>The PHY Status control register is at address 0x1 and in this case the PHY Address is 0x10.</p>
<p>Perform an MII Mgmt read cycle of Status Register.</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (2) and Register address (2) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY] = 0,</p> <p>read the MIIMSTAT register and check bit 10 (AN Done)</p> <p>MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p> <p>Other information about the link is also returned. (Extend Status, No pre, Remote Fault, An Ability, Link status, extend Ability)</p>
<p>Perform an MII Mgmt read cycle of AN Expansion Register.</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0110]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x10) and Register address (6) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY] = 0,</p> <p>read the MII Mgmt AN Expansion register and check bits 13 and 14 (NP Able and Page Rx'd)</p> <p>MII Mgmt AN Expansion ---> [0000_0000_0000_0000_0000_0000_0000_0110]</p>
<p>Perform an MII Mgmt read cycle of AN Link Partner Base Page Ability Register. (Optional)</p> <p>Setup MIIMADD[0000_0000_0000_0000_0001_0000_0000_0101]</p> <p>Clear MIIMCOM[Read Cycle]</p> <p>Set MIIMCOM[Read Cycle]</p> <p>(Uses the PHY address (0x10) and Register address (5) placed in MIIMADD register),</p> <p>When MIIMIND[BUSY] = 0,</p> <p>read the MII Mgmt AN Link Partner Base Page Ability register and check bits 9 and 10. (Half and Full Duplex)</p> <p>MII Mgmt AN Link Partner Base Page Ability ---> [0000_0000_0000_0000_0000_000x_1110_0000]</p>
<p>Clear IEVENT register,</p> <p>IEVENT[0000_0000_0000_0000_0000_0000_0000_0000]</p>

Table continues on the next page...

Table 14-941. SGMII Mode Register Initialization Steps (continued)

Initialize IMASK (Optional) IMASK[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize MACnADDR1/2 (Optional) MACnADDR1/2[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize GADDRn (Optional) GADDRn[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize RCTRL (Optional) RCTRL[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize DMACTRL (Optional) DMACTRL[0000_0000_0000_0000_0000_0000_0000_0000]
Initialize (Empty) Transmit Descriptor ring and fill buffers with Data Initialize TBASE0-TBASE7, TBASE0-TBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Initialize (Empty) Receive Descriptor ring and fill with empty buffers Initialize RBASE0-RBASE7, RBASE0-RBASE7[LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_LLLL_L000]
Enable Transmit Queues Initialize TQUEUE
Enable Receive Queues Initialize RQUEUE
Enable Rx and Tx, MACCFG1[0000_0000_0000_0000_0000_0000_0000_0101]

Chapter 15

serial RapidIO

The serial RapidIO (SRIO) controller consists of a RapidIO endpoint and the RapidIO messaging unit (RMU).

15.1 Serial RapidIO overview

The serial RapidIO (SRIO) controller consists of a RapidIO endpoint and the RapidIO messaging unit (RMU). Both portions are compliant with the *RapidIO Interconnect Specification, Revision 1.2*.

The serial RapidIO interface provides a RapidIO port and message unit to communicate with other RapidIO devices. This figure shows the RapidIO port and message unit as they interface with OCeaN and with each other.

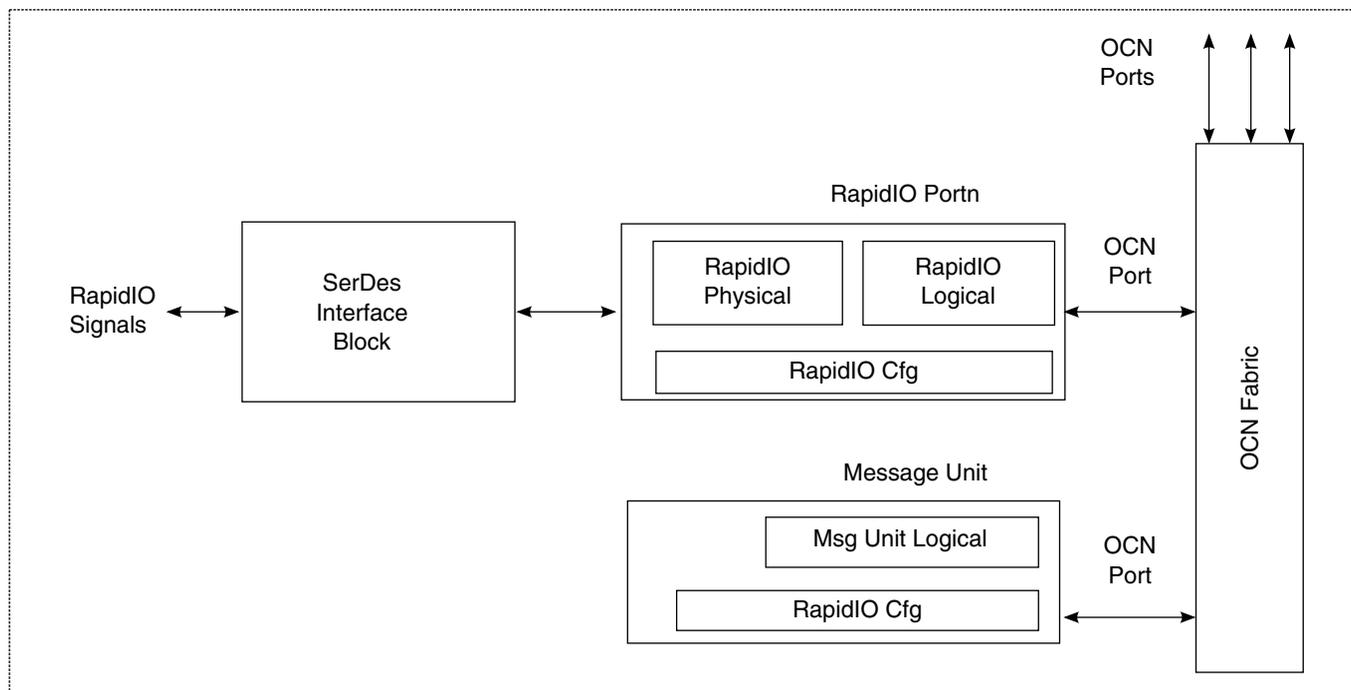


Figure 15-1. RapidIO endpoint and RMU

15.2 Serial RapidIO features summary

The RapidIO port supports the following features of the *RapidIO Interconnect Specification, Revision 1.2*:

- Small or large size transport information field
- 34-bit addressing
- Up to 256-byte data payload
- Up to eight outstanding unacknowledged RapidIO transactions
- Hardware recovery only
- All transaction flows and all priorities
- Register and register bit extensions as described in the *RapidIO Interconnect Specification, Revision 1.2, Part VIII: Error Management Extensions Specification*.
- Hot swap
- ATOMIC set/clr/inc/dec for read-modify-write operations
- IO_READ_HOME and FLUSH with data for accessing cache-coherent data from a remote memory system
- Only supports receiver-controlled flow control
- Inbound transactions to the configuration registers are limited to 32-bit accesses only.
- Outbound maintenance transactions can be any valid size.

The RMU supports the following features of the *RapidIO Interconnect Specification, Revision 1.2*:

- Two outbound message controllers
- Two inbound message controllers
- One outbound doorbell controller
- One inbound doorbell controller
- One inbound port-write controller

RapidIO endpoint supports the following user-defined features:

- Nine outbound ATMU windows with each window having up to 32 subwindows (except for the default window 0)
- Five inbound ATMU windows
- Logical outbound packet time-to-live counter to prevent local processor from hanging when the RIO interface fails
- Accept-all mode of operation for failover support
- RapidIO random bit error injection
- Performance monitor interface

The RMU supports the following user-defined features:

- Performance monitor interface

RapidIO endpoint does not support or has limited support of the following features of the *RapidIO Interconnect Specification, Revision 1.2*:

- No support for 50- and 66-bit addressing
- No support for software assisted error recovery
- No support for ATOMIC test-and-swap transaction
- No support for coherent (CC-NUMA) transactions with the exception of IO_READ_HOME and FLUSH with data transactions
- No support for transmitter-controlled flow control
- No decrementing of a maintenance packet hop count
- No support for multicast event control symbols

RapidIO endpoint supports the following features of RapidIO LP-Serial:

- 1x and 4x LP-Serial link interfaces
- Transmission rates of 1.25, 2.5, 3.125 Gbaud (data rates of 1.0, 2.0, 2.5 Gbps) per lane
- Auto detection of 1x and 4x mode operation during port initialization
- Error detection for packets and control symbols
- Support for link initialization, synchronization, error recovery, and time-out

RapidIO endpoint does not support the following features of RapidIO LP-Serial:

- RapidIO endpoint cannot be configured as four 1x ports

15.3 Serial RapidIO modes of operation

This section provides information on the Serial RapidIO modes of operation.

15.3.1 RapidIO port

The RapidIO port's primary operating modes are the following:

- 1x or 4x LP-Serial link interfaces
- Transmission rates of 1.25, 2.5, 3.125 Gbaud (data rates of 1.0, 2.0, 2.5 Gbps) per lane
- Small or large size transport information field
- Accept-all mode of operation-all packets are accepted regardless of the target ID

15.3.2 Message unit

The message unit's primary operating modes are the following:

- Outbound message controller
 - Direct mode. No descriptors are involved. Software must initialize the required fields before starting a transfer.
 - Chaining mode. Software must initialize descriptors in memory and the required fields before starting a transfer.
 - Multicast mode. Single segment messages can be transferred to up to 32 destinations.

15.4 LP-serial signal descriptions

The high-speed interface signals used for the serial RapidIO interface are shared with other I/O ports and must be configured at power-on reset for use with the serial RapidIO controller.

See [Signal Descriptions](#) for specific configuration details.

The following sections describe the serial RapidIO signal functionality. Refer to the *RapidIO Interconnect Specification, Revision 1.2, Part VI: Physical Layer 1x/ , Chapter 8, Electrical Specifications*, for electrical characteristic details.

15.4.1 Serial Rapid I/O overview

The serial Rapid I/O interface is compatible with the *RapidIO Interconnect Specification, Revision 1.2, Part VI: Physical Layer 1x/4x LP-Serial Specifications* .

15.4.2 Serial Rapid I/O detailed signal descriptions

15.4.2.1 SD_TXn/SD_TX n_B-outputs

These are the serial data outputs. They are differential pairs, one for each lane in 4x mode of operation. Please refer to the Reset and Signals chapters of the device reference manual for external signal mapping and configuration information. In the case of 1x mode of operation on the serial RapidIO interface, only the first SRIO lane is used. For information on selection of SRIO x1 lanes, see [Table 4-14](#).

See *Part VI: Physical Layer 1x/4x LP-Serial Specification, RapidIO Interconnect Specification, Revision 1.2*, for complete details.

These outputs are asynchronous, as described in Part VI of the *RapidIO Interconnect Specification, Revision 1.2*. This implementation supports data rates of 1.25, 2.5, 3.125 Gbaud.

15.4.2.2 SD_RXn/SD_RX n_B-inputs

These are the serial data input pads. They are differential pairs, one for each lane in 4x mode of operation.

In the case of 1x mode of operation on the serial RapidIO interface, only the first SRIO lane is used. For information on selection of SRIO x1 lanes, see [Table 4-14](#).

See *Part VI: Physical Layer 1x/4x LP-Serial Specification, RapidIO Interconnect Specification, Revision 1.2*. for complete details.

These inputs are asynchronous, as described in Part VI of the *RapidIO Interconnect Specification, Revision 1.2*. This implementation supports data rates of 1.25, 2.5, 3.125 Gbaud.

15.5 Serial RapidIO Memory Map/Register Definition

The following table is the memory map of the RapidIO configuration registers .

SRIO memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C_0000	Device identity capability register (SRIO_DIDCAR)	32	R	See section	15.5.1/1040
C_0004	Device information capability register (SRIO_DICAR)	32	R	See section	15.5.2/1041

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C_0008	Assembly identity capability register (SRIO_AIDCAR)	32	R/W	0000_0000h	15.5.3/1041
C_000C	Assembly information capability register (SRIO_AICAR)	32	R/W	0000_0100h	15.5.4/1042
C_0010	Processing element features capability register (SRIO_PEF CAR)	32	R	See section	15.5.5/1043
C_0018	Source operations capability register (SRIO_SOCAR)	32	R	0600_FCF0h	15.5.6/1045
C_001C	Destination operations capability register (SRIO_DOCAR)	32	R	0000_FCF4h	15.5.7/1048
C_0040	Mailbox command and status register (SRIO_MCSR)	32	R	2020_0000h	15.5.8/1051
C_0044	Port -Write and doorbell command and status register (SRIO_PWDCSR)	32	R	2000_0020h	15.5.9/1053
C_004C	Processing element logical layer control command and status register (SRIO_PELLCCSR)	32	R	0000_0001h	15.5.10/1055
C_005C	Local configuration space base address 1 command and status register (SRIO_LCSBA1CSR)	32	R/W	0000_0000h	15.5.11/1056
C_0060	Base device ID command and status register (SRIO_BDIDCSR)	32	R/W	See section	15.5.12/1057
C_0068	Host base device ID lock command and status register (SRIO_HBDIDLCSR)	32	R/W	0000_FFFFh	15.5.13/1058
C_006C	Component tag command and status register (SRIO_CTCSR)	32	R/W	0000_0000h	15.5.14/1058
C_0100	Port maintenance block header 0 (SRIO_PMBH0)	32	R	0600_0001h	15.5.15/1059
C_0120	Port link time-out control command and status register (SRIO_PLTOCCSR)	32	R/W	FFFF_FF00h	15.5.16/1059
C_0124	Port response time-out control command and status register (SRIO_PRTOCCSR)	32	R/W	FFFF_FF00h	15.5.17/1060
C_013C	Port General control command and status register (SRIO_PGCCSR)	32	R/W	See section	15.5.18/1061
C_0140	Port 1 Link maintenance request command and status register (SRIO_P1LMREQCSR)	32	R/W	0000_0000h	15.5.19/1062
C_0144	Port 1 Link maintenance response command and status register (SRIO_P1LMRESPCSR)	32	R	0000_0000h	15.5.20/1062
C_0148	Port 1 Local ackID status command and status register (SRIO_P1LASCSR)	32	R/W	0000_0000h	15.5.21/1063
C_0158	Port 1 Error and status command and status register (SRIO_P1ESCSR)	32	R/W	0000_0001h	15.5.22/1065
C_015C	Port 1 Control command and status register (SRIO_P1CCSR)	32	R/W	5060_0001h	15.5.23/1068
C_0160	Port 2 Link maintenance request command and status register (SRIO_P2LMREQCSR)	32	R/W	0000_0000h	15.5.24/1070
C_0164	Port 2 Link maintenance response command and status register (SRIO_P2LMRESPCSR)	32	R	0000_0000h	15.5.25/1070
C_0168	Port 2 Local ackID status command and status register (SRIO_P2LASCSR)	32	R/W	0000_0000h	15.5.26/1071

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C_0178	Port 2 Error and status command and status register (SRIO_P2ESCSR)	32	R/W	0000_0001h	15.5.27/1073
C_017C	Port 2 Control command and status register (SRIO_P2CCSR)	32	R/W	5060_0001h	15.5.28/1076
C_0600	Error reporting block header (SRIO_ERBH)	32	R	0000_0007h	15.5.29/1078
C_0608	Logical/Transport layer error detect command and status register (SRIO_LTLDCSR)	32	R/W	0000_0000h	15.5.30/1078
C_060C	Logical/Transport layer error enable command and status register (SRIO_LTLECSR)	32	R/W	0000_0000h	15.5.31/1081
C_0614	Logical/Transport layer address capture command and status register (SRIO_LTLACCSR)	32	R/W	0000_0000h	15.5.32/1082
C_0618	Logical/Transport layer device ID capture command and status register (SRIO_LTLIDCCSR)	32	R/W	0000_0000h	15.5.33/1083
C_061C	Logical/Transport layer control capture command and status register (SRIO_LTLCCSR)	32	R/W	0000_0000h	15.5.34/1084
C_0640	Port 1 Error detect command and status register (SRIO_P1EDCSR)	32	R/W	0000_0000h	15.5.35/1085
C_0644	Port 1 Error rate enable command and status register (SRIO_P1ERCSR)	32	R/W	0000_0000h	15.5.36/1086
C_0648	Port 1 Error capture attributes command and status register (SRIO_P1ECACSR)	32	R/W	0000_0000h	15.5.37/1088
C_064C	Port 1 Packet/control symbol error capture command and status register 0 (SRIO_P1PCSECCSR0)	32	R/W	0000_0000h	15.5.38/1089
C_0650	Port 1 Packet error capture command and status register 1 (SRIO_P1PECCSR1)	32	R/W	0000_0000h	15.5.39/1090
C_0654	Port 1 Packet error capture command and status register 2 (SRIO_P1PECCSR2)	32	R/W	0000_0000h	15.5.40/1090
C_0658	Port 1 Packet error capture command and status register 3 (SRIO_P1PECCSR3)	32	R/W	0000_0000h	15.5.41/1091
C_0668	Port 1 Error rate command and status register (SRIO_P1ERCSR)	32	R/W	8000_0000h	15.5.42/1091
C_066C	Port 1 Error rate threshold command and status register (SRIO_P1ERTCSR)	32	R/W	FFFF_0000h	15.5.43/1092
C_0680	Port 2 Error detect command and status register (SRIO_P2EDCSR)	32	R/W	0000_0000h	15.5.44/1093
C_0684	Port 2 Error rate enable command and status register (SRIO_P2ERCSR)	32	R/W	0000_0000h	15.5.45/1095
C_0688	Port 2 Error capture attributes command and status register (SRIO_P2ECACSR)	32	R/W	0000_0000h	15.5.46/1096
C_068C	Port 2 Packet/control symbol error capture command and status register 0 (SRIO_P2PCSECCSR0)	32	R/W	0000_0000h	15.5.47/1097
C_0690	Port 2 Packet error capture command and status register 1 (SRIO_P2PECCSR1)	32	R/W	0000_0000h	15.5.48/1098

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
C_0694	Port 2 Packet error capture command and status register 2 (SRIO_P2PECCSR2)	32	R/W	0000_0000h	15.5.49/ 1098
C_0698	Port 2 Packet error capture command and status register 3 (SRIO_P2PECCSR3)	32	R/W	0000_0000h	15.5.50/ 1099
C_06A8	Port 2 Error rate command and status register (SRIO_P2ERCSR)	32	R/W	8000_0000h	15.5.51/ 1099
C_06AC	Port 2 Error rate threshold command and status register (SRIO_P2ERTCSR)	32	R/W	FFFF_0000h	15.5.52/ 1100
D_0004	Logical layer configuration register (SRIO_LLCR)	32	R/W	0000_0000h	15.5.53/ 1101
D_0010	Error / port-write interrupt status register (SRIO_EPWISR)	32	R	0000_0000h	15.5.54/ 1103
D_0020	Logical retry error threshold configuration register (SRIO_LRETCR)	32	R/W	0000_00FFh	15.5.55/ 1104
D_0080	Physical retry error threshold configuration register (SRIO_PRETCR)	32	R/W	0000_00FFh	15.5.56/ 1105
D_0100	Port 1 Alternate device ID command and status register (SRIO_P1ADIDCSR)	32	R/W	0000_0000h	15.5.57/ 1105
D_0120	Port 1 accept-all configuration register (SRIO_P1AACR)	32	R/W	1000_0000h	15.5.58/ 1106
D_0124	Port 1 Logical Outbound Packet time-to-live configuration register (SRIO_P1LOPTTLCR)	32	R/W	0000_0000h	15.5.59/ 1107
D_0130	Port 1 Implementation error command and status register (SRIO_P1IECSR)	32	w1c	0000_0000h	15.5.60/ 1108
D_0140	Port 1 Physical configuration register (SRIO_P1PCR)	32	R/W	0000_8010h	15.5.61/ 1109
D_0158	Port 1 Serial link command and status register (SRIO_P1SLCSR)	32	w1c	0000_0000h	15.5.62/ 1110
D_0160	Port 1 Serial link error injection configuration register (SRIO_P1SLEICR)	32	R/W	0000_0000h	15.5.63/ 1111
D_0180	Port 2 Alternate device ID command and status register (SRIO_P2ADIDCSR)	32	R/W	0000_0000h	15.5.64/ 1112
D_01A0	Port 2 accept-all configuration register (SRIO_P2AACR)	32	R/W	1000_0000h	15.5.65/ 1113
D_01A4	Port 2 Logical Outbound Packet time-to-live configuration register (SRIO_P2LOPTTLCR)	32	R/W	0000_0000h	15.5.66/ 1113
D_01B0	Port 2 Implementation error command and status register (SRIO_P2IECSR)	32	w1c	0000_0000h	15.5.67/ 1114
D_01C0	Port 2 Physical configuration register (SRIO_P2PCR)	32	R/W	0000_8010h	15.5.68/ 1115
D_01D8	Port 2 Serial link command and status register (SRIO_P2SLCSR)	32	w1c	0000_0000h	15.5.69/ 1116
D_01E0	Port 2 Serial link error injection configuration register (SRIO_P2SLEICR)	32	R/W	0000_0000h	15.5.70/ 1117

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_0BF8	IP Block Revision Register 1 (SRIO_IPBRR1)	32	R	01C0_0102h	15.5.71/1118
D_0BFC	IP Block Revision Register 2 (SRIO_IPBRR2)	32	R	0000_0001h	15.5.72/1119
D_0C00	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR0)	32	R/W	FF80_0000h	15.5.73/1119
D_0C04	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR0)	32	R/W	0000_003Fh	15.5.74/1121
D_0C10	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR0)	32	R/W	8004_4023h	15.5.75/1122
D_0C20	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR1)	32	R/W	0000_0000h	15.5.76/1125
D_0C24	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR1)	32	R/W	0000_003Fh	15.5.74/1121
D_0C28	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR1)	32	R/W	0000_0000h	15.5.77/1126
D_0C30	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR1)	32	R/W	8004_4023h	15.5.75/1122
D_0C34	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R1)	32	R/W	0044_0000h	15.5.78/1127
D_0C38	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R1)	32	R/W	0044_0000h	15.5.79/1129
D_0C3C	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R1)	32	R/W	0044_0000h	15.5.80/1131
D_0C40	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR2)	32	R/W	0000_0000h	15.5.76/1125
D_0C44	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR2)	32	R/W	0000_003Fh	15.5.74/1121
D_0C48	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR2)	32	R/W	0000_0000h	15.5.77/1126
D_0C50	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR2)	32	R/W	8004_4023h	15.5.75/1122
D_0C54	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R2)	32	R/W	0044_0000h	15.5.78/1127
D_0C58	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R2)	32	R/W	0044_0000h	15.5.79/1129
D_0C5C	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R2)	32	R/W	0044_0000h	15.5.80/1131
D_0C60	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR3)	32	R/W	0000_0000h	15.5.76/1125
D_0C64	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR3)	32	R/W	0000_003Fh	15.5.74/1121
D_0C68	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR3)	32	R/W	0000_0000h	15.5.77/1126

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_0C70	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR3)	32	R/W	8004_4023h	15.5.75/ 1122
D_0C74	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R3)	32	R/W	0044_0000h	15.5.78/ 1127
D_0C78	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R3)	32	R/W	0044_0000h	15.5.79/ 1129
D_0C7C	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R3)	32	R/W	0044_0000h	15.5.80/ 1131
D_0C80	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR4)	32	R/W	0000_0000h	15.5.76/ 1125
D_0C84	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR4)	32	R/W	0000_003Fh	15.5.74/ 1121
D_0C88	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR4)	32	R/W	0000_0000h	15.5.77/ 1126
D_0C90	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR4)	32	R/W	8004_4023h	15.5.75/ 1122
D_0C94	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R4)	32	R/W	0044_0000h	15.5.78/ 1127
D_0C98	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R4)	32	R/W	0044_0000h	15.5.79/ 1129
D_0C9C	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R4)	32	R/W	0044_0000h	15.5.80/ 1131
D_0CA0	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR5)	32	R/W	0000_0000h	15.5.76/ 1125
D_0CA4	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR5)	32	R/W	0000_003Fh	15.5.74/ 1121
D_0CA8	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR5)	32	R/W	0000_0000h	15.5.77/ 1126
D_0CB0	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR5)	32	R/W	8004_4023h	15.5.75/ 1122
D_0CB4	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R5)	32	R/W	0044_0000h	15.5.78/ 1127
D_0CB8	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R5)	32	R/W	0044_0000h	15.5.79/ 1129
D_0CBC	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R5)	32	R/W	0044_0000h	15.5.80/ 1131
D_0CC0	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR6)	32	R/W	0000_0000h	15.5.76/ 1125
D_0CC4	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR6)	32	R/W	0000_003Fh	15.5.74/ 1121
D_0CC8	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR6)	32	R/W	0000_0000h	15.5.77/ 1126
D_0CD0	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR6)	32	R/W	8004_4023h	15.5.75/ 1122

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_0CD4	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R6)	32	R/W	0044_0000h	15.5.78/ 1127
D_0CD8	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R6)	32	R/W	0044_0000h	15.5.79/ 1129
D_0CDC	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R6)	32	R/W	0044_0000h	15.5.80/ 1131
D_0CE0	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR7)	32	R/W	0000_0000h	15.5.76/ 1125
D_0CE4	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR7)	32	R/W	0000_003Fh	15.5.74/ 1121
D_0CE8	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR7)	32	R/W	0000_0000h	15.5.77/ 1126
D_0CF0	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR7)	32	R/W	8004_4023h	15.5.75/ 1122
D_0CF4	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R7)	32	R/W	0044_0000h	15.5.78/ 1127
D_0CF8	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R7)	32	R/W	0044_0000h	15.5.79/ 1129
D_0CFC	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R7)	32	R/W	0044_0000h	15.5.80/ 1131
D_0D00	Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR8)	32	R/W	0000_0000h	15.5.76/ 1125
D_0D04	Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR8)	32	R/W	0000_003Fh	15.5.74/ 1121
D_0D08	Port 1 RapidIO outbound window base address register n (SRIO_P1ROWBAR8)	32	R/W	0000_0000h	15.5.77/ 1126
D_0D10	Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWAR8)	32	R/W	8004_4023h	15.5.75/ 1122
D_0D14	Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1R8)	32	R/W	0044_0000h	15.5.78/ 1127
D_0D18	Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2R8)	32	R/W	0044_0000h	15.5.79/ 1129
D_0D1C	Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3R8)	32	R/W	0044_0000h	15.5.80/ 1131
D_0D60	Port 1 RapidIO Inbound window translation address register n (SRIO_P1RIWTAR4)	32	R/W	0000_0000h	15.5.81/ 1133
D_0D68	Port 1 RapidIO Inbound window base address register n (SRIO_P1RIWBAR4)	32	R/W	0000_0000h	15.5.82/ 1134
D_0D70	Port 1 RapidIO inbound window attributes register n (SRIO_P1RIWAR4)	32	R/W	0004_4021h	15.5.83/ 1135
D_0D80	Port 1 RapidIO Inbound window translation address register n (SRIO_P1RIWTAR3)	32	R/W	0000_0000h	15.5.81/ 1133
D_0D88	Port 1 RapidIO Inbound window base address register n (SRIO_P1RIWBAR3)	32	R/W	0000_0000h	15.5.82/ 1134

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_0D90	Port 1 RapidIO inbound window attributes register n (SRIO_P1RIWAR3)	32	R/W	0004_4021h	15.5.83/ 1135
D_0DA0	Port 1 RapidIO Inbound window translation address register n (SRIO_P1RIWTAR2)	32	R/W	0000_0000h	15.5.81/ 1133
D_0DA8	Port 1 RapidIO Inbound window base address register n (SRIO_P1RIWBAR2)	32	R/W	0000_0000h	15.5.82/ 1134
D_0DB0	Port 1 RapidIO inbound window attributes register n (SRIO_P1RIWAR2)	32	R/W	0004_4021h	15.5.83/ 1135
D_0DC0	Port 1 RapidIO Inbound window translation address register n (SRIO_P1RIWTAR1)	32	R/W	0000_0000h	15.5.81/ 1133
D_0DC8	Port 1 RapidIO Inbound window base address register n (SRIO_P1RIWBAR1)	32	R/W	0000_0000h	15.5.82/ 1134
D_0DD0	Port 1 RapidIO inbound window attributes register n (SRIO_P1RIWAR1)	32	R/W	0004_4021h	15.5.83/ 1135
D_0DE0	Port 1 RapidIO Inbound window translation address register n (SRIO_P1RIWTAR0)	32	R/W	0000_0000h	15.5.81/ 1133
D_0DF0	Port 1 RapidIO inbound window attributes register 0 (SRIO_P1RIWAR0)	32	R/W	8004_4021h	15.5.84/ 1137
D_0E00	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR0)	32	R/W	FF80_0000h	15.5.85/ 1140
D_0E04	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR0)	32	R/W	0000_003Fh	15.5.86/ 1141
D_0E10	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR0)	32	R/W	8004_4023h	15.5.87/ 1142
D_0E20	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR1)	32	R/W	FF80_0000h	15.5.88/ 1145
D_0E24	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR1)	32	R/W	0000_003Fh	15.5.86/ 1141
D_0E28	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR1)	32	R/W	0000_0000h	15.5.89/ 1146
D_0E30	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR1)	32	R/W	8004_4023h	15.5.87/ 1142
D_0E34	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R1)	32	R/W	0044_0000h	15.5.90/ 1147
D_0E38	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R1)	32	R/W	0044_0000h	15.5.91/ 1149
D_0E3C	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R1)	32	R/W	0044_0000h	15.5.92/ 1151
D_0E40	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR2)	32	R/W	FF80_0000h	15.5.88/ 1145
D_0E44	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR2)	32	R/W	0000_003Fh	15.5.86/ 1141
D_0E48	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR2)	32	R/W	0000_0000h	15.5.89/ 1146

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_0E50	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR2)	32	R/W	8004_4023h	15.5.87/1142
D_0E54	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R2)	32	R/W	0044_0000h	15.5.90/1147
D_0E58	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R2)	32	R/W	0044_0000h	15.5.91/1149
D_0E5C	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R2)	32	R/W	0044_0000h	15.5.92/1151
D_0E60	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR3)	32	R/W	FF80_0000h	15.5.88/1145
D_0E64	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR3)	32	R/W	0000_003Fh	15.5.86/1141
D_0E68	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR3)	32	R/W	0000_0000h	15.5.89/1146
D_0E70	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR3)	32	R/W	8004_4023h	15.5.87/1142
D_0E74	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R3)	32	R/W	0044_0000h	15.5.90/1147
D_0E78	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R3)	32	R/W	0044_0000h	15.5.91/1149
D_0E7C	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R3)	32	R/W	0044_0000h	15.5.92/1151
D_0E80	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR4)	32	R/W	FF80_0000h	15.5.88/1145
D_0E84	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR4)	32	R/W	0000_003Fh	15.5.86/1141
D_0E88	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR4)	32	R/W	0000_0000h	15.5.89/1146
D_0E90	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR4)	32	R/W	8004_4023h	15.5.87/1142
D_0E94	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R4)	32	R/W	0044_0000h	15.5.90/1147
D_0E98	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R4)	32	R/W	0044_0000h	15.5.91/1149
D_0E9C	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R4)	32	R/W	0044_0000h	15.5.92/1151
D_0EA0	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR5)	32	R/W	FF80_0000h	15.5.88/1145
D_0EA4	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR5)	32	R/W	0000_003Fh	15.5.86/1141
D_0EA8	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR5)	32	R/W	0000_0000h	15.5.89/1146
D_0EB0	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR5)	32	R/W	8004_4023h	15.5.87/1142

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_0EB4	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R5)	32	R/W	0044_0000h	15.5.90/ 1147
D_0EB8	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R5)	32	R/W	0044_0000h	15.5.91/ 1149
D_0EBC	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R5)	32	R/W	0044_0000h	15.5.92/ 1151
D_0EC0	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR6)	32	R/W	FF80_0000h	15.5.88/ 1145
D_0EC4	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR6)	32	R/W	0000_003Fh	15.5.86/ 1141
D_0EC8	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR6)	32	R/W	0000_0000h	15.5.89/ 1146
D_0ED0	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR6)	32	R/W	8004_4023h	15.5.87/ 1142
D_0ED4	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R6)	32	R/W	0044_0000h	15.5.90/ 1147
D_0ED8	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R6)	32	R/W	0044_0000h	15.5.91/ 1149
D_0EDC	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R6)	32	R/W	0044_0000h	15.5.92/ 1151
D_0EE0	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR7)	32	R/W	FF80_0000h	15.5.88/ 1145
D_0EE4	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR7)	32	R/W	0000_003Fh	15.5.86/ 1141
D_0EE8	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR7)	32	R/W	0000_0000h	15.5.89/ 1146
D_0EF0	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR7)	32	R/W	8004_4023h	15.5.87/ 1142
D_0EF4	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R7)	32	R/W	0044_0000h	15.5.90/ 1147
D_0EF8	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R7)	32	R/W	0044_0000h	15.5.91/ 1149
D_0EFC	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R7)	32	R/W	0044_0000h	15.5.92/ 1151
D_0F00	Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR8)	32	R/W	FF80_0000h	15.5.88/ 1145
D_0F04	Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR8)	32	R/W	0000_003Fh	15.5.86/ 1141
D_0F08	Port 2 RapidIO outbound window base address register n (SRIO_P2ROWBAR8)	32	R/W	0000_0000h	15.5.89/ 1146
D_0F10	Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR8)	32	R/W	8004_4023h	15.5.87/ 1142
D_0F14	Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1R8)	32	R/W	0044_0000h	15.5.90/ 1147

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_0F18	Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2R8)	32	R/W	0044_0000h	15.5.91/ 1149
D_0F1C	Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3R8)	32	R/W	0044_0000h	15.5.92/ 1151
D_0F60	Port 2 RapidIO Inbound window translation address register n (SRIO_P2RIWTAR4)	32	R/W	0000_0000h	15.5.93/ 1153
D_0F68	Port 2 RapidIO Inbound window base address register n (SRIO_P2RIWBAR4)	32	R/W	0000_0000h	15.5.94/ 1154
D_0F70	Port 2 RapidIO inbound window attributes register n (SRIO_P2RIWAR4)	32	R/W	0004_4021h	15.5.95/ 1155
D_0F80	Port 2 RapidIO Inbound window translation address register n (SRIO_P2RIWTAR3)	32	R/W	0000_0000h	15.5.93/ 1153
D_0F88	Port 2 RapidIO Inbound window base address register n (SRIO_P2RIWBAR3)	32	R/W	0000_0000h	15.5.94/ 1154
D_0F90	Port 2 RapidIO inbound window attributes register n (SRIO_P2RIWAR3)	32	R/W	0004_4021h	15.5.95/ 1155
D_0FA0	Port 2 RapidIO Inbound window translation address register n (SRIO_P2RIWTAR2)	32	R/W	0000_0000h	15.5.93/ 1153
D_0FA8	Port 2 RapidIO Inbound window base address register n (SRIO_P2RIWBAR2)	32	R/W	0000_0000h	15.5.94/ 1154
D_0FB0	Port 2 RapidIO inbound window attributes register n (SRIO_P2RIWAR2)	32	R/W	0004_4021h	15.5.95/ 1155
D_0FC0	Port 2 RapidIO Inbound window translation address register n (SRIO_P2RIWTAR1)	32	R/W	0000_0000h	15.5.93/ 1153
D_0FC8	Port 2 RapidIO Inbound window base address register n (SRIO_P2RIWBAR1)	32	R/W	0000_0000h	15.5.94/ 1154
D_0FD0	Port 2 RapidIO inbound window attributes register n (SRIO_P2RIWAR1)	32	R/W	0004_4021h	15.5.95/ 1155
D_0FE0	Port 2 RapidIO Inbound window translation address register n (SRIO_P2RIWTAR0)	32	R/W	0000_0000h	15.5.93/ 1153
D_0FF0	Port 2 RapidIO inbound window attributes register 0 (SRIO_P2RIWAR0)	32	R/W	8004_4021h	15.5.96/ 1157
D_3000	Outbound message n mode register (SRIO_OM0MR)	32	R/W	0000_0000h	15.5.97/ 1159
D_3004	Outbound message n status register (SRIO_OM0SR)	32	R/W	0000_0000h	15.5.98/ 1163
D_3008	Extended outbound message n descriptor queue dequeue pointer address register (SRIO_EOM0DQDPAR)	32	R/W	0000_0000h	15.5.99/ 1165
D_300C	Outbound message n descriptor queue dequeue pointer address register (SRIO_OM0DQDPAR)	32	R/W	0000_0000h	15.5.100/ 1166
D_3010	Extended outbound message n source address register (SRIO_EOM0SAR)	32	R/W	0000_0000h	15.5.101/ 1167
D_3014	Outbound message n source address register (SRIO_OM0SAR)	32	R/W	0000_0000h	15.5.102/ 1167

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_3018	Outbound message n destination port register (SRIO_OM0DPR)	32	R/W	0000_0000h	15.5.103/ 1168
D_301C	Outbound message n destination attributes Register (SRIO_OM0DATR)	32	R/W	0000_0000h	15.5.104/ 1169
D_3020	Outbound message n double-word count register (SRIO_OM0DCR)	32	R/W	0000_0000h	15.5.105/ 1170
D_3024	Extended outbound message n descriptor queue enqueue pointer address register (SRIO_EOM0DQEPAR)	32	R/W	0000_0000h	15.5.106/ 1171
D_3028	Outbound message n descriptor queue enqueue pointer address register (SRIO_OM0DQEPAR)	32	R/W	0000_0000h	15.5.107/ 1171
D_302C	Outbound message n retry error threshold configuration register (SRIO_OM0RETCR)	32	R/W	0000_0000h	15.5.108/ 1173
D_3030	Outbound message n multicast group register (SRIO_OM0MGR)	32	R/W	0000_0000h	15.5.109/ 1174
D_3034	Outbound message n multicast list register (SRIO_OM0MLR)	32	R/W	0000_0000h	15.5.110/ 1175
D_3060	Inbound message n mode register (SRIO_IM0MR)	32	R/W	0000_0000h	15.5.111/ 1176
D_3064	Inbound message n status register (SRIO_IM0SR)	32	R/W	0000_0002h	15.5.112/ 1179
D_3068	Extended inbound message n frame queue dequeue pointer address register (SRIO_EIM0FQDPPAR)	32	R/W	0000_0000h	15.5.113/ 1181
D_306C	Inbound message n frame queue dequeue pointer address register (SRIO_IM0FQDPPAR)	32	R/W	0000_0000h	15.5.114/ 1181
D_3070	Extended inbound message n frame queue enqueue pointer address register (SRIO_EIM0FQEPAR)	32	R/W	0000_0000h	15.5.115/ 1182
D_3074	Inbound message n frame queue enqueue pointer address register (SRIO_IM0FQEPAR)	32	R/W	0000_0000h	15.5.116/ 1183
D_3078	Inbound message n maximum interrupt report interval register (SRIO_IM0MIRIR)	32	R/W	FFFF_FF00h	15.5.117/ 1184
D_3100	Outbound message n mode register (SRIO_OM1MR)	32	R/W	0000_0000h	15.5.97/ 1159
D_3104	Outbound message n status register (SRIO_OM1SR)	32	R/W	0000_0000h	15.5.98/ 1163
D_3108	Extended outbound message n descriptor queue dequeue pointer address register (SRIO_EOM1DQDPPAR)	32	R/W	0000_0000h	15.5.99/ 1165
D_310C	Outbound message n descriptor queue dequeue pointer address register (SRIO_OM1DQDPPAR)	32	R/W	0000_0000h	15.5.100/ 1166
D_3110	Extended outbound message n source address register (SRIO_EOM1SAR)	32	R/W	0000_0000h	15.5.101/ 1167
D_3114	Outbound message n source address register (SRIO_OM1SAR)	32	R/W	0000_0000h	15.5.102/ 1167
D_3118	Outbound message n destination port register (SRIO_OM1DPR)	32	R/W	0000_0000h	15.5.103/ 1168

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_311C	Outbound message n destination attributes Register (SRIO_OM1DATR)	32	R/W	0000_0000h	15.5.104/ 1169
D_3120	Outbound message n double-word count register (SRIO_OM1DCR)	32	R/W	0000_0000h	15.5.105/ 1170
D_3124	Extended outbound message n descriptor queue enqueue pointer address register (SRIO_EOM1DQEPAR)	32	R/W	0000_0000h	15.5.106/ 1171
D_3128	Outbound message n descriptor queue enqueue pointer address register (SRIO_OM1DQEPAR)	32	R/W	0000_0000h	15.5.107/ 1171
D_312C	Outbound message n retry error threshold configuration register (SRIO_OM1RETCR)	32	R/W	0000_0000h	15.5.108/ 1173
D_3130	Outbound message n multicast group register (SRIO_OM1MGR)	32	R/W	0000_0000h	15.5.109/ 1174
D_3134	Outbound message n multicast list register (SRIO_OM1MLR)	32	R/W	0000_0000h	15.5.110/ 1175
D_3160	Inbound message n mode register (SRIO_IM1MR)	32	R/W	0000_0000h	15.5.111/ 1176
D_3164	Inbound message n status register (SRIO_IM1SR)	32	R/W	0000_0002h	15.5.112/ 1179
D_3168	Extended inbound message n frame queue dequeue pointer address register (SRIO_EIM1FQDPAR)	32	R/W	0000_0000h	15.5.113/ 1181
D_316C	Inbound message n frame queue dequeue pointer address register (SRIO_IM1FQDPAR)	32	R/W	0000_0000h	15.5.114/ 1181
D_3170	Extended inbound message n frame queue enqueue pointer address register (SRIO_EIM1FQEPAR)	32	R/W	0000_0000h	15.5.115/ 1182
D_3174	Inbound message n frame queue enqueue pointer address register (SRIO_IM1FQEPAR)	32	R/W	0000_0000h	15.5.116/ 1183
D_3178	Inbound message n maximum interrupt report interval register (SRIO_IM1MIRIR)	32	R/W	FFFF_FF00h	15.5.117/ 1184
D_3400	Outbound doorbell mode register (SRIO_ODMR)	32	R/W	0000_0000h	15.5.118/ 1184
D_3404	Outbound doorbell status register (SRIO_ODSR)	32	R/W	0000_0000h	15.5.119/ 1186
D_3418	Outbound doorbell n destination port register (SRIO_ODDPR)	32	R/W	0000_0000h	15.5.120/ 1188
D_341C	Outbound doorbell n destination attributes register (SRIO_ODDATR)	32	R/W	0000_0000h	15.5.121/ 1189
D_342C	Outbound doorbell n retry error threshold configuration register (SRIO_ODRETCR)	32	R/W	0000_0000h	15.5.122/ 1190
D_3460	Inbound doorbell n mode register (SRIO_IDMR)	32	R/W	0000_0000h	15.5.123/ 1191
D_3464	Inbound doorbell n status register (SRIO_IDSR)	32	R/W	0000_0002h	15.5.124/ 1194
D_3468	Extended inbound doorbell n queue dequeue pointer address register (SRIO_EIDQDPAR)	32	R/W	0000_0000h	15.5.125/ 1195

Table continues on the next page...

SRIO memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
D_346C	Inbound doorbell n queue dequeue Pointer address register (SRIO_IDQDPAR)	32	R/W	0000_0000h	15.5.126/ 1196
D_3470	Extended inbound doorbell n queue enqueue pointer address register (SRIO_EIDQEPAR)	32	R/W	0000_0000h	15.5.127/ 1197
D_3474	Inbound doorbell n Queue enqueue pointer address register (SRIO_IDQEPAR)	32	R/W	0000_0000h	15.5.128/ 1197
D_3478	Inbound doorbell n maximum interrupt report interval register (SRIO_IDMIRIR)	32	R/W	FFFF_FF00h	15.5.129/ 1198
D_34E0	Inbound port-write n mode register (SRIO_IPWMR)	32	R/W	0000_0000h	15.5.130/ 1199
D_34E4	Inbound port-write n status register (SRIO_IPWSR)	32	R/W	0000_0000h	15.5.131/ 1200
D_34E8	Extended inbound port-write n queue base address register (SRIO_EIPWQBAR)	32	R/W	0000_0000h	15.5.132/ 1201
D_34EC	Inbound port-write n queue base address register (SRIO_IPWQBAR)	32	R/W	0000_0000h	15.5.133/ 1201

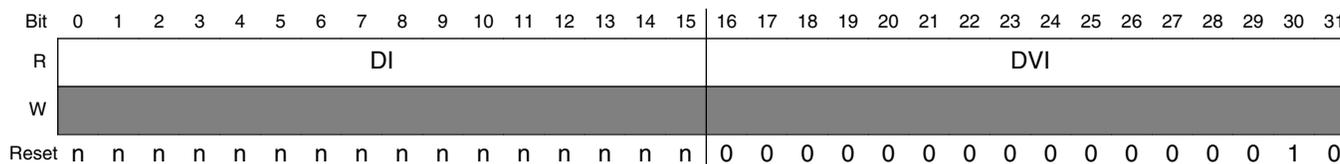
15.5.1 Device identity capability register (SRIO_DIDCAR)

The device vendor identity field (DVI) identifies the vendor that manufactured the device containing the processing element. A value for DVI is uniquely assigned to a device vendor by the registration authority of the RapidIO Trade Association.

The device identity field (DI) is intended to uniquely identify the type of device from the vendor specified by DVI. The values for DI are assigned and managed by the respective vendor.

DIDCAR is a read-only register.

Address: C_0000h base + 0h offset = C_0000h



SRIO_DIDCAR field descriptions

Field	Description
0–15 DI	Device identity 0x0070 = P2020 with security 0x0071 = P2020

Table continues on the next page...

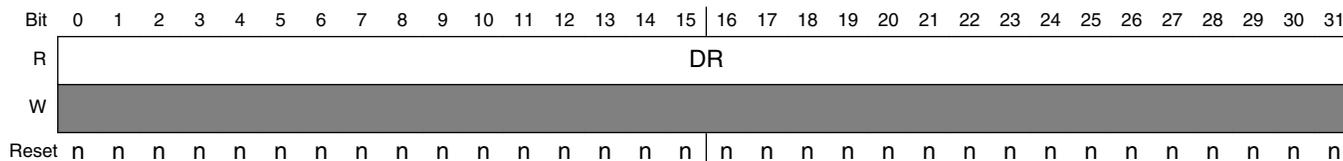
SRIO_DIDCAR field descriptions (continued)

Field	Description
	0x0078 = P2010 with security 0x0079 = P2010
16–31 DVI	Device vendor identity (Freescale = 0x0002)

15.5.2 Device information capability register (SRIO_DICAR)

The device revision field (DR) is intended to identify the revision level of the device. The value for DR is assigned and managed by the vendor specified by DIDCAR[DVI]. DICAR represents a copy of the device's system version register (SVR). See System Version Register (SVR) for more information.

Address: C_0000h base + 4h offset = C_0004h



SRIO_DICAR field descriptions

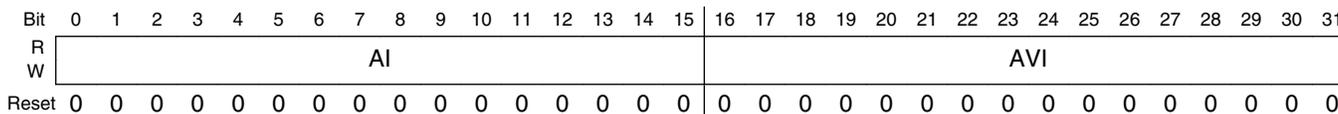
Field	Description
0–31 DR	Device revision. This is a copy of the device's system version register (SVR). See System version register (GUTS_SVR) for additional information. 0x80EA_0021 for P2020 with security 0x80E2_0021 for P2020 without security 0x80EB_0021 for P2010 with security 0x80E3_0021 for P2010 without security

15.5.3 Assembly identity capability register (SRIO_AIDCAR)

The assembly vendor identity field (AVI) identifies the vendor that manufactured the assembly or subsystem containing the device. A value for AVI is uniquely assigned to an assembly vendor by the registration authority of the RapidIO Trade Association.

The assembly identity field (AI) is intended to uniquely identify the type of assembly from the vendor specified by AVI. The values for AI are assigned and managed by the respective vendor.

Address: C_0000h base + 8h offset = C_0008h



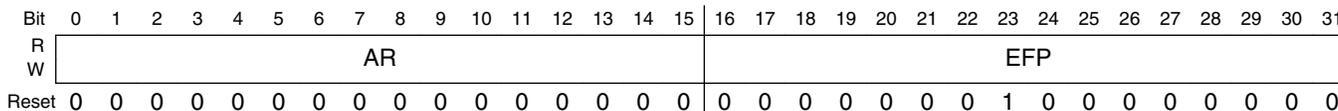
SRIO_AIDCAR field descriptions

Field	Description
0–15 AI	Assembly identity (all zeros)
16–31 AVI	Assembly vendor identity (all zeros)

15.5.4 Assembly information capability register (SRIO_AICAR)

AICAR contains additional information about the assembly and the pointer to the first entry in the extended features list.

Address: C_0000h base + Ch offset = C_000Ch



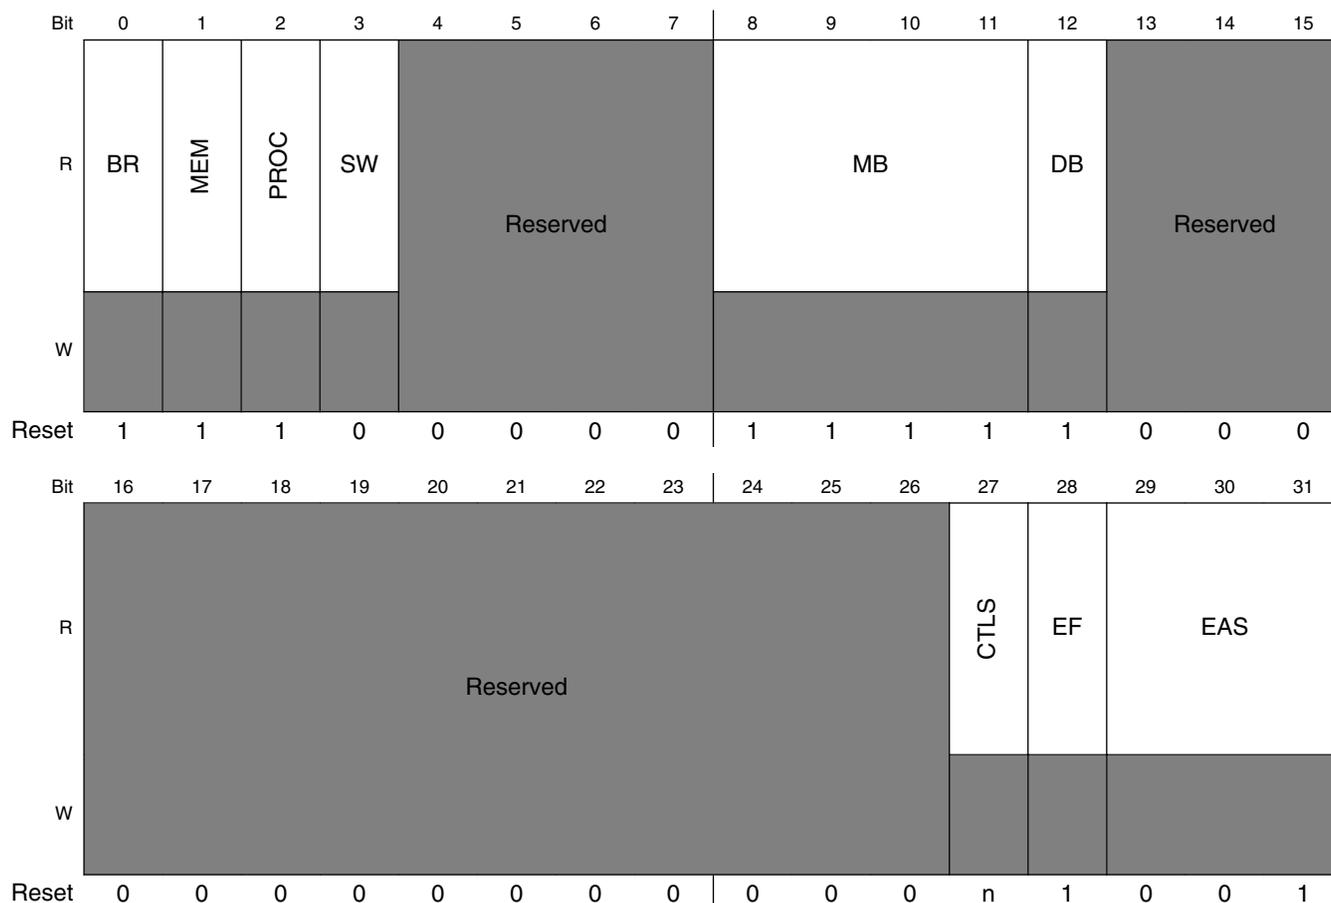
SRIO_AICAR field descriptions

Field	Description
0–15 AR	AssyRev field (all zeros)
16–31 EFP	ExtendedFeaturesPtr field (0x100)

15.5.5 Processing element features capability register (SRIO_PEF CAR)

The processing element features capability register (PEFCAR) identifies the major functionality provided by the processing element. PEFCAR is a read-only register.

Address: C_0000h base + 10h offset = C_0010h



SRIO_PEF CAR field descriptions

Field	Description
0 BR	Bridge. PE can bridge to another interface. (BR = 1)
1 MEM	Memory. PE has physically addressable local address space and can be accessed as an endpoint through non-maintenance operations. (MEM = 1)
2 PROC	Processor. PE physically contains a local processor that executes code. (PROC = 1)
3 SW	Switch. PE does not bridge to another external RapidIO interface. (SW = 0)

Table continues on the next page...

SRIO_PFCAR field descriptions (continued)

Field	Description
4–7 -	This field is reserved. Reserved
8–11 MB	Mailbox 0-3. (MB = 1111) 0 indicates PE supports inbound mailbox 0. 1 indicates PE supports inbound mailbox 1. 2 indicates PE supports inbound mailbox 2. 3 indicates PE supports inbound mailbox 3.
12 DB	Doorbell. The RapidIO controller supports inbound doorbells. (DB = 1)
13–26 -	This field is reserved. Reserved
27 CTLS	This bit reflects the selected RapidIO common transport system size. The RapidIO common transport system size is determined at power-on reset. 0 PE only supports common transport small systems (up to 256 devices). 1 PE supports common transport large systems (up to 65,536 devices).
28 EF	Extended features pointer is valid. (EF = 1)
29–31 EAS	Indicates the number of address bits supported by the PE both as a source and target of an operation. EAS = 001(34-bit addresses)

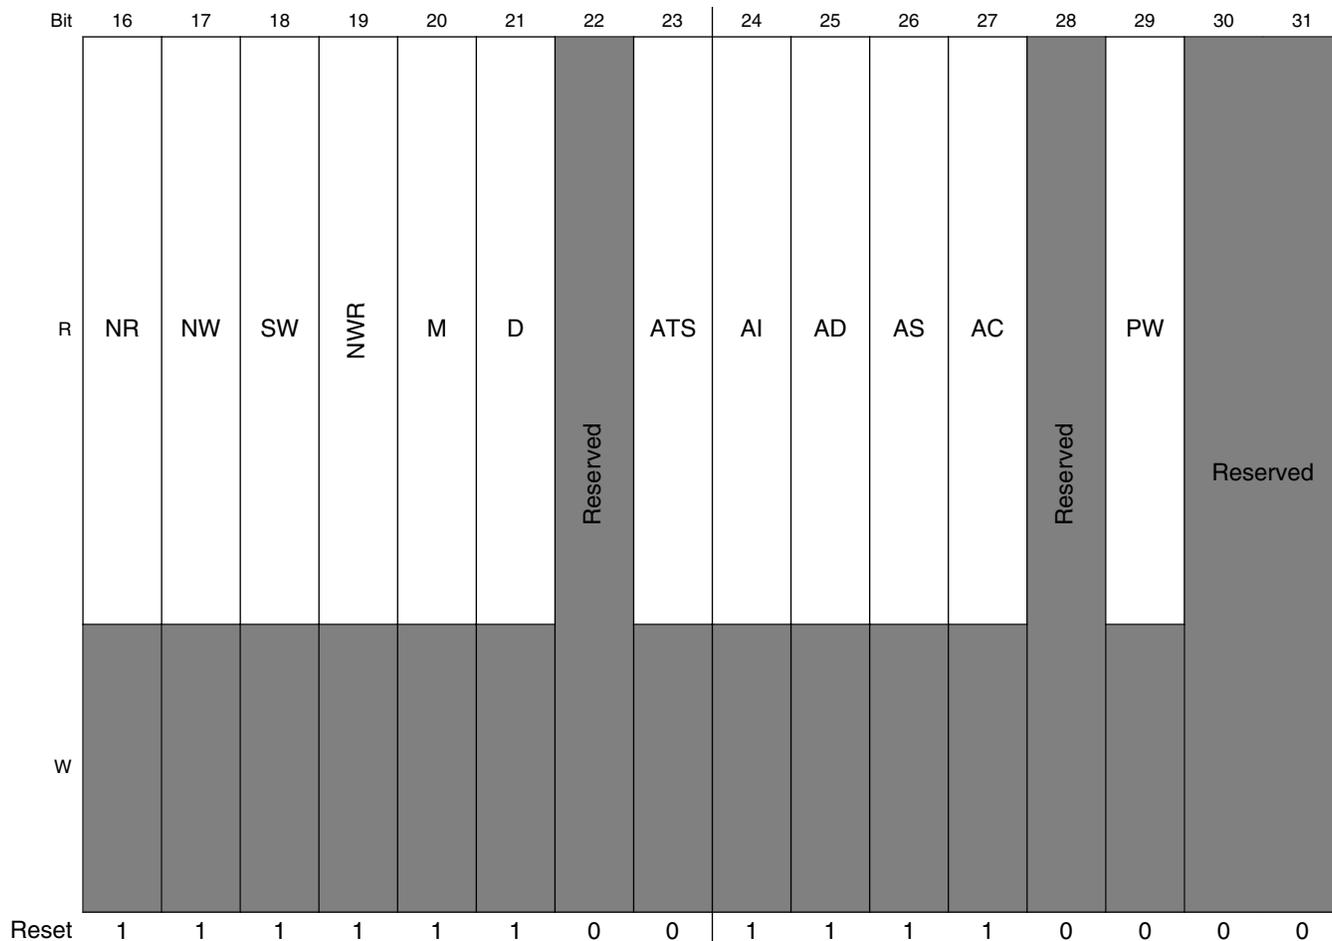
15.5.6 Source operations capability register (SRIO_SOCAR)

The source operations capability register (SOCAR) defines the set of RapidIO I/O logical operations that can be issued by this processing element. SOCAR is a read-only register.

Address: C_0000h base + 18h offset = C_0018h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	R	IR	RO	DCI	C	F	IOR	ICI	TIE	TIES	Reserved					
W	Reserved															
Reset	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0

Serial RapidIO Memory Map/Register Definition



SRIO_SOCAR field descriptions

Field	Description
0 R	PE does not support a Read operation. (R = 0)
1 IR	PE does not support an IRead operation. (IR = 0)
2 RO	PE does not support a Read_To_Own operation. (RO = 0)
3 DCI	PE does not support a Data Cache Invalidate operation. (DCI = 0)
4 C	PE does not support a Castout operation. (C = 0)
5 F	PE supports a Flush operation. (F = 1)
6 IOR	PE supports an I/O-Read operation. (IOR = 1)
7 ICI	PE does not support an Instruction Cache Invalidate operation. (ICI = 0)
8 TIE	PE does not support a TLBIE operation. (TIE = 0)

Table continues on the next page...

SRIO_SOCAR field descriptions (continued)

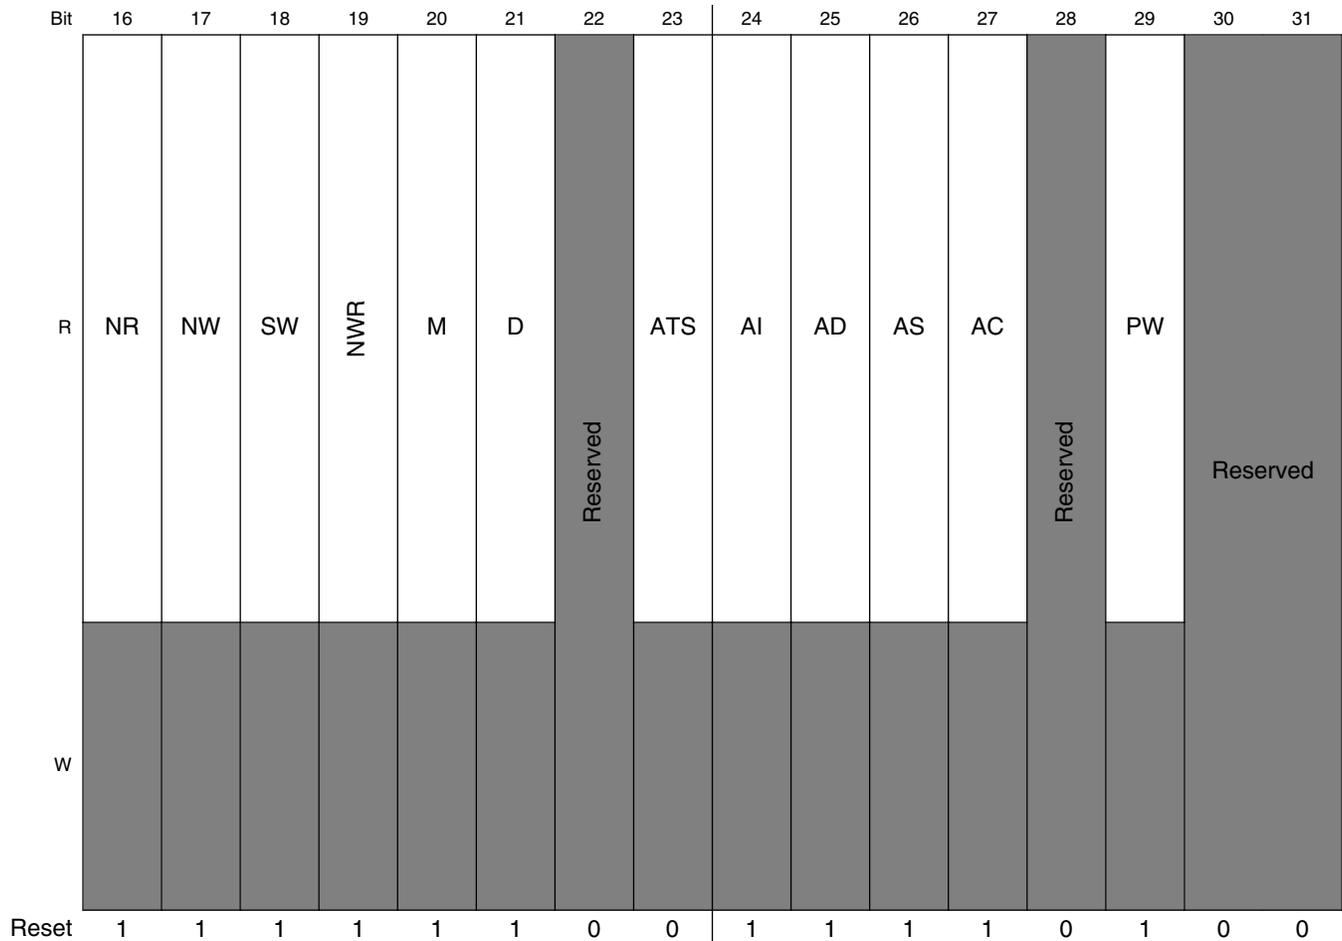
Field	Description
9 TIES	PE does not support a TLBSYNC operation. (TIES = 0)
10–15 -	This field is reserved. Reserved
16 NR	PE supports a Nread operation. (NR = 1)
17 NW	PE supports a Nwrite operation. (NW = 1)
18 SW	PE supports an Swrite operation. (SW = 1)
19 NWR	PE supports a Nwrite_R operation. (NWR = 1)
20 M	PE supports a Message operation. (M = 1)
21 D	PE supports a Doorbell operation. (D = 1)
22 -	This field is reserved. Reserved
23 ATS	PE does not support an Atomic-Test-and-Swap operation. (ATS = 0)
24 AI	PE supports an Atomic-Inc operation. (AI = 1)
25 AD	PE supports an Atomic-Dec operation. (AD = 1)
26 AS	PE supports an Atomic-Set operation. (AS = 1)
27 AC	PE supports an Atomic-Clr operation. (AC = 1)
28 -	This field is reserved. Reserved
29 PW	PE does not support a Port-Write operation. (PW = 0)
30–31 -	This field is reserved. Reserved

15.5.7 Destination operations capability register (SRIO_DOCAR)

The destination operations capability register (DOCAR) defines the set of RapidIO I/O operations that can be supported by this processing element. DOCAR is a read-only register.

Address: C_0000h base + 1Ch offset = C_001Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	R	IR	RO	DCI	C	F	IOR	ICI	TIE	TIES	Reserved					
W	Reserved										Reserved					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0



SRIO_DOCAR field descriptions

Field	Description
0 R	PE does not support a Read operation. (R = 0)
1 IR	PE does not support an IRead operation. (IR = 0)
2 RO	PE does not support a Read_To_Own operation. (RO = 0)
3 DCI	PE does not support a Data Cache Invalidate operation. (DCI = 0)
4 C	PE does not support a Castout operation. (C = 0)
5 F	PE does not support a Flush operation. (F = 0)
6 IOR	PE does not support an I/O-Read operation. (IOR = 0)
7 ICI	PE does not support an Instruction Cache Invalidate operation. (ICI = 0)
8 TIE	PE does not support a TLBIE operation. (TIE = 0)

Table continues on the next page...

SRIO_DOCAR field descriptions (continued)

Field	Description
9 TIES	PE does not support a TLBSYNC operation. (TIES = 0)
10–15 -	This field is reserved. Reserved
16 NR	PE supports a Nread operation. (NR = 1)
17 NW	PE supports a Nwrite operation. (NW = 1)
18 SW	PE supports an Swrite operation. (SW = 1)
19 NWR	PE supports a Nwrite_R operation. (NWR = 1)
20 M	PE supports a Message operation. (M = 1)
21 D	PE supports a Doorbell operation. (D = 1)
22 -	This field is reserved. Reserved
23 ATS	PE does not support an Atomic-Test-and-Swap operation. (ATS = 0)
24 AI	PE supports an Atomic-Inc operation. (AI = 1)
25 AD	PE supports an Atomic-Dec operation. (AD = 1)
26 AS	PE supports an Atomic-Set operation. (AS = 1)
27 AC	PE supports an Atomic-Clr operation. (AC = 1)
28 -	This field is reserved. Reserved
29 PW	PE supports a Port-Write operation. (PW = 1)
30–31 -	This field is reserved. Reserved

15.5.8 Mailbox command and status register (SRIO_MCSR)

The MCSR reflects the status of the RapidIO message controllers on this device.

Address: C_0000h base + 40h offset = C_0040h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	A0	FU0	EM0	B0	FA0	ERR0	Reserved		A1	FU1	EM1	B1	FA1	ERR1	Reserved	
W																
Reset	0	0	1	0	0	0	0	0	0	0	1	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_MCSR field descriptions

Field	Description
0 A0	Available 0 Message Controller 0 is not ready to accept messages. All incoming message transactions return error responses. 1 Message Controller 0 is initialized and ready to accept messages.
1 FU0	Full 0 Message Controller 0 is not full. 1 Message Controller 0 is full. New messages are retried.
2 EM0	Empty 0 Message Controller 0 contains outstanding messages. 1 Message Controller 0 contains no outstanding messages.
3 B0	Busy 0 Message Controller 0 is not busy processing a message. 1 Message Controller 0 is busy processing a message. New message operations return retry responses.
4 FA0	Failure 0 Message Controller 0 has not encountered an internal error. 1 Message Controller 0 had an internal fault or error condition and is waiting for assistance. All incoming message transactions return error responses.
5 ERR0	Error This field always returns a 0.
6–7 -	This field is reserved. Reserved

Table continues on the next page...

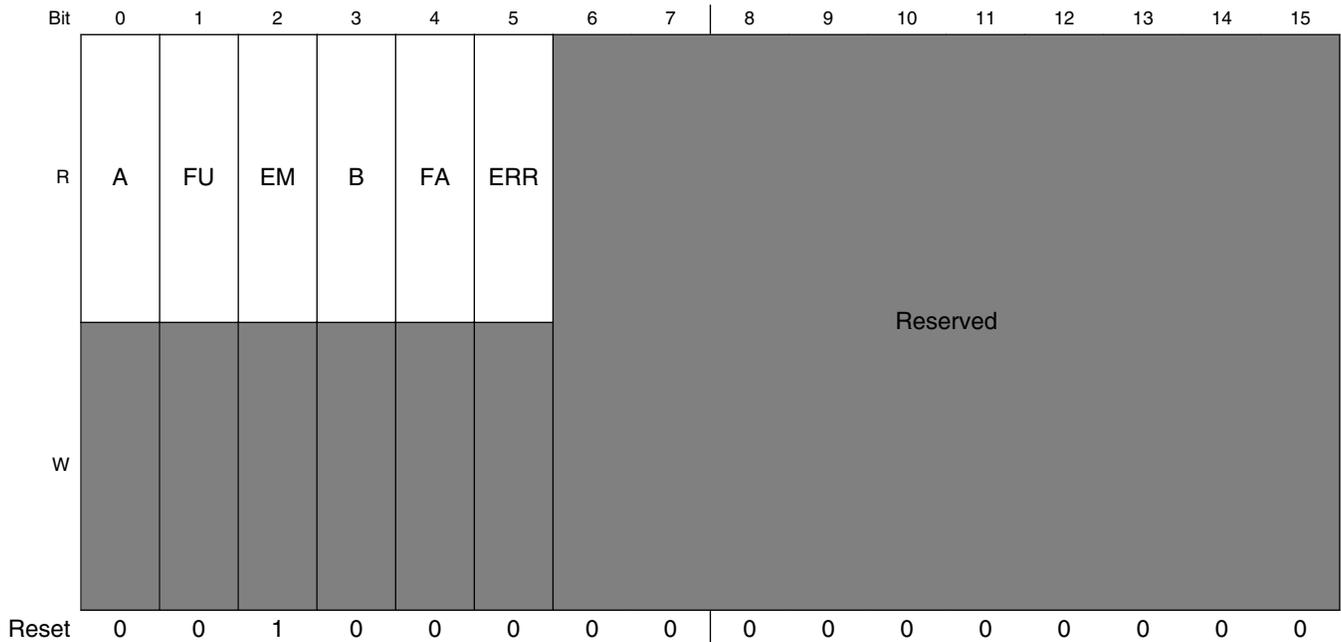
SRIO_MCSR field descriptions (continued)

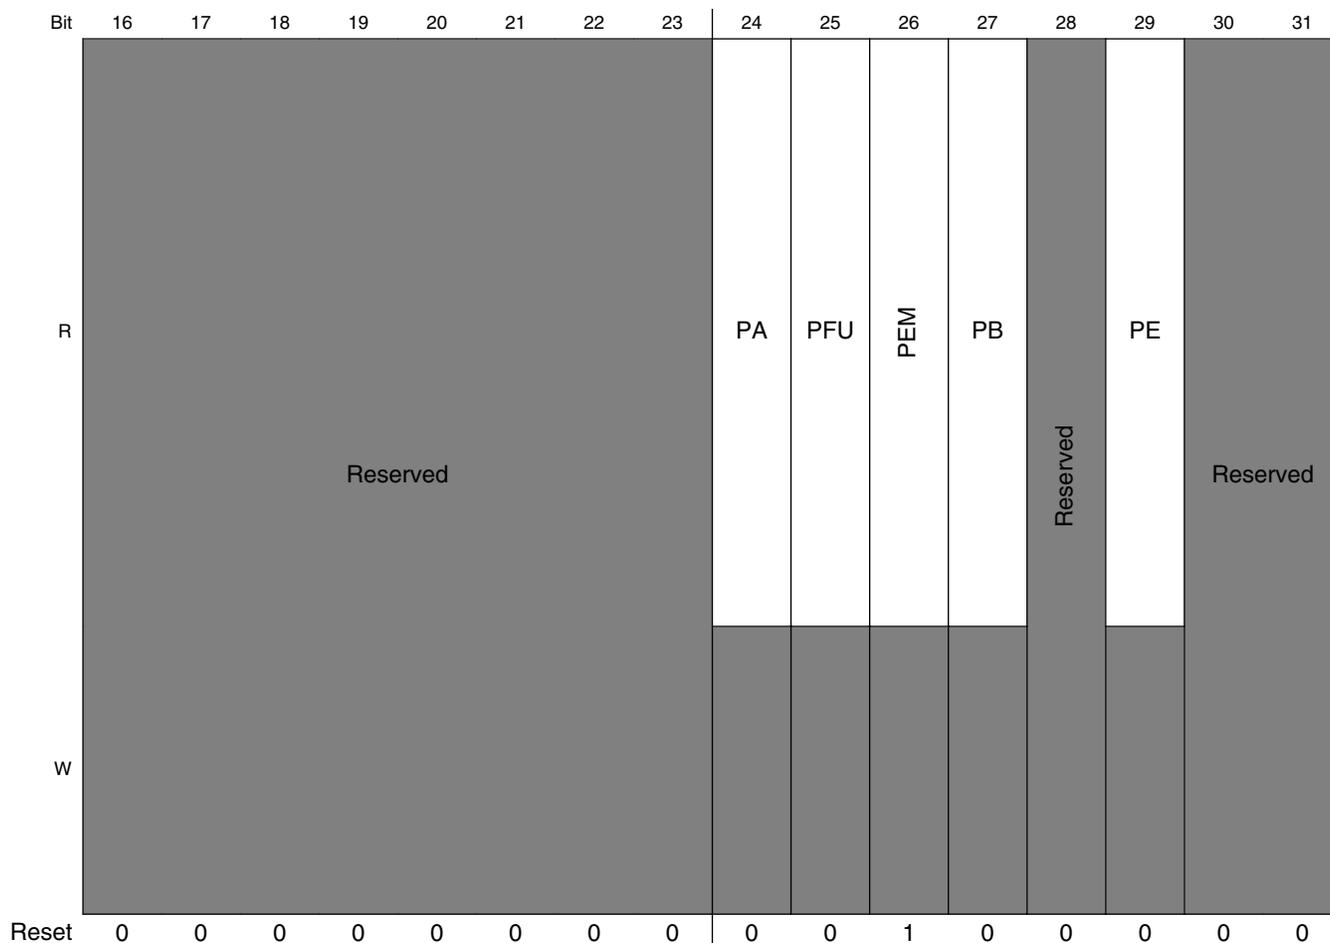
Field	Description
8 A1	Available 0 Message Controller 1 is not ready to accept messages. All incoming message transactions return error responses. 1 Message Controller 1 is initialized and ready to accept messages.
9 FU1	Full 0 Message Controller 1 is not full. 1 Message Controller 1 is full. New messages are retried.
10 EM1	Empty 0 Message Controller 1 contains outstanding messages. 1 Message Controller 1 contains no outstanding messages.
11 B1	Busy 0 Message Controller 1 is not busy processing a message. 1 Message Controller 1 is busy processing a message. New message operations return retry responses.
12 FA1	Failure 0 Message Controller 1 has not encountered an internal error. 1 Message Controller 1 had an internal fault or error condition and is waiting for assistance. All incoming message transactions return error responses.
13 ERR1	Error This field always returns a 0.
14–31 -	This field is reserved. Reserved

15.5.9 Port -Write and doorbell command and status register (SRIO_PWDCSR)

The PWDCSR reflects the status of the RapidIO doorbell and port-write hardware on this device. Additional details can be found in the *RapidIO Interconnect Specification, Revision 1.2*, in sections "Doorbell CSR" and "Write Port CSR."

Address: C_0000h base + 44h offset = C_0044h





SRIO_PWDCSR field descriptions

Field	Description
0 A	Available 0 Doorbell unit is not ready to accept doorbell messages. All incoming doorbell transactions return error responses. 1 Doorbell unit is initialized and ready to accept doorbell messages.
1 FU	Full 0 Doorbell unit is not full. 1 Doorbell unit is full. New doorbell messages are retried.
2 EM	Empty 0 Doorbell unit has outstanding doorbell messages. 1 Doorbell unit has no outstanding doorbell messages.
3 B	Busy 0 Doorbell unit is not busy processing a doorbell message. 1 Doorbell unit is busy processing a doorbell message. Incoming transactions are not retried.
4 FA	Failure

Table continues on the next page...

SRIO_PWDCSR field descriptions (continued)

Field	Description
	0 Doorbell unit has not encountered an internal error. 1 Doorbell unit had an internal fault or error condition and is waiting for assistance. All incoming doorbell transactions return error responses.
5 ERR	Error This field always returns a 0.
6–23 -	This field is reserved. Reserved
24 PA	Port-write unit available. 0 Port-write unit is not available. All incoming port-write packets are discarded. 1 Port-write unit is initialized and ready to accept a port-write packet.
25 PFU	Port-write unit full 0 Port-write unit is not full. 1 Port-write unit is full. All incoming port-write packets are discarded.
26 PEM	Port-write unit empty This field always returns a 1.
27 PB	Port-write unit busy 0 Port-write unit is not busy. 1 Port-write unit is busy processing a port-write packet. Incoming port-write packets are discarded.
28 -	This field is reserved. Reserved
29 PE	Error This field always returns a 0.
30–31 -	This field is reserved. Reserved

15.5.10 Processing element logical layer control command and status register (SRIO_PELLCCSR)

The processing element logical layer control command and status register (PELLCCSR) controls the extended addressing abilities. The RapidIO endpoint only supports 34-bit addressing. PELLCCSR is a read-only register.

Address: C_0000h base + 4Ch offset = C_004Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															Reserved											EAC					
W	Reserved															Reserved																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

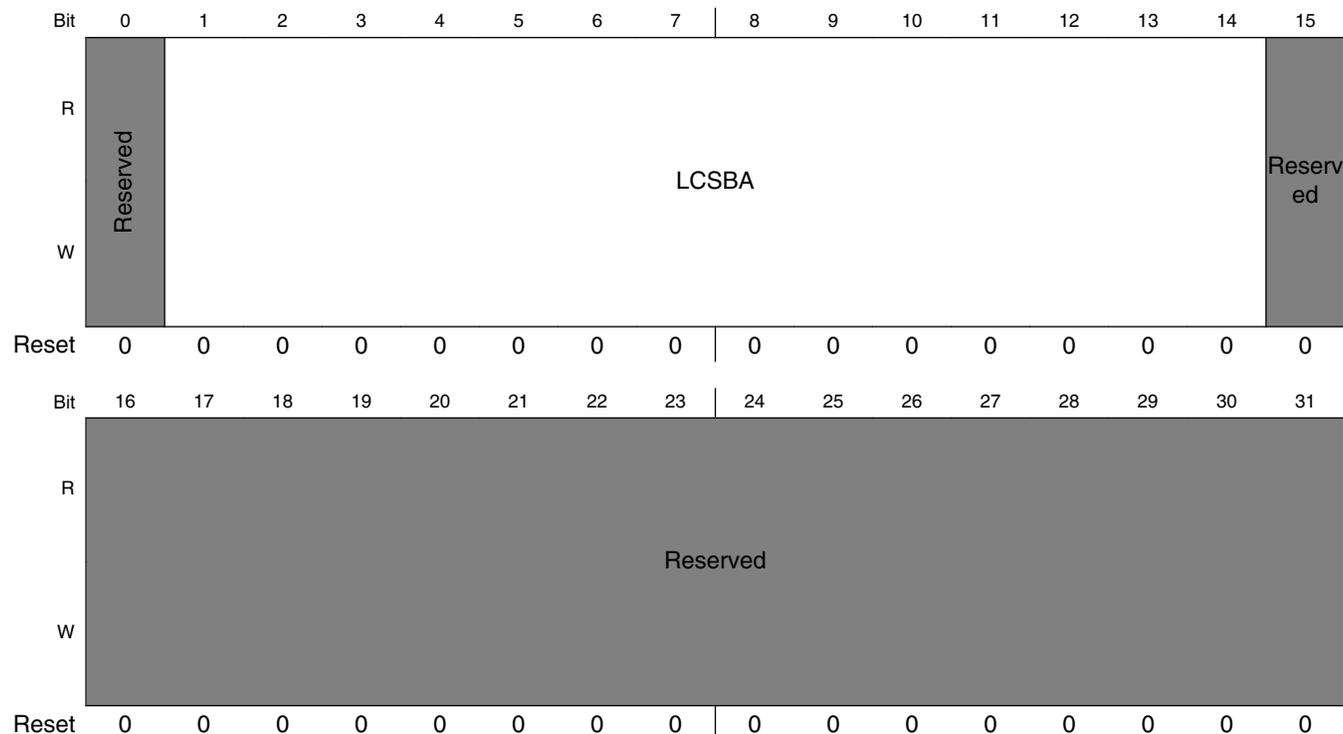
SRIO_PELLCSR field descriptions

Field	Description
0–28 -	This field is reserved. Reserved
29–31 EAC	Extended addressing control. Read-only value is 0b001.

15.5.11 Local configuration space base address 1 command and status register (SRIO_LCSBA1CSR)

The local configuration space base address 1 command and status register (LCSBA1CSR) specifies the least-significant bits of the local physical address double-word offset for the processing element's configuration register space, allowing the configuration register space to be physically mapped in the processing element. This register allows configuration and maintenance of a processing element through regular read and write operations rather than maintenance operations. The double-word offset is right-justified in the register. This window has priority over all ATMU windows. As is the case with all registers, an external processor writing to LCSBA1CSR should not assume it has been written until a response has been received.

Address: C_0000h base + 5Ch offset = C_005Ch



SRIO_LCSBA1CSR field descriptions

Field	Description
0 -	This field is reserved. Reserved
1–14 LCSBA	Local configuration space base address. These bits correspond to the highest 14 bits of the 34-bit RapidIO address space.
15–31 -	This field is reserved. Reserved

15.5.12 Base device ID command and status register (SRIO_BDIDCSR)

BDIDCSR contains the base device ID values for the processing element.

Address: C_0000h base + 60h offset = C_0060h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	

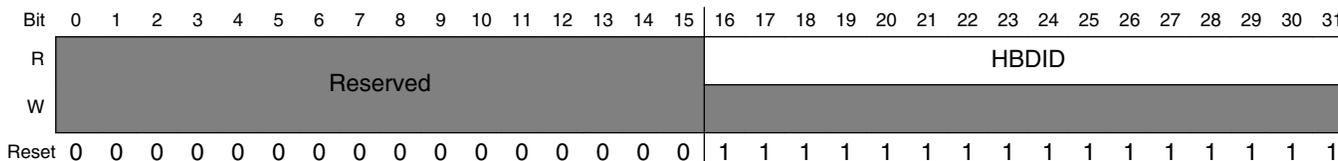
SRIO_BDIDCSR field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–15 BDID	Base device ID of the device in a small common transport system (RapidIO device ID). If RapidIO is configured as a host, BDID = {0b0000_0, cfg_device_ID[5:7]}. If RapidIO is configured as an agent, BDID = {0b1111_111, cfg_device_ID[7]}. (See Host/agent configuration and RapidIO device ID for details.)
16–31 LBDID	Large base device ID of the device in a large common transport system. This field is valid only if PEFCAR[CTLS] is set. If RapidIO is configured as a host, LBDID = 0b0000_0000_0000_0, cfg_device_ID[5:7]. If RapidIO is configured as an agent, LBDID = 0b1111_1111_1111_111, cfg_device_ID[7]. If RapidIO is configured as a host, LBDID = 0b0000_0000. If RapidIO is configured as an agent, LBDID = 0b1111_1111. (See Host/agent configuration for details.)

15.5.13 Host base device ID lock command and status register (SRIO_HBDIDLCSR)

The host base device ID lock command and status register (HBDIDLCSR) contains the base device ID value for the processing element in the system that is responsible for initializing this processing element. The HBDID field is a write-once/resettable field which provides a lock function. Once HBDID is written, all subsequent writes to the field are ignored, except in the case that the value written matches the value contained in the field. In this case, the register is re-initialized to 0xFFFF. After writing HBDID, a processing element must then read the host base device ID lock CSR to verify that it owns the lock before attempting to initialize this processing element.

Address: C_0000h base + 68h offset = C_0068h



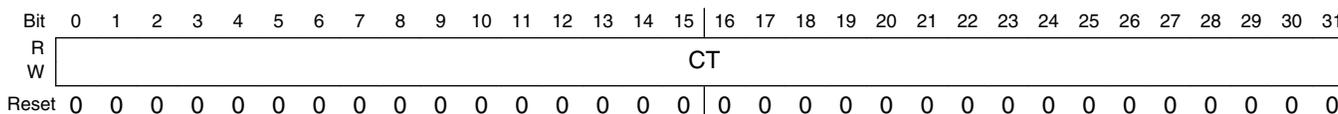
SRIO_HBDIDLCSR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16–31 HBDID	This is the host base device ID for the processing element that is responsible for initializing this device. Only the first write to this field is accepted; all other writes are ignored, except in the case that the value written matches the value contained in the field. In this case, the register is re-initialized to 0xFFFF.

15.5.14 Component tag command and status register (SRIO_CTCSR)

The component tag command and status register (CTCSR) contains a component tag value for the processing element and can be assigned by software when the device is initialized. It is unused internally in the RapidIO endpoint.

Address: C_0000h base + 6Ch offset = C_006Ch



SRIO_CTCSR field descriptions

Field	Description
0–31 CT	Component tag

15.5.15 Port maintenance block header 0 (SRIO_PMBH0)

The port maintenance block header 0 register (PMBH0) contains a pointer (EF_PTR) to the next EF_BLK (Extended Features Space, Error Management) and the EF_ID that identifies this as the generic end point port maintenance block header. Note that while registers defined by software-assisted error recovery are supported, software-assisted error recovery is not (these registers are included for hot insertion only); therefore, the RapidIO endpoint is defined here as not supporting software-assisted error recovery. PMBH0 is a read-only register.

Address: C_0000h base + 100h offset = C_0100h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	EF_PTR															EF_ID																
W																																
Reset	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

SRIO_PMBH0 field descriptions

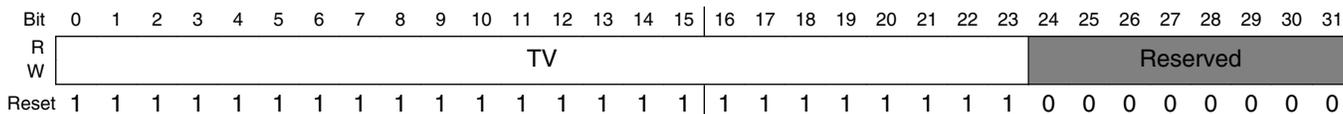
Field	Description
0–15 EF_PTR	Extended features pointer
16–31 EF_ID	Extended features ID

15.5.16 Port link time-out control command and status register (SRIO_PLTOCCSR)

The port link time-out control command and status register (PLTOCCSR) contains the link time-out value for all ports on a device. This time-out is for link events such as sending a packet to receiving the corresponding acknowledge and sending a link-request to receiving the corresponding link-response. The reset, or default, value of this counter is the maximum time-out interval . Minimal programmable TV value supported is 0x200.

The resolution of this timer is $168/(\text{platform frequency})$. For example, at a platform frequency of 400 MHz, the maximum timeout value is $0xFF_FFFF * 168/400\text{MHz} = 7.05$ seconds.

Address: C_0000h base + 120h offset = C_0120h



SRIO_PLTOCCSR field descriptions

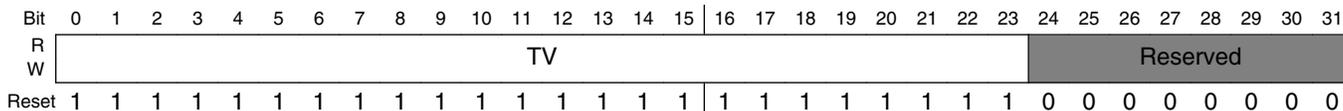
Field	Description
0–23 TV	Time-out value. Setting to all zeros disables the link time-out timer. This value is loaded each time the link time-out timer starts.
24–31 -	This field is reserved. Reserved

15.5.17 Port response time-out control command and status register (SRIO_PRTOCCSR)

The port response time-out control command and status register (PRTOCCSR) contains the time-out timer count for all ports on a device. This time-out is for sending a request packet to receiving the corresponding response packet. Note that this applies to the RapidIO endpoint and the messaging unit. The reset, or default, value of this counter is the maximum time-out interval .

The resolution of this timer is $168/(\text{platform frequency})$. For example, at a platform frequency of 400 MHz, the maximum timeout value is $0xFF_FFFF * 168/400\text{MHz} = 7.05$ seconds.

Address: C_0000h base + 124h offset = C_0124h



SRIO_PRTOCCSR field descriptions

Field	Description
0–23 TV	Time-out value. Setting to all zeros disables the response time-out timer. This value is loaded each time the response time-out timer starts.
24–31 -	This field is reserved. Reserved

15.5.18 Port General control command and status register (SRIO_PGCCSR)

The port general control command and status register (PGCCSR) contains control register bits applicable to the RapidIO interface .

Address: C_0000h base + 13Ch offset = C_013Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	H	M	D	Reserved												
W																
Reset	n	n	n	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W																
Reset	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*	0*

* Notes:

- Values indicated with n are read from configuration signals at reset.

SRIO_PGCCSR field descriptions

Field	Description
0 H	<p>Host. On reset the value of this bit is assigned by power-on reset configuration signals, $\overline{LWEn} / \overline{LBSn}$, which determine host/agent configuration for the device. See Host/agent configuration for more information.</p> <p>0 Agent device 1 Host device</p>
1 M	<p>Master. The value of this bit is identical to PGCCSR[H] which is assigned by power-on reset configuration signals, $\overline{LWEn} / \overline{LBSn}$, which determine host/agent configuration for the device. See Host/agent configuration for more information. For agent devices, this bit may be modified manually by software after device enumeration is complete.</p> <p>This bit controls whether or not a device is allowed to issue requests into the system. If M = 0, the device may only respond to requests.</p> <p>Software must ensure no requests are sent to the device by disabling any local access windows and disabling inbound ATMU windows from other devices targeting SRIO while M=0.</p> <p>0 Device is not enabled to issue requests into the system. 1 Device is enabled to issue requests into the system.</p>
2 D	<p>Discovered. The value of this bit is identical to PGCCSR[H] which is assigned by power-on reset configuration signals, $\overline{LWEn} / \overline{LBSn}$, which determine host/agent configuration for the device. See Host/agent configuration for more information. For agent devices, this bit should be modified by the host device to indicate this device has been discovered and enumerated.</p> <p>0 Device has not been discovered by system host. 1 Device has been discovered by system host.</p>
3–31 -	<p>This field is reserved. Reserved</p>

15.5.19 Port 1 Link maintenance request command and status register (SRIO_P1LMREQCSR)

The port 1 link maintenance request command and status register (P1LMREQCSR) is accessible both by the local processor and an external device. A write to this register generates a link-request control symbol on the corresponding RapidIO endpoint port interface.

Address: C_0000h base + 140h offset = C_0140h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															C																
W	Reserved															C																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1LMREQCSR field descriptions

Field	Description
0–28 -	This field is reserved. Reserved
29–31 C	LINK_REQUEST command to send. If read, this field returns the last written value. If written with a value other than 0b011 (reset-device) or 0b100 (input-status), the resulting operation is undefined, as all other values are reserved in the RapidIO specification.

15.5.20 Port 1 Link maintenance response command and status register (SRIO_P1LMRESPCSR)

The port 1 link maintenance response command and status register (P1LMRESPCSR) is accessible both by the local processor and an external device. A read to this register returns the status received in a link-response control symbol. This register is read only.

Address: C_0000h base + 144h offset = C_0144h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	RV	Reserved														
W	RV	Reserved														
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved							AS					LS			
W	Reserved							AS					LS			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

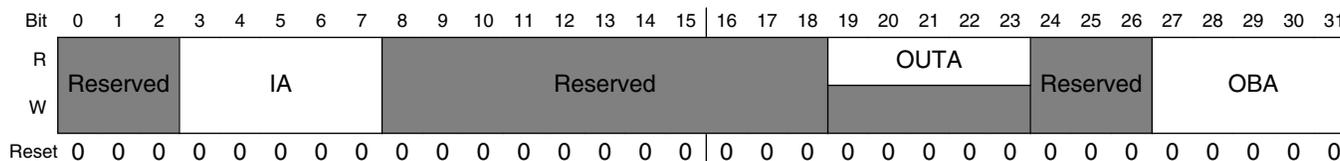
SRIO_P1LMRESPCSR field descriptions

Field	Description
0 RV	Response valid. If the link-request causes a link-response, this bit indicates that the link-response has been received and the status fields are valid. If the link-request does not cause a link-response, this bit indicates that the link-request has been transmitted. This bit clears on read, so be careful when polling its value.
1–21 -	This field is reserved. Reserved
22–26 AS	AckID_status field from LINK_RESPONSE
27–31 LS	Link_status field from LINK_RESPONSE

15.5.21 Port 1 Local ackID status command and status register (SRIO_P1LASCSSR)

The port 1 local ackID status command and status register (P1LASCSSR) is accessible both by the local processor and an external device. A read to this register returns the local ackID status for both the output and input ports of the device.

Address: C_0000h base + 148h offset = C_0148h



SRIO_P1LASCSSR field descriptions

Field	Description
0–2 -	This field is reserved. Reserved
3–7 IA	Input port next expected ackID value.
8–18 -	This field is reserved. Reserved
19–23 OUTA	Outstanding port unacknowledged ackID status. Next expected acknowledge control symbol ackID field that indicates the ackID value expected in the next received acknowledge control symbol. Note that this value is read only even though RapidIO specification allows for it to be writable.
24–26 -	This field is reserved. Reserved

Table continues on the next page...

SRIO_P1LASCSR field descriptions (continued)

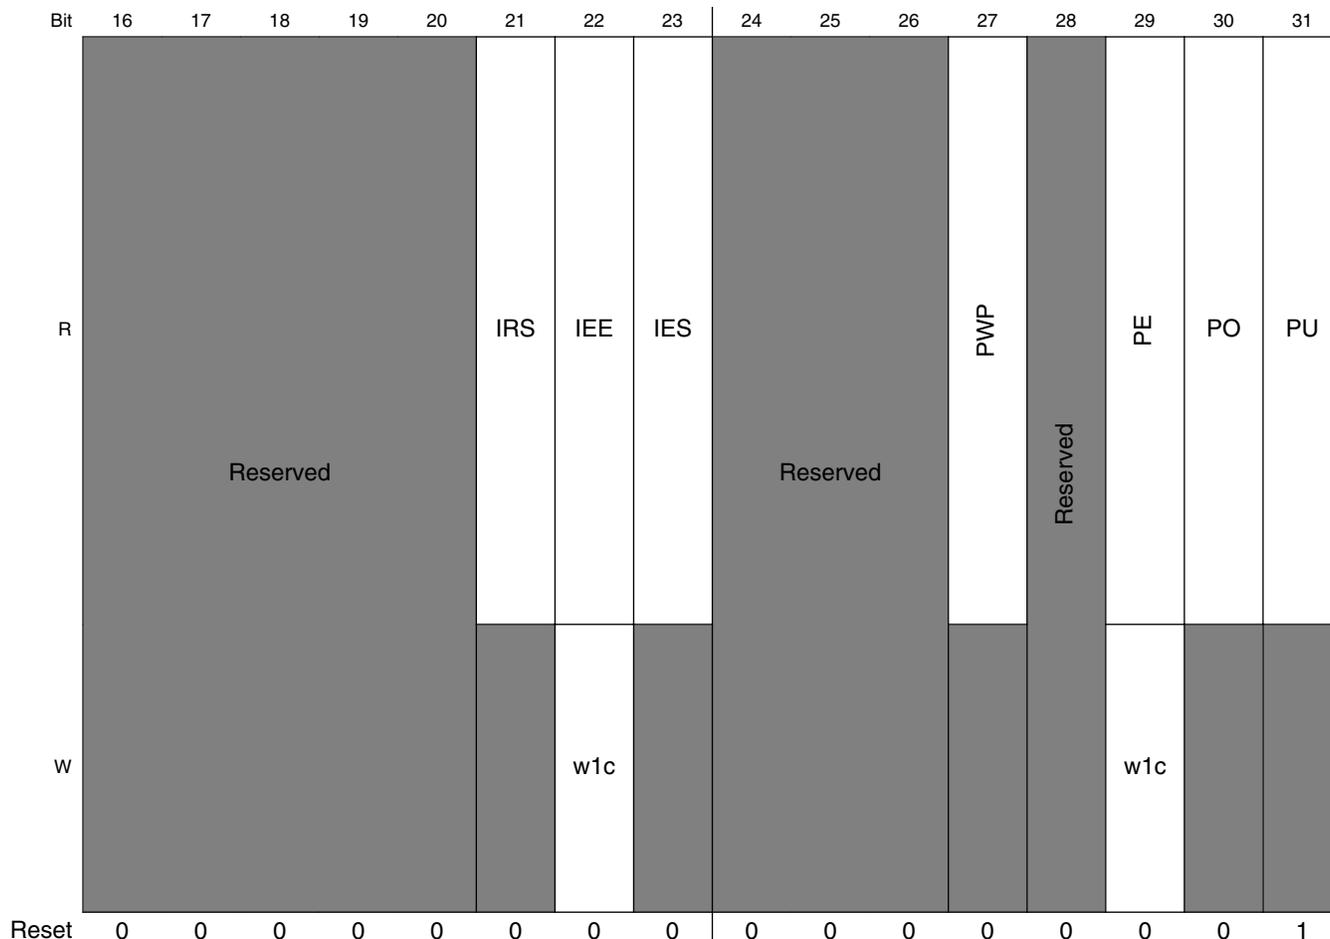
Field	Description
27–31 OBA	Outbound ackID output port next transmitted ackID value. This can be written by software but only if there are no outstanding unacknowledged packets. If there are, the newly-written value is ignored.

15.5.22 Port 1 Error and status command and status register (SRIO_P1ESCSR)

The port 1 error and status command and status register (P1ESCSR) is accessed when the local processor or an external device wishes to examine the port error and status information.

Address: C_0000h base + 158h offset = C_0158h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved					OPD	OFE	ODE	Reserved			ORE	OR	ORS	OEE	OES
W						w1c	w1c	w1c				w1c			w1c	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0



SRIO_P1ESCSR field descriptions

Field	Description
0-4 -	This field is reserved. Reserved
5 OPD	Output Packet-dropped. Output port has discarded a packet. A packet is discarded if: <ul style="list-style-type: none"> It is received while OFE is set and PmCCSR[DPE] (drop packet enable) is set and PmCCSR[SPF] (stop on port failed) is set. It is received while PmPCR[OBDEN] (output buffer drain enable) is set. It is not-accepted by the link-partner while PmERCSR[ERFTT] (error rate failed threshold trigger) is met or exceeded and PmCCSR[DPE] is set and PmCCSR[SPF] is not set (and link-response returns expected ackID). Once OPD is set, it remains set until written with a logic 1 to clear.
6 OFE	Output Failed-encountered. Output port has encountered a failed condition, meaning that the Error Rate Counter has met or exceeded the port's failed error threshold (ERFTT) Once set, remains set until written with a logic 1 to clear. Once cleared, does not assert again unless the Error Rate Counter dips below the port's failed error threshold and then meets or exceeds it again.
7 ODE	Output port has encountered a degraded condition, meaning that the Error Rate Counter has met or exceeded the port's degraded error threshold. Once set, remains set until written with a logic 1 to clear.

Table continues on the next page...

SRIO_P1ESCSR field descriptions (continued)

Field	Description
	Once cleared, does not assert again unless the Error Rate Counter dips below the port's degraded error threshold and then meets or exceeds it again.
8–10 -	This field is reserved. Reserved
11 ORE	Output port has encountered a retry condition. This bit is set when bit 13 is set. Once set, remains set until written with a logic 1 to clear.
12 OR	Output port has received a packet retry control symbol and cannot make forward progress. This bit is set when bit 13 is set and cleared when a packet-accepted or packet-not-accepted control symbol is received. (read only)
13 ORS	Output port is stopped due to a retry (read only)
14 OEE	Output port has encountered (and possibly recovered from) a transmission error. This bit is set when bit 15 is set. Once set, remains set until written with a logic 1 to clear.
15 OES	Output port is stopped due to a transmission error (read only)
16–20 -	This field is reserved. Reserved
21 IRS	Input port is stopped due to a retry (read only)
22 IEE	Input port has encountered (and possibly recovered from) a transmission error. This bit is set when bit 23 is set. Once set, remains set until written with a logic 1 to clear.
23 IES	Input port is stopped due to a transmission error (read-only)
24–26 -	This field is reserved. Reserved
27 PWP	Port has encountered a condition which required it to initiate a maintenance port-write operation. This bit is only valid if the device is capable of issuing a maintenance port-write transaction. The RapidIO endpoint is not capable of issuing port-writes. This bit is hardwired to 0.
28 -	This field is reserved. Reserved
29 PE	Input or output port has encountered an error from which hardware was unable to recover. Once set, remains set until written with a logic 1 to clear. This bit indicates that OFE is set while PmCCSR[SPF] is set; in other words, the failed threshold has been reached which has caused the output port to stop transmitting packets.
30 PO	The input and output ports are initialized and the port is exchanging error-free control symbols with the attached device. (read-only).
31 PU	Input and output ports are not initialized . This bit and bit 30 are mutually exclusive (read-only).

15.5.23 Port 1 Control command and status register (SRIO_P1CCSR)

The port 1 control command and status register (P1CCSR) contains control register bits for the RapidIO port .

Address: C_0000h base + 15Ch offset = C_015Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	PW		IPW			PWO			PD	OPE	IPE	ECD	MEP	Reserved		
W																
Reset	0	1	0	1	0	0	0	0	0	1	1	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved											SPF	DPE	PL	PT	
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

SRIO_P1CCSR field descriptions

Field	Description
0-1 PW	Hardware width of the port (read-only). 00 Single-lane port 01 Four-lane port 10-11 Reserved
2-4 IPW	Width of the ports after initialized (read-only). 000 Single-lane port, lane 0 001 Single-lane port, lane 2 010 Four-lane port 011-111 Reserved
5-7 PWO	Soft port configuration to override the hardware size. This field should be changed only when the port is uninitialized. To achieve this, first disable the RapidIO port. Then change PWO to any valid value. Finally, re-enable the RapidIO port. To achieve this, first set PmCCSR[PD] (port disabled). Then change PWO to any legal value. Finally, clear PmCCSR[PD] (enabled). 000 No override 001 Reserved 010 Force single lane, lane 0 011 Force single lane, lane 2 100-111 Reserved, causes undefined operation
8 PD	Port disable.

Table continues on the next page...

SRIO_P1CCSR field descriptions (continued)

Field	Description
	0 Input error state machine operates normally 1 Input error state machine is forced to normal state
9 OPE	Output port transmit enable. OPE is ignored by RapidIO endpoints. It is expected that if OPE = 0, software will not send packets out of outbound. If packets are sent by OCN to outbound, they are sent out of RapidIO endpoints, regardless of the value of OPE. Initial value read from configuration pins. 0 Port is stopped and not enabled to issue any packets except to route or respond to I/O logical MAINTENANCE packets. Control symbols are not affected and are sent normally. 1 Port is enabled to issue any packets.
10 IPE	Input port receive enable. This bit value must equal the value of the output port enable (OPE) bit in order for the RapidIO controller to function properly. Initial value read from configuration pins. 0 Port is stopped and only enabled to route or respond to I/O logical MAINTENANCE packets. Other packets generate packet-not-accepted control symbols to force an error condition to be signaled by the sending device. Control symbols are not affected and are received and handled normally. 1 Port is enabled to respond to any packet.
11 ECD	Error checking disable. This bit is hardwired to 0. This bit disables all RapidIO transmission error checking 0 Error checking and recovery is enabled. 1 Error checking and recovery is disabled.
12 MEP	Multicast-event participant. This bit is hard-wired to 0.
13–27 -	This field is reserved. Reserved
28 SPF	Stop on port failed-encountered enable. This bit is used with the drop packet enable bit to force certain behavior when the error rate failed threshold has been met or exceeded.
29 DPE	Drop packet enable. This bit is used with the stop on port failed-encountered enable bit to force certain behavior when the error rate failed threshold has been met or exceeded.
30 PL	Port lockout. 0 The packets that may be received and issued are controlled by the state of the OPE and IPE bits. 1 This port is stopped and is not enabled to issue or receive any packets; the input port can still follow the training procedure and can still send and respond to link-requests; all received packets return packet-not-accepted control symbols to force an error condition to be signaled by the sending device.
31 PT	Port type (read-only). 0 Reserved 1 Serial port.

15.5.24 Port 2 Link maintenance request command and status register (SRIO_P2LMREQCSR)

The port 2 link maintenance request command and status register (P2LMREQCSR) is accessible both by the local processor and an external device. A write to this register generates a link-request control symbol on the corresponding RapidIO endpoint port interface.

Address: C_0000h base + 160h offset = C_0160h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																															
W	Reserved															C																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2LMREQCSR field descriptions

Field	Description
0–28 -	This field is reserved. Reserved
29–31 C	LINK_REQUEST command to send. If read, this field returns the last written value. If written with a value other than 0b011 (reset-device) or 0b100 (input-status), the resulting operation is undefined, as all other values are reserved in the RapidIO specification.

15.5.25 Port 2 Link maintenance response command and status register (SRIO_P2LMRESPCSR)

The port 2 link maintenance response command and status register (P2LMRESPCSR) is accessible both by the local processor and an external device. A read to this register returns the status received in a link-response control symbol. This register is read only.

Address: C_0000h base + 164h offset = C_0164h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	RV		Reserved													
W	Reserved		Reserved													
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved							AS					LS			
W	Reserved							Reserved					Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

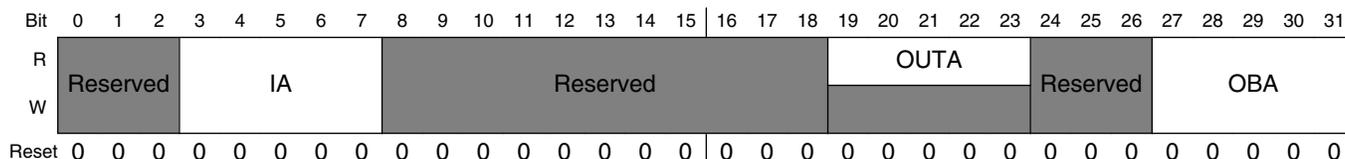
SRIO_P2LMRESPCSR field descriptions

Field	Description
0 RV	Response valid. If the link-request causes a link-response, this bit indicates that the link-response has been received and the status fields are valid. If the link-request does not cause a link-response, this bit indicates that the link-request has been transmitted. This bit clears on read, so be careful when polling its value.
1–21 -	This field is reserved. Reserved
22–26 AS	AckID_status field from LINK_RESPONSE
27–31 LS	Link_status field from LINK_RESPONSE

15.5.26 Port 2 Local ackID status command and status register (SRIO_P2LASCSR)

The port 2 local ackID status command and status register (P2LASCSR) is accessible both by the local processor and an external device. A read to this register returns the local ackID status for both the output and input ports of the device.

Address: C_0000h base + 168h offset = C_0168h



SRIO_P2LASCSR field descriptions

Field	Description
0–2 -	This field is reserved. Reserved
3–7 IA	Input port next expected ackID value.
8–18 -	This field is reserved. Reserved
19–23 OUTA	Outstanding port unacknowledged ackID status. Next expected acknowledge control symbol ackID field that indicates the ackID value expected in the next received acknowledge control symbol. Note that this value is read only even though RapidIO specification allows for it to be writable.
24–26 -	This field is reserved. Reserved

Table continues on the next page...

SRIO_P2LASCSR field descriptions (continued)

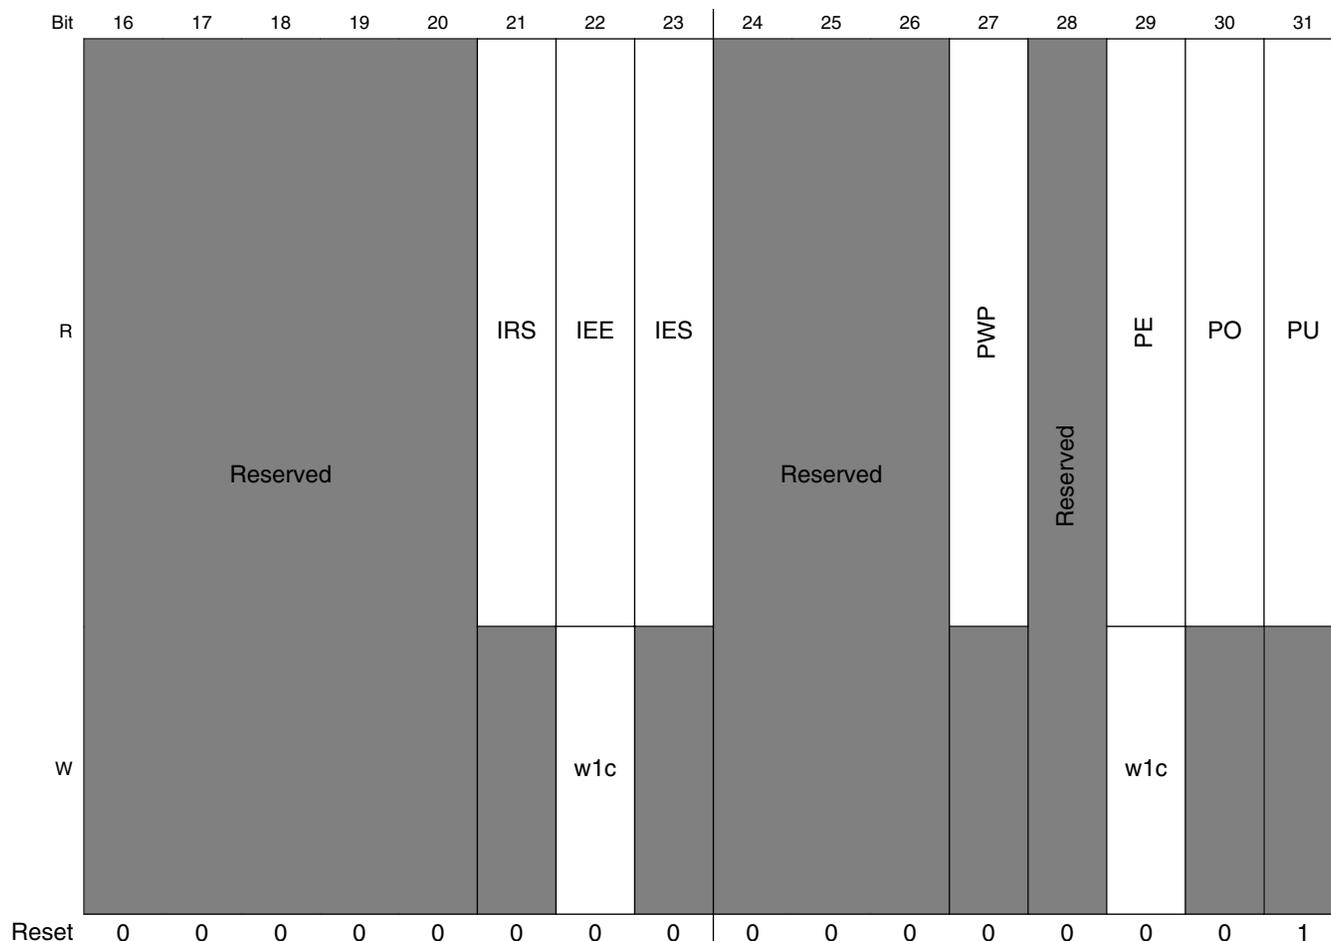
Field	Description
27–31 OBA	Outbound ackID output port next transmitted ackID value. This can be written by software but only if there are no outstanding unacknowledged packets. If there are, the newly-written value is ignored.

15.5.27 Port 2 Error and status command and status register (SRIO_P2ESCSR)

The port 2 error and status command and status register (P2ESCSR) is accessed when the local processor or an external device wishes to examine the port error and status information.

Address: C_0000h base + 178h offset = C_0178h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved					OPD	OFE	ODE	Reserved			ORE	OR	ORS	OEE	OES
W						w1c	w1c	w1c				w1c			w1c	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0



SRIO_P2ESCSR field descriptions

Field	Description
0-4 -	This field is reserved. Reserved
5 OPD	Output Packet-dropped. Output port has discarded a packet. A packet is discarded if: <ul style="list-style-type: none"> It is received while OFE is set and PmCCSR[DPE] (drop packet enable) is set and PmCCSR[SPF] (stop on port failed) is set. It is received while PmPCR[OBDEN] (output buffer drain enable) is set. It is not-accepted by the link-partner while PmERCSR[ERFTT] (error rate failed threshold trigger) is met or exceeded and PmCCSR[DPE] is set and PmCCSR[SPF] is not set (and link-response returns expected ackID). Once OPD is set, it remains set until written with a logic 1 to clear.
6 OFE	Output Failed-encountered. Output port has encountered a failed condition, meaning that the Error Rate Counter has met or exceeded the port's failed error threshold (ERFTT) Once set, remains set until written with a logic 1 to clear. Once cleared, does not assert again unless the Error Rate Counter dips below the port's failed error threshold and then meets or exceeds it again.
7 ODE	Output port has encountered a degraded condition, meaning that the Error Rate Counter has met or exceeded the port's degraded error threshold. Once set, remains set until written with a logic 1 to clear.

Table continues on the next page...

SRIO_P2ESCSR field descriptions (continued)

Field	Description
	Once cleared, does not assert again unless the Error Rate Counter dips below the port's degraded error threshold and then meets or exceeds it again.
8–10 -	This field is reserved. Reserved
11 ORE	Output port has encountered a retry condition. This bit is set when bit 13 is set. Once set, remains set until written with a logic 1 to clear.
12 OR	Output port has received a packet retry control symbol and cannot make forward progress. This bit is set when bit 13 is set and cleared when a packet-accepted or packet-not-accepted control symbol is received. (read only)
13 ORS	Output port is stopped due to a retry (read only)
14 OEE	Output port has encountered (and possibly recovered from) a transmission error. This bit is set when bit 15 is set. Once set, remains set until written with a logic 1 to clear.
15 OES	Output port is stopped due to a transmission error (read only)
16–20 -	This field is reserved. Reserved
21 IRS	Input port is stopped due to a retry (read only)
22 IEE	Input port has encountered (and possibly recovered from) a transmission error. This bit is set when bit 23 is set. Once set, remains set until written with a logic 1 to clear.
23 IES	Input port is stopped due to a transmission error (read-only)
24–26 -	This field is reserved. Reserved
27 PWP	Port has encountered a condition which required it to initiate a maintenance port-write operation. This bit is only valid if the device is capable of issuing a maintenance port-write transaction. The RapidIO endpoint is not capable of issuing port-writes. This bit is hardwired to 0.
28 -	This field is reserved. Reserved
29 PE	Input or output port has encountered an error from which hardware was unable to recover. Once set, remains set until written with a logic 1 to clear. This bit indicates that OFE is set while PmCCSR[SPF] is set; in other words, the failed threshold has been reached which has caused the output port to stop transmitting packets.
30 PO	The input and output ports are initialized and the port is exchanging error-free control symbols with the attached device. (read-only).
31 PU	Input and output ports are not initialized . This bit and bit 30 are mutually exclusive (read-only).

15.5.28 Port 2 Control command and status register (SRIO_P2CCSR)

The port 2 control command and status register (P2CCSR) contains control register bits for the RapidIO port .

Address: C_0000h base + 17Ch offset = C_017Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	PW		IPW			PWO			PD	OPE	IPE	ECD	MEP	Reserved		
W																
Reset	0	1	0	1	0	0	0	0	0	1	1	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved											SPF	DPE	PL	PT	
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

SRIO_P2CCSR field descriptions

Field	Description
0-1 PW	Hardware width of the port (read-only). 00 Single-lane port 01 Four-lane port 10-11 Reserved
2-4 IPW	Width of the ports after initialized (read-only). 000 Single-lane port, lane 0 001 Single-lane port, lane 2 010 Four-lane port 011-111 Reserved
5-7 PWO	Soft port configuration to override the hardware size. This field should be changed only when the port is uninitialized. To achieve this, first disable the RapidIO port. Then change PWO to any valid value. Finally, re-enable the RapidIO port. To achieve this, first set PmCCSR[PD] (port disabled). Then change PWO to any legal value. Finally, clear PmCCSR[PD] (enabled). 000 No override 001 Reserved 010 Force single lane, lane 0 011 Force single lane, lane 2 100-111 Reserved, causes undefined operation
8 PD	Port disable.

Table continues on the next page...

SRIO_P2CCSR field descriptions (continued)

Field	Description
	0 Input error state machine operates normally 1 Input error state machine is forced to normal state
9 OPE	Output port transmit enable. OPE is ignored by RapidIO endpoints. It is expected that if OPE = 0, software will not send packets out of outbound. If packets are sent by OCN to outbound, they are sent out of RapidIO endpoints, regardless of the value of OPE. Initial value read from configuration pins. 0 Port is stopped and not enabled to issue any packets except to route or respond to I/O logical MAINTENANCE packets. Control symbols are not affected and are sent normally. 1 Port is enabled to issue any packets.
10 IPE	Input port receive enable. This bit value must equal the value of the output port enable (OPE) bit in order for the RapidIO controller to function properly. Initial value read from configuration pins. 0 Port is stopped and only enabled to route or respond to I/O logical MAINTENANCE packets. Other packets generate packet-not-accepted control symbols to force an error condition to be signaled by the sending device. Control symbols are not affected and are received and handled normally. 1 Port is enabled to respond to any packet.
11 ECD	Error checking disable. This bit is hardwired to 0. This bit disables all RapidIO transmission error checking 0 Error checking and recovery is enabled. 1 Error checking and recovery is disabled.
12 MEP	Multicast-event participant. This bit is hard-wired to 0.
13–27 -	This field is reserved. Reserved
28 SPF	Stop on port failed-encountered enable. This bit is used with the drop packet enable bit to force certain behavior when the error rate failed threshold has been met or exceeded.
29 DPE	Drop packet enable. This bit is used with the stop on port failed-encountered enable bit to force certain behavior when the error rate failed threshold has been met or exceeded.
30 PL	Port lockout. 0 The packets that may be received and issued are controlled by the state of the OPE and IPE bits. 1 This port is stopped and is not enabled to issue or receive any packets; the input port can still follow the training procedure and can still send and respond to link-requests; all received packets return packet-not-accepted control symbols to force an error condition to be signaled by the sending device.
31 PT	Port type (read-only). 0 Reserved 1 Serial port.

15.5.29 Error reporting block header (SRIO_ERBH)

The error reporting block header register contains the EF_PTR to the next EF_BLK (the next EF_PTR is 0x0000 since this is the last set of registers in the extended features space) and the EF_ID that identifies this as the error reporting block header. ERBH is a read-only register.

Address: C_0000h base + 600h offset = C_0600h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
R	EF_PTR															EF_ID																			
W	[Shaded]																																		
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1

SRIO_ERBH field descriptions

Field	Description
0–15 EF_PTR	Extended features pointer
16–31 EF_ID	Extended features ID

15.5.30 Logical/Transport layer error detect command and status register (SRIO_LTLEDCSR)

This register indicates the error that was detected by the Logical or Transport logic layer. Software should write this register with all 0s to clear any detected errors and unlock the capture registers.

Error information that corresponds to two or more different error events are not captured on the same clock cycle. However, one error event (such as a corrupted inbound packet) can cause multiple bits to be set. The priority of errors is PRT and then all other errors. OACB error results in PRT. When PRT bit is set with OACB bit, error capture is for an OCN transaction for which an Outbound ATMU window boundary was crossed or a segment or subsegment boundary was crossed.

An error that is not enabled will set the detect bit in this register as long as a capture has not yet occurred.

Note that fields in this register can be set by writing to this register. This can be used to emulate a hardware error during software development. However, undefined results occur if fields in this register are set while an actual Logical/Transport Layer error is being detected. Also, note that this register can be written with an invalid combination of bits set and care should be taken to avoid this.

Address: C_0000h base + 608h offset = C_0608h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R											Reserved					
W											Reserved					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								IACB	OACB	DMAMER	RETE	TSE	PTTL	Reserved	
W	Reserved														Reserved	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_LTLEDCSR field descriptions

Field	Description
0 IER	IO error response. Received a response of ERROR for an IO logical layer request.
1 MER	Reserved for message error response. Received a response of ERROR for an MSG logical layer request. Error detected and captured in the message unit, if one exists.
2 GER	GSM error response. Received a response of ERROR for a GSM logical layer request.
3 MFE	Reserved for message format error. Received MESSAGE packet data payload with an invalid size or segment. Error detected and captured in the message unit, if one exists.
4 ITD	Illegal transaction decode. Received illegal fields in the request/response packet for a supported transaction (IO/MSG/GSM logical)
5 ITTE	Illegal transaction target error. Received a packet that contained a destination ID that is not defined for this end point . Endpoints with multiple ports and a built-in switch function may not report this as an error (transport)
6 MRT	Reserved for message request time-out. A required message request has not been received within the specified time-out interval. Error detected and captured in the message unit, if one exists.

Table continues on the next page...

SRIO_LTLEDCSR field descriptions (continued)

Field	Description
7 PRT	Packet response time-out. A required response has not been received within the specified time out interval (IO/MSG/GSM logical)
8 UR	Unsolicited response. An unsolicited/unexpected response packet was received (IO/MSG/GSM logical; only maintenance response for switches)
9 UT	Unsupported transaction. A transaction is received that is not supported in the destination operations CAR (IO/MSG/GSM logical; only maintenance port-write for switches)
10–23 -	This field is reserved. Reserved
24 IACB	Inbound ATMU crossed boundary. A transaction is received that crosses an inbound ATMU boundary.
25 OACB	Outbound ATMU crossed boundary. A transaction is being sent that crosses an outbound ATMU boundary, a segment boundary, or a subsegment boundary.
26 DMAMER	DMA message error response. An error response was received for a DMA message (detected in the RapidIO endpoint, not the message unit).
27 RETE	Retry error threshold exceeded. The allowed number of logical retries (given by LRETCR[RET]) has been exceeded. This bit is also driven by the Message Unit when the allowed number of message retries has been exceeded.
28 TSE	Transport size error. The tt field is not consistent with bit 27 of the processing element features CAR (that is, the tt value is reserved or indicates a common transport system that is unsupported by this device).
29 PTTL	Packet time-to-live error. A packet time-to-live error occurred (a packet could not be successfully transmitted before the packet time-to-live counter expired).
30–31 -	This field is reserved. Reserved

15.5.31 Logical/Transport layer error enable command and status register (SRIO_LTLEECR)

This register contains the bits that control whether an error condition locks the logical/transport layer error detect and capture registers and is reported to the system host. LTLEEDCSR is stored in all ports and the message unit.

Address: C_0000h base + 60Ch offset = C_060Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R											Reserved					
W	IER	MER	GER	MFE	ITD	ITTE	MRT	PRT	UR	UT						
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								IACB	OACB	DMAMER	RETE	TSE	PTTL	Reserved	
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_LTLEECR field descriptions

Field	Description
0 IER	IO error response enable.Enable reporting of an IO error response. Capture and lock the error.
1 MER	Message error response enable. Enable reporting of a Message error response. Capture and lock the error (capture done in Message Unit, if one exists).
2 GER	GSM error response enable. Enable reporting of a GSM error response. Capture and lock the error.
3 MFE	Message format error enable. Enable reporting of a message format error. Capture and lock the error.(capture done in Message Unit, if one exists).
4 ITD	Illegal transaction decode enable. Enable reporting of an illegal transaction decode error. Capture and lock the error.
5 ITTE	Illegal transaction target error enable. Enable reporting of an illegal transaction target error. Capture and lock the error.
6 MRT	Message request time-out enable. Enable reporting of a Message Request time-out error. Capture and lock the error. (capture done in Message Unit, if one exists)

Table continues on the next page...

SRIO_LTLECSR field descriptions (continued)

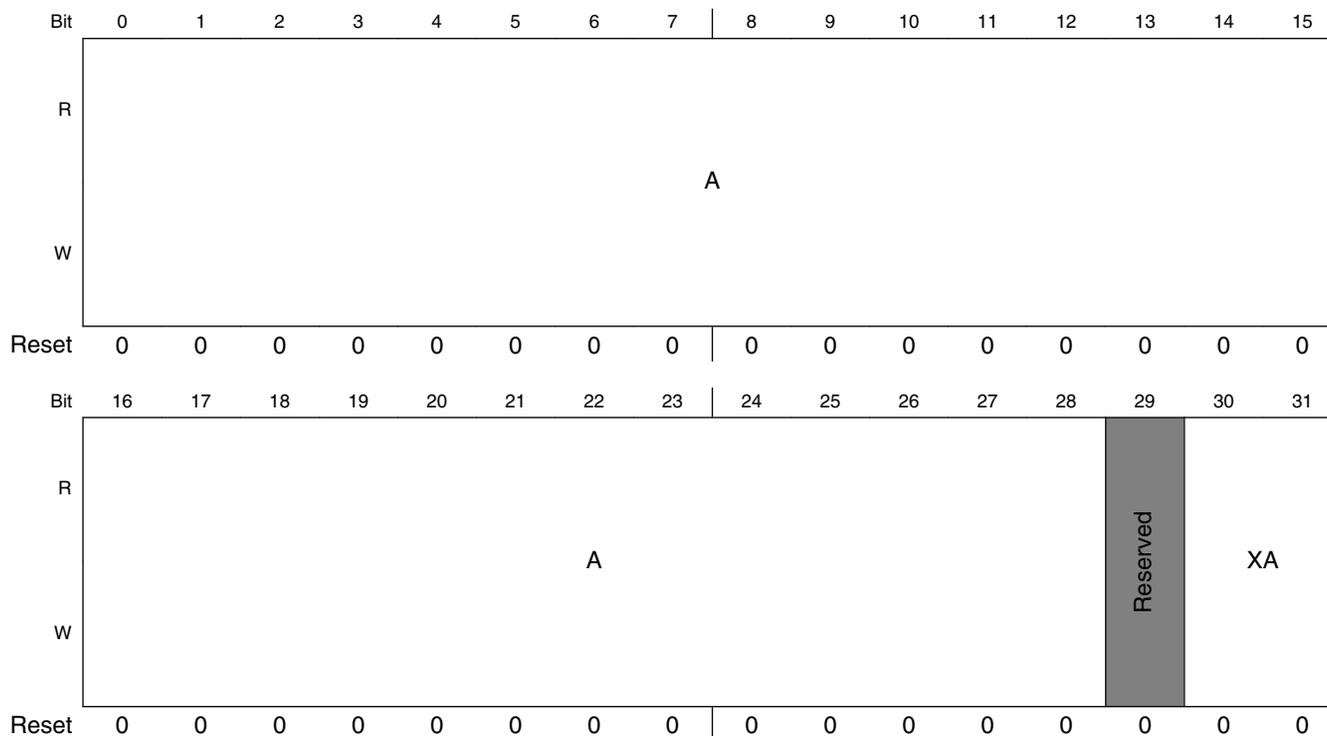
Field	Description
7 PRT	Packet response time-out error enable. Enable reporting of a packet response time-out error. Capture and lock the error.
8 UR	Unsolicited response error enable. Enable reporting of an unsolicited response error. Capture and lock the error.
9 UT	Unsupported transaction error enable. Enable reporting of an unsupported transaction error. Capture and lock the error.
10–23 -	This field is reserved. Reserved
24 IACB	Inbound ATMU crossed boundary error enable. Enable reporting of a received transaction that crosses an inbound ATMU boundary. Capture and lock the error.
25 OACB	Outbound ATMU crossed boundary error enable. Enable reporting of a transaction that crosses an outbound ATMU boundary, a segment boundary, or a subsegment boundary. Capture and lock the error.
26 DMAMER	DMA message error response. Enable error reporting of an error response for a DMA message. Capture and lock the error.
27 RETE	Retry error threshold exceeded. Enable error reporting when the allowed number of logical retries has been exceeded.
28 TSE	Transport size error. Enable error reporting when the tt field is not consistent with bit 27 of the Processing Element Features CAR (that is, the tt value is reserved or indicates a common transport system that is unsupported by this device).
29 PTTL	Packet time-to-live error. Enable reporting of a packet time-to-live time-out error. Capture and lock the error.
30–31 -	This field is reserved. Reserved

15.5.32 Logical/Transport layer address capture command and status register (SRIO_LTLACCSR)

The LTLACCSR provides error information. It is locked when a logical/transport error is detected, and the corresponding enable bit is set. LTLACCSR is stored in each port and in the message unit, although the values in this register can differ between each port and message unit. The message unit LTLACCSR cannot lock if the message unit or any other port has locked.

Note that fields in this register can be set by to writing this register. This can be used to emulate a hardware error during software development. However, undefined results occur if fields in this register are set while an actual logical/transport layer error is being detected.

Address: C_0000h base + 614h offset = C_0614h



SRIO_LTLACCSR field descriptions

Field	Description
0–28 A	Normally the least significant 29 bits of the address associated with the error (for requests, for responses if available). See Logical layer errors and error handling for details.
29 -	This field is reserved. Reserved
30–31 XA	xamsbs. Normally the extended address bits of the address associated with the error (for requests, responses, if available). See Logical layer errors and error handling for details.

15.5.33 Logical/Transport layer device ID capture command and status register (SRIO_LTLDIDCCSR)

LTLDIDCCSR contains error information. It is locked when a logical/transport error is detected, and the corresponding enable bit is set. LTLDIDCCSR is stored in each port and in the message unit, although the values in this register can differ between each port and message unit. The message unit LTLDIDCCSR cannot lock if the message unit or any other port has locked.

Note that fields in this register can be set by writing to this register, in order to emulate a hardware error during software development. However, undefined results occur if fields in this register are set while an actual logical/transport layer error is being detected.

Serial RapidIO Memory Map/Register Definition

Address: C_0000h base + 618h offset = C_0618h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	DIDMSB							DID							SIDMSB							SID										
W	DIDMSB							DID							SIDMSB							SID										
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_LTLIDCCSR field descriptions

Field	Description
0–7 DIDMSB	Normally the most significant byte of the destinationID associated with the error. This field is valid only if bit 27 of the Processing Element Features CAR is set (large transport systems only). See Logical layer errors and error handling for details.
8–15 DID	Normally the destinationID (or least significant byte of the destination ID if large transport system) associated with the error. See Logical layer errors and error handling for details.
16–23 SIDMSB	Normally the most significant byte of the sourceID associated with the error. This field is valid only if bit 27 of the Processing Element Features CAR is set (large transport systems only). See Logical layer errors and error handling for details.
24–31 SID	Normally the sourceID (or least significant byte of the source ID if large transport system) associated with the error. See Logical layer errors and error handling for details.

15.5.34 Logical/Transport layer control capture command and status register (SRIO_LTLCCCSR)

LTLCCCS contains error information. It is stored in each port and in the message unit, although the values in this register can differ between each port and message unit. The message unit LTLCCCSR cannot lock if the message unit or any other port has locked.

Note that fields in this register can be set by writing to this register, in order to emulate a hardware error during software development. However, undefined results occur if fields in this register are set while an actual logical/transport layer error is being detected.

Address: C_0000h base + 61Ch offset = C_061Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	FT			TT			MI									Reserved																
W	FT			TT			MI									Reserved																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_LTLCCCSR field descriptions

Field	Description
0–3 FT	Normally the format type associated with the error. See Logical layer errors and error handling for details.
4–7 TT	Normally the transaction type associated with the error. See Logical layer errors and error handling for details.
8–15 MI	Normally the message Information: letter, mbox, and msgseg for the last message request received for the mailbox that had an error (message errors only). See Logical layer errors and error handling for details.
16–31 -	This field is reserved. Reserved

15.5.35 Port 1 Error detect command and status register (SRIO_P1EDCSR)

The port *m* error detect command and status register (*Pm*EDCSR) indicates transmission errors that are detected by the hardware. Software can write bits in this register with 1 to cause the Error Rate Counter to increment. Undefined results occur if this register is written while actual physical layer errors are being detected by the port.

Address: C_0000h base + 640h offset = C_0640h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved									CCS	AUA	PNA	UA	CRC	EM	Reserv
W	Reserved															ed
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										NOA	PE	-	DE	UCS	LTO
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1EDCSR field descriptions

Field	Description
0–8 -	This field is reserved. Reserved
9 CCS	Received a control symbol with a bad CRC value.
10 AUA	Received acknowledge control symbol with unexpected ackID (packet-accepted or packet-retry).
11 PNA	Received packet-not-accepted acknowledge control symbol
12 UA	Received packet with unexpected ackID value.
13 CRC	Received a packet with a bad CRC value
14 EM	Received packet which exceed the maximum allowed size (276 bytes).
15–25 -	This field is reserved. Reserved
26 NOA	Link-response received with an ackID that is not outstanding.
27 PE	Protocol Error: An unexpected packet or control symbol was received.
28 -	Reserved

Table continues on the next page...

SRIO_P1EDCSR field descriptions (continued)

Field	Description
29 DE	Received unaligned /SC/ or /PD/ or undefined code-group.
30 UCS	An unexpected acknowledge control symbol was received.
31 LTO	An acknowledge or link-response control symbol is not received within the specified time-out interval.

15.5.36 Port 1 Error rate enable command and status register (SRIO_P1ERECSR)

The P_m ERECSR contains the bits that control when an error condition is allowed to increment the error rate counter in the port m error rate threshold register and lock the port m error capture registers.

Address: C_0000h base + 644h offset = C_0644h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved									CCS	AUA	PNA	UA	CRC	EM	Reserved
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										NOA	PE	-	DE	UCS	LTO
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1ERECSR field descriptions

Field	Description
0–8 -	This field is reserved. Reserved
9 CCS	Enable error rate counting of a corrupt control symbol
10 AUA	Enable error rate counting of an acknowledge control symbol with an unexpected ackID
11 PNA	Enable error rate counting of received packet-not-accepted control symbols.
12 UA	Enable error rate counting of packet with unexpected ackID value.
13 CRC	Enable error rate counting of packet with a bad CRC value.
14 EM	Enable error rate counting of packet which exceeds the maximum allowed size

Table continues on the next page...

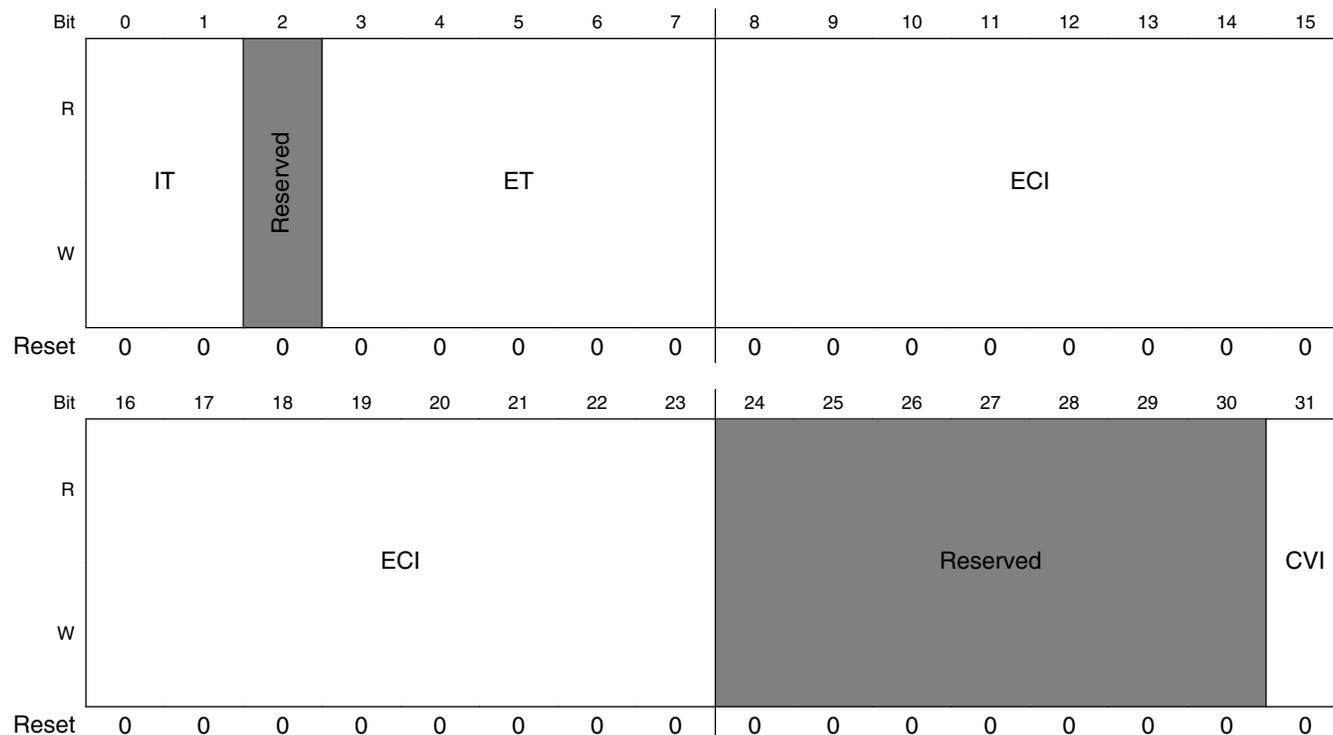
SRIO_P1ERECSR field descriptions (continued)

Field	Description
15–25 -	This field is reserved. Reserved
26 NOA	Enable error rate counting of link-responses received with an ackID that is not outstanding.
27 PE	Enable error rate counting of protocol errors
28 -	Reserved
29 DE	Enable error rate counting of delineation errors.
30 UCS	Enable error rate counting of unsolicited acknowledge control symbol errors.
31 LTO	Enable error rate counting of link time-out errors.

15.5.37 Port 1 Error capture attributes command and status register (SRIO_P1ECACSR)

The error capture attribute register indicates the type of information contained in the port *m* error capture registers. In the case of multiple detected errors during the same clock cycle one of the errors must be reflected in the Error type field. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P *m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 648h offset = C_0648h



SRIO_P1ECACSR field descriptions

Field	Description
0-1 IT	Type of information logged: 00 Packet (error capture registers hold the first 4 words of the packet, or the entire packet if it is less than 4 words long). 01 Control symbol (only error capture register 0 is valid) 10 Reserved 11 Undefined (not clearly a control symbol or packet error. Error capture registers hold the symbol that caused the error and the next 3 symbols.)

Table continues on the next page...

SRIO_P1ECACSR field descriptions (continued)

Field	Description
2 -	This field is reserved. Reserved
3–7 ET	The encoded value of the bit in the port <i>m</i> error detect CSR that describes the error captured in the port <i>m</i> error capture CSRs
8–23 ECI	Extended capture information [0:15]. ECI contains the control/data character signal corresponding to each byte of captured data.
24–30 -	This field is reserved. Reserved
31 CVI	This bit is set by hardware to indicate that the Packet/control symbol capture registers contain valid information. For control symbols, only capture register 0 contains meaningful information.

15.5.38 Port 1 Packet/control symbol error capture command and status register 0 (SRIO_P1PCSECCSR0)

P m ECACSR contains the first four bytes of captured packet symbol information or a control character and control symbol. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the *P m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 64Ch offset = C_064Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	C0																															
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

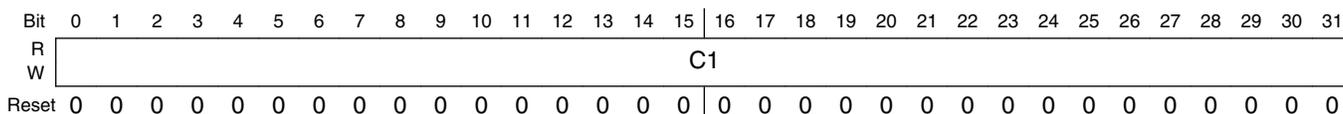
SRIO_P1PCSECCSR0 field descriptions

Field	Description
0–31 C0	Capture 0: Control Character and control symbol or bytes 0 to 3 of packet header.

15.5.39 Port 1 Packet error capture command and status register 1 (SRIO_P1PECCSR1)

Error capture register 1 contains bytes 4 through 7 of the packet header. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P m ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 650h offset = C_0650h



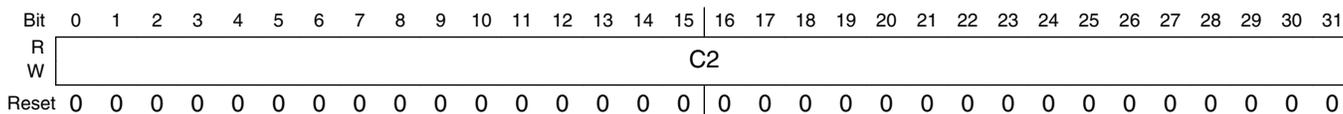
SRIO_P1PECCSR1 field descriptions

Field	Description
0–31 C1	Capture 1. Bytes 4 to 7 of the packet header

15.5.40 Port 1 Packet error capture command and status register 2 (SRIO_P1PECCSR2)

Error capture register 2 contains bytes 8 through 11 of the packet header. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P m ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 654h offset = C_0654h



SRIO_P1PECCSR2 field descriptions

Field	Description
0–31 C2	Capture 2. Bytes 8 to 11 of the packet header

15.5.41 Port 1 Packet error capture command and status register 3 (SRIO_P1PECCSR3)

Error capture register 3 contains bytes 12 through 15 of the packet header. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P *m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 658h offset = C_0658h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	C3																															
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1PECCSR3 field descriptions

Field	Description
0–31 C3	Capture 3. Bytes 12 to 15 of the packet header

15.5.42 Port 1 Error rate command and status register (SRIO_P1ERCSR)

The port *m* error rate register is a 32-bit register used with the port *m* error rate threshold register to monitor and control the reporting of transmission errors.

Address: C_0000h base + 668h offset = C_0668h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	ERB							Reserved								ERR				PER					ERC							
W																																
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1ERCSR field descriptions

Field	Description
0–7 ERB	These bits provide the error rate bias value. Other values are reserved and cause undefined operation. 0x00 Do not decrement the error rate counter 0x01 Decrement every 1 ms (±34%) 0x02 Decrement every 10 ms (±34%) 0x04 Decrement every 100 ms (±34%) 0x08 Decrement every 1 s (±34%)

Table continues on the next page...

SRIO_P1ERCSR field descriptions (continued)

Field	Description
	0x10 Decrement every 10 s ($\pm 34\%$) 0x20 Decrement every 100 s ($\pm 34\%$) 0x40 Decrement every 1000 s ($\pm 34\%$) 0x80 Decrement every 10000 s ($\pm 34\%$)
8–13 -	This field is reserved. Reserved
14–15 ERR	These bits limit the incrementing of the error rate counter above the failed threshold trigger. Note that the Error Rate Counter never increments above 0xFF, even if the combination of the settings of ERR and the failed threshold trigger imply that it might. 00 Only count 2 errors above 01 Only count 4 errors above 10 Only count 16 error above 11 Do not limit incrementing the error rate count
16–23 PER	Peak error rate. Contains the peak value attained by the error rate counter
24–31 ERC	Error rate counter. These bits maintain a count of the number of transmission errors that have been detected by the port, decremented by the Error Rate Bias mechanism, to create an indication of the link error rate. Software should not attempt to write this field to a value higher than failed threshold trigger plus the number of errors specified in the ERR field (the maximum ERC value).

15.5.43 Port 1 Error rate threshold command and status register (SRIO_P1ERTCSR)

The port *m* error rate threshold register is a 32-bit register used to control the reporting of the link status to the system host.

Address: C_0000h base + 66Ch offset = C_066Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	ERFTT															ERDTT											Reserved					
W	1															1											0					
Reset	1															1											0					

SRIO_P1ERTCSR field descriptions

Field	Description
0–7 ERFTT	Error rate failed threshold trigger. These bits provide the threshold value for reporting an error condition due to a possibly broken link. The P <i>m</i> ESCSR[OFE] bit is not set if ERFTT is written to a value lower than or equal to the P <i>m</i> ERCSR[ERC]. See Table 15-324 for more details.

Table continues on the next page...

SRIO_P1ERTCSR field descriptions (continued)

Field	Description
	0x00 Disable the Error Rate Failed Threshold Trigger. 0x01 Set the error reporting threshold to 1. 0x02 Set the error reporting threshold to 2. ... 0xFF Set the error reporting threshold to 255.
8–15 ERDTT	Error rate degraded threshold trigger. These bits provide the threshold value for reporting an error condition due to a degrading link. The P <i>m</i> ESCSR[ODE] bit is not set if ERDTT is written to a value lower than or equal to the P <i>m</i> ERCSR[ERC]. See Table 15-324 for more details. 0x00 Disable the Error Rate Degraded Threshold Trigger. 0x01 Set the error reporting threshold to 1. 0x02 Set the error reporting threshold to 2. ... 0xFF Set the error reporting threshold to 255.
16–31 -	This field is reserved. Reserved

15.5.44 Port 2 Error detect command and status register (SRIO_P2EDCSR)

The port *m* error detect command and status register (P *m* EDCSR) indicates transmission errors that are detected by the hardware. Software can write bits in this register with 1 to cause the Error Rate Counter to increment. Undefined results occur if this register is written while actual physical layer errors are being detected by the port.

Address: C_0000h base + 680h offset = C_0680h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved										CCS	AUA	PNA	UA	CRC	EM	Reserv	
W	Reserved																ed	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved											NOA	PE	-	DE	UCS	LTO	
W	Reserved																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	

SRIO_P2EDCSR field descriptions

Field	Description
0–8 -	This field is reserved. Reserved
9 CCS	Received a control symbol with a bad CRC value.
10 AUA	Received acknowledge control symbol with unexpected ackID (packet-accepted or packet-retry).
11 PNA	Received packet-not-accepted acknowledge control symbol
12 UA	Received packet with unexpected ackID value.
13 CRC	Received a packet with a bad CRC value
14 EM	Received packet which exceed the maximum allowed size (276 bytes).
15–25 -	This field is reserved. Reserved
26 NOA	Link-response received with an ackID that is not outstanding.
27 PE	Protocol Error: An unexpected packet or control symbol was received.
28 -	Reserved
29 DE	Received unaligned /SC/ or /PD/ or undefined code-group.
30 UCS	An unexpected acknowledge control symbol was received.
31 LTO	An acknowledge or link-response control symbol is not received within the specified time-out interval.

15.5.45 Port 2 Error rate enable command and status register (SRIO_P2ERECSR)

The P_m ERECSR contains the bits that control when an error condition is allowed to increment the error rate counter in the port m error rate threshold register and lock the port m error capture registers.

Address: C_0000h base + 684h offset = C_0684h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved								CCS	AUA	PNA	UA	CRC	EM	Reserved	
W	Reserved								CCS	AUA	PNA	UA	CRC	EM	Reserved	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										NOA	PE	-	DE	UCS	LTO
W	Reserved										NOA	PE	-	DE	UCS	LTO
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2ERECSR field descriptions

Field	Description
0–8 -	This field is reserved. Reserved
9 CCS	Enable error rate counting of a corrupt control symbol
10 AUA	Enable error rate counting of an acknowledge control symbol with an unexpected ackID
11 PNA	Enable error rate counting of received packet-not-accepted control symbols.
12 UA	Enable error rate counting of packet with unexpected ackID value.
13 CRC	Enable error rate counting of packet with a bad CRC value.
14 EM	Enable error rate counting of packet which exceeds the maximum allowed size
15–25 -	This field is reserved. Reserved
26 NOA	Enable error rate counting of link-responses received with an ackID that is not outstanding.
27 PE	Enable error rate counting of protocol errors
28 -	Reserved
29 DE	Enable error rate counting of delineation errors.

Table continues on the next page...

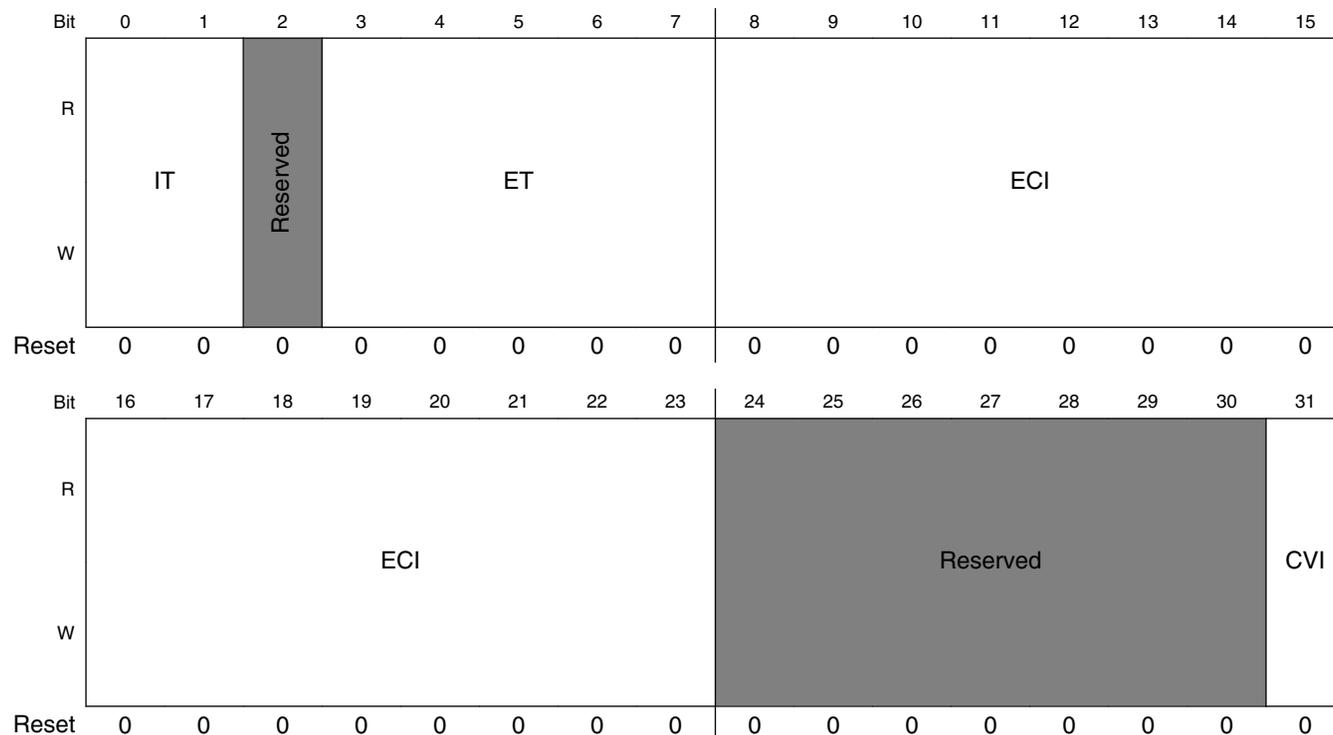
SRIO_P2ERECSR field descriptions (continued)

Field	Description
30 UCS	Enable error rate counting of unsolicited acknowledge control symbol errors.
31 LTO	Enable error rate counting of link time-out errors.

15.5.46 Port 2 Error capture attributes command and status register (SRIO_P2ECACSR)

The error capture attribute register indicates the type of information contained in the port *m* error capture registers. In the case of multiple detected errors during the same clock cycle one of the errors must be reflected in the Error type field. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P *m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 688h offset = C_0688h



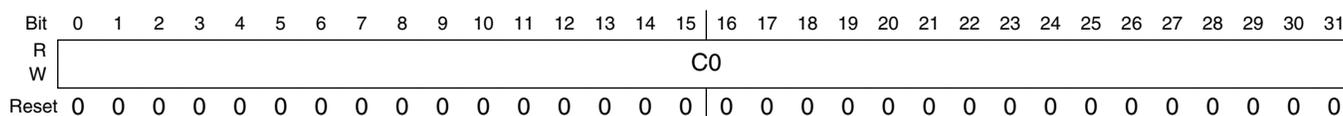
SRIO_P2ECACSR field descriptions

Field	Description
0–1 IT	Type of information logged: 00 Packet (error capture registers hold the first 4 words of the packet, or the entire packet if it is less than 4 words long). 01 Control symbol (only error capture register 0 is valid) 10 Reserved 11 Undefined (not clearly a control symbol or packet error. Error capture registers hold the symbol that caused the error and the next 3 symbols.)
2 -	This field is reserved. Reserved
3–7 ET	The encoded value of the bit in the port <i>m</i> error detect CSR that describes the error captured in the port <i>m</i> error capture CSRs
8–23 ECI	Extended capture information [0:15]. ECI contains the control/data character signal corresponding to each byte of captured data.
24–30 -	This field is reserved. Reserved
31 CVI	This bit is set by hardware to indicate that the Packet/control symbol capture registers contain valid information. For control symbols, only capture register 0 contains meaningful information.

15.5.47 Port 2 Packet/control symbol error capture command and status register 0 (SRIO_P2PCSECCSR0)

P m ECACSR contains the first four bytes of captured packet symbol information or a control character and control symbol. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the *P m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 68Ch offset = C_068Ch



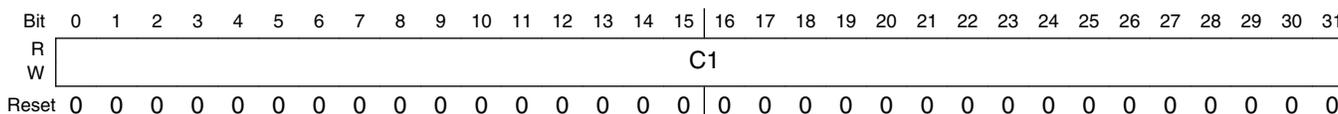
SRIO_P2PCSECCSR0 field descriptions

Field	Description
0–31 C0	Capture 0: Control Character and control symbol or bytes 0 to 3 of packet header.

15.5.48 Port 2 Packet error capture command and status register 1 (SRIO_P2PECCSR1)

Error capture register 1 contains bytes 4 through 7 of the packet header. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P *m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 690h offset = C_0690h



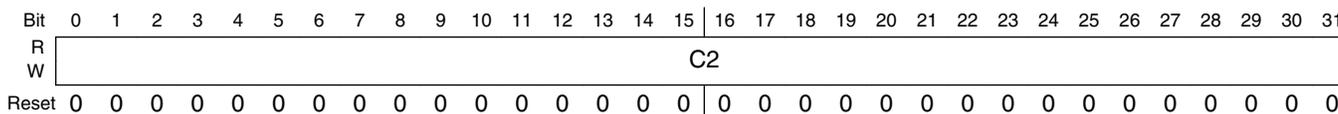
SRIO_P2PECCSR1 field descriptions

Field	Description
0–31 C1	Capture 1. Bytes 4 to 7 of the packet header

15.5.49 Port 2 Packet error capture command and status register 2 (SRIO_P2PECCSR2)

Error capture register 2 contains bytes 8 through 11 of the packet header. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P *m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 694h offset = C_0694h



SRIO_P2PECCSR2 field descriptions

Field	Description
0–31 C2	Capture 2. Bytes 8 to 11 of the packet header

15.5.50 Port 2 Packet error capture command and status register 3 (SRIO_P2PECCSR3)

Error capture register 3 contains bytes 12 through 15 of the packet header. Undefined results occur if this register is written while actual physical layer errors are being detected by the port. Software should check that the P *m* ECACSR[CVI] bit is set before reading the capture registers to ensure that the error has been properly captured.

Address: C_0000h base + 698h offset = C_0698h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	C3																															
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2PECCSR3 field descriptions

Field	Description
0–31 C3	Capture 3. Bytes 12 to 15 of the packet header

15.5.51 Port 2 Error rate command and status register (SRIO_P2ERCSR)

The port *m* error rate register is a 32-bit register used with the port *m* error rate threshold register to monitor and control the reporting of transmission errors.

Address: C_0000h base + 6A8h offset = C_06A8h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	ERB							Reserved								ERR			PER						ERC							
W																																
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2ERCSR field descriptions

Field	Description
0–7 ERB	These bits provide the error rate bias value. Other values are reserved and cause undefined operation. 0x00 Do not decrement the error rate counter 0x01 Decrement every 1 ms (±34%) 0x02 Decrement every 10 ms (±34%) 0x04 Decrement every 100 ms (±34%) 0x08 Decrement every 1 s (±34%)

Table continues on the next page...

SRIO_P2ERCSR field descriptions (continued)

Field	Description
	0x10 Decrement every 10 s ($\pm 34\%$) 0x20 Decrement every 100 s ($\pm 34\%$) 0x40 Decrement every 1000 s ($\pm 34\%$) 0x80 Decrement every 10000 s ($\pm 34\%$)
8–13 -	This field is reserved. Reserved
14–15 ERR	These bits limit the incrementing of the error rate counter above the failed threshold trigger. Note that the Error Rate Counter never increments above 0xFF, even if the combination of the settings of ERR and the failed threshold trigger imply that it might. 00 Only count 2 errors above 01 Only count 4 errors above 10 Only count 16 error above 11 Do not limit incrementing the error rate count
16–23 PER	Peak error rate. Contains the peak value attained by the error rate counter
24–31 ERC	Error rate counter. These bits maintain a count of the number of transmission errors that have been detected by the port, decremented by the Error Rate Bias mechanism, to create an indication of the link error rate. Software should not attempt to write this field to a value higher than failed threshold trigger plus the number of errors specified in the ERR field (the maximum ERC value).

15.5.52 Port 2 Error rate threshold command and status register (SRIO_P2ERTCSR)

The port *m* error rate threshold register is a 32-bit register used to control the reporting of the link status to the system host.

Address: C_0000h base + 6ACh offset = C_06ACh

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2ERTCSR field descriptions

Field	Description
0–7 ERFTT	Error rate failed threshold trigger. These bits provide the threshold value for reporting an error condition due to a possibly broken link. The P <i>m</i> ESCSR[OFE] bit is not set if ERFTT is written to a value lower than or equal to the P <i>m</i> ERCSR[ERC]. See Table 15-324 for more details.

Table continues on the next page...

SRIO_P2ERTCSR field descriptions (continued)

Field	Description
	0x00 Disable the Error Rate Failed Threshold Trigger. 0x01 Set the error reporting threshold to 1. 0x02 Set the error reporting threshold to 2. ... 0xFF Set the error reporting threshold to 255.
8–15 ERDTT	Error rate degraded threshold trigger. These bits provide the threshold value for reporting an error condition due to a degrading link. The P m ESCSR[ODE] bit is not set if ERDTT is written to a value lower than or equal to the P m ERCSR[ERC]. See Physical layer RapidIO errors detected for more details. 0x00 Disable the Error Rate Degraded Threshold Trigger. 0x01 Set the error reporting threshold to 1. 0x02 Set the error reporting threshold to 2. ... 0xFF Set the error reporting threshold to 255.
16–31 -	This field is reserved. Reserved

15.5.53 Logical layer configuration register (SRIO_LLCR)

The logical layer configuration register contains general port-common logical layer mode enables.

Address: C_0000h base + 1_0004h offset = D_0004h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R				Reserved												
W	Unused	DMAME	ECRAB	Reserved												
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

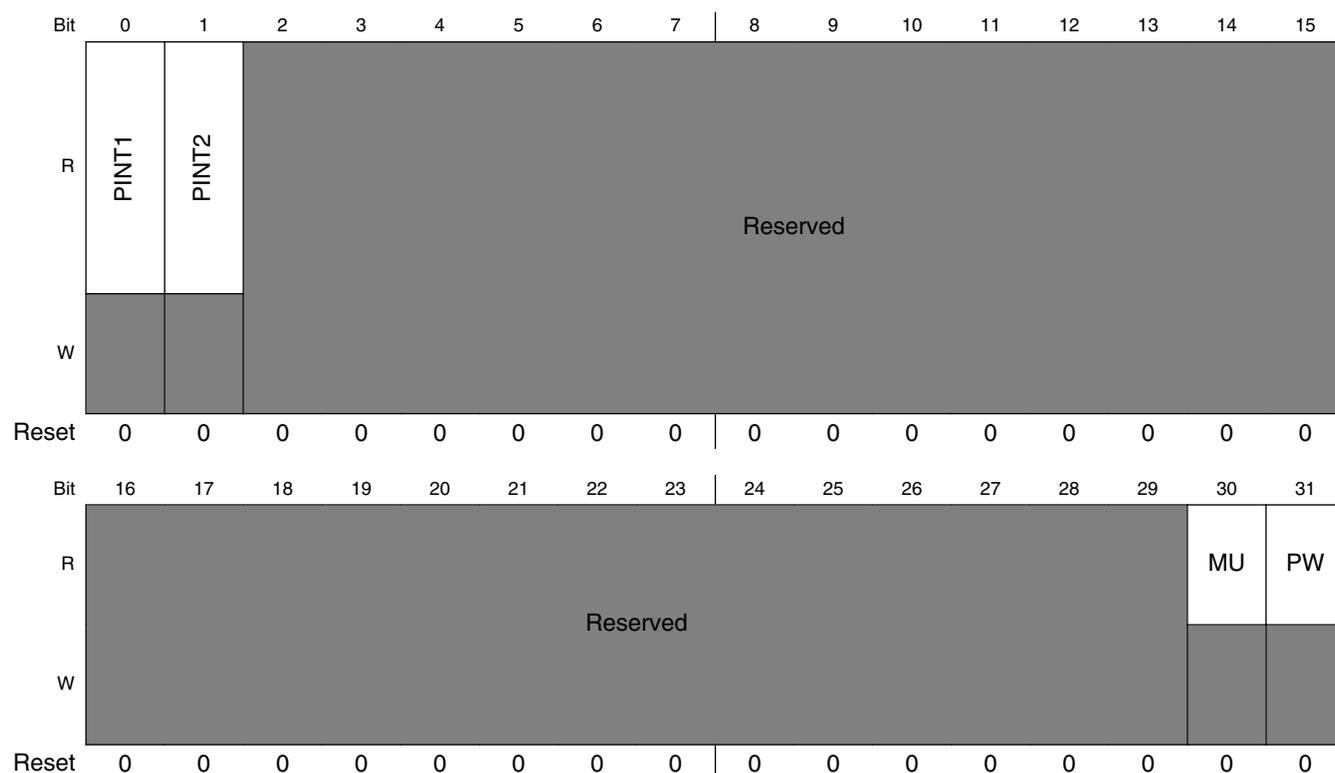
SRIO_LLCR field descriptions

Field	Description
0 Unused	This bit is unused. It is readable and writable.
1 DMAME	DMA message enable (for outbound DMA messages). 0 Message Unit messages can be assigned to letter 0,1,2,3. No DMA Messages. 1 DMA messages are assigned letter 3, Message Unit messages 0,1,2
2 ECRAB	External configuration register access block. When set, all maintenance requests and accesses that hit LCSBA1CSR are blocked; reads return all 0's, and writes are ignored (both return done response). When clear, any external RapidIO device can access registers.
3–31 -	This field is reserved. Reserved

15.5.54 Error / port-write interrupt status register (SRIO_EPWISR)

The EPWISR register contains status bits of the interrupts that have been generated by any port or the message unit for physical or logical/transport layer errors or inbound port-writes. Because errors from all ports are reported to the core with one interrupt signal, this register provides the core with quick access to where the error occurred. This register is read only and is stored in each port and the message unit as a logically equivalent copy.

Address: C_0000h base + 1_0010h offset = D_0010h



SRIO_EPWISR field descriptions

Field	Description
0 PINT1	A physical or logical/transport error interrupt was generated for port 1 . This bit is also set for outbound doorbell packet response time-out (PRT) errors.
1 PINT2	A physical or logical/transport error interrupt was generated for port 2. This bit is also set for outbound doorbell packet response time-out (PRT) errors.
2–29 -	This field is reserved. Reserved
30 MU	A logical/transport layer error interrupt was generated in the message unit.

Table continues on the next page...

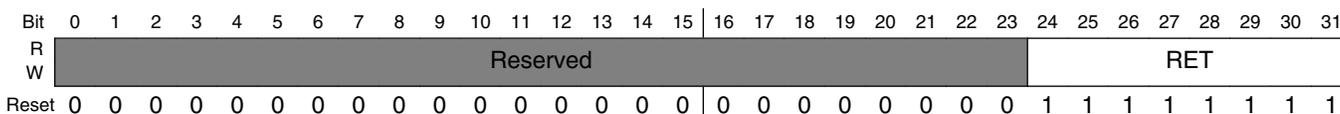
SRIO_EPWISR field descriptions (continued)

Field	Description
31 PW	An inbound port-write was received.

15.5.55 Logical retry error threshold configuration register (SRIO_LRETCR)

The LRETCR register contains the retry error threshold for the logical layer. When the number of consecutive logical retries for a given packet is greater than this value, an error interrupt is generated. Note that the number of retries must be greater than this value unlike other registers that define a retry threshold.

Address: C_0000h base + 1_0020h offset = D_0020h



SRIO_LRETCR field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24–31 RET	<p>Retry error threshold.</p> <p>These bits provide the threshold value for the number of consecutive logical retries (for GSM responses) received for a given packet that causes RAPIDIO ENDPOINT to report an error condition.</p> <p>0x00 Disable the RET 0x01 Set the error reporting threshold to 1 ... 0xFF Set the error reporting threshold to 255</p>

15.5.56 Physical retry error threshold configuration register (SRIO_PRETCR)

The PRETCR register contains the retry error threshold for the physical layer. When the number of consecutive ACK-retries is greater than or equal to this value, an error interrupt is generated.

Address: C_0000h base + 1_0080h offset = D_0080h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R																												RET					
W																												RET					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1

SRIO_PRETCR field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24–31 RET	Retry error threshold. These bits provide the threshold value for the number of consecutive ACK retries received that causes RAPIDIO ENDPOINT to report an error condition. 0x00 Disable the RET 0x01 Set the error reporting threshold to 1 ... 0xFF Set the error reporting threshold to 255

15.5.57 Port 1 Alternate device ID command and status register (SRIO_P1ADIDCSR)

The port *m* alternate device id CSR contains an alternate deviceID . This register should be used on a multi-port device to enable separate deviceIDs for each port. It is intended that this register should be enabled before the master enabled bit of the P GCCSR is set, such that when it is enabled, all other devices in the RapidIO system (including switches) send packets to and receive packets from the deviceID contained in this register, instead of the deviceID contained in BDIDCSR.

When the alternate deviceID is enabled, the inbound RapidIO endpoint only accepts packets sent with the deviceID contained in P *m* ADIDCSR or with the deviceID contained in BDIDCSR (except during Accept All mode, during which the inbound RapidIO endpoint accepts packets using the same common transport system). In addition,

Serial RapidIO Memory Map/Register Definition

the outbound RapidIO endpoint only generates requests using the deviceID contained in P *m* ADIDCSR; it generates responses with the deviceID given in the original request packet (either from P *m* ADIDCSR or BDIDCSR).

Address: C_0000h base + 1_0100h offset = D_0100h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R																	
W	ADE	Reserved						ADID									
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R																	
W	LADID																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

SRIO_P1ADIDCSR field descriptions

Field	Description
0 ADE	Alternate device ID enable When set, this bit causes the port to use the deviceID specified in this register instead of the deviceid specified in BDIDCSR
1–7 -	This field is reserved. Reserved
8–15 ADID	Alternate device ID for the device in a small transport system
16–31 LADID	Alternate device ID for the device in a large transport system

15.5.58 Port 1 accept-all configuration register (SRIO_P1AACR)

The port *m* accept-all configuration register contains information on accept-all mode.

Address: C_0000h base + 1_0120h offset = D_0120h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	Reserved															
Reset	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																AA
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1AACR field descriptions

Field	Description
0–30 -	This field is reserved. Reserved
31 AA	Accept all. 1 All packets are accepted without checking the target ID. However, the tt field must be consistent with the common transport system specified by bit 27 of the processing element features CAR (PEFCAR). 0 Normal RapidIO acceptance based on target ID.

15.5.59 Port 1 Logical Outbound Packet time-to-live configuration register (SRIO_P1LOPTTLCR)

The port *m* logical outbound packet time-to-live configuration register contains the time-to-live count for all ports on a device. This packet time-to-live counter starts when a packet is ready to be transmitted. If the packet is not successfully transmitted before the timer expires, the packet is discarded. Successfully transmitted means that a packet accept was received for the packet on the RIO interface. If the packet requires a response, an internal error response is returned after the response time-out occurs (PRTOCCSR). The packet time-to-live counter prevents the local processor from being stalled when packets cannot be successfully transmitted (acknowledged with an accept by the link partner at the physical level). The value of this register should always be larger than the link time-out value (PLTOCCSR). When the packet time-to-live counter expires, P *m* PCR[OB DEN] is automatically set. P *m* PCR[OB DEN] must be cleared by software. By default, this time-out value is disabled (all zeros).

The resolution of this timer is $168/(\text{platform frequency})$. For example, at a platform frequency of 400 MHz, the maximum timeout value is $0xFF_FFFF * 168/400\text{MHz} = 7.05$ seconds.

When the packet time-to-live counter expires, P *m* PCR[OB DEN] is automatically set. P *m* PCR[OB DEN] must be cleared by software.

Address: C_0000h base + 1_0124h offset = D_0124h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																TV												Reserved				
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1LOPTTLCR field descriptions

Field	Description
0–23 TV	Time-out value. Setting to all zeros disables the time-to-live time-out timer. This value is loaded each time the time-to-live time-out timer starts.
24–31 -	This field is reserved. Reserved

15.5.60 Port 1 Implementation error command and status register (SRIO_P1IECSR)

The *P m* IECSR register contains status bits that are asserted whenever an implementation-defined error occurs.

Address: C_0000h base + 1_0130h offset = D_0130h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R																	
W	RETE	Reserved															
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

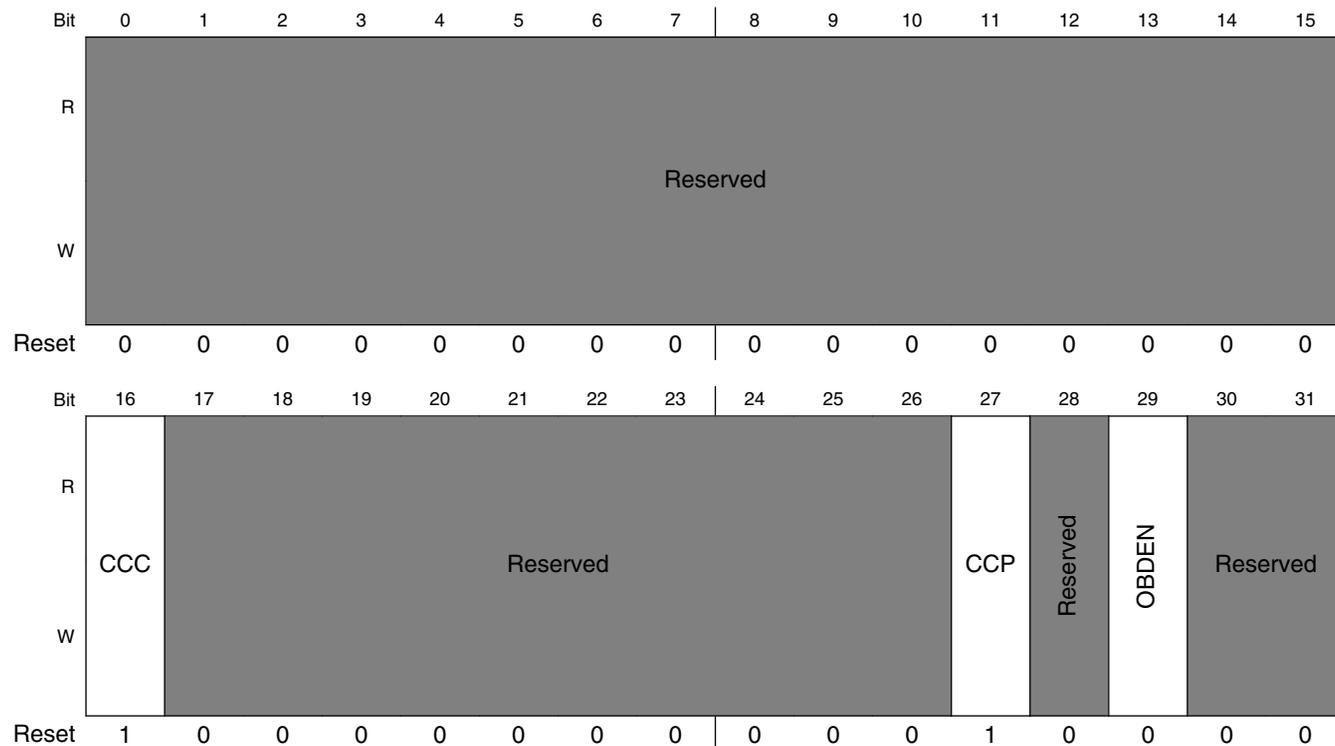
SRIO_P1IECSR field descriptions

Field	Description
0 RETE	<p>Retry error threshold exceeded.</p> <p>This bit is asserted when the number of consecutive retries has reached retry error threshold in the retry error threshold register. This bit is cleared by writing a 1 to it.</p> <p>This bit sets again if another retry is received and the number of consecutive retries continues to exceed the retry error threshold.</p>
1–31 -	This field is reserved. Reserved

15.5.61 Port 1 Physical configuration register (SRIO_P1PCR)

The P m PCR contains general physical layer protocol and link mode enables.

Address: C_0000h base + 1_0140h offset = D_0140h



SRIO_P1PCR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16 CCC	CRC checking enable-control symbol. When set, CRC is checked on received control symbols. When cleared, no CRC is checked on received control symbols.
17–26 -	This field is reserved. Reserved
27 CCP	CRC checking enable-packet. When set, CRC is checked on received packets. When cleared, no CRC is checked on received packets
28 -	This field is reserved. Reserved
29 OB DEN	Output buffer drain enable.

Table continues on the next page...

SRIO_P1PCR field descriptions (continued)

Field	Description
	When set, the output drains packets from the outbound buffer and does not send them out. This intentionally causes the inbound to time-out (when a response on a drained request was expected) and send an error response to OCN. A packet time-to-live time-out causes this bit to be set (See Inbound port-write n queue base address register). This bit is also used during software retraining and error recovery. See Outbound drain mode and Software assisted error recovery register support .
30–31 -	This field is reserved. Reserved

15.5.62 Port 1 Serial link command and status register (SRIO_P1SLCSR)

The *P m* SLCSR contains status of the of the serial physical link.

Address: C_0000h base + 1_0158h offset = D_0158h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	LS0	LS1	LS2	LS3	Reserved				LA	Reserved						
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1SLCSR field descriptions

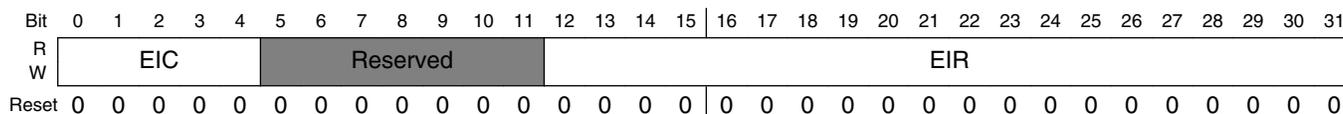
Field	Description
0 LS0	Lane sync achieved for lane 0. Write with 1 to clear
1 LS1	Lane sync achieved for lane 1. Write with 1 to clear
2 LS2	Lane sync achieved for lane 2. Write with 1 to clear
3 LS3	Lane sync achieved for lane 3. Write with 1 to clear.
4–7 -	This field is reserved. Reserved
8 LA	Lane alignment achieved. This bit is set to 1 only if all four lanes achieve sync. LA will be 0 if a x1 or x2 connection is achieved. Write with 1 to clear.
9–31 -	This field is reserved. Reserved

15.5.63 Port 1 Serial link error injection configuration register (SRIO_P1SLEICR)

The *PmSLEICR* is used to control the injection of bit errors into the transmit bit stream .

The *PmSLEICR* register is used to generate pseudo-random errors into the outbound serial RapidIO data stream. If the EIC field is any of the allowable non-zero values (as shown in the table above), then error injection is enabled for one lane or all four lanes, as selected by the EIC value. When enabled, at pseudo-random intervals, an error is injected by inverting a single bit in the outgoing data stream. This occurs only in the lane(s) that have error injection enabled. The range of the pseudo-random value (delay between injected errors) is controlled by the EIR field, as described above. That is, the value of EIR, multiplied by 32, determines the maximum number of character times between injected errors.

Address: C_0000h base + 1_0160h offset = D_0160h



SRIO_P1SLEICR field descriptions

Field	Description
0–4 EIC	Error injection control. Enables and controls serial link error injection as follows: All other values are reserved. 00000 Error injection is disabled. 10000 Error injection, lane 0 only 01000 Error injection, lane 1 only 00100 Error injection, lane 2 only 00010 Error injection, lane 3 only 11110 Error injection, all 4 lanes simultaneously 11111 Error injection, randomly distributed over all 4 lanes
5–11 -	This field is reserved. Reserved
12–31 EIR	Error injection range. The value of EIR x 32 determines the maximum value of the pseudo-random delay between errors. For example, a value of 0x1 would indicate a maximum delay of 32 character times. Value within this register should be right-justified.

15.5.64 Port 2 Alternate device ID command and status register (SRIO_P2ADIDCSR)

The port *m* alternate device id CSR contains an alternate deviceID . This register should be used on a multi-port device to enable separate deviceIDs for each port. It is intended that this register should be enabled before the master enabled bit of the P GCCSR is set, such that when it is enabled, all other devices in the RapidIO system (including switches) send packets to and receive packets from the deviceID contained in this register, instead of the deviceID contained in BDIDCSR.

When the alternate deviceID is enabled, the inbound RapidIO endpoint only accepts packets sent with the deviceID contained in P *m* ADIDCSR or with the deviceID contained in BDIDCSR (except during Accept All mode, during which the inbound RapidIO endpoint accepts packets using the same common transport system). In addition, the outbound RapidIO endpoint only generates requests using the deviceID contained in P *m* ADIDCSR; it generates responses with the deviceID given in the original request packet (either from P *m* ADIDCSR or BDIDCSR).

Address: C_0000h base + 1_0180h offset = D_0180h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R																	
W	ADE	Reserved							ADID								
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R																	
W	LADID																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

SRIO_P2ADIDCSR field descriptions

Field	Description
0 ADE	Alternate device ID enable When set, this bit causes the port to use the deviceID specified in this register instead of the deviceid specified in BDIDCSR
1-7 -	This field is reserved. Reserved
8-15 ADID	Alternate device ID for the device in a small transport system
16-31 LADID	Alternate device ID for the device in a large transport system

15.5.65 Port 2 accept-all configuration register (SRIO_P2AACR)

The port *m* accept-all configuration register contains information on accept-all mode.

Address: C_0000h base + 1_01A0h offset = D_01A0h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved																	
W	Reserved																	
Reset	0	0	0	1	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved																AA	
W	Reserved																AA	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	

SRIO_P2AACR field descriptions

Field	Description
0–30 -	This field is reserved. Reserved
31 AA	Accept all. 1 All packets are accepted without checking the target ID. However, the tt field must be consistent with the common transport system specified by bit 27 of the processing element features CAR (PEFCAR). 0 Normal RapidIO acceptance based on target ID.

15.5.66 Port 2 Logical Outbound Packet time-to-live configuration register (SRIO_P2LOPTTLCR)

The port *m* logical outbound packet time-to-live configuration register contains the time-to-live count for all ports on a device. This packet time-to-live counter starts when a packet is ready to be transmitted. If the packet is not successfully transmitted before the timer expires, the packet is discarded. Successfully transmitted means that a packet accept was received for the packet on the RIO interface. If the packet requires a response, an internal error response is returned after the response time-out occurs (PRTOCCSR). The packet time-to-live counter prevents the local processor from being stalled when packets cannot be successfully transmitted (acknowledged with an accept by the link partner at the physical level). The value of this register should always be larger than the link time-out value (PLTOCCSR). When the packet time-to-live counter expires, P *m* PCR[OB DEN] is automatically set. P *m* PCR[OB DEN] must be cleared by software. By default, this time-out value is disabled (all zeros).

The resolution of this timer is $168/(\text{platform frequency})$. For example, at a platform frequency of 400 MHz, the maximum timeout value is $0xFF_FFFF * 168/400\text{MHz} = 7.05$ seconds.

When the packet time-to-live counter expires, P_m PCR[OB DEN] is automatically set. P_m PCR[OB DEN] must be cleared by software.

Address: C_0000h base + 1_01A4h offset = D_01A4h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																												Reserved				
W																												Reserved				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2LOPTTLCR field descriptions

Field	Description
0–23 TV	Time-out value. Setting to all zeros disables the time-to-live time-out timer. This value is loaded each time the time-to-live time-out timer starts.
24–31 -	This field is reserved. Reserved

15.5.67 Port 2 Implementation error command and status register (SRIO_P2IECSR)

The P_m IECSR register contains status bits that are asserted whenever an implementation-defined error occurs.

Address: C_0000h base + 1_01B0h offset = D_01B0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	RETE	Reserved														
W	RETE	Reserved														
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2IECSR field descriptions

Field	Description
0 RETE	<p>Retry error threshold exceeded.</p> <p>This bit is asserted when the number of consecutive retries has reached retry error threshold in the retry error threshold register. This bit is cleared by writing a 1 to it.</p> <p>This bit sets again if another retry is received and the number of consecutive retries continues to exceed the retry error threshold.</p>

Table continues on the next page...

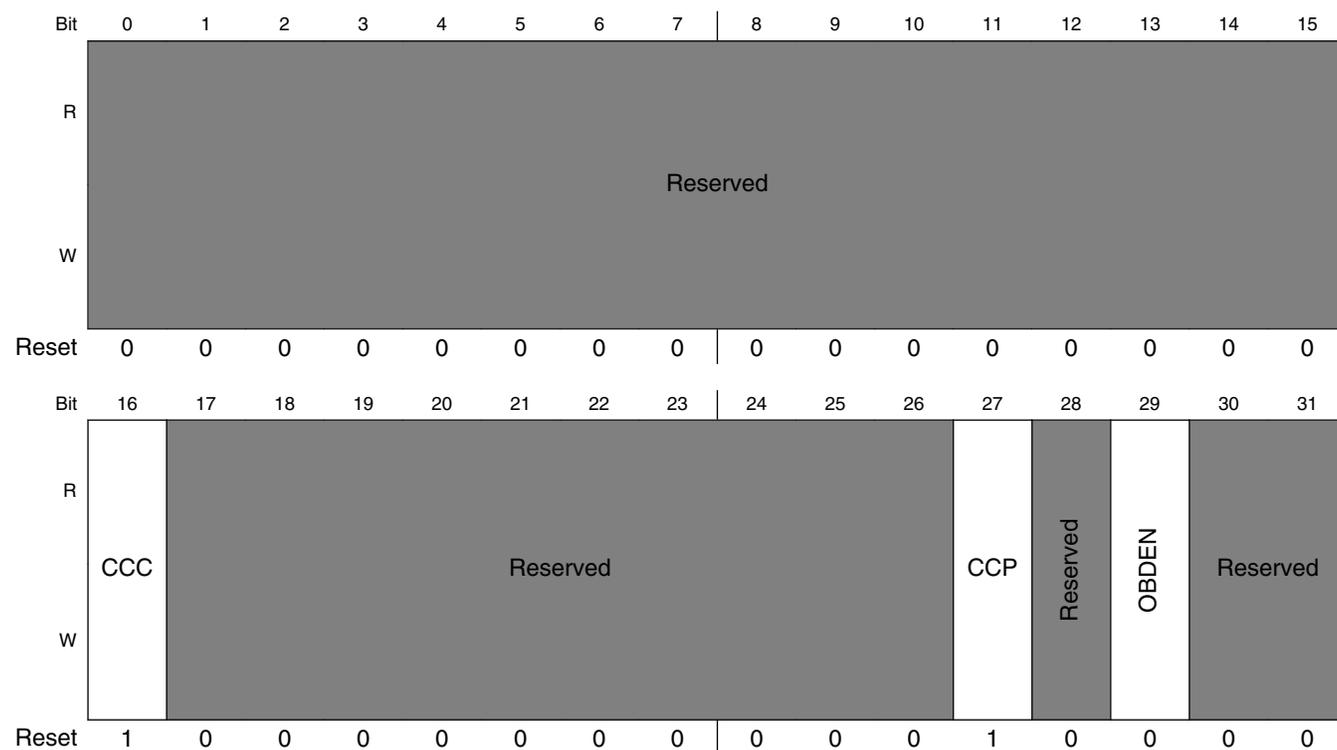
SRIO_P2IECSR field descriptions (continued)

Field	Description
1–31 -	This field is reserved. Reserved

15.5.68 Port 2 Physical configuration register (SRIO_P2PCR)

The P m PCR contains general physical layer protocol and link mode enables.

Address: C_0000h base + 1_01C0h offset = D_01C0h


SRIO_P2PCR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16 CCC	CRC checking enable-control symbol. When set, CRC is checked on received control symbols. When cleared, no CRC is checked on received control symbols.
17–26 -	This field is reserved. Reserved

Table continues on the next page...

SRIO_P2PCR field descriptions (continued)

Field	Description
27 CCP	CRC checking enable-packet. When set, CRC is checked on received packets. When cleared, no CRC is checked on received packets
28 -	This field is reserved. Reserved
29 OBDEN	Output buffer drain enable. When set, the output drains packets from the outbound buffer and does not send them out. This intentionally causes the inbound to time-out (when a response on a drained request was expected) and send an error response to OCN. A packet time-to-live time-out causes this bit to be set. (See Inbound port-write n queue base address register ."). This bit is also used during software retraining and error recovery. See Outbound drain mode and Software assisted error recovery register support .
30-31 -	This field is reserved. Reserved

15.5.69 Port 2 Serial link command and status register (SRIO_P2SLCSR)

The P m SLCSR contains status of the of the serial physical link.

Address: C_0000h base + 1_01D8h offset = D_01D8h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	LS0	LS1	LS2	LS3	Reserved				LA	Reserved						
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2SLCSR field descriptions

Field	Description
0 LS0	Lane sync achieved for lane 0. Write with 1 to clear
1 LS1	Lane sync achieved for lane 1. Write with 1 to clear
2 LS2	Lane sync achieved for lane 2. Write with 1 to clear
3 LS3	Lane sync achieved for lane 3. Write with 1 to clear.

Table continues on the next page...

SRIO_P2SLCSR field descriptions (continued)

Field	Description
4–7 -	This field is reserved. Reserved
8 LA	Lane alignment achieved. This bit is set to 1 only if all four lanes achieve sync. LA will be 0 if a x1 or x2 connection is achieved. Write with 1 to clear.
9–31 -	This field is reserved. Reserved

15.5.70 Port 2 Serial link error injection configuration register (SRIO_P2SLEICR)

The P m SLEICR is used to control the injection of bit errors into the transmit bit stream .

The P m SLEICR register is used to generate pseudo-random errors into the outbound serial RapidIO data stream. If the EIC field is any of the allowable non-zero values (as shown in the table above), then error injection is enabled for one lane or all four lanes, as selected by the EIC value. When enabled, at pseudo-random intervals, an error is injected by inverting a single bit in the outgoing data stream. This occurs only in the lane(s) that have error injection enabled. The range of the pseudo-random value (delay between injected errors) is controlled by the EIR field, as described above. That is, the value of EIR, multiplied by 32, determines the maximum number of character times between injected errors.

Address: C_0000h base + 1_01E0h offset = D_01E0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2SLEICR field descriptions

Field	Description
0–4 EIC	Error injection control. Enables and controls serial link error injection as follows: All other values are reserved. 00000 Error injection is disabled. 10000 Error injection, lane 0 only 01000 Error injection, lane 1 only 00100 Error injection, lane 2 only 00010 Error injection, lane 3 only

Table continues on the next page...

SRIO_P2SLEICR field descriptions (continued)

Field	Description
	11110 Error injection, all 4 lanes simultaneously 11111 Error injection, randomly distributed over all 4 lanes
5–11 -	This field is reserved. Reserved
12–31 EIR	Error injection range. The value of EIR x 32 determines the maximum value of the pseudo-random delay between errors. For example, a value of 0x1 would indicate a maximum delay of 32 character times. Value within this register should be right-justified.

15.5.71 IP Block Revision Register 1 (SRIO_IPBRR1)

IP block revision register 1 is used to track changes and revisions of the RapidIO endpoint. It is a read-only register.

Address: C_0000h base + 1_0BF8h offset = D_0BF8h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	IPID															IPMJ								IPMN									
W	[Shaded]																																
Reset	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0

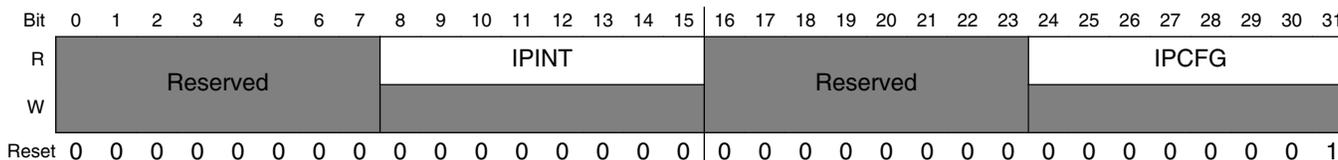
SRIO_IPBRR1 field descriptions

Field	Description
0–15 IPID	IP block ID = 0x01C0
16–23 IPMJ	Major revision of the IP block = 0x01
24–31 IPMN	Minor revision of the IP block = 0x 2

15.5.72 IP Block Revision Register 2 (SRIO_IPBRR2)

IP block revision register 2 is used to track changes and revisions of the RapidIO endpoint. It is a read-only register.

Address: C_0000h base + 1_0BFCh offset = D_0BFCh



SRIO_IPBRR2 field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–15 IPINT	IP block Integration options = 0x0
16–23 -	This field is reserved. Reserved
24–31 IPCFG	IP block Configuration options = 0x 1

15.5.73 Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR0)

ATMU registers are used for outbound and inbound transactions. Their purpose is to translate RapidIO packets to OCN packets on inbound and to translate OCN packets to RapidIO packets on outbound. ATMU window misses use the window 0 register set by default, and overlapping window hits results in the use of the lowest number window register set hit. For both inbound and outbound translation, the smallest window size is 4K and the largest window size is 16G for inbound translation and 64G for outbound translation. The default window register set causes no translation of the transaction address for inbound transactions since the RapidIO address space has 34 bits and the OCN address space has 36 bits. For outbound transactions, the default window maps each of the four 16G chunks to the RapidIO 16G address space. The inbound and outbound translation windows must be aligned based on the granularity selected by the size fields. The packet device ID fields are not used in the inbound translation process, only the address field.

As is the case with all registers, an external processor writing the ATMU registers should not assume that the write has completed until a response is received.

Note that when booting from serial RapidIO, outbound ATMU window 0 must be used.

See [Segmented Outbound Window Description](#) .

The port *m* RapidIO outbound window translation address registers select the starting addresses in the external address space for window hits within the outbound translation windows. The new translated address is created by concatenating the transaction offset to this translation address.

Address: C_0000h base + 1_0C00h offset = D_0C00h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W																
Reset	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0
Field	LTGTID		TREXAD						TRAD							
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Field	TRAD															

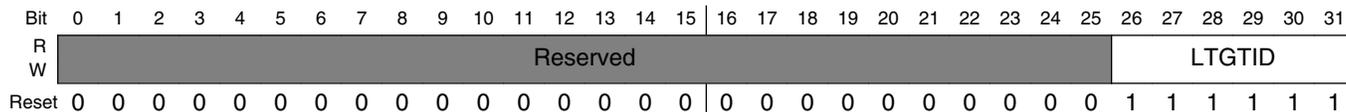
SRIO_P1ROWTAR0 field descriptions

Field	Description
0-1 LTGTID	LTGTID correspond to bits 6-7 of the target ID for a large transport system. This field is valid only if bit 27 of PEFCAR is set. Bits 0-5 of the target ID are specified in the window's port <i>m</i> RapidIO outbound window translation extended address register.
2-11 TREXAD	Translation extended address. TREXAD[0-7] correspond to the target ID for a small transport system or the least significant byte (bits 8-15) of the target ID for a large transport system. TREXAD[8-9] corresponds to bits [0-1] of a 34-bit RapidIO translation address. For maintenance transactions and default window 0, TREXAD[8-9] is reserved.
12-31 TRAD	Translation address. System address which represents the starting point of the outbound translated address. The translation address must be aligned based on the size field. This corresponds to bits [2-21] of the 34-bit RapidIO translation address. For maintenance transactions, the hop count is formed from TRAD[0-7], and the upper 12 bits of the maintenance offset is formed from TRAD[8-19]; the rest of the maintenance offset is formed from the untranslated address. This field is reserved for default window 0.

15.5.74 Port 1 RapidIO outbound window translation extended address register 0 (SRIO_P1ROWTEAR_n)

The port *m* RapidIO outbound window translation extended address registers contain bits 0-5 of the target ID for a common transport large system.

Address: C_0000h base + 1_0C04h offset + (32d × i), where i=0d to 8d



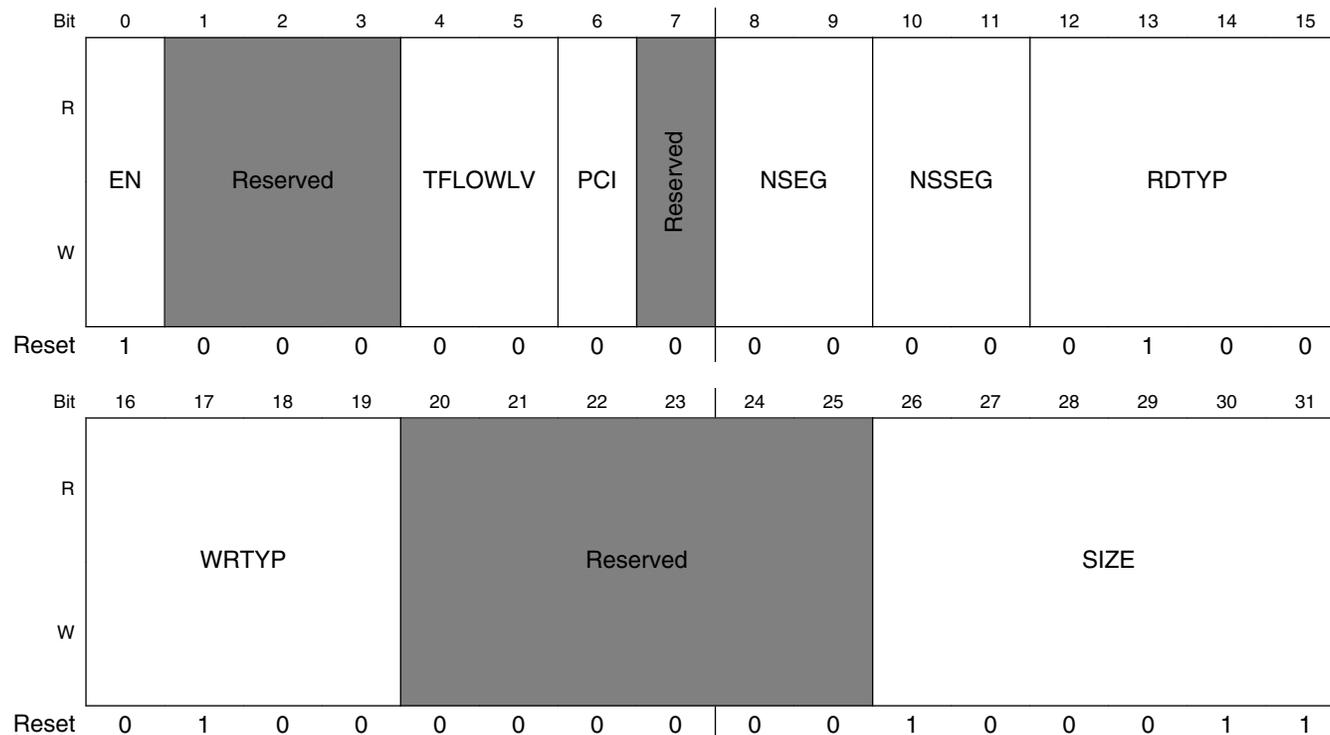
SRIO_P1ROWTEAR_n field descriptions

Field	Description
0–25 -	This field is reserved. Reserved
26–31 LTGTID	LTGTID correspond to bits 0-5 of the target ID for a large transport system. This field is valid only if bit 27 of PEFCAR is set. Bits 6-7 of the target ID are specified in the window's port <i>m</i> RapidIO outbound window translation address register.

15.5.75 Port 1 RapidIO outbound window attributes register n (SRIO_P1ROWARn)

The port *m* RapidIO outbound window attributes registers define the window sizes to translate and other attributes for the translations. 64G is the largest window size allowed. For a segmented window, these attributes are used for segment 0. The PCI_Window bit applies for all segments.

Address: C_0000h base + 1_0C10h offset + (32d × i), where i=0d to 8d



SRIO_P1ROWARn field descriptions

Field	Description
0 EN	This field enables this address translation window. It is set to 1 and is read only for default window 0.
1-3 -	This field is reserved. Reserved
4-5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI bit is set. Also, the RapidIO priority given by this field must always be greater than or equal to the OCN priority or a deadlock can occur. Normally, the OCN priority of all packets is 0 except for packets coming from a PCI type interface. Read type packets coming from a PCI type interface are priority 0 and write type packets are priority 1. Packets coming from a PCI type interface should set this field to 0 and set the PCI field to 1.

Table continues on the next page...

SRIO_P1ROWAR n field descriptions (continued)

Field	Description
	00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6 PCI	PCI_Window. This window follows PCI ordering rules as defined in the RapidIO Inter-operability specification The TFLOWLV field must be set to 00 if this bit is set causing reads to have RapidIO priority 0 and writes to have RapidIO priority 1.
7 -	This field is reserved. Reserved
8–9 NSEG	Number of segments for this window. This field is reserved for default window 0. 00 One segment (normal window) 01 Two segments (half size aliasing window) 10 Four segments (quarter size aliasing window) 11 Reserved
10–11 NSSEG	Number of subsegments for this segment. This field is reserved for default window 0. Note that this field is valid only when ROWAR n [NSEG] contains a non-zero value (1 or 2). 00 One target deviceID for this segment 01 Two target deviceIDs for this segment 10 Four target deviceIDs for this segment 11 Eight target deviceIDs for this segment
12–15 RDTYP	Transaction type to run on RapidIO interface if access is a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved 0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
16–19 WRTYP	Transaction type to run on RapidIO interface if access was a write. Write-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead.

Table continues on the next page...

SRIO_P1ROWAR n field descriptions (continued)

Field	Description
	0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1111 Reserved
20–25 -	This field is reserved. Reserved
26–31 SIZE	Outbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window size is 4 Kbytes. This field is read only for default window 0. 000000 Reserved ... 001011 4-Kbyte window size 001100 8-Kbyte window size ... 011111 4-Gbyte window size 100000 8-Gbyte window size 100001 16-Gbyte window size 100010 32-Gbyte window size 100011 64-Gbyte window size 100100 Reserved ... 111111 Reserved

15.5.76 Port 1 RapidIO outbound window translation address register n (SRIO_P1ROWTAR_n)

ATMU registers are used for outbound and inbound transactions. Their purpose is to translate RapidIO packets to OCN packets on inbound and to translate OCN packets to RapidIO packets on outbound. ATMU window misses use the window 0 register set by default, and overlapping window hits results in the use of the lowest number window register set hit. For both inbound and outbound translation, the smallest window size is 4K and the largest window size is 16G for inbound translation and 64G for outbound translation. The default window register set causes no translation of the transaction address for inbound transactions since the RapidIO address space has 34 bits and the OCN address space has 36 bits. For outbound transactions, the default window maps each of the four 16G chunks to the RapidIO 16G address space. The inbound and outbound translation windows must be aligned based on the granularity selected by the size fields. The packet device ID fields are not used in the inbound translation process, only the address field.

As is the case with all registers, an external processor writing the ATMU registers should not assume that the write has completed until a response is received.

Note that when booting from serial RapidIO, outbound ATMU window 0 must be used.

See [Segmented Outbound Window Description](#) .

The port *m* RapidIO outbound window translation address registers select the starting addresses in the external address space for window hits within the outbound translation windows. The new translated address is created by concatenating the transaction offset to this translation address.

Address: C_0000h base + 1_0C20h offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	LTGTID		TRESAD										TRAD			
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	TRAD															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1ROWTAR_n field descriptions

Field	Description
0-1 LTGTID	LTGTID correspond to bits 6-7 of the target ID for a large transport system. This field is valid only if bit 27 of PEFCAR is set.

Table continues on the next page...

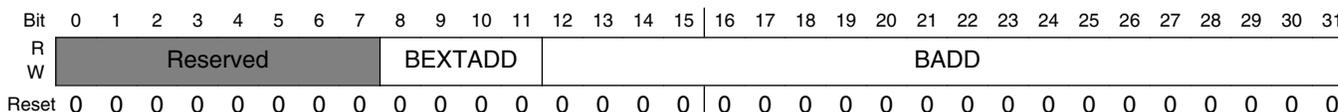
SRIO_P1ROWTAR_n field descriptions (continued)

Field	Description
	Bits 0-5 of the target ID are specified in the window's port <i>m</i> RapidIO outbound window translation extended address register.
2–11 TREXAD	Translation extended address. TREXAD[0-7] correspond to the target ID for a small transport system or the least significant byte (bits 8-15) of the target ID for a large transport system. TREXAD[8-9] corresponds to bits [0-1] of a 34-bit RapidIO translation address. For maintenance transactions and default window 0, TREXAD[8-9] is reserved.
12–31 TRAD	Translation address. System address which represents the starting point of the outbound translated address. The translation address must be aligned based on the size field. This corresponds to bits [2-21] of the 34-bit RapidIO translation address. For maintenance transactions, the hop count is formed from TRAD[0-7], and the upper 12 bits of the maintenance offset is formed from TRAD[8-19]; the rest of the maintenance offset is formed from the untranslated address. This field is reserved for default window 0.

15.5.77 Port 1 RapidIO outbound window base address register *n* (SRIO_P1ROWBAR_n)

The port *m* RapidIO outbound window base address registers select the base address for the windows which are translated to an alternate system address space. Addresses for outbound transactions are compared to these windows. If such a transaction does not fall within one of these spaces the transaction is forwarded through the default register set. For transactions that cross more than one window, see [Outbound window boundary crossing errors](#) ."

Address: C_0000h base + 1_0C28h offset + (32d × *i*), where *i*=0d to 7d



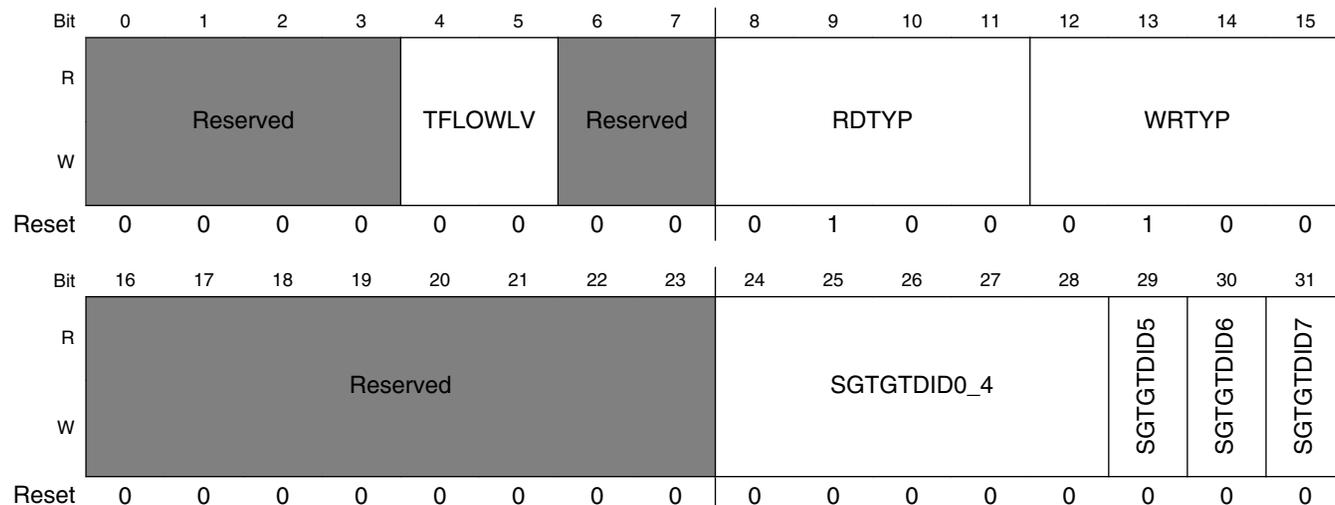
SRIO_P1ROWBAR_n field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–11 BEXTADD	Window base extended address. Corresponds to bits [0-3] of the 36-bit OCN base address.
12–31 BADD	Window base address. Source address which is the starting point for the outbound translation window. The window must be aligned based on the size selected in the window size bits. This corresponds to bits [4-23] of the 36-bit OCN base address.

15.5.78 Port 1 RapidIO outbound window segment 1 register n (SRIO_P1ROWS1Rn)

The port *m* RapidIO outbound window segment registers define the attributes and target device ID that are to be used for a transaction that hits in that segment rather than the primary attributes and target deviceID. There is one of these registers for each segment (except segment 0).

Address: C_0000h base + 1_0C34h offset + (32d × i), where i=0d to 7d



SRIO_P1ROWS1Rn field descriptions

Field	Description
0-3 -	This field is reserved. Reserved
4-5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI_Window bit is set 00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6-7 -	This field is reserved. Reserved
8-11 RDTYP	Transaction type to run on RapidIO interface if access was a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved

Table continues on the next page...

SRIO_P1ROWS1Rn field descriptions (continued)

Field	Description
	0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
12–15 WRTYP	Transaction type to run on RapidIO interface if access was a write. Writes-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead. 0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1110 Reserved 1111 Reserved
16–23 -	This field is reserved. Reserved
24–28 SGTGTID0_4	Bits 0-4 (or bits 8-12 if large transport system) of the target device ID for this segment.
29 SGTGTID5	Bit 5 (or bit 13 if large transport system) of the target deviceID; this bit is reserved if 8 target subsegments are selected.
30 SGTGTID6	Bit 6 (or bit 14 if large transport system) of the target deviceID; this bit is reserved if 8 or 4 target subsegments are selected.
31 SGTGTID7	Bit 7 (or bit 15 if large transport system) of the target deviceID; this bit is reserved if 8, 4, or 2 target subsegments are selected.

15.5.79 Port 1 RapidIO outbound window segment 2 register n (SRIO_P1ROWS2Rn)

The port *m* RapidIO outbound window segment registers define the attributes and target device ID that are to be used for a transaction that hits in that segment rather than the primary attributes and target deviceID. There is one of these registers for each segment (except segment 0).

Address: C_0000h base + 1_0C38h offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP				
W	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP				
Reset	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								SGTGTID0_4				SGTGTID5	SGTGTID6	SGTGTID7	
W	Reserved								SGTGTID0_4				SGTGTID5	SGTGTID6	SGTGTID7	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1ROWS2Rn field descriptions

Field	Description
0–3 -	This field is reserved. Reserved
4–5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI_Window bit is set 00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6–7 -	This field is reserved. Reserved
8–11 RDTYP	Transaction type to run on RapidIO interface if access was a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved

Table continues on the next page...

SRIO_P1ROWS2Rn field descriptions (continued)

Field	Description
	0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
12–15 WRTYP	Transaction type to run on RapidIO interface if access was a write. Writes-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead. 0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1110 Reserved 1111 Reserved
16–23 -	This field is reserved. Reserved
24–28 SGTGTDID0_4	Bits 0-4 (or bits 8-12 if large transport system) of the target device ID for this segment.
29 SGTGTDID5	Bit 5 (or bit 13 if large transport system) of the target deviceID; this bit is reserved if 8 target subsegments are selected.
30 SGTGTDID6	Bit 6 (or bit 14 if large transport system) of the target deviceID; this bit is reserved if 8 or 4 target subsegments are selected.
31 SGTGTDID7	Bit 7 (or bit 15 if large transport system) of the target deviceID; this bit is reserved if 8, 4, or 2 target subsegments are selected.

15.5.80 Port 1 RapidIO outbound window segment 3 register n (SRIO_P1ROWS3Rn)

The port *m* RapidIO outbound window segment registers define the attributes and target device ID that are to be used for a transaction that hits in that segment rather than the primary attributes and target deviceID. There is one of these registers for each segment (except segment 0).

Address: C_0000h base + 1_0C3Ch offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15			
R	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP							
W	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP							
Reset	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	0			
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
R	Reserved								SGTGTID0_4								SGTGTID5	SGTGTID6	SGTGTID7
W	Reserved								SGTGTID0_4								SGTGTID5	SGTGTID6	SGTGTID7
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			

SRIO_P1ROWS3Rn field descriptions

Field	Description
0–3 -	This field is reserved. Reserved
4–5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI_Window bit is set 00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6–7 -	This field is reserved. Reserved
8–11 RDTYP	Transaction type to run on RapidIO interface if access was a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved

Table continues on the next page...

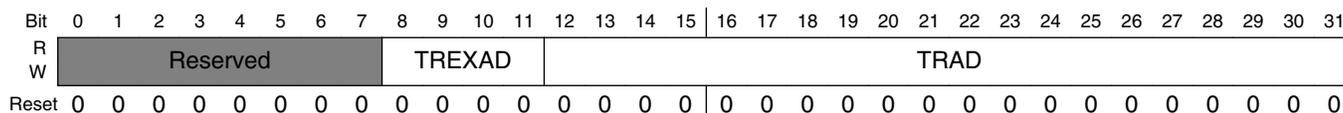
SRIO_P1ROWS3Rn field descriptions (continued)

Field	Description
	0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
12–15 WRTYP	Transaction type to run on RapidIO interface if access was a write. Writes-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead. 0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1110 Reserved 1111 Reserved
16–23 -	This field is reserved. Reserved
24–28 SGTGTDID0_4	Bits 0-4 (or bits 8-12 if large transport system) of the target device ID for this segment.
29 SGTGTDID5	Bit 5 (or bit 13 if large transport system) of the target deviceID; this bit is reserved if 8 target subsegments are selected.
30 SGTGTDID6	Bit 6 (or bit 14 if large transport system) of the target deviceID; this bit is reserved if 8 or 4 target subsegments are selected.
31 SGTGTDID7	Bit 7 (or bit 15 if large transport system) of the target deviceID; this bit is reserved if 8, 4, or 2 target subsegments are selected.

15.5.81 Port 1 RapidIO Inbound window translation address register n (SRIO_P1RIWTAR_n)

The port *m* RapidIO inbound window translation address registers point to the starting addresses in local address space for window hits within the inbound translation windows. The new translated address is created by concatenating the transaction offset to this translation address.

Address: C_0000h base + 1_0D60h offset + (32d × i), where i=0d to 4d



SRIO_P1RIWTAR_n field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–11 TREXAD	Translation extended address. Corresponds to bits 0-3 of the 36-bit OCN translation address. TREXAD[2-3] are reserved for default window 0.
12–31 TRAD	Translation address. Target address which indicates the starting point of the inbound translated address. The translation address must be aligned based on the size field. This corresponds to bits 4-23 of the 36-bit OCN translation address. TRAD is reserved for default window 0.

15.5.82 Port 1 RapidIO Inbound window base address register n (SRIO_P1RIWBARn)

The port *m* RapidIO inbound window *n* base address registers select the base address for the windows which are translated to an alternate target address space. Addresses for inbound transactions are compared to these windows. If such a transaction does not fall within one of these spaces, then the transaction is forwarded to the interior of the chip using the default window. For transactions that cross more than one window, see [Outbound window boundary crossing errors](#) . Note that the LCSBA1CSR register (See [Inbound port-write n queue base address register](#)) has priority over all ATMU windows if both are configured for the same address space.

Address: C_0000h base + 1_0D68h offset + (32d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved									BEXAD		BADD				
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	BADD															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P1RIWBARn field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–11 BEXAD	Base extended address. BEXAD represents bits 0-1 of the 34-bit RapidIO address.
12–31 BADD	Base address. System address which is the starting point for the inbound translation window. The window must be aligned based on the size selected in the window size bits. This corresponds to bits 2-21 of the 34-bit RapidIO base address.

15.5.83 Port 1 RapidIO inbound window attributes register n (SRIO_P1RIWAR_n)

The port *m* RapidIO inbound window attributes registers define the window sizes to translate and other attributes for the translations. 16 Gbytes is the largest window size allowed. The RDTYP and WRTYP fields are used for attributes on the request, but do not actually change the transaction type of the transaction. In other words, P *m* RIWAR *n* does not modify whether or not the request requires a response or if the request was atomic; this type of attribute remains unchanged on the request as it is translated through the ATMU.

Address: C_0000h base + 1_0D70h offset + (32d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	EN	PW	Reserved					TGINT				RDTYP				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W	WRTYP				Reserved						SIZE					
Reset	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	1

SRIO_P1RIWAR_n field descriptions

Field	Description
0 EN	This bit enables this address translation. This field is set to 1 and read-only for default window 0.
1 PW	Protected Window. This bit indicates that this window is protected. Writes requiring response and reads to this window generate an error response. Writes not requiring response are silently discarded.
2–7 -	This field is reserved. Reserved
8–11 TGINT	Target interface. If this field is set to anything other than local address space, the attributes for the transaction must be assigned in a corresponding outbound window at the target. This is the field definition for the P2020: 0000 PCI Express interface 3 0001 PCI Express interface 2 0010 PCI Express Interface 1 0011-1011 Reserved 1100 Serial RapidIO interface 1 1101 Serial RapidIO interface 2 1110 Reserved 1111 Local memory (DDRs, local bus controller, memory-mapped registers)
12–15 RDTYP	Transaction type to run on the I/O interface if access is a read. 0000 Reserved

Table continues on the next page...

SRIO_P1RIWAR_n field descriptions (continued)

Field	Description
	<p>...</p> <p>0100 Read</p> <p>0101 Reserved</p> <p>...</p> <p>1111 Reserved</p> <p>Transaction type to run on local memory if access is a read.</p> <p>0000 Reserved</p> <p>...</p> <p>0100 Read, do not snoop local processor</p> <p>0101 Read, snoop local processor</p> <p>0110 Reserved</p> <p>0111 Read, unlock L2 cache line</p> <p>1000 Reserved</p> <p>...</p> <p>1111 Reserved</p>
16–19 WRTYP	<p>Transaction type to run on I/O interface if access is a write.</p> <p>0000 Reserved</p> <p>...</p> <p>0100 write</p> <p>0101 Reserved</p> <p>...</p> <p>1111 Reserved</p> <p>Transaction type to run on local memory if access is a write.</p> <p>0000 Reserved</p> <p>...</p> <p>0011 Reserved</p> <p>0100 Write, do not snoop local processor</p> <p>0101 Write, snoop local processor</p> <p>0110 Write, allocate cache line</p> <p>0111 Write, allocate and lock cache line</p> <p>1000 Reserved</p> <p>...</p> <p>1111 Reserved</p>
20–25 -	<p>This field is reserved. Reserved</p>
26–31 SIZE	<p>Inbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window size is 4 Kbytes.</p> <p>This field is read only for default window 0.</p>

Table continues on the next page...

SRIO_P1RIWAR n field descriptions (continued)

Field	Description
000000	Reserved
...	
001011	4-Kbyte window size
001100	8-Kbyte window size
...	
011111	4-Gbyte window size
100000	8-Gbyte window size
100001	16-Gbyte window size
100010	Reserved
...	
111111	Reserved

15.5.84 Port 1 RapidIO inbound window attributes register 0 (SRIO_P1RIWAR0)

The port m RapidIO inbound window attributes registers define the window sizes to translate and other attributes for the translations. 16 Gbytes is the largest window size allowed. The RDTYP and WRTYP fields are used for attributes on the request, but do not actually change the transaction type of the transaction. In other words, P m RIWAR n does not modify whether or not the request requires a response or if the request was atomic; this type of attribute remains unchanged on the request as it is translated through the ATMU.

Address: C_0000h base + 1_0DF0h offset = D_0DF0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	EN		Reserved						TGINT				RDTYP			
W		PW	Reserved						TGINT				RDTYP			
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	WRTYP				Reserved								SIZE			
W	WRTYP				Reserved								SIZE			
Reset	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	1

SRIO_P1RIWAR0 field descriptions

Field	Description
0 EN	This bit enables this address translation.

Table continues on the next page...

SRIO_P1RIWAR0 field descriptions (continued)

Field	Description
	This field is set to 1 and read-only for default window 0.
1 PW	Protected Window. This bit indicates that this window is protected. Writes requiring response and reads to this window generate an error response. Writes not requiring response are silently discarded.
2-7 -	This field is reserved. Reserved
8-11 TGINT	Target interface. If this field is set to anything other than local address space, the attributes for the transaction must be assigned in a corresponding outbound window at the target. This is the field definition for the P2020: 0000 PCI Express interface 3 0001 PCI Express interface 2 0010 PCI Express Interface 1 0011-1011 Reserved 1100 Serial RapidIO interface 1 1101 Serial RapidIO interface 2 1110 Reserved 1111 Local memory (DDRs, local bus controller, memory-mapped registers)
12-15 RDTYP	Transaction type to run on the I/O interface if access is a read. 0000 Reserved ... 0100 Read 0101 Reserved ... 1111 Reserved Transaction type to run on local memory if access is a read. 0000 Reserved ... 0100 Read, do not snoop local processor 0101 Read, snoop local processor 0110 Reserved 0111 Read, unlock L2 cache line 1000 Reserved ... 1111 Reserved
16-19 WRTYP	Transaction type to run on I/O interface if access is a write. 0000 Reserved ... 0100 write 0101 Reserved ... 1111 Reserved

Table continues on the next page...

SRIO_P1RIWAR0 field descriptions (continued)

Field	Description
	Transaction type to run on local memory if access is a write. 0000 Reserved ... 0011 Reserved 0100 Write, do not snoop local processor 0101 Write, snoop local processor 0110 Write, allocate cache line 0111 Write, allocate and lock cache line 1000 Reserved ... 1111 Reserved
20–25 -	This field is reserved. Reserved
26–31 SIZE	Inbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window size is 4 Kbytes. This field is read only for default window 0. 000000 Reserved ... 001011 4-Kbyte window size 001100 8-Kbyte window size ... 011111 4-Gbyte window size 100000 8-Gbyte window size 100001 16-Gbyte window size 100010 Reserved ... 111111 Reserved

15.5.85 Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTAR0)

ATMU registers are used for outbound and inbound transactions. Their purpose is to translate RapidIO packets to OCN packets on inbound and to translate OCN packets to RapidIO packets on outbound. ATMU window misses use the window 0 register set by default, and overlapping window hits results in the use of the lowest number window register set hit. For both inbound and outbound translation, the smallest window size is 4K and the largest window size is 16G for inbound translation and 64G for outbound translation. The default window register set causes no translation of the transaction address for inbound transactions since the RapidIO address space has 34 bits and the OCN address space has 36 bits. For outbound transactions, the default window maps each of the four 16G chunks to the RapidIO 16G address space. The inbound and outbound translation windows must be aligned based on the granularity selected by the size fields. The packet device ID fields are not used in the inbound translation process, only the address field.

As is the case with all registers, an external processor writing the ATMU registers should not assume that the write has completed until a response is received.

Note that when booting from serial RapidIO, outbound ATMU window 0 must be used.

See [Segmented Outbound Window Description](#) .

The port *m* RapidIO outbound window translation address registers select the starting addresses in the external address space for window hits within the outbound translation windows. The new translated address is created by concatenating the transaction offset to this translation address.

Address: C_0000h base + 1_0E00h offset = D_0E00h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	LTGTID		TREXAD						TRAD							
Reset	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W	TRAD															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2ROWTAR0 field descriptions

Field	Description
0-1 LTGTID	LTGTID correspond to bits 6-7 of the target ID for a large transport system. This field is valid only if bit 27 of PEFCAR is set.

Table continues on the next page...

SRIO_P2ROWTAR0 field descriptions (continued)

Field	Description
	Bits 0-5 of the target ID are specified in the window's port <i>m</i> RapidIO outbound window translation extended address register.
2–11 TRESAD	Translation extended address. TRESAD[0-7] correspond to the target ID for a small transport system or the least significant byte (bits 8-15) of the target ID for a large transport system. TRESAD[8-9] corresponds to bits [0-1] of a 34-bit RapidIO translation address. For maintenance transactions and default window 0, TRESAD[8-9] is reserved.
12–31 TRAD	Translation address. System address which represents the starting point of the outbound translated address. The translation address must be aligned based on the size field. This corresponds to bits [2-21] of the 34-bit RapidIO translation address. For maintenance transactions, the hop count is formed from TRAD[0-7], and the upper 12 bits of the maintenance offset is formed from TRAD[8-19]; the rest of the maintenance offset is formed from the untranslated address. This field is reserved for default window 0.

15.5.86 Port 2 RapidIO outbound window translation extended address register 0 (SRIO_P2ROWTEAR_n)

The port *m* RapidIO outbound window translation extended address registers contain bits 0-5 of the target ID for a common transport large system.

Address: C_0000h base + 1_0E04h offset + (32d × *i*), where *i*=0d to 8d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	Reserved															LTGTID																	
W	0																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1

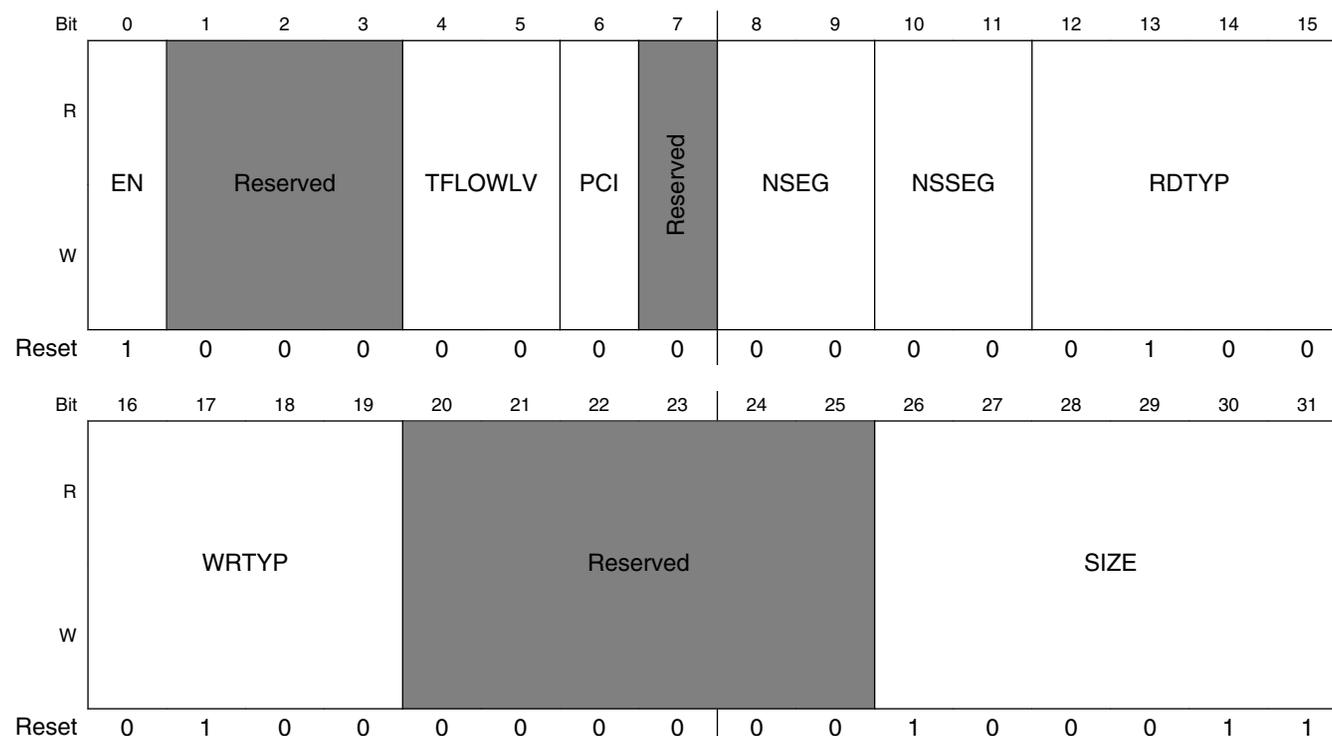
SRIO_P2ROWTEAR_n field descriptions

Field	Description
0–25 -	This field is reserved. Reserved
26–31 LTGTID	LTGTID correspond to bits 0-5 of the target ID for a large transport system. This field is valid only if bit 27 of PEFCAR is set. Bits 6-7 of the target ID are specified in the window's port <i>m</i> RapidIO outbound window translation address register.

15.5.87 Port 2 RapidIO outbound window attributes register n (SRIO_P2ROWAR_n)

The port *m* RapidIO outbound window attributes registers define the window sizes to translate and other attributes for the translations. 64G is the largest window size allowed. For a segmented window, these attributes are used for segment 0. The PCI_Window bit applies for all segments.

Address: C_0000h base + 1_0E10h offset + (32d × i), where i=0d to 8d



SRIO_P2ROWAR_n field descriptions

Field	Description
0 EN	This field enables this address translation window. It is set to 1 and is read only for default window 0.
1–3 -	This field is reserved. Reserved
4–5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI bit is set. Also, the RapidIO priority given by this field must always be greater than or equal to the OCN priority or a deadlock can occur. Normally, the OCN priority of all packets is 0 except for packets coming from a PCI type interface. Read type packets coming from a PCI type interface are priority 0 and write type packets are priority 1. Packets coming from a PCI type interface should set this field to 0 and set the PCI field to 1.

Table continues on the next page...

SRIO_P2ROWAR n field descriptions (continued)

Field	Description
	00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6 PCI	PCI_Window. This window follows PCI ordering rules as defined in the RapidIO Inter-operability specification The TFLOWLV field must be set to 00 if this bit is set causing reads to have RapidIO priority 0 and writes to have RapidIO priority 1.
7 -	This field is reserved. Reserved
8–9 NSEG	Number of segments for this window. This field is reserved for default window 0. 00 One segment (normal window) 01 Two segments (half size aliasing window) 10 Four segments (quarter size aliasing window) 11 Reserved
10–11 NSSEG	Number of subsegments for this segment. This field is reserved for default window 0. Note that this field is valid only when ROWAR n [NSEG] contains a non-zero value (1 or 2). 00 One target deviceID for this segment 01 Two target deviceIDs for this segment 10 Four target deviceIDs for this segment 11 Eight target deviceIDs for this segment
12–15 RDTYP	Transaction type to run on RapidIO interface if access is a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved 0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
16–19 WRTYP	Transaction type to run on RapidIO interface if access was a write. Write-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead.

Table continues on the next page...

SRIO_P2ROWAR n field descriptions (continued)

Field	Description
	0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1111 Reserved
20–25 -	This field is reserved. Reserved
26–31 SIZE	Outbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window size is 4 Kbytes. This field is read only for default window 0. 000000 Reserved ... 001011 4-Kbyte window size 001100 8-Kbyte window size ... 011111 4-Gbyte window size 100000 8-Gbyte window size 100001 16-Gbyte window size 100010 32-Gbyte window size 100011 64-Gbyte window size 100100 Reserved ... 111111 Reserved

15.5.88 Port 2 RapidIO outbound window translation address register n (SRIO_P2ROWTARn)

ATMU registers are used for outbound and inbound transactions. Their purpose is to translate RapidIO packets to OCN packets on inbound and to translate OCN packets to RapidIO packets on outbound. ATMU window misses use the window 0 register set by default, and overlapping window hits results in the use of the lowest number window register set hit. For both inbound and outbound translation, the smallest window size is 4K and the largest window size is 16G for inbound translation and 64G for outbound translation. The default window register set causes no translation of the transaction address for inbound transactions since the RapidIO address space has 34 bits and the OCN address space has 36 bits. For outbound transactions, the default window maps each of the four 16G chunks to the RapidIO 16G address space. The inbound and outbound translation windows must be aligned based on the granularity selected by the size fields. The packet device ID fields are not used in the inbound translation process, only the address field.

As is the case with all registers, an external processor writing the ATMU registers should not assume that the write has completed until a response is received.

Note that when booting from serial RapidIO, outbound ATMU window 0 must be used.

See [Segmented Outbound Window Description](#) .

The port m RapidIO outbound window translation address registers select the starting addresses in the external address space for window hits within the outbound translation windows. The new translated address is created by concatenating the transaction offset to this translation address.

Address: C_0000h base + 1_0E20h offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	LTGTID		TRESAD										TRAD			
W																
Reset	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	TRAD															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2ROWTARn field descriptions

Field	Description
0-1 LTGTID	LTGTID correspond to bits 6-7 of the target ID for a large transport system. This field is valid only if bit 27 of PEFCAR is set.

Table continues on the next page...

SRIO_P2ROWTAR_n field descriptions (continued)

Field	Description
	Bits 0-5 of the target ID are specified in the window's port <i>m</i> RapidIO outbound window translation extended address register.
2–11 TREXAD	Translation extended address. TREXAD[0-7] correspond to the target ID for a small transport system or the least significant byte (bits 8-15) of the target ID for a large transport system. TREXAD[8-9] corresponds to bits [0-1] of a 34-bit RapidIO translation address. For maintenance transactions and default window 0, TREXAD[8-9] is reserved.
12–31 TRAD	Translation address. System address which represents the starting point of the outbound translated address. The translation address must be aligned based on the size field. This corresponds to bits [2-21] of the 34-bit RapidIO translation address. For maintenance transactions, the hop count is formed from TRAD[0-7], and the upper 12 bits of the maintenance offset is formed from TRAD[8-19]; the rest of the maintenance offset is formed from the untranslated address. This field is reserved for default window 0.

15.5.89 Port 2 RapidIO outbound window base address register *n* (SRIO_P2ROWBAR_n)

The port *m* RapidIO outbound window base address registers select the base address for the windows which are translated to an alternate system address space. Addresses for outbound transactions are compared to these windows. If such a transaction does not fall within one of these spaces the transaction is forwarded through the default register set. For transactions that cross more than one window, see Section 15.9.6.2, "Window Boundary Crossing Errors ."

Address: C_0000h base + 1_0E28h offset + (32d × *i*), where *i*=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2ROWBAR_n field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–11 BEXTADD	Window base extended address. Corresponds to bits [0-3] of the 36-bit OCN base address.
12–31 BADD	Window base address. Source address which is the starting point for the outbound translation window. The window must be aligned based on the size selected in the window size bits. This corresponds to bits [4-23] of the 36-bit OCN base address.

15.5.90 Port 2 RapidIO outbound window segment 1 register n (SRIO_P2ROWS1Rn)

The port *m* RapidIO outbound window segment registers define the attributes and target device ID that are to be used for a transaction that hits in that segment rather than the primary attributes and target deviceID. There is one of these registers for each segment (except segment 0).

Address: C_0000h base + 1_0E34h offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15			
R	Reserved				TFLOWLV		Reserved		RDTYP				WRTYP						
W	Reserved				TFLOWLV		Reserved		RDTYP				WRTYP						
Reset	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	0			
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
R	Reserved								SGTGTID0_4								SGTGTID5	SGTGTID6	SGTGTID7
W	Reserved								SGTGTID0_4								SGTGTID5	SGTGTID6	SGTGTID7
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			

SRIO_P2ROWS1Rn field descriptions

Field	Description
0–3 -	This field is reserved. Reserved
4–5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI_Window bit is set 00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6–7 -	This field is reserved. Reserved
8–11 RDTYP	Transaction type to run on RapidIO interface if access was a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved

Table continues on the next page...

SRIO_P2ROWS1Rn field descriptions (continued)

Field	Description
	0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
12–15 WRTYP	Transaction type to run on RapidIO interface if access was a write. Writes-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead. 0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1110 Reserved 1111 Reserved
16–23 -	This field is reserved. Reserved
24–28 SGTGTDID0_4	Bits 0-4 (or bits 8-12 if large transport system) of the target device ID for this segment.
29 SGTGTDID5	Bit 5 (or bit 13 if large transport system) of the target deviceID; this bit is reserved if 8 target subsegments are selected.
30 SGTGTDID6	Bit 6 (or bit 14 if large transport system) of the target deviceID; this bit is reserved if 8 or 4 target subsegments are selected.
31 SGTGTDID7	Bit 7 (or bit 15 if large transport system) of the target deviceID; this bit is reserved if 8, 4, or 2 target subsegments are selected.

15.5.91 Port 2 RapidIO outbound window segment 2 register n (SRIO_P2ROWS2Rn)

The port *m* RapidIO outbound window segment registers define the attributes and target device ID that are to be used for a transaction that hits in that segment rather than the primary attributes and target deviceID. There is one of these registers for each segment (except segment 0).

Address: C_0000h base + 1_0E38h offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP				
W	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP				
Reset	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								SGTGTID0_4				SGTGTID5	SGTGTID6	SGTGTID7	
W	Reserved								SGTGTID0_4				SGTGTID5	SGTGTID6	SGTGTID7	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2ROWS2Rn field descriptions

Field	Description
0–3 -	This field is reserved. Reserved
4–5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI_Window bit is set 00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6–7 -	This field is reserved. Reserved
8–11 RDTYP	Transaction type to run on RapidIO interface if access was a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved

Table continues on the next page...

SRIO_P2ROWS2Rn field descriptions (continued)

Field	Description
	0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
12–15 WRTYP	Transaction type to run on RapidIO interface if access was a write. Writes-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead. 0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1110 Reserved 1111 Reserved
16–23 -	This field is reserved. Reserved
24–28 SGTGTDID0_4	Bits 0-4 (or bits 8-12 if large transport system) of the target device ID for this segment.
29 SGTGTDID5	Bit 5 (or bit 13 if large transport system) of the target deviceID; this bit is reserved if 8 target subsegments are selected.
30 SGTGTDID6	Bit 6 (or bit 14 if large transport system) of the target deviceID; this bit is reserved if 8 or 4 target subsegments are selected.
31 SGTGTDID7	Bit 7 (or bit 15 if large transport system) of the target deviceID; this bit is reserved if 8, 4, or 2 target subsegments are selected.

15.5.92 Port 2 RapidIO outbound window segment 3 register n (SRIO_P2ROWS3Rn)

The port *m* RapidIO outbound window segment registers define the attributes and target device ID that are to be used for a transaction that hits in that segment rather than the primary attributes and target deviceID. There is one of these registers for each segment (except segment 0).

Address: C_0000h base + 1_0E3Ch offset + (32d × i), where i=0d to 7d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15			
R	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP							
W	Reserved			TFLOWLV		Reserved		RDTYP				WRTYP							
Reset	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	0			
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
R	Reserved								SGTGTID0_4								SGTGTID5	SGTGTID6	SGTGTID7
W	Reserved								SGTGTID0_4								SGTGTID5	SGTGTID6	SGTGTID7
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			

SRIO_P2ROWS3Rn field descriptions

Field	Description
0–3 -	This field is reserved. Reserved
4–5 TFLOWLV	Transaction flow level. This field must be set to 00 if the PCI_Window bit is set 00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority level transaction request flow 11 Reserved
6–7 -	This field is reserved. Reserved
8–11 RDTYP	Transaction type to run on RapidIO interface if access was a read. 0000 Reserved 0001 Reserved 0010 IO_READ_HOME 0011 Reserved

Table continues on the next page...

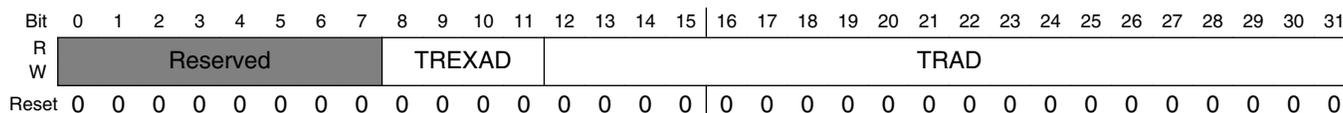
SRIO_P2ROWS3Rn field descriptions (continued)

Field	Description
	0100 NREAD 0101 Reserved 0110 Reserved 0111 MAINTENANCE read 1000 Reserved ... 1100 ATOMIC increment 1101 ATOMIC decrement 1110 ATOMIC set 1111 ATOMIC clear
12–15 WRTYP	Transaction type to run on RapidIO interface if access was a write. Writes-requiring-response sent from OCN must generate a write-requiring-response to RapidIO. Therefore, if an OCN write-requiring-response request hits a window with WRTYP = SWRITE or NWRITE, the RapidIO endpoint generates a NWRITE_R instead. 0000 Reserved 0001 FLUSH 0010 Reserved 0011 SWRITE 0100 NWRITE 0101 NWRITE_R 0110 Reserved 0111 MAINTENANCE write 1000 Reserved ... 1110 Reserved 1111 Reserved
16–23 -	This field is reserved. Reserved
24–28 SGTGTID0_4	Bits 0-4 (or bits 8-12 if large transport system) of the target device ID for this segment.
29 SGTGTID5	Bit 5 (or bit 13 if large transport system) of the target deviceID; this bit is reserved if 8 target subsegments are selected.
30 SGTGTID6	Bit 6 (or bit 14 if large transport system) of the target deviceID; this bit is reserved if 8 or 4 target subsegments are selected.
31 SGTGTID7	Bit 7 (or bit 15 if large transport system) of the target deviceID; this bit is reserved if 8, 4, or 2 target subsegments are selected.

15.5.93 Port 2 RapidIO Inbound window translation address register n (SRIO_P2RIWTAR_n)

The port *m* RapidIO inbound window translation address registers point to the starting addresses in local address space for window hits within the inbound translation windows. The new translated address is created by concatenating the transaction offset to this translation address.

Address: C_0000h base + 1_0F60h offset + (32d × i), where i=0d to 4d



SRIO_P2RIWTAR_n field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–11 TRESAD	Translation extended address. Corresponds to bits 0-3 of the 36-bit OCN translation address. TRESAD[2-3] are reserved for default window 0.
12–31 TRAD	Translation address. Target address which indicates the starting point of the inbound translated address. The translation address must be aligned based on the size field. This corresponds to bits 4-23 of the 36-bit OCN translation address. TRAD is reserved for default window 0.

15.5.94 Port 2 RapidIO Inbound window base address register n (SRIO_P2RIWBARn)

The port *m* RapidIO inbound window *n* base address registers select the base address for the windows which are translated to an alternate target address space. Addresses for inbound transactions are compared to these windows. If such a transaction does not fall within one of these spaces, then the transaction is forwarded to the interior of the chip using the default window. For transactions that cross more than one window, see Section 15.9.7.2, "Window Boundary Crossing Errors ." Note that the LCSBA1CSR register (See [Inbound port-write n queue base address register](#)) has priority over all ATMU windows if both are configured for the same address space.

Address: C_0000h base + 1_0F68h offset + (32d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved									BEXAD		BADD				
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	BADD															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_P2RIWBARn field descriptions

Field	Description
0–9 -	This field is reserved. Reserved
10–11 BEXAD	Base extended address. BEXAD represents bits 0-1 of the 34-bit RapidIO address.
12–31 BADD	Base address. System address which is the starting point for the inbound translation window. The window must be aligned based on the size selected in the window size bits. This corresponds to bits 2-21 of the 34-bit RapidIO base address.

15.5.95 Port 2 RapidIO inbound window attributes register n (SRIO_P2RIWAR_n)

The port *m* RapidIO inbound window attributes registers define the window sizes to translate and other attributes for the translations. 16 Gbytes is the largest window size allowed. The RDTYP and WRTYP fields are used for attributes on the request, but do not actually change the transaction type of the transaction. In other words, P *m* RIWAR *n* does not modify whether or not the request requires a response or if the request was atomic; this type of attribute remains unchanged on the request as it is translated through the ATMU.

Address: C_0000h base + 1_0F70h offset + (32d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	EN	PW	Reserved					TGINT				RDTYP				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W	WRTYP				Reserved						SIZE					
Reset	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	1

SRIO_P2RIWAR_n field descriptions

Field	Description
0 EN	This bit enables this address translation. This field is set to 1 and read-only for default window 0.
1 PW	Protected Window. This bit indicates that this window is protected. Writes requiring response and reads to this window generate an error response. Writes not requiring response are silently discarded.
2–7 -	This field is reserved. Reserved
8–11 TGINT	Target interface. If this field is set to anything other than local address space, the attributes for the transaction must be assigned in a corresponding outbound window at the target. This is the field definition for the P2020: 0000 PCI Express interface 3 0001 PCI Express interface 2 0010 PCI Express Interface 1 0011-1011 Reserved 1100 Serial RapidIO interface 1 1101 Serial RapidIO interface 2 1110 Reserved 1111 Local memory (DDRs, local bus controller, memory-mapped registers)
12–15 RDTYP	Transaction type to run on the I/O interface if access is a read. 0000 Reserved

Table continues on the next page...

SRIO_P2RIWAR_n field descriptions (continued)

Field	Description
	<p>...</p> <p>0100 Read</p> <p>0101 Reserved</p> <p>...</p> <p>1111 Reserved</p> <p>Transaction type to run on local memory if access is a read.</p> <p>0000 Reserved</p> <p>...</p> <p>0100 Read, do not snoop local processor</p> <p>0101 Read, snoop local processor</p> <p>0110 Reserved</p> <p>0111 Read, unlock L2 cache line</p> <p>1000 Reserved</p> <p>...</p> <p>1111 Reserved</p>
16–19 WRTYP	<p>Transaction type to run on I/O interface if access is a write.</p> <p>0000 Reserved</p> <p>...</p> <p>0100 write</p> <p>0101 Reserved</p> <p>...</p> <p>1111 Reserved</p> <p>Transaction type to run on local memory if access is a write.</p> <p>0000 Reserved</p> <p>...</p> <p>0011 Reserved</p> <p>0100 Write, do not snoop local processor</p> <p>0101 Write, snoop local processor</p> <p>0110 Write, allocate cache line</p> <p>0111 Write, allocate and lock cache line</p> <p>1000 Reserved</p> <p>...</p> <p>1111 Reserved</p>
20–25 -	<p>This field is reserved. Reserved</p>
26–31 SIZE	<p>Inbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window size is 4 Kbytes.</p> <p>This field is read only for default window 0.</p>

Table continues on the next page...

SRIO_P2RIWAR n field descriptions (continued)

Field	Description
000000	Reserved
...	
001011	4-Kbyte window size
001100	8-Kbyte window size
...	
011111	4-Gbyte window size
100000	8-Gbyte window size
100001	16-Gbyte window size
100010	Reserved
...	
111111	Reserved

15.5.96 Port 2 RapidIO inbound window attributes register 0 (SRIO_P2RIWAR0)

The port m RapidIO inbound window attributes registers define the window sizes to translate and other attributes for the translations. 16 Gbytes is the largest window size allowed. The RDTYP and WRTYP fields are used for attributes on the request, but do not actually change the transaction type of the transaction. In other words, P m RIWAR n does not modify whether or not the request requires a response or if the request was atomic; this type of attribute remains unchanged on the request as it is translated through the ATMU.

Address: C_0000h base + 1_0FF0h offset = D_0FF0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	EN		Reserved						TGINT				RDTYP			
W		PW														
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	WRTYP				Reserved								SIZE			
W																
Reset	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	1

SRIO_P2RIWAR0 field descriptions

Field	Description
0 EN	This bit enables this address translation.

Table continues on the next page...

SRIO_P2RIWAR0 field descriptions (continued)

Field	Description
	This field is set to 1 and read-only for default window 0.
1 PW	Protected Window. This bit indicates that this window is protected. Writes requiring response and reads to this window generate an error response. Writes not requiring response are silently discarded.
2-7 -	This field is reserved. Reserved
8-11 TGINT	Target interface. If this field is set to anything other than local address space, the attributes for the transaction must be assigned in a corresponding outbound window at the target. This is the field definition for the P2020: 0000 PCI Express interface 3 0001 PCI Express interface 2 0010 PCI Express Interface 1 0011-1011 Reserved 1100 Serial RapidIO interface 1 1101 Serial RapidIO interface 2 1110 Reserved 1111 Local memory (DDRs, local bus controller, memory-mapped registers)
12-15 RDTYP	Transaction type to run on the I/O interface if access is a read. 0000 Reserved ... 0100 Read 0101 Reserved ... 1111 Reserved Transaction type to run on local memory if access is a read. 0000 Reserved ... 0100 Read, do not snoop local processor 0101 Read, snoop local processor 0110 Reserved 0111 Read, unlock L2 cache line 1000 Reserved ... 1111 Reserved
16-19 WRTYP	Transaction type to run on I/O interface if access is a write. 0000 Reserved ... 0100 write 0101 Reserved ... 1111 Reserved

Table continues on the next page...

SRIO_P2RIWAR0 field descriptions (continued)

Field	Description
	Transaction type to run on local memory if access is a write. 0000 Reserved ... 0011 Reserved 0100 Write, do not snoop local processor 0101 Write, snoop local processor 0110 Write, allocate cache line 0111 Write, allocate and lock cache line 1000 Reserved ... 1111 Reserved
20–25 -	This field is reserved. Reserved
26–31 SIZE	Inbound translation window size N which is the encoded $2^{(N+1)}$ bytes window size. The smallest window size is 4 Kbytes. This field is read only for default window 0. 000000 Reserved ... 001011 4-Kbyte window size 001100 8-Kbyte window size ... 011111 4-Gbyte window size 100000 8-Gbyte window size 100001 16-Gbyte window size 100010 Reserved ... 111111 Reserved

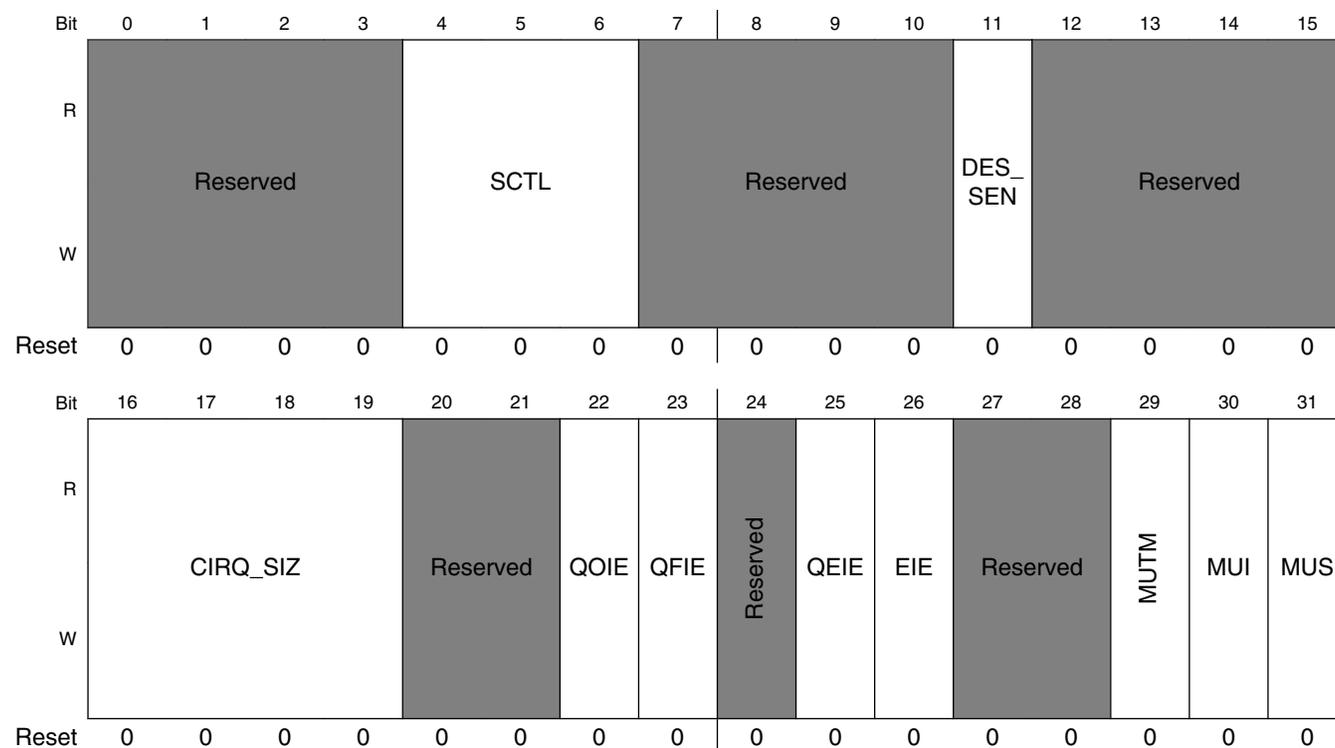
15.5.97 Outbound message n mode register (SRIO_OMnMR)

RapidIO Outbound Message 0 Registers control RapidIO outbound messages. The following provide partial descriptions of these registers.

The outbound message mode register allows software to start a message operation and to control various message operation characteristics.

Serial RapidIO Memory Map/Register Definition

Address: C_0000h base + 1_3000h offset + (256d × i), where i=0d to 1d



SRIO_OMnMR field descriptions

Field	Description
0–3 -	This field is reserved. Reserved
4–6 SCTL	Service control. Determines the number of descriptors to process before servicing the next queue. Note that if one queue has SCTL set to fixed priority, all SCTL values in other queues are ignored. For example, of the two outbound message units, if unit 1's service control value is set to 0b000, fixed priority results with unit 0 the highest priority results. If a queue is in direct mode only, one message operation is serviced before servicing the next queue. For proper operation, this field should only be modified when the outbound message controller is not enabled. The value of this field cannot be changed unless both units are disabled. 000 Fixed priority based on outbound message unit number 001 1 descriptors 010 2 descriptors 011 4 descriptors 100 8 descriptors 101 16 descriptors 110 32 descriptors 111 64 descriptors
7–10 -	This field is reserved. Reserved
11 DES_SEN	Descriptor snoop enable. When set enables snooping of the local processor when reading descriptors from memory. For proper operation, this field should only be modified when the outbound message controller is not enabled

Table continues on the next page...

SRIO_OM n MR field descriptions (continued)

Field	Description																																		
12–15 -	This field is reserved. Reserved																																		
16–19 CIRQ_SIZ	<p>Circular descriptor queue size. Determines the number of descriptors that can be placed on the circular queue without overflow.</p> <p>For proper operation, this field should only be modified when the outbound message controller is not enabled</p> <table> <tr><td>0000</td><td>2</td></tr> <tr><td>0001</td><td>4</td></tr> <tr><td>0010</td><td>8</td></tr> <tr><td>0011</td><td>16</td></tr> <tr><td>0100</td><td>32</td></tr> <tr><td>0101</td><td>64</td></tr> <tr><td>0110</td><td>128</td></tr> <tr><td>0111</td><td>256</td></tr> <tr><td>1000</td><td>512</td></tr> <tr><td>1001</td><td>1024</td></tr> <tr><td>1010</td><td>2048</td></tr> <tr><td>1011</td><td>Reserved</td></tr> <tr><td>...</td><td></td></tr> <tr><td>1100</td><td>Reserved</td></tr> <tr><td>1101</td><td>Reserved</td></tr> <tr><td>1110</td><td>Reserved</td></tr> <tr><td>1111</td><td>Reserved</td></tr> </table>	0000	2	0001	4	0010	8	0011	16	0100	32	0101	64	0110	128	0111	256	1000	512	1001	1024	1010	2048	1011	Reserved	...		1100	Reserved	1101	Reserved	1110	Reserved	1111	Reserved
0000	2																																		
0001	4																																		
0010	8																																		
0011	16																																		
0100	32																																		
0101	64																																		
0110	128																																		
0111	256																																		
1000	512																																		
1001	1024																																		
1010	2048																																		
1011	Reserved																																		
...																																			
1100	Reserved																																		
1101	Reserved																																		
1110	Reserved																																		
1111	Reserved																																		
20–21 -	This field is reserved. Reserved																																		
22 QOIE	Queue overflow interrupt enable. When set, bit generates an interrupt on detection of a queue overflow (that is, the enqueue and dequeue pointers are no longer equal after being incremented by the processor and the queue was full). No queue overflow interrupt is generated if this bit is cleared. (only applicable to chaining mode). If this bit is not set and the queue overflows, the result is undefined.																																		
23 QFIE	Queue full interrupt enable. When set will generate an interrupt when the queue transitions to full (that is, the enqueue and dequeue pointers are equal after being incremented by the processor). No QF interrupt is generated if this bit is cleared. If this bit is set and OM n SR[QF] is a 1, OM n SR[QFI] will assert.																																		
24 -	This field is reserved. Reserved																																		
25 QEIE	Queue empty interrupt enable. When set, bit generates an interrupt at the completion of all outstanding message operations (that is, the enqueue and dequeue pointers are equal after an increment by the message unit controller). No QE interrupt is generated if this bit is cleared. For proper operation, this field should only be modified when the outbound message controller is not enabled																																		
26 EIE	Error interrupt enable. When set, bit generates the port-write/error interrupt when a transfer error (OM n SR[TE]), a message error response (OM n SR[MER]), a packet response time-out (OM n SR[PRT]), or a retry threshold event exceeded (OM n SR[RETE]) event occurs. No port-write/error interrupt is generated if this bit is cleared.																																		
27–28 -	This field is reserved. Reserved																																		
29 MUTM	Message unit transfer mode. Setting this bit puts the message unit into direct mode. This means that software is responsible for placing all the required parameters into necessary registers to start the message transmission. Clearing this bit configures the message unit in chaining mode.																																		

Table continues on the next page...

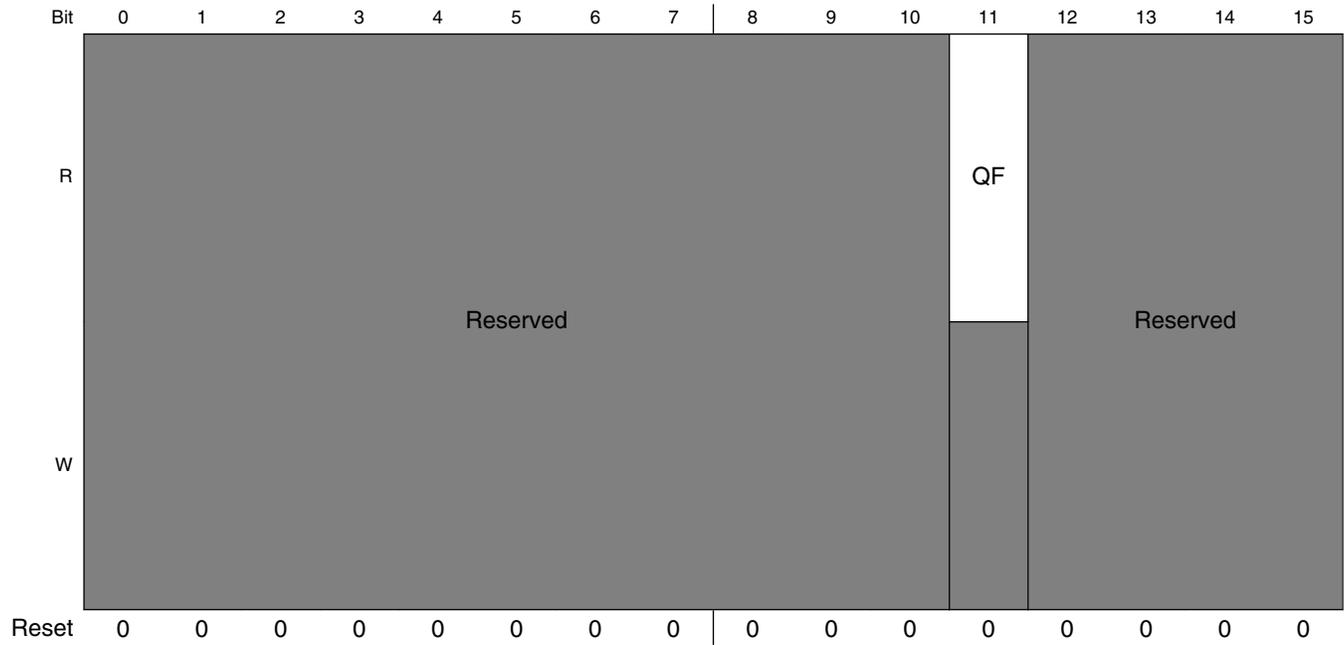
SRIO_OM n MR field descriptions (continued)

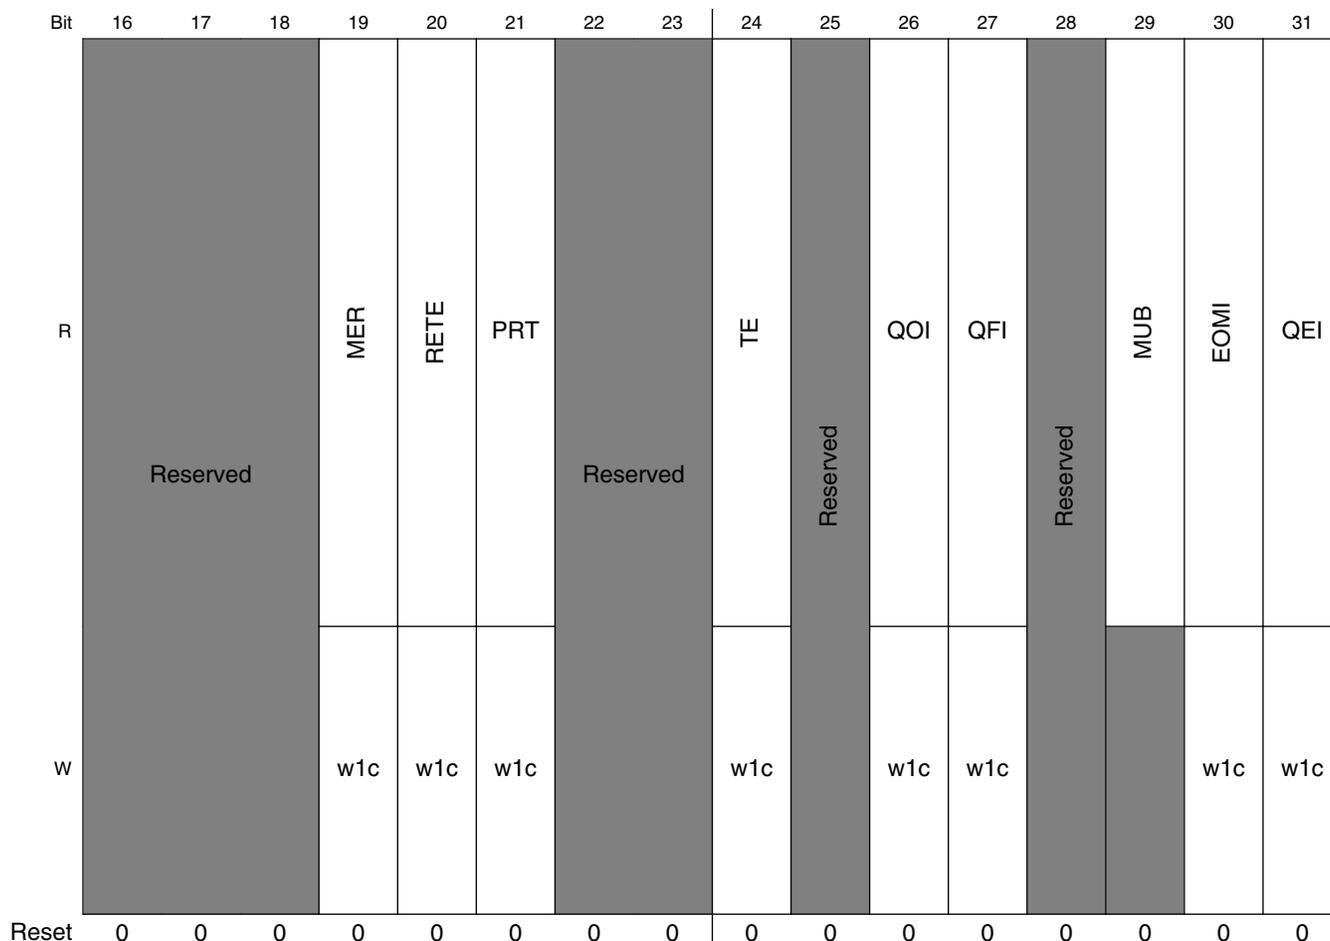
Field	Description
30 MUI	Message unit increment. Software sets this bit after writing a descriptor to memory. Hardware then increments the OM n DQEPAR and clear this bit. Always reads as 0 when MUS is set.
31 MUS	<p>Message unit start.</p> <p>Direct mode-A 0 to 1 transition when the message unit is not busy (MUB bit is 0) starts the message unit. A 1 to 0 transition will have no effect.</p> <p>Chaining mode-If this bit is set, the message unit starts whenever the enqueue and dequeue pointers are not equal.</p>

15.5.98 Outbound message n status register (SRIO_OMnSR)

The outbound message status register reports various message unit conditions during and after a message operation. Writing a 1 to the corresponding set bit clears the bit.

Address: C_0000h base + 1_3004h offset + (256d × i), where i=0d to 1d





SRIO_OMnSR field descriptions

Field	Description
0–10 -	This field is reserved. Reserved
11 QF	Queue full-If the queue becomes full, then this bit is set. (Read-only)
12–18 -	This field is reserved. Reserved
19 MER	Message error response. This bit is set when an ERROR response is received from the message target. The Error response received field indicates value of the error response status bits when an error response is received. This bit is cleared by writing a 1.
20 RETE	Retry error threshold exceeded. This bit is set when the message unit has been unable to complete a message operation because the retry error threshold value has been exceeded due to RapidIO retry response. This bit is cleared by writing a 1.
21 PRT	Packet response time-out. This bit is set when the message unit has been unable to complete a message operation and a packet response time-out occurred. This bit is cleared by writing a 1.
22–23 -	This field is reserved. Reserved

Table continues on the next page...

SRIO_OMnSR field descriptions (continued)

Field	Description
24 TE	Transaction error. This bit is set when an internal error condition occurs during the message operation. This bit is cleared by writing a 1. For proper operation, this field should only be modified when the outbound message controller is not enabled
25 -	This field is reserved. Reserved
26 QOI	Queue overflow interrupt. This bit is set when a queue overflow condition is detected. This bit is cleared by writing a 1. (only applicable to chaining mode)
27 QFI	Queue full interrupt. If the queue becomes full, if the QFIE bit in the Mode Register is set, then this bit is set and an interrupt generated. This bit is cleared by writing a 1.
28 -	This field is reserved. Reserved
29 MUB	Message unit busy. When set indicates that a message operation is currently in progress. This bit is cleared as a result of an error or the message operation is finished. (Read-only)
30 EOMI	End-Of-Message interrupt. After finishing this message operation, if the EOMIE bit in the Destination Attributes Register is set, then this bit is set and an interrupt generated. This bit is cleared by writing a 1.
31 QEI	Queue Empty Interrupt. When the last message operation in the outbound descriptor queue is finished, if the QEIE bit in the Mode Register is set, then this bit is set and an interrupt is generated. Otherwise, no interrupt is generated. This bit is cleared by writing a 1.

15.5.99 Extended outbound message n descriptor queue dequeue pointer address register (SRIO_EOMnDQDPA)

Used with OM *n* DQDPA; see description in [Outbound message n descriptor queue dequeue pointer address register](#) .

Address: C_0000h base + 1_3008h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															EDQDPA																
W	Reserved															EDQDPA																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

SRIO_EOMnDQDPA field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EDQDPA	Extended descriptor queue dequeue pointer address bits. These are the highest order address bits. For proper operation, this field should only be modified when the outbound message controller is not enabled.

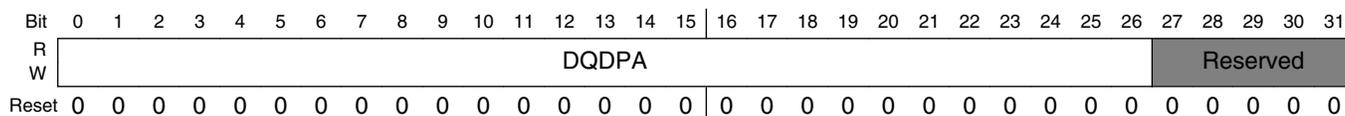
15.5.100 Outbound message n descriptor queue dequeue pointer address register (SRIO_OMnDQDPA)

The outbound message descriptor queue dequeue pointer address registers contain the address of the first descriptor in memory to be processed. Software must initialize these registers to point to the first descriptor in memory. After processing this descriptor, the message unit controller increments the outbound message descriptor queue dequeue pointer address in OM n DQDPA and EOM n DQDPA to point to the next descriptor. If the outbound message descriptor queue enqueue pointer and the outbound message descriptor queue dequeue pointer are not equal (indicating that the queue is not empty), the message unit controller reads the next descriptor from memory for processing. If the enqueue and dequeue pointers are equal after the dequeue pointer has been incremented by the message unit controller, the queue is empty and the message unit halts until the enqueue pointer is incremented by the processor. Incrementing the pointer indicates that a new descriptor has been added to the queue and is ready for processing. If the queue becomes empty and OM n MR[QEIE] is set, OM n SR[QEI] is set and an interrupt is generated.

When software initializes these registers, the queue to which they point must be aligned on a boundary equal to the number of queue entries x 32 bytes (size of each queue descriptor). For example, if there are eight entries in the queue, the queue must be 256-byte aligned.

The number of queue entries is set in OM n MR[CIRQ_SIZ]; see [Outbound message n mode register](#) .

Address: C_0000h base + 1_300Ch offset + (256d × i), where i=0d to 1d



SRIO_OMnDQDPA field descriptions

Field	Description
0–26 DQDPA	Descriptor queue dequeue pointer address. Contains the address of the first descriptor in memory to process. Each descriptor must be size aligned to a 32 byte boundary as each descriptor is 32 bytes in size. For proper operation, this field should only be modified when the outbound message controller is not enabled. When writing to this register, it is not necessary to shift the address, since the lower 5 bits are merely ignored to ensure 32 byte alignment.
27–31 -	This field is reserved. Reserved

15.5.101 Extended outbound message n source address register (SRIO_EOMnSAR)

Address: C_0000h base + 1_3010h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															ESAD																
W	Reserved															ESAD																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_EOMnSAR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 ESAD	Extended source address bits. These are the highest order address bits. For proper operation, this field should only be modified when the outbound message controller is not enabled

15.5.102 Outbound message n source address register (SRIO_OMnSAR)

The outbound message unit source address registers indicate the address from which the message unit controller is to read data. Software must ensure that this is a valid local memory address. The source address must be aligned to a double-word boundary, so the least significant three bits are reserved.

Address: C_0000h base + 1_3014h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	SAD															
W	SAD															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	SAD												SNEN	Reserved		
W	SAD													SNEN	Reserved	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0		0	0

SRIO_OMnSAR field descriptions

Field	Description
0–28 SAD	Source address. This is the source address of the message operation. The contents are updated after every memory read operation. For proper operation, this field should only be modified when the outbound message controller is not enabled
29 SNEN	Snoop enable. This bit enables snooping of the local processor caches for the data reads from local memory. For proper operation, this field should only be modified when the outbound message controller is not enabled
30–31 -	This field is reserved. Reserved

15.5.103 Outbound message n destination port register (SRIO_OMnDPR)

The destination port register indicates the RapidIO destination ID and mailbox to which the message unit controller is to send data. The software must ensure that this is a valid port in the receiving device.

Address: C_0000h base + 1_3018h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	EDTGTRROUTE								DTGTRROUTE							
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										XMAILBOX			MAILBOX		
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

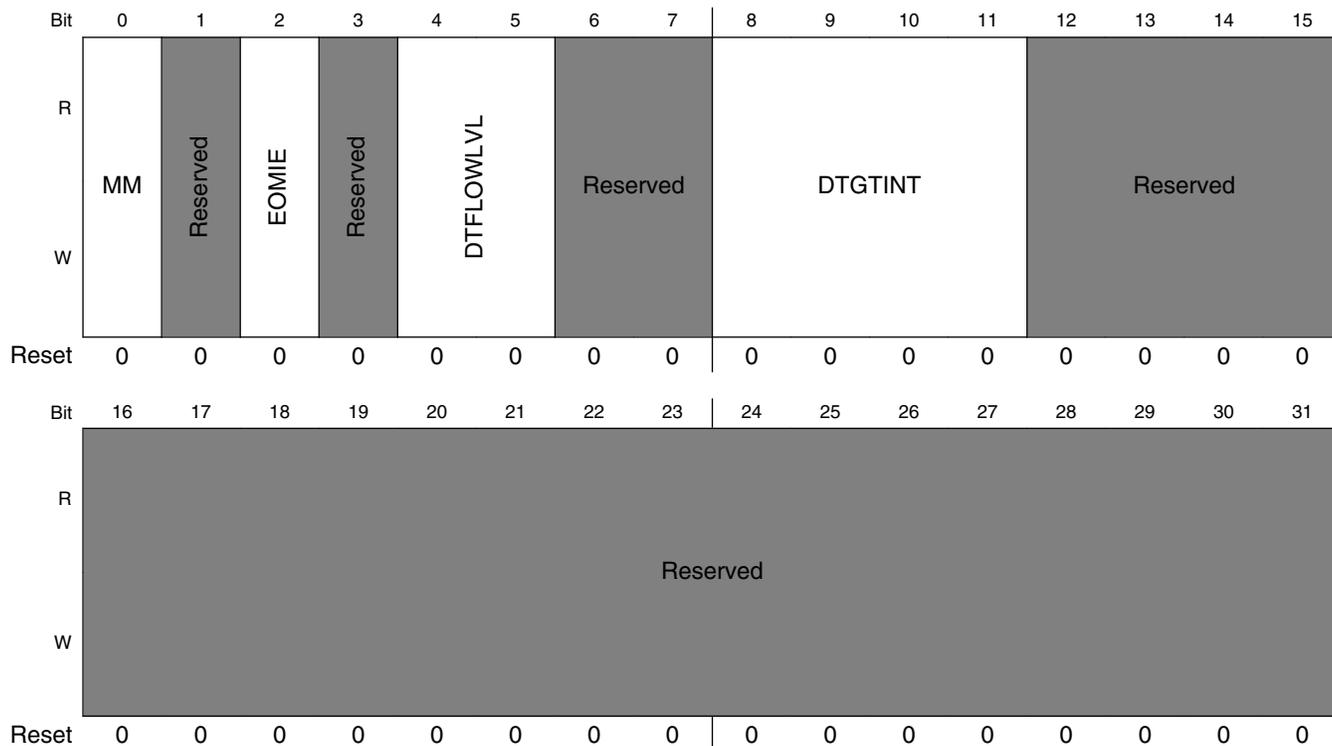
SRIO_OMnDPR field descriptions

Field	Description
0–7 EDTGTRROUTE	Extended destination target route. Most significant byte of a 16-bit target route when operating in large transport mode. Reserved when operating in small transport mode. For proper operation, this field should only be modified when the outbound message controller is not enabled
8–15 DTGTRROUTE	Destination target route. Contains the target route field of the transaction (Device ID of the target). This value is overridden by the multicast group and list if multicast mode is enabled. On error, if multicast mode is enabled, this field is loaded with the destination of the failed operation. For proper operation, this field should only be modified when the outbound message controller is not enabled
16–25 -	This field is reserved. Reserved
26–29 XMAILBOX	Value for 'xmbx' field in MESSAGE packet. This field is only used when OM n DATR[MM] is set. For proper operation, this field should only be modified when the outbound message controller is not enabled
30–31 MAILBOX	Value for 'mbx' field in MESSAGE packet. For proper operation, this field should only be modified when the outbound message controller is not enabled

15.5.104 Outbound message n destination attributes Register (SRIO_OMnDATR)

The outbound message destination attributes register contains the transaction attributes to be used for the message operation.

Address: C_0000h base + 1_301Ch offset + (256d × i), where i=0d to 1d



SRIO_OMnDATR field descriptions

Field	Description
0 MM	Multicast mode. When set, the message operation is sent to all of the targets indicated by the multicast group and list. Messages are limited to one segment and 256 bytes or less when this mode is enabled.
1 -	This field is reserved. Reserved
2 EOMIE	End-of-Message interrupt enable. When set, generates an interrupt when the current message operation is finished. For proper operation, this field should only be modified when the outbound message controller is not enabled.
3 -	This field is reserved. Reserved
4–5 DTFLOWLVL	Transaction flow level. For proper operation, this field should only be modified when the outbound message controller is not enabled

Table continues on the next page...

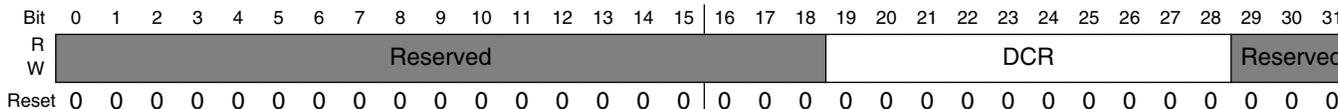
SRIO_OMnDATR field descriptions (continued)

Field	Description
	00 Lowest priority transaction request flow 01 Next highest priority transaction request flow 10 Highest priority transaction request flow 11 Reserved
6–7 -	This field is reserved. Reserved
8–11 DTGTINT	Target interface. The value of this field should be set according to the following: 0000 RapidIO Port 1 0001 RapidIO Port 2
12–31 -	This field is reserved. Reserved

15.5.105 Outbound message n double-word count register (SRIO_OMnDCR)

The outbound message double-word count register contains the number of double-words for the message operation. The maximum message operation size is 4 Kbytes and the minimum is 8 bytes.

Address: C_0000h base + 1_3020h offset + (256d × i), where i=0d to 1d



SRIO_OMnDCR field descriptions

Field	Description
0–18 -	This field is reserved. Reserved
19–28 DCR	Transfer count register. Contains the number of bytes for the message operation. All other values yield undefined behavior. For proper operation, this field should only be modified when the outbound message controller is not enabled. 00 0000 0000 Reserved 00 0000 0001 8 bytes 00 0000 0010 16 bytes 00 0000 0100 32 bytes 00 0000 1000 64 bytes 00 0001 0000 128 bytes

Table continues on the next page...

SRIO_OMnDCR field descriptions (continued)

Field	Description
	00 0010 0000 256 bytes 00 0100 0000 512 bytes (multi segment mode only) 00 1000 0000 1024 bytes (multi segment mode only) 01 0000 0000 2048 bytes (multi segment mode only) 10 0000 0000 4096 bytes (multi segment mode only)
29–31 -	This field is reserved. Reserved

15.5.106 Extended outbound message n descriptor queue enqueue pointer address register (SRIO_EOMnDQEPAR)

Used with OM *n* DQEPAR; see description in [Outbound message n descriptor queue enqueue pointer address register](#) .

Address: C_0000h base + 1_3024h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															EDQEPA																
W	Reserved															EDQEPA																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_EOMnDQEPAR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EDQEPA	Extended descriptor queue enqueue pointer address. These are the highest-order address bits.

15.5.107 Outbound message n descriptor queue enqueue pointer address register (SRIO_OMnDQEPAR)

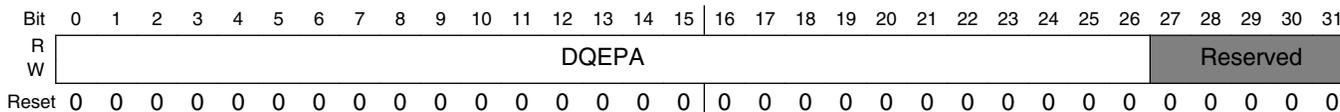
The outbound message descriptor queue enqueue pointer address registers contain the address for the next descriptor in memory to be added to the queue. Software must initialize these registers to match the outbound message descriptor queue dequeue pointer address. When a message is ready to be sent, the processor writes a descriptor to the next location in the queue (indicated by the address in OM *n* DQEPAR and EOM *n* DQEPAR), and then either writes the OM *n* DQEPAR to point to the next descriptor location in memory or sets OM *n* MR[MUI]. This can result in a number of actions:

- If the enqueue and dequeue pointers match, the queue is now full. If the OM *n* MR[QFIE] bit is set, then the OM *n* SR[QFI] bit is set, and an interrupt is generated.
- If the enqueue and dequeue pointer no longer match after the enqueue pointer has been incremented and the queue is full, then the queue overflows, and the message unit stops. If OM *n* MR[QOIE] is set, OM *n* SR[QOI] is set and an interrupt is generated. OM *n* MR[MUS] must transition to a 0 to clear this error condition. If the enqueue pointer is directly written, the queue overflow condition is not detected.
- If the enqueue and dequeue pointers were the same before reading the register, the message unit controller will start (if enabled).

When software initializes these registers, the queue to which they point must be aligned on a boundary equal to the number of queue entries x 32 bytes (size of each queue descriptor). For example, if there are eight entries in the queue, the queue must be 256-byte aligned.

The number of queue entries is set in OM *n* MR[CIRQ_SIZ]; see [Outbound message *n* mode register](#) .

Address: C_0000h base + 1_3028h offset + (256d x i), where i=0d to 1d



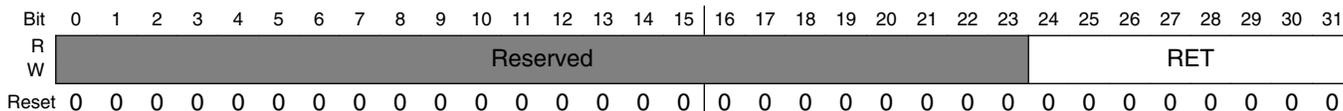
SRIO_OMnDQEPAR field descriptions

Field	Description
0–26 DQEPA	Descriptor queue enqueue pointer address. Contains the address of the last descriptor in memory to process. The descriptor must be size aligned to a 32 byte boundary. When writing to this register it is not necessary to shift the address, since the lower 5 bits are merely ignored to ensure 32 byte alignment .
27–31 -	This field is reserved. Reserved

15.5.108 Outbound message n retry error threshold configuration register (SRIO_OMnRETCR)

The outbound message retry error threshold configuration register controls the number of times that the message unit attempts to transfer a message segment to a particular destination before reporting an error. A message segment is retransmitted if a RETRY response is received from the target.

Address: C_0000h base + 1_302Ch offset + (256d × i), where i=0d to 1d



SRIO_OMnRETCR field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24–31 RET	<p>Retry error threshold. This value is the number of times that the message unit attempts to transmit a message segment to a particular target.</p> <p>For proper operation, this field should only be modified when the outbound message controller is not enabled</p> <p>0x00 Disabled 0x01 Message segment transmitted only 1 time 0x02 Message segment transmitted up to 2 times ... 0xFF Message segment transmitted up to 255 times</p>

15.5.109 Outbound message n multicast group register (SRIO_OMnMGR)

The multicast group register contains a multicast group (MG) value and extended multicast group number (EMG) which, in combination with the multicast list register and the multicast enable (OM *n* DATR[MM]), indicates which deviceIDs are targets of a multicast operation. The multicast group represents the most significant three bits (bits[0-2]) of the RapidIO deviceIDs that are destinations of the message operation, whereas the multicast list indicates a list of targets within that group. Each individual set bit, when encoded, determines the least significant five bits of the RapidIO deviceIDs (bits[3-7]) that are targets of the message operation. Therefore, multicast group 0 (MG = 0) contains target deviceIDs (0,1,...,31), multicast group 1 (MG = 1) contains target deviceIDs (32,33,...,63), and so on.

If in large transport mode, the extended multicast group represents the eight most significant bits (bits[0-7]), the multicast group represents the next three bits (bits[8-10]) and the multicast list indicates a list of targets within that group.

If multicast is enabled, this information in the multicast group and mask register is used to determine the target of the message operation instead of the DTGTROUTE and EDTGTROUTE fields in the outbound message *n* destination attributes register.

Address: C_0000h base + 1_3030h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															EMG								MG								
W	Reserved															EMG								MG								
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

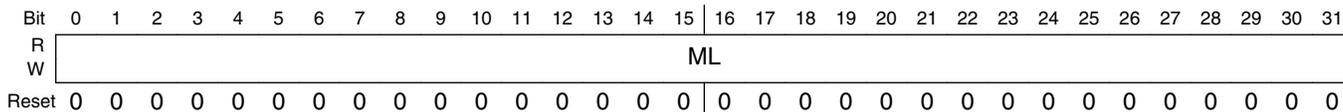
SRIO_OMnMGR field descriptions

Field	Description
0–20 -	This field is reserved. Reserved
21–28 EMG	Extended multicast group. This is the most significant eight bits of the target deviceIDs for the multicast operation when operating in large transport mode. For proper operation, this field should only be modified when the outbound message controller is not enabled
29–31 MG	Multicast group. This is the most significant three bits of the target deviceIDs for the multicast operation. For proper operation, this field should only be modified when the outbound message controller is not enabled

15.5.110 Outbound message n multicast list register (SRIO_OMnMLR)

See the multicast group register description for more information.

Address: C_0000h base + 1_3034h offset + (256d × i), where i=0d to 1d



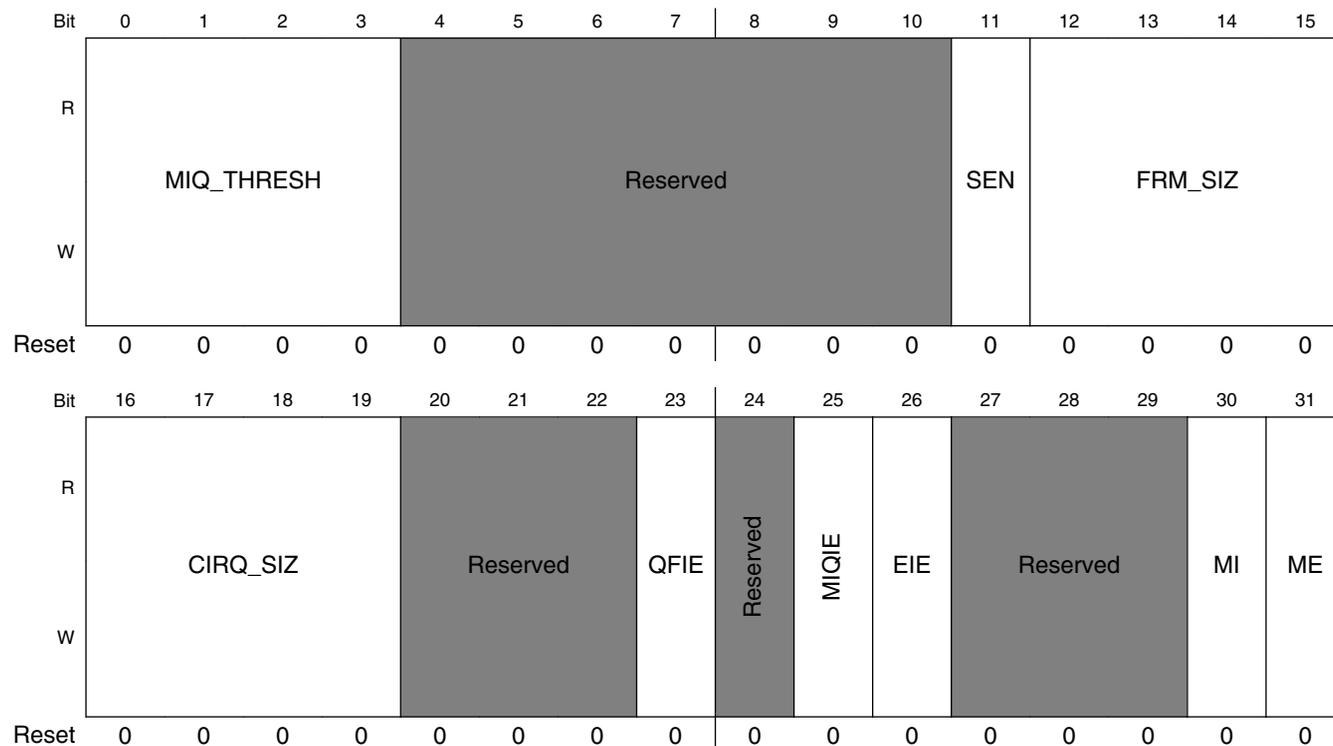
SRIO_OMnMLR field descriptions

Field	Description
0–31 ML	Multicast list. This is the group target list for the message operation. Depending on the value of the multicast group value, bit 0 corresponds to deviceID 0, 32, 64, 96, and so on; bit 1 corresponds to deviceID 1, 33, 65, 97, and so on. If none of the bits are set, bit 0 is assumed to be set. For proper operation, this field should only be modified when the outbound message controller is not enabled.

15.5.111 Inbound message n mode register (SRIO_IMnMR)

The inbound message mode register allows software to enable the mailbox controller and to control various message operation characteristics.

Address: C_0000h base + 1_3060h offset + (256d × i), where i=0d to 1d



SRIO_IMnMR field descriptions

Field	Description
0-3 MIQ_THRESH	<p>Message in queue threshold. Determines the number of message frames to be accumulated in the frame queue before Message In Queue is signaled. Undefined operation results if the message in queue threshold is set greater than or equal to the message queue size (IM n MR[CIRQ_SIZ]).</p> <p>For proper operation, this field should only be modified when the inbound message controller is not enabled.</p> <p>0000 1 0001 2 0010 4 0011 8 0100 16 0101 32 0110 64 0111 128 1000 256</p>

Table continues on the next page...

SRIO_IMnMR field descriptions (continued)

Field	Description
	1001 512 1010 1024 1011 Reserved ... 1111 Reserved
4–10 -	This field is reserved. Reserved
11 SEN	Snoop enable. When set, enables snooping the local processor when writing messages into memory. For proper operation, this field should only be modified when the inbound message controller is not enabled.
12–15 FRM_SIZ	Message frame size. Determines the maximum message size that can be accepted by this mailbox without error. This parameter combined with CIRQ_SIZ determine the maximum contiguous memory space allocated to the mailbox. For proper operation, this field should only be modified when the inbound message controller is not enabled. 0000 Reserved ... 0010 8 bytes 0011 16 bytes 0100 32 bytes 0101 64 bytes 0110 128 bytes 0111 256 bytes 1000 512 bytes 1001 1024 bytes 1010 2048 bytes 1011 4096 bytes 1100 Reserved ... 1111 Reserved
16–19 CIRQ_SIZ	Circular frame queue size. Determines the number of message frames that can be place on the circular queue without overflow. This parameter combined with FRM_SIZ determine the maximum contiguous memory space allocated to the mailbox. For proper operation, this field should only be modified when the inbound message controller is not enabled. 0000 2 0001 4 0010 8 0011 16 0100 32 0101 64 0110 128 0111 256 1000 512 1001 1024 1010 2048

Table continues on the next page...

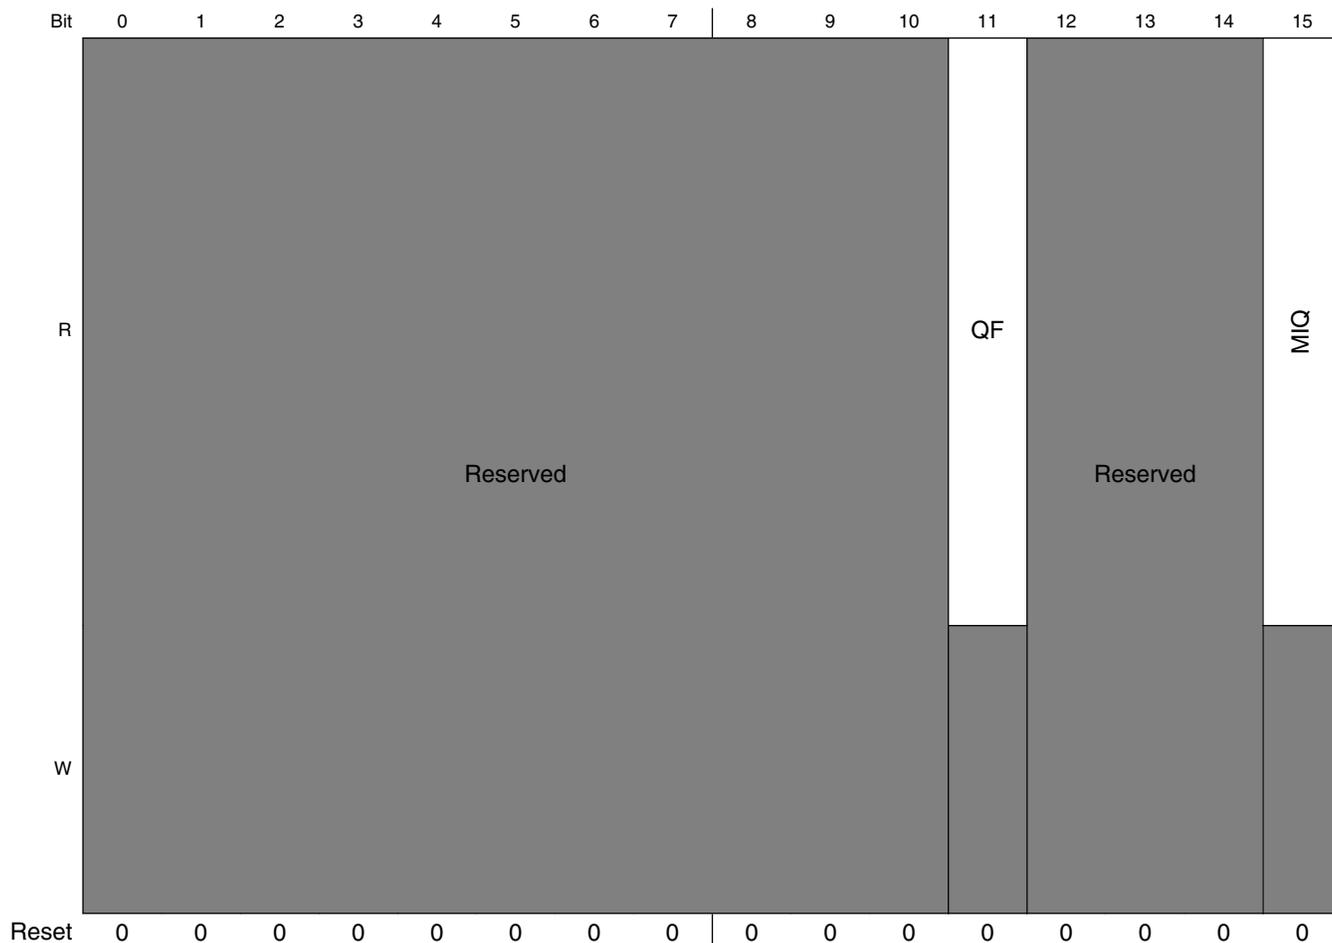
SRIO_IMnMR field descriptions (continued)

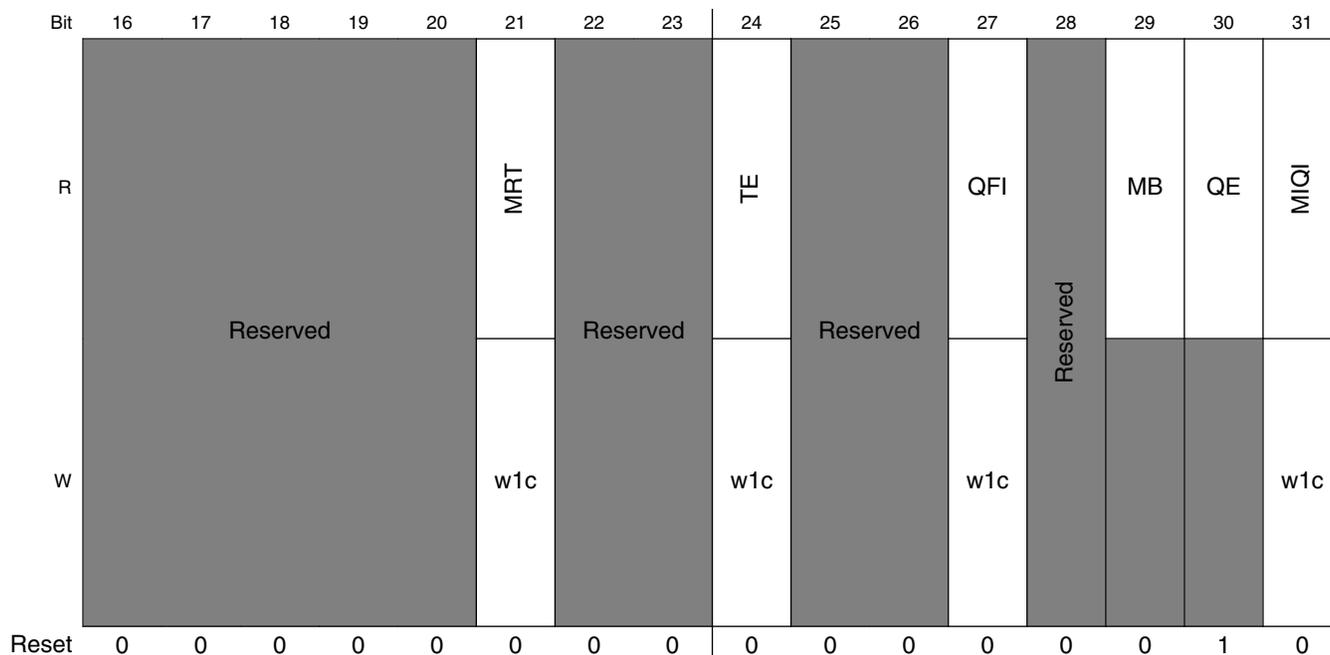
Field	Description
	1011 Reserved ... 1111 Reserved
20–22 -	This field is reserved. Reserved
23 QFIE	Queue full interrupt enable. When set, generates an interrupt when the queue is full (that is, the enqueue and dequeue pointers are equal after the dequeue pointer was incremented by the mailbox controller). No QFI interrupt is generated if this if this bit is cleared. If this bit is set and IM n SR[QF] is a 1, IM n SR[QFI] asserts.
24 -	This field is reserved. Reserved
25 MIQIE	Message in queue interrupt enable. When set, generates an interrupt when the queue has accumulated the number of messages specified by the IM n MR[MIQ_THRESH]. No MIQ interrupt is generated if this bit is cleared. If this bit is set and IM n SR[MIQ] is a 1, IM n SR[MIQI] asserts. If this bit is set and IM n MR[MI] is also set simultaneously, IM n SR[MIQI] indicates reflects the value of MIQ after the increment.
26 EIE	Error interrupt enable. When set, generates the port-write/error interrupt when a transfer error (IM n SR[TE]) or a message request time-out (IM n SR[MRT]) event occurs. No port-write/error interrupt is generated if this bit is cleared.
27–29 -	This field is reserved. Reserved
30 MI	Mailbox increment. Software sets this bit after processing an inbound message. Hardware increments the IM n FQDPAR and clears this bit. Always reads as 0.
31 ME	Mailbox enable. If this bit is set, the mailbox is initialized and can service incoming message operations. If this bit is cleared after the first segment of a multi-segment message has arrived, a message request time-out results (IM n SR[MRT]) and the busy bit (IM n SR[MB]) clears as long as the port response timer value (PRTOCCSR[TV]) is not set to the disabled value. If the port response timer value is set to the disabled value, the busy bit will not clear.

15.5.112 Inbound message n status register (SRIO_IMnSR)

The inbound message status register reports various mailbox conditions during and after a message operation. Writing a 1 to the corresponding set bit clears the bit.

Address: C_0000h base + 1_3064h offset + (256d × i), where i=0d to 1d





SRIO_IMnSR field descriptions

Field	Description
0–10 -	This field is reserved. Reserved
11 QF	Queue full. If the queue becomes full, then this bit will be set. This bit is cleared if the message controller is disabled. (Read-only)
12–14 -	This field is reserved. Reserved
15 MIQ	Message-In-Queue. If the queue has accumulated the number of messages specified by the IM <i>n</i> MR[MIQ_THRESH], then this bit will be set. This bit is cleared if the message controller is disabled. (Read-only)
16–20 -	This field is reserved. Reserved
21 MRT	Message request time-out. This bit is set when the message unit has not received another message segment for a multi segment message and a time-out occurred. This bit is cleared by writing a 1.
22–23 -	This field is reserved. Reserved
24 TE	Transaction error. This bit is set when there is an internal error condition occurs during the message operation. This bit is cleared by writing a 1. For proper operation, this field should only be modified when the inbound message controller is not enabled.
25–26 -	This field is reserved. Reserved
27 QFI	Queue full interrupt. If the queue is full and if the QFIE bit in the mode register is set, then this bit will be set and an interrupt generated. This bit is cleared by writing a 1.
28 -	This field is reserved. Reserved
29 MB	Mailbox busy. When set, indicates that a message operation is currently in progress. This bit is cleared as a result of an error or the message operation is finished. (Read-only)

Table continues on the next page...

SRIO_IMnSR field descriptions (continued)

Field	Description
30 QE	Queue empty. If the queue is empty, then this bit is set. This bit is also set when the message controller is disabled. (Read-only)
31 MIQI	Message-In-Queue interrupt. If the queue has accumulated the number of messages specified by the IM <i>n</i> MR[MIQ_THRESH] and if the MIQIE bit in the Mode Register is set, then this bit is set and an interrupt generated. This bit is cleared by writing a 1.

15.5.113 Extended inbound message *n* frame queue dequeue pointer address register (SRIO_EIMnFQDPAR)

Used with IM *n* FQDPAR; see description in [Inbound message *n* frame queue dequeue pointer address register](#) .

Address: C_0000h base + 1_3068h offset + (256d × *i*), where *i*=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															EFQDPA																
W	Reserved															EFQDPA																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_EIMnFQDPAR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EFQDPA	Extended frame queue dequeue pointer address bits. These are the highest order address bits.

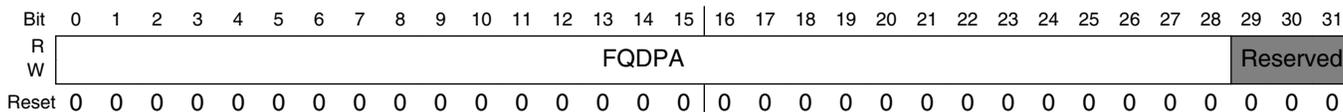
15.5.114 Inbound message *n* frame queue dequeue pointer address register (SRIO_IMnFQDPAR)

The inbound message frame queue dequeue pointer address registers contain the address for the first message in memory to be processed. Software must initialize these registers to the first frame location in memory. When a message has been processed, the processor sets IM *n* MR[MI]. The mailbox hardware then increments IM *n* FQDPAR and EIM *n* FQDPAR to point to the next frame in memory and clears the IM *n* MR[MI] bit. If the inbound message frame queue enqueue pointer and the inbound message frame queue dequeue pointer are not equal (indicating that the queue is not empty), the processor can read the next message frame from memory for processing. If the enqueue and dequeue pointers are equal after being incremented by the processor, the queue is empty and all outstanding messages have been processed.

When software initializes these registers, the queue to which they point must be aligned on a boundary equal to the number of queue entries × frame size. For example, if there are eight entries in the queue and the frame size is 128 bytes, the queue must be 1024-byte aligned.

The number of queue entries is set in IM *n* MR[CIRQ_SIZ], and the frame size is set in IM *n* MR[FRM_SIZ]. See [Inbound message n mode register](#) .

Address: C_0000h base + 1_306Ch offset + (256d × i), where i=0d to 1d



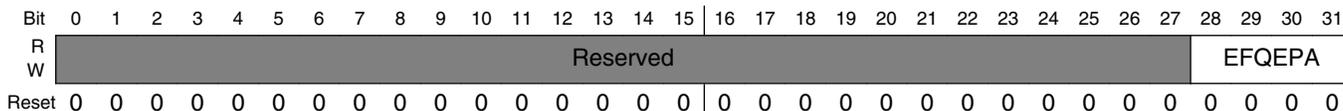
SRIO_IMnFQDPA field descriptions

Field	Description
0–28 FQDPA	Frame queue dequeue pointer address. Contains the address of the first message in memory to process.
29–31 -	This field is reserved. Reserved

15.5.115 Extended inbound message n frame queue enqueue pointer address register (SRIO_EIMnFQEPA)

Used with IM *n* FQEPA; see description in [Inbound message n frame queue enqueue pointer address register](#) .

Address: C_0000h base + 1_3070h offset + (256d × i), where i=0d to 1d



SRIO_EIMnFQEPA field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EFQEPA	Extended frame queue enqueue pointer address bits. These are the highest order address bits. For proper operation, this field should only be modified when the inbound message controller is not enabled.

15.5.116 Inbound message n frame queue enqueue pointer address register (SRIO_IMnFQEPAR)

The inbound message frame queue enqueue pointer address registers contain the address for the next message frame in memory to be added to the queue. Software must initialize these registers to match the frame queue dequeue pointer address. When a message is received by the mailbox controller, it writes the message data to the next location in the queue (indicated by the address in IM *n* FQEPAR and EIM *n* FQEPAR) and then increments IM *n* FQEPAR and EIM *n* FQEPAR to point to the next frame location in memory. This can result in a number of actions:

- If the enqueue and dequeue pointers match, the queue is now full and the mailbox controller does not accept any more incoming messages, returning RETRY responses to the sending devices until the queue is no longer full. If the IM *n* MR[QFIE] bit is set, then the IM *n* MR[QFI] bit is set and an interrupt is generated.
- If the enqueue and dequeue pointers were the same before the register was read, the queue has transitioned from empty to not empty. If IM *n* MR[MIQIE] is set, IM *n* SR[MIQI] is set, and an interrupt is generated.

When software initializes these registers, the queue to which they point must be aligned on a boundary equal to the number of queue entries x frame size. For example, if there are eight entries in the queue and the frame size is 128 bytes, the queue must be 1024-byte aligned.

The number of queue entries is set in IM *n* MR[CIRQ_SIZ], and the frame size is set in IM *n* MR[FRM_SIZ]. See [Inbound message n mode register](#) .

Address: C_0000h base + 1_3074h offset + (256d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	FQEPA															Reserved																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_IMnFQEPAR field descriptions

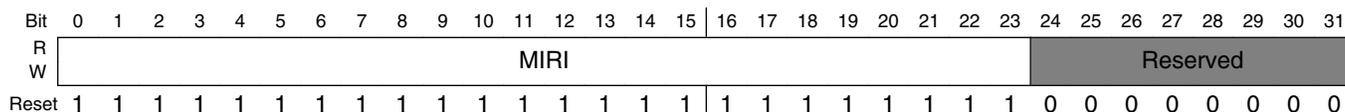
Field	Description
0–28 FQEPA	Frame queue enqueue pointer address. Contains the address of the next message frame to be added to the queue. For proper operation, this field should only be modified when the inbound message controller is not enabled.
29–31 -	This field is reserved. Reserved

15.5.117 Inbound message n maximum interrupt report interval register (SRIO_IMnMIRIR)

The maximum interrupt interval register contains a time-out timer value to define the maximum amount of time that the mailbox controller is to wait from transitioning from not empty to signaling an interrupt (if enabled) to the local processor if the IM *n* MR[MIQ_THRESH] limit is not reached. The reset, or default, value of this counter is the maximum interrupt interval .

The resolution of this timer is 168/(platform frequency). For example, at a platform frequency of 400 MHz, the maximum interval value is 0xFF_FFFF * 168/400MHz = 7.05 seconds.

Address: C_0000h base + 1_3078h offset + (256d × i), where i=0d to 1d



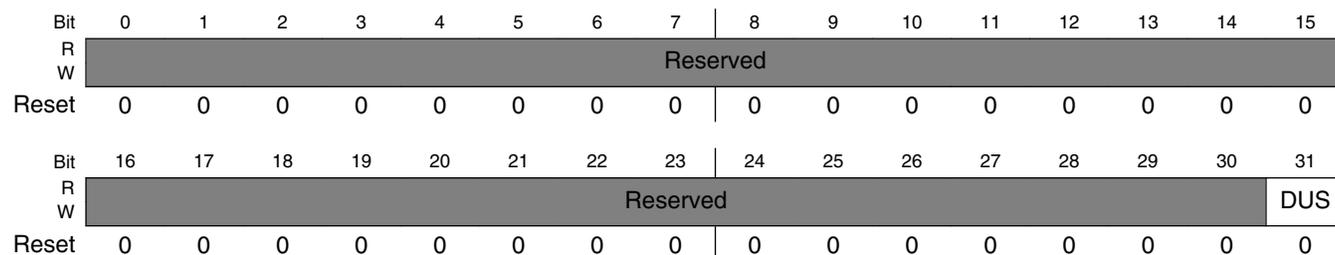
SRIO_IMnMIRIR field descriptions

Field	Description
0–23 MIRI	Maximum interrupt report interval. Maximum message-in-queue to interrupt generation time. A value of 0 disables the time-out timer. For proper operation, this field should only be modified when the inbound message controller is not enabled.
24–31 -	This field is reserved. Reserved

15.5.118 Outbound doorbell mode register (SRIO_ODMR)

The outbound mode register allows software to start a doorbell operation and to control various doorbell operation characteristics.

Address: C_0000h base + 1_3400h offset = D_3400h



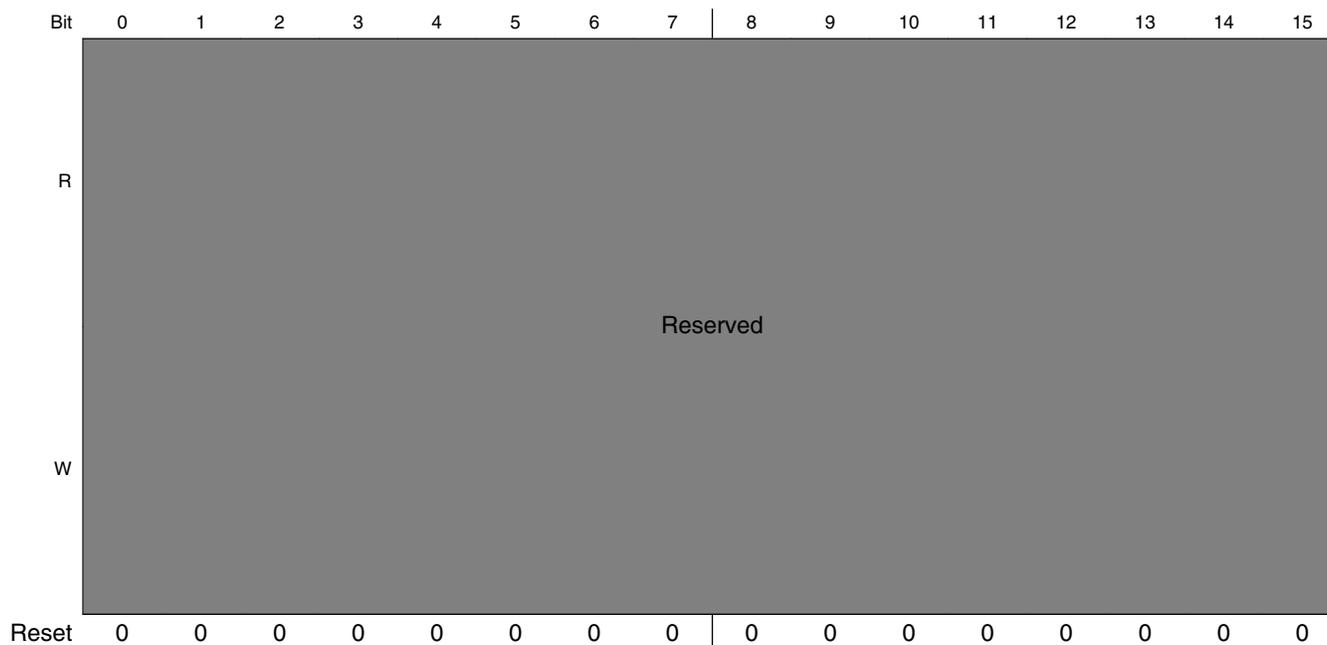
SRIO_ODMR field descriptions

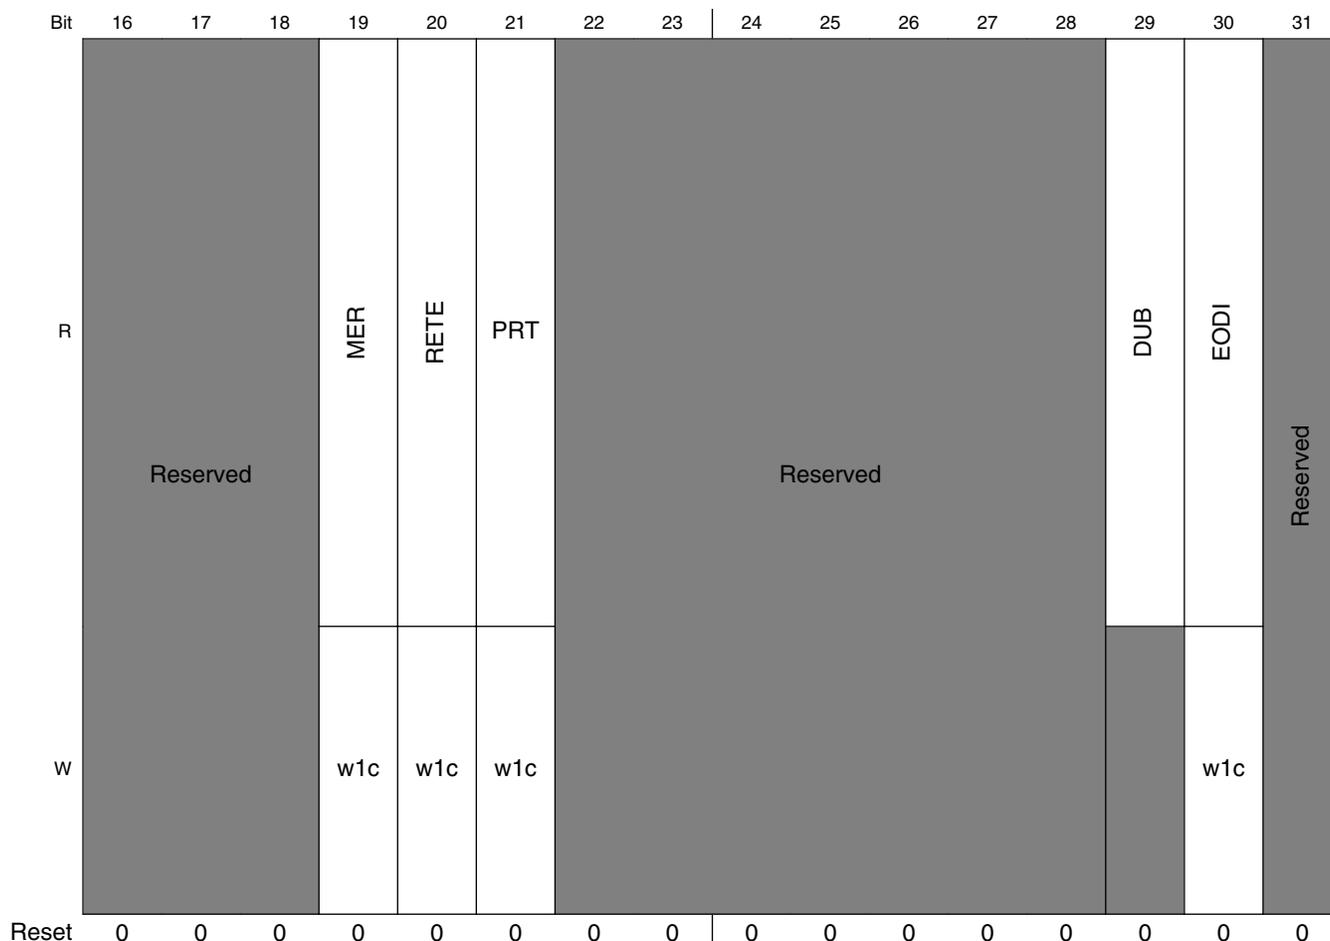
Field	Description
0–30 -	This field is reserved. Reserved
31 DUS	Doorbell unit start. Direct mode - A 0 to 1 transition when the doorbell unit is not busy (DUB bit is 0) starts the doorbell unit. A 1 to 0 transition has no effect.

15.5.119 Outbound doorbell status register (SRIO_ODSR)

The outbound status register reports various doorbell unit conditions during and after a doorbell operation. For some bits, writing a 1 clears the bit.

Address: C_0000h base + 1_3404h offset = D_3404h





SRIO_ODSR field descriptions

Field	Description
0–18 -	This field is reserved. Reserved
19 MER	Message error response. This bit is set when an ERROR response is received from the doorbell target. The Error response received field indicates value of the error response status bits when an error response is received. This bit is cleared by writing a 1. For proper operation, this bit should only be cleared when a doorbell operation is not in progress.
20 RETE	Retry error threshold exceeded. This bit is set when the doorbell unit has been unable to complete a doorbell operation because the retry error threshold value has been exceeded due to RapidIO retry response. This bit is cleared by writing a 1. For proper operation, this bit should only be cleared when a doorbell operation is not in progress.
21 PRT	Packet response time-out. This bit is set when the doorbell unit has been unable to complete a doorbell operation and a packet response time-out occurred. This bit is cleared by writing a 1. For proper operation, this bit should only be cleared when a doorbell operation is not in progress.
22–28 -	This field is reserved. Reserved
29 DUB	Doorbell unit busy. When set, indicates that a doorbell operation is currently in progress. This bit is cleared as a result of an error or when the doorbell operation is finished. (Read-only)

Table continues on the next page...

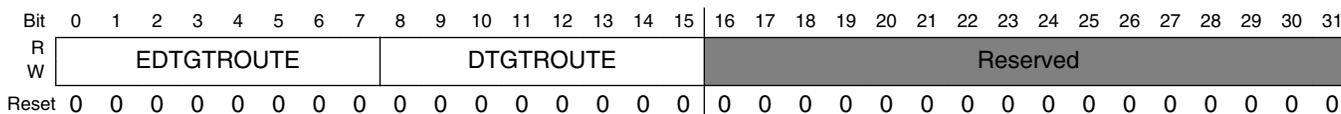
SRIO_ODSR field descriptions (continued)

Field	Description
30 EODI	End-of-doorbell interrupt. After finishing this doorbell operation, if the EODIE bit in the destination attributes register is set, then this bit will be set and an interrupt generated. This bit is cleared by writing a 1. For proper operation, this bit should only be cleared when a doorbell operation is not in progress.
31 -	This field is reserved. Reserved

15.5.120 Outbound doorbell n destination port register (SRIO_ODDPR)

The destination port register indicates the RapidIO Destination ID and mailbox to which the doorbell unit controller is to send data. The software must ensure that this is a valid port in the receiving device.

Address: C_0000h base + 1_3418h offset = D_3418h



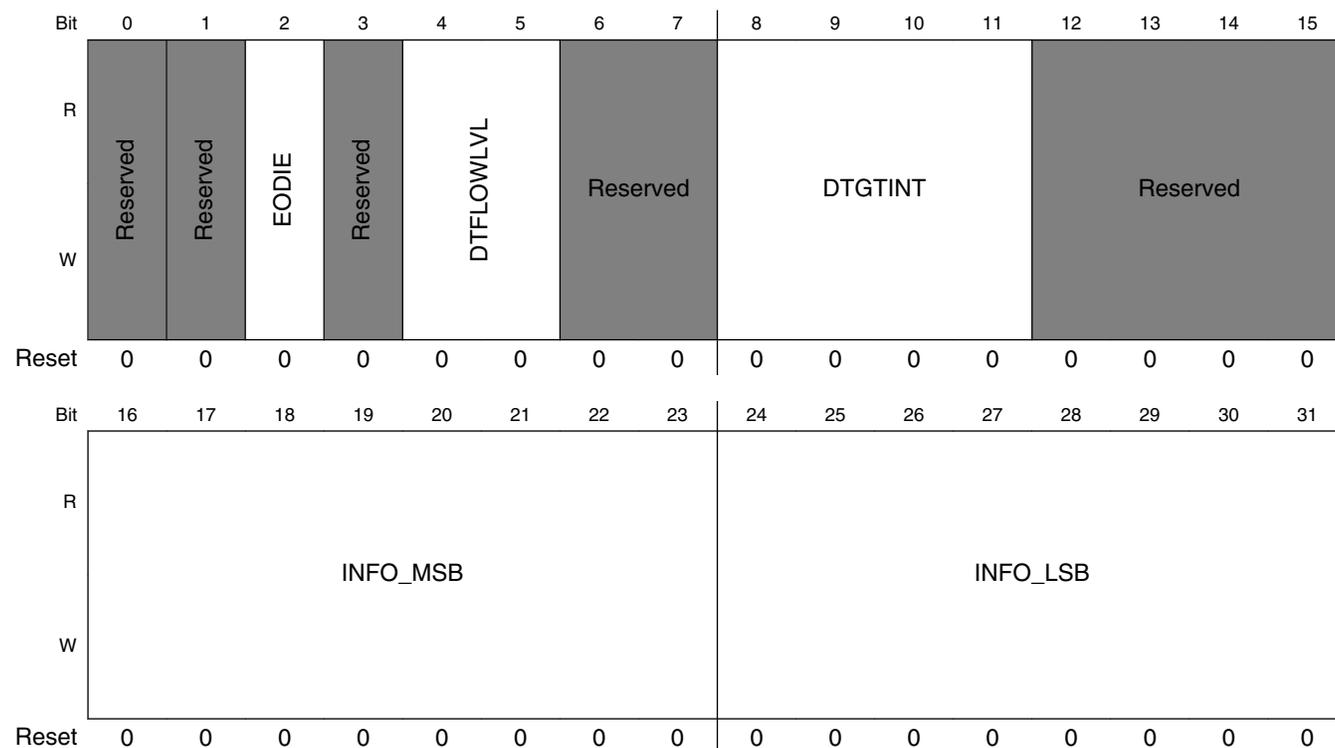
SRIO_ODDPR field descriptions

Field	Description
0–7 EDTGTRROUTE	Extended destination target route. Most significant byte of a 16-bit target route when operating in large transport mode. Reserved when operating in small transport mode. For proper operation, this field should only be modified when a doorbell operation is not in progress.
8–15 DTGTRROUTE	Destination target route. Contains the target route field of the transaction (Device ID of the target). For proper operation, this field should only be modified when a doorbell operation is not in progress.
16–31 -	This field is reserved. Reserved

15.5.121 Outbound doorbell n destination attributes register (SRIO_ODDATR)

The outbound destination attributes register contains the transaction attributes to be used for the doorbell operation.

Address: C_0000h base + 1_341Ch offset = D_341Ch



SRIO_ODDATR field descriptions

Field	Description
0 -	This field is reserved. Reserved
1 -	This field is reserved. Reserved, hard wired to 0
2 EODIE	End-of-Doorbell interrupt enable. When set, generates an interrupt when the current doorbell operation is finished. For proper operation, this field should only be modified when a doorbell operation is not in progress.
3 -	This field is reserved. Reserved
4-5 DTFLOWLVL	Transaction flow level. For proper operation, this field should only be modified when a doorbell operation is not in progress.

Table continues on the next page...

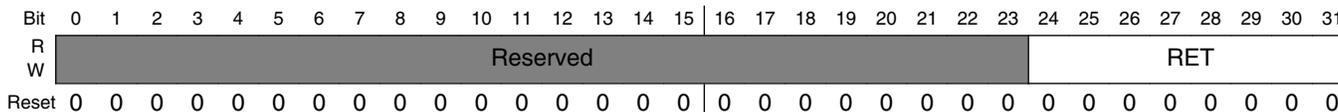
SRIO_ODDATR field descriptions (continued)

Field	Description
	00 Lowest priority transaction flow 01 Next highest priority transaction flow 10 Highest priority transaction flow 11 Reserved
6-7 -	This field is reserved. Reserved
8-11 DTGTINT	Target interface. The value of this field should be set according to the following: 0000 RapidIO Port 1 0001 RapidIO Port 2
12-15 -	This field is reserved. Reserved
16-23 INFO_MSB	Most significant byte of the doorbell "info" field. For proper operation, this field should only be modified when a doorbell operation is not in progress.
24-31 INFO_LSB	Least significant byte of the doorbell "info" field. For proper operation, this field should only be modified when a doorbell operation is not in progress.

15.5.122 Outbound doorbell n retry error threshold configuration register (SRIO_ODRETCR)

The retry error threshold configuration register controls the number of times that the doorbell unit attempts to complete a doorbell operation before reporting an error and moving on to the next task, if one is available. Failures to complete an operation are indicated by receiving a RapidIO logical layer RETRY response from the target. If the programmed count is exceeded, the ODSR[RETE] bit is set.

Address: C_0000h base + 1_342Ch offset = D_342Ch



SRIO_ODRETCR field descriptions

Field	Description
0-23 -	This field is reserved. Reserved
24-31 RET	Retry error threshold. This value is the number of times that the doorbell unit attempts to transmit a doorbell. For proper operation, this field should only be modified when a doorbell operation is not in progress. 0x00 Disabled

Table continues on the next page...

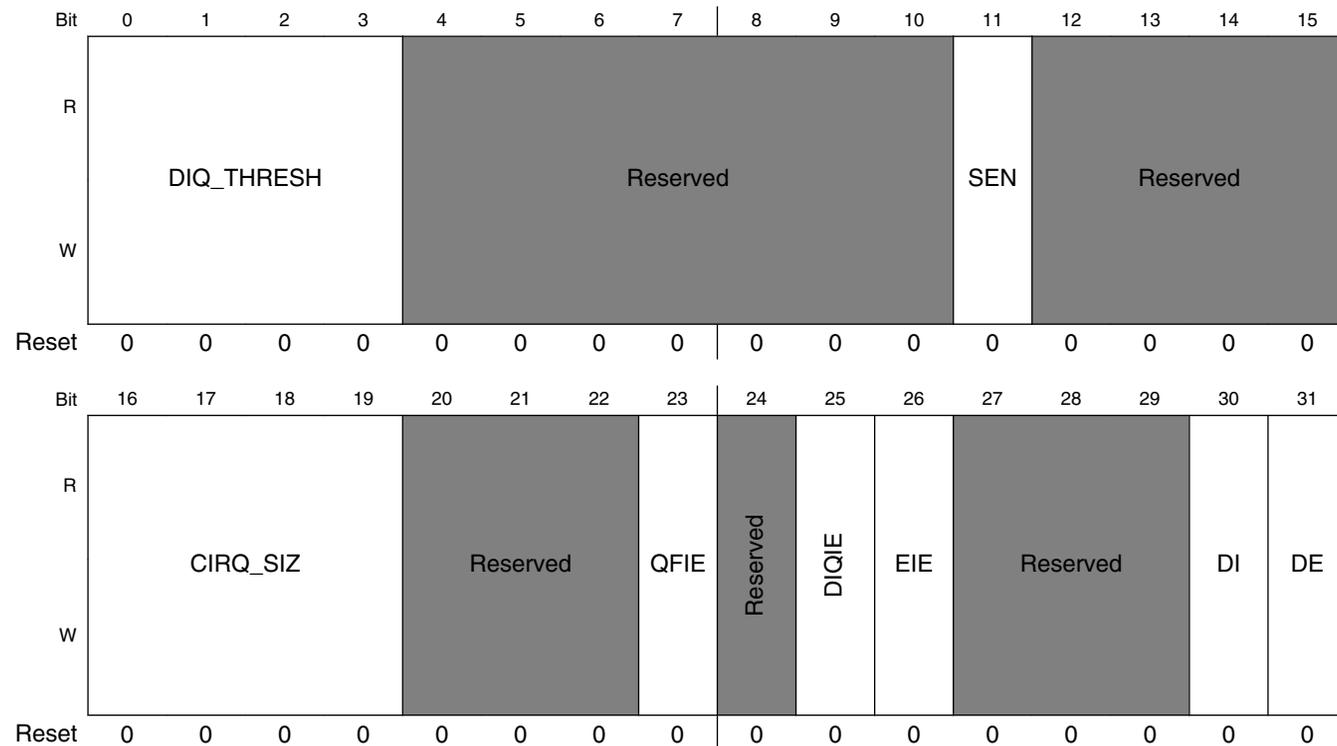
SRIO_ODRETCR field descriptions (continued)

Field	Description
0x01	Doorbell transmitted only 1 time
0x02	Doorbell transmitted up to 2 times
...	
0xFF	Doorbell transmitted up to 255 times

15.5.123 Inbound doorbell n mode register (SRIO_IDMR)

The doorbell mode register allows software to enable the doorbell controller to control various doorbell operation characteristics.

Address: C_0000h base + 1_3460h offset = D_3460h



SRIO_IDMR field descriptions

Field	Description
0–3 DIQ_THRESH	Doorbell-in-queue threshold. Determines the number of doorbells to be accumulated in the doorbell queue before the Doorbell-in-queue bit is set (IDSR[DIQ]). Undefined operation results if the actual number of entries defined by the doorbell-in-queue threshold is set greater than or equal to the actual size of the doorbell queue. For proper operation, this field should only be modified when the doorbell controller is not enabled.

Table continues on the next page...

SRIO_IDMR field descriptions (continued)

Field	Description
	0000 1 0001 2 0010 4 0011 8 0100 16 0101 32 0110 64 0111 128 1000 256 1001 512 1010 1024 1011 reserved ... 1111 reserved
4–10 -	This field is reserved. Reserved
11 SEN	Snoop enable. When set enables snooping the local processor when writing messages into memory. For proper operation, this field should only be modified when the doorbell controller is not enabled.
12–15 -	This field is reserved. Reserved
16–19 CIRQ_SIZ	Circular doorbell queue size. Determines the number of doorbell entries that can be place on the circular queue. CIRQ_SIZ x 8 bytes determine the maximum contiguous memory space allocated to the doorbell unit. For proper operation, this field should only be modified when the doorbell controller is not enabled. For proper operation, this field should only be modified when the doorbell controller is not enabled. 0000 2 0001 4 0010 8 0011 16 0100 32 0101 64 0110 128 0111 256 1000 512 (4-kbyte page boundary) 1001 1024 1010 2048 1011 reserved ... 1111 reserved
20–22 -	This field is reserved. Reserved
23 QFIE	Queue full interrupt enable. If this bit is set and ID SR[QF] is a 1, ID SR[QFI] is asserted. 0 No QF interrupt is generated 1 Generates an interrupt when the queue is full (that is, the enqueue and dequeue pointers are equal after the dequeue pointer was incremented by the doorbell controller).

Table continues on the next page...

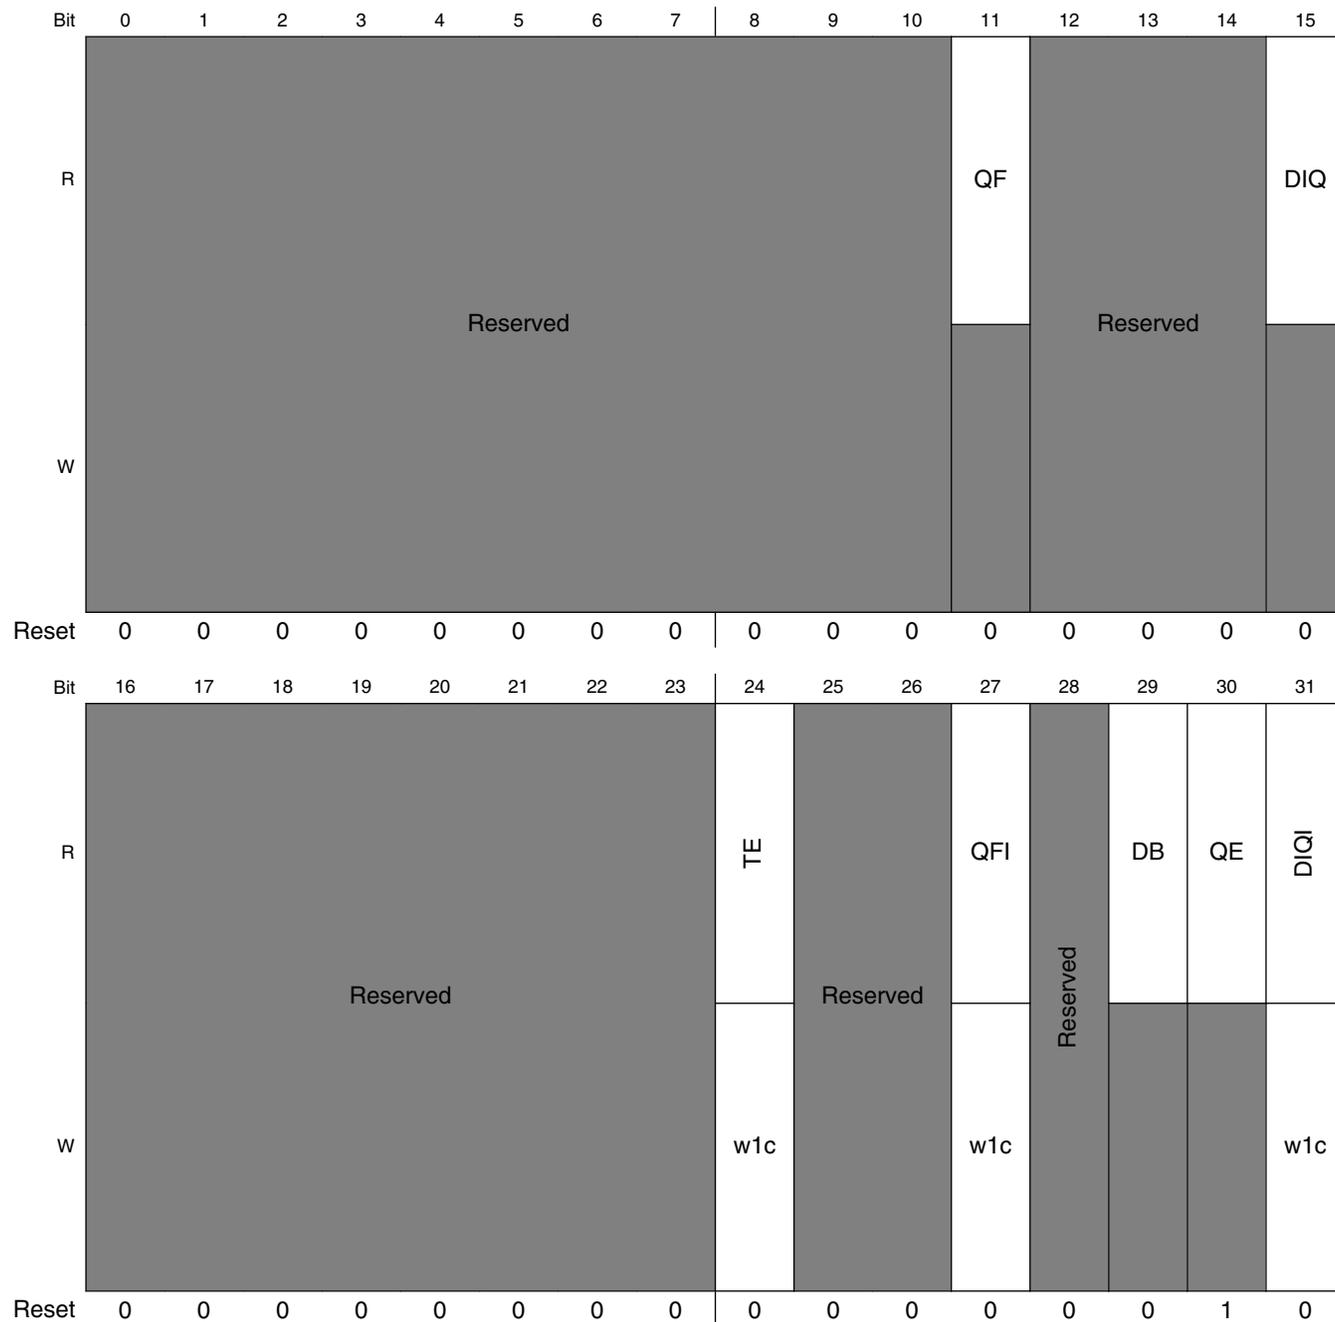
SRIO_IDMR field descriptions (continued)

Field	Description
24 -	This field is reserved. Reserved
25 DIQIE	Doorbell in queue interrupt enable. When set, allows the doorbell in queue interrupt to assert. The doorbell in queue interrupt cannot be asserted if this bit is cleared. Doorbell in queue interrupt enable. If this bit is set and ID SR[DIQ] is a 1, ID SR[DIQ] will assert. If this bit is set and ID MR[DI] is also set simultaneously, ID SR[DIQ] indicates reflects the value of DIQ after the increment. 0 No DIQ interrupt is generated 1 Generates an interrupt when the DIQ bit is set (ID SR[DIQ])
26 EIE	Error interrupt enable. When set, generates the port-write/error interrupt when a transfer error (ID SR[TE]) event occurs. No port-write/error interrupt is generated if this bit is cleared.
27–29 -	This field is reserved. Reserved
30 DI	Doorbell increment. Software sets this bit after processing an inbound doorbell. Hardware then increments the OD QEPAR and clears this bit. Always reads as 0.
31 DE	Doorbell enable. 0 Doorbell disabled 1 The doorbell has been initialized and can service incoming doorbell operations.

15.5.124 Inbound doorbell n status register (SRIO_IDSR)

The doorbell status register reports various doorbell conditions after a doorbell operation. Writing a 1 to the corresponding set bit will clear the bit.

Address: C_0000h base + 1_3464h offset = D_3464h



SRIO_IDSR field descriptions

Field	Description
0–10 -	This field is reserved. Reserved
11 QF	Queue full. If the queue is full, then this bit is set. This bit clears if the doorbell controller is disabled. (Read-only)
12–14 -	This field is reserved. Reserved
15 DIQ	Doorbell-In-Queue. If the queue has accumulated the number of doorbells specified by DIQ_THRESH, then this bit is set. Also, if a valid entry pointed to by the dequeue address pointers has not been serviced within the configured maximum interval, this bit is set. This bit clears if the doorbell controller is disabled. (Read-only)
16–23 -	This field is reserved. Reserved
24 TE	Transaction error. This bit is set when an internal error occurs during the doorbell operation. (Bit reset, write 1 clear). For proper operation, this field should only be modified when the doorbell controller is not enabled.
25–26 -	This field is reserved. Reserved
27 QFI	Queue full interrupt. If the queue becomes full and the QFIE bit in the Mode Register is set, then this bit is set and an interrupt generated. This bit is cleared by writing a 1.
28 -	This field is reserved. Reserved
29 DB	Doorbell busy. When set, indicates that a doorbell has been received and the doorbell queue is currently being written. This bit clears when the write to memory has finished. Disabling the doorbell controller does not affect this bit. (Read-only)
30 QE	Queue empty. If the queue is empty, then this bit is set. This bit is also set if the doorbell controller is disabled. (Read-only)
31 DIQI	Doorbell-In-Queue interrupt. If DIQ is set, then this bit is set and an interrupt generated. This bit is cleared by writing a 1.

15.5.125 Extended inbound doorbell n queue dequeue pointer address register (SRIO_EIDQDPAR)

Used with ID QDPAR; see description in [Inbound doorbell n queue dequeue Pointer address register](#) .

Address: C_0000h base + 1_3468h offset = D_3468h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															EDQDPA																
W	Reserved															EDQDPA																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_EIDQDPAR field descriptions

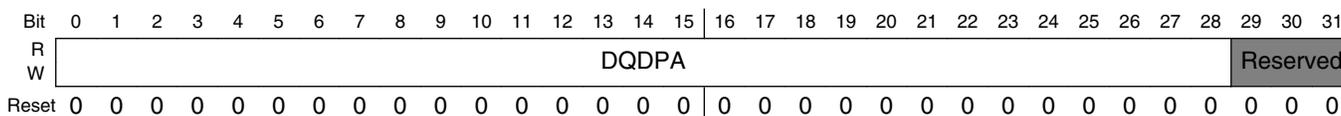
Field	Description
0–27 -	This field is reserved. Reserved
28–31 EDQDPA	Extended doorbell queue dequeue pointer address bits. These are the highest order address bits.

15.5.126 Inbound doorbell n queue dequeue Pointer address register (SRIO_IDQDPAR)

The inbound doorbell queue dequeue pointer address registers (EID QDPAR and ID QDPAR) contain the double-word address for the first doorbell in memory to be processed. Software must initialize these registers to the first doorbell entry location in memory. When the current doorbell has been processed, the processor sets ID MR[DI]. The doorbell hardware then increments the doorbell queue dequeue pointer address register to point to the next doorbell entry in memory and clears ID MR[DI]. When processing multiple doorbells, instead of setting ID MR[DI] for each doorbell, the processor can write these registers directly.

If the doorbell queue enqueue pointer and the doorbell queue dequeue pointer are not equal (indicating that the queue is not empty), the processor can read the next doorbell from memory for processing. If the enqueue and dequeue pointers are equal after being incremented by the processor, the queue is empty and all outstanding doorbells have been processed.

Address: C_0000h base + 1_346Ch offset = D_346Ch



SRIO_IDQDPAR field descriptions

Field	Description
0–28 DQDPA	Doorbell queue dequeue pointer address. Contains the double-word address of the first doorbell in memory to process. Note that this base address must be queue-size aligned.
29–31 -	This field is reserved. Reserved

15.5.127 Extended inbound doorbell n queue enqueue pointer address register (SRIO_EIDQEPAR)

Used with ID QEPAR; see description in [Inbound doorbell n Queue enqueue pointer address register](#) .

Address: C_0000h base + 1_3470h offset = D_3470h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved																			EDQEPA												
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_EIDQEPAR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EDQEPA	Doorbell queue enqueue pointer address bits. These are the highest order address bits. This field can only be written when the doorbell controller is disabled or undefined operation results.

15.5.128 Inbound doorbell n Queue enqueue pointer address register (SRIO_IDQEPAR)

The doorbell queue enqueue pointer address registers (EID QEPAR and ID QEPAR) contain the double-word address for the next doorbell entry in memory to be added to the queue. Software must initialize ID QEPAR and EID QEPAR to match the doorbell queue dequeue pointer address. When a doorbell packet is received by the doorbell controller, it writes the doorbell information to the next location in the queue (indicated by the address in ID QEPAR and EID QEPAR) and then increments ID QEPAR and EID QEPAR to point to the next doorbell location in memory. This can result in a number of actions:

- If the enqueue and dequeue pointers match, then the queue is now full and the doorbell controller does not accept any more incoming doorbell packets, returning RETRY responses to the sending devices until the queue is no longer full. If the ID MR[QFIE] bit is set, then the ID SR[QFI] is set and the interrupt is generated.
- If the enqueue and dequeue pointers were the same before receiving the doorbell, the queue has transitioned from empty to not empty. When the number of doorbells received matches the configured threshold, the ID SR[DIQ] bit is set. If the DRM[DIQIE] bit is set, then the ID SR[DIQI] bit is also set and the inbound doorbell interrupt is generated.

Serial RapidIO Memory Map/Register Definition

Address: C_0000h base + 1_3474h offset = D_3474h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	DQEPA																Reserved															
W	DQEPA																Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_IDQEPAR field descriptions

Field	Description
0–28 DQEPA	Doorbell queue enqueue pointer address. Contains the double-word address of the next doorbell location to be added to the queue. This field can only be written when the doorbell controller is disabled or undefined operation results. Note that this base address must be queue-size aligned.
29–31 -	This field is reserved. Reserved

15.5.129 Inbound doorbell n maximum interrupt report interval register (SRIO_IDMIRIR)

The maximum interrupt interval register contains a time-out timer value to define the maximum amount of time that a doorbell entry can be at the head of the doorbell queue before generating an interrupt (if enabled) if the DIQ_THRESH limit is not reached. The reset, or default, value of this counter is the maximum interrupt interval .

The resolution of this timer is $168 / (\text{platform frequency})$. For example, at a platform frequency of 400 MHz, the maximum interval value is $0xFF_FFFF * 168 / 400\text{MHz} = 7.05$ seconds.

Address: C_0000h base + 1_3478h offset = D_3478h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	MIRI																Reserved															
W	MIRI																Reserved															
Reset	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0

SRIO_IDMIRIR field descriptions

Field	Description
0–23 MIRI	Maximum interrupt report interval-Maximum doorbell-in-queue to interrupt generation time. A value of 0 disables the timer. This field can only be written when the doorbell controller is disabled or undefined operation results.
24–31 -	This field is reserved. Reserved

15.5.130 Inbound port-write n mode register (SRIO_IPWMR)

The port-write mode register allows software to enable the port-write controller and to control various port-write operation characteristics.

Address: C_0000h base + 1_34E0h offset = D_34E0h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved										SEN	Reserved				
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved							QFIE	Reserved		EIE	Reserved			CQ	PWE
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRIO_IPWMR field descriptions

Field	Description
0–10 -	This field is reserved. Reserved
11 SEN	Snoop enable. When set enables snooping the local processor when writing port-write data payload into memory. For proper operation, this field should only be modified when the port-write controller is not enabled.
12–22 -	This field is reserved. Reserved
23 QFIE	Queue full interrupt enable. When set, generates the error/port-write interrupt if the queue is full (that is, the controller has written the port-write data payload into memory). No error/port-write interrupt is generated if this bit is cleared. For proper operation, this field should only be modified when the port-write controller is not enabled.
24–25 -	This field is reserved. Reserved
26 EIE	Error Interrupt enable. When set, generates the port-write/error interrupt when a transfer error (IPW0SR[TE]) event occurs. No port-write/error interrupt is generated if this bit is cleared.
27–29 -	This field is reserved. Reserved
30 CQ	Clear queue. Software sets this bit after processing an inbound port-write operation. Hardware clears the queue full bit (IPWSR[QF]), clears this bit, and allows another port-write to be received. This bit is always read as a 0.
31 PWE	Port write enable. If this bit is set the port-write controller has been initialized and can service an incoming port-write operation.

15.5.131 Inbound port-write n status register (SRIO_IPWSR)

The port-write status register reports various port-write conditions.

Address: C_0000h base + 1_34E4h offset = D_34E4h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
R	Reserved										QF	Reserved					
W	Reserved											Reserved					
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	Reserved								TE	Reserved			QFI	PWD	PWB	Reserved	
W	Reserved								w1c	Reserved			w1c	w1c		Reserved	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

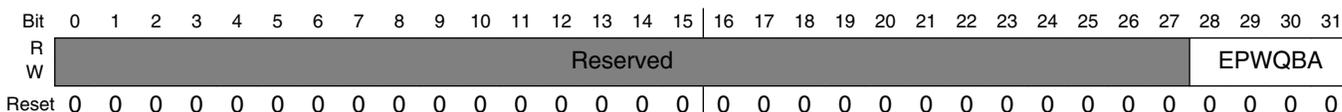
SRIO_IPWSR field descriptions

Field	Description
0–10 -	This field is reserved. Reserved
11 QF	Queue full. If the queue becomes full, then this bit is set. This bit is cleared when the clear queue bit is set (IPWMR[CQ]) and the queue is not full. This bit is also cleared if the port-write controller is disabled. (Read-only)
12–23 -	This field is reserved. Reserved
24 TE	Transaction error. This bit is set when there an internal error condition occurs during the port-write operation. (Bit reset, write 1 clear). For proper operation, this field should only be modified when the port-write controller is not enabled.
25–26 -	This field is reserved. Reserved
27 QFI	Queue full interrupt. If the queue becomes full and the QFIE bit in the Mode Register is set, then this bit is set and an interrupt generated. This bit is cleared by writing a 1.
28 PWD	Port-write discarded. This bit is set if a port-write is discarded while the port-write controller is enabled but busy. This bit is cleared by writing a 1.
29 PWB	Port-write busy (Read-only). 0 This bit is cleared after the port-write payload has been written to memory. 1 Indicates that a port-write has been received and the port-write payload is currently being written to memory. Disabling the port-write controller does not affect this bit.
30–31 -	This field is reserved. Reserved

15.5.132 Extended inbound port-write n queue base address register (SRIO_EIPWQBAR)

Used with IPW QBAR; see description in [Inbound port-write n queue base address register](#) .

Address: C_0000h base + 1_34E8h offset = D_34E8h



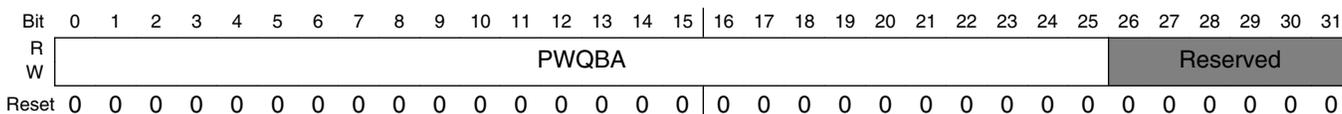
SRIO_EIPWQBAR field descriptions

Field	Description
0–27 -	This field is reserved. Reserved
28–31 EPWQBA	Extended port-write queue base address bits. These are the highest order address bits. This field can only be written when the port-write controller is disabled or undefined operation results.

15.5.133 Inbound port-write n queue base address register (SRIO_IPWQBAR)

The port-write queue base address registers (EIPW QBAR and IPW QBAR) contain the 64-byte cache line address for the port-write data payload. Software must initialize these registers to the desired location in memory.

Address: C_0000h base + 1_34ECh offset = D_34ECh



SRIO_IPWQBAR field descriptions

Field	Description
0–25 PWQBA	Port-write queue base address. Contains the cache line address of the port-write data payload. This field can only be written when the port-write controller is disabled or undefined operation results. When populating the address in this register it is not necessary to shift the address as the lower 6 bits are merely ignored to ensure cache line alignment.
26–31 -	This field is reserved. Reserved

SRIO_IPWQBAR field descriptions (continued)

Field	Description
-------	-------------

15.6 Serial RapidIO functional description

15.6.1 Segmented Outbound Window Description

All outbound windows have the capability to have 2 or 4 segments, all of which are equal in size, numbered 0 and 1, or 0 through 3, respectively (the standard 8540 unsegmented window definition becomes segment 0). Each segment assigns attributes and the target deviceID for an outbound transaction. All segments of a window translate to the same translation address in the target.

Additionally, each segment can be set up with 2, 4, or 8 subsegments, all of which are equal in size. These subsegments allow a single segment to target a number of numerically adjacent target device IDs, and, again, they all translate to the same translation address in the targets. For example, a segment with 8 subsegments can be configured to generate a transaction with the same set of attributes to target deviceIDs 0, 1, 2, 3, 4, 5, 6, or 7, depending on which subsegment is addressed.

Note that subsegments are only supported when multiple segments are chosen.

This allows a window to be configured so that aliases can be created to the same offset within the target device so that a single window can be used to generate different transaction types. Without segmented windows, achieving the equivalent behavior would require multiple windows. The figure below shows an example of this capability. A window is defined to be 4kB in size, and is defined to have 4 segments and no subsegments. Each segment is assigned to target deviceID 0x05, and each segment is given a different write transaction type attribute - segment 0 is assigned NWRITE, segment 1 is assigned SWRITE, segment 2 is assigned NWRITE_R, and segment 3 is assigned FLUSH. Since all of the segments are assigned to target the same device, by writing to the same offset in each segment, a different write transaction can be generated to the target to the same offset in the target.

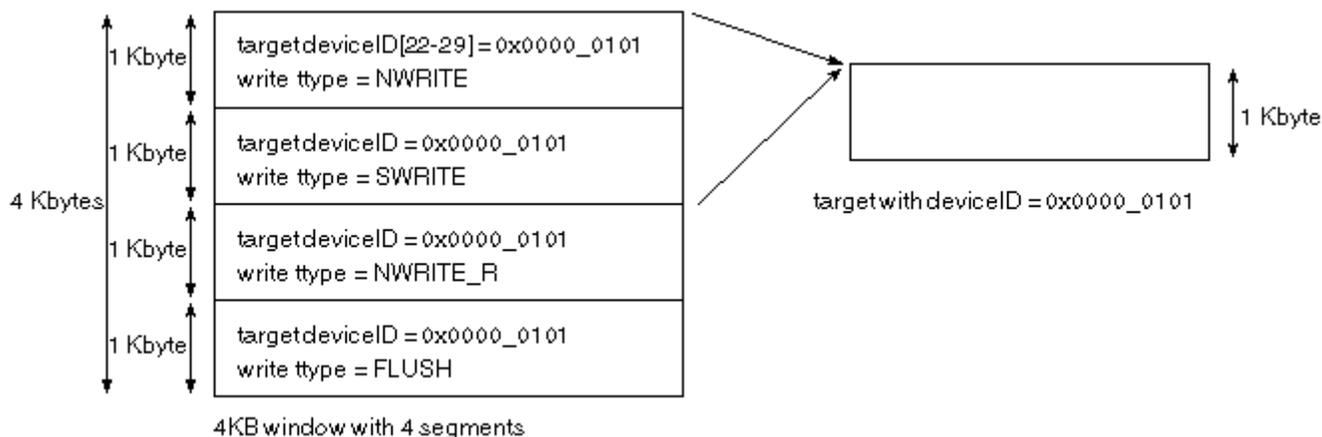


Figure 15-319. Example of Attribute Aliasing

So, writing to offset 0x0 in segment 0 is translated (as defined in the translation address registers) and generates a NWRITE transaction to offset 0x0 in a 1KB window in the target with deviceID = 0x05. A write to offset 0x0 in segment 2 is also translated, to the same offset in the target device as the write to segment 0, but this time a NWRITE_R transaction is generated.

Another use is that the same window can be used to target multiple devices with the same translation offset. Without segmented, (and subsegmented) windows, achieving the equivalent behavior would require multiple windows. The following figure shows an example of this multi-targeting. For example, a 4kB window is set up with 2 segments of 2 subsegments. Each segment is assigned a write ttype of NWRITE, but each segment and subsegment has a different target deviceID. Segments 0 and 1 are assigned target deviceIDs 4 and 5, and 8 and 9, respectively.

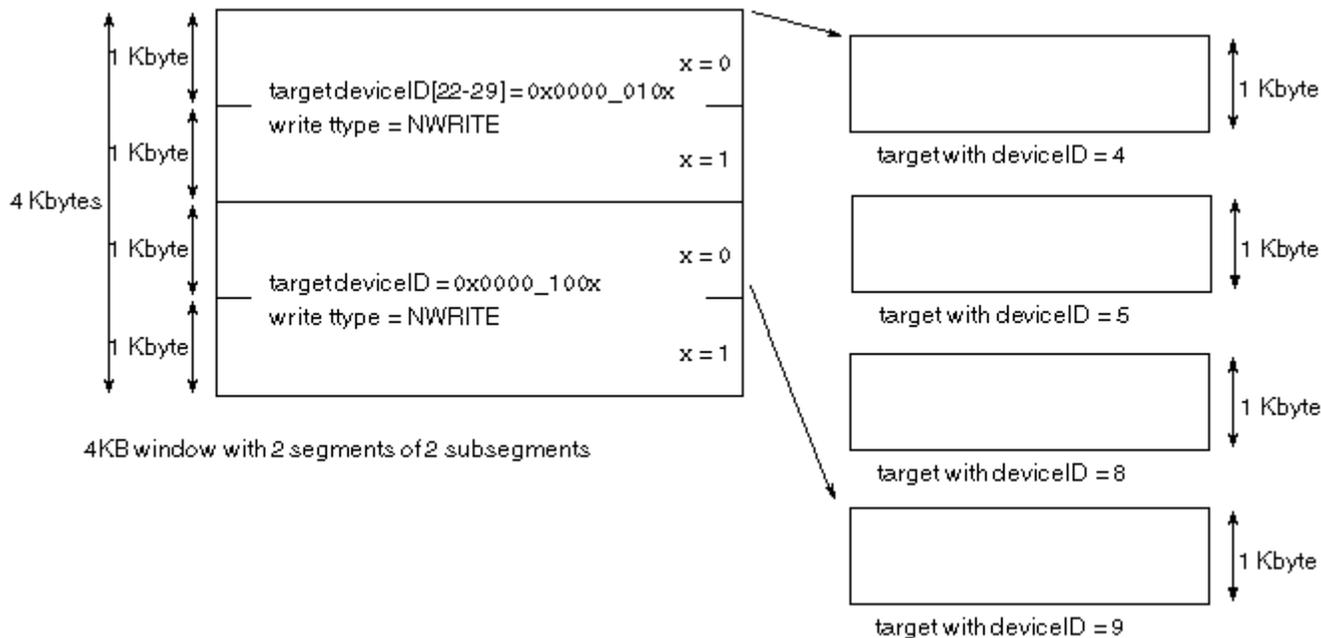


Figure 15-320. Example of Multi-Targeting

In this example, a write to offset 0x0 in segment 0 is translated as defined, and a NWRITE transaction is generated targeted to deviceID 4. A corresponding write to segment 1 to offset 0x400 is also translated but also using the assigned deviceID instead of the translation address bits [22-29]. The generated NWRITE transaction has the same target device offset as the write to segment 0, but is instead targeted to deviceID 9.

Combinations of aliasing and multi-targeting are also possible for a window.

15.6.2 RapidIO transaction summary

The RapidIO endpoint on this device supports all RapidIO I/O transactions, all RapidIO message passing transactions, and a few RapidIO GSM transactions.

Table 15-318. RapidIO I/O transactions

IO Transaction	ftype	ttype	Status	Description	
NREAD	0010	0100	NA	Read	
ATOMIC inc		1100		Read and post-increment with response	
ATOMIC dec		1101		Read and post-decrement with response	
ATOMIC set		1110		Read and set to all-1s with response	
ATOMIC clr		1111		Read and set to all-0s with response	
NWRITE	0101	0100	NA	Write with no response	
NWRITE_R		0101		Write with response	
SWRITE	0110	N/A		Streaming-Write	
MAINT read	1000	0000	NA	Maintenance read	
MAINT write		0001		Maintenance write	
MAINT read response		0010		0000	Done maintenance read response
				0111	Error response
MAINT write response		0011		0000	Done maintenance write response
				0111	Error response
MAINT port-write		0100		NA	Maintenance port-write ¹
RESPONSE without data	1101	0000	0000	I/O done response	
			0111	I/O error response	
RESPONSE with data		1000	0000	I/O done response with data	

1. Limited to inbound RapidIO packets only

Table 15-319. RapidIO message passing transactions

MSG Transaction	ftype	ttype	status	Description
DOORBELL	1010	NA	NA	Doorbell
MESSAGE	1011			Message
RESPONSE without data	1101	0000	0000	Doorbell done response
			0011	Doorbell retry response
			0111	Doorbell error response
		0001	0000	Message done response
			0011	Message retry response
			0111	Message error response

Table 15-320. RapidIO GSM transactions

GSM Transaction	ftype	ttype	status	Description
IO_READ_HOME	0010	0010	NA	I/O Read Home ¹
FLUSH with data	0101	0001		GSM Flush with data ¹
RESPONSE without data	1101	0000	0000	GSM done response
			0011	GSM retry response
			0101	GSM done intervention response
			0111	GSM error response
RESPONSE with data		1000	0000	GSM done response
			0001	GSM data only response

1. Limited to RapidIO packet generation only

15.6.3 RapidIO packet format summary

Table 15-321 summarizes the small transport field packet formats for supported RapidIO transaction types for LP-Serial operation.

Note that this RapidIO endpoint limits configuration read and write requests to 32-bit data accesses.

Table 15-321. RapidIO small transport field packet format

Transaction	Bits														
	32							32					16		64
	5	3	2	2	4	8	8	4	4	8	8	8	1	1	
NREAD, ATOMIC (inc/dec/inc/dec), IO_READ_HOME	ackl D	rsv = 0	prio	tt	ftype	dest ID	src ID	ttype	rdsz	src TID	addr		w d p tr	xam bs	NA
NWRITE_R, FLUSH with data	ackl D	rsv = 0	prio	tt	ftype	dest ID	src ID	ttype	wrsz	src TID	addr		w d p tr	xam bs	dword 0 -> dword n
NWRITE	ackl D	rsv = 0	prio	tt	ftype	dest ID	src ID	ttype	wrsz	don't care	addr		w d p tr	xam bs	dword 0 -> dword n
SWRITE	ackl D	rsv = 0	prio	tt	ftype	dest ID	src ID	addr(29), rsv(1) = 0, xambs(2)				dword 0 -> dword n			

Table continues on the next page...

Table 15-321. RapidIO small transport field packet format (continued)

Transaction	Bits															
	32							32					16			64
	5	3	2	2	4	8	8	4	4	8	8	8	1	1	2	
MAINT read	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	rd/ wrsiz	src TID	hop cnt	cfg offs et	w d p tr	rsv = 0	NA	
MAINT write	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	rd/ wrsiz	src TID	hop cnt	cfg offs et	w d p tr	rsv = 0	dword 0 (32-bit)	
MAINT port-write	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	rd/ wrsiz	rsv = 0	hop cnt	rsv = 0	w d p tr	rsv = 0	dword 0 -> dword n	
MAINT response without data	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	status	tar TID	hop cnt	rsv = 0			NA	
MAINT response with data	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	status	tar TID	hop cnt	rsv = 0			dword 0 (32-bit)	
RESPONSE without data	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	status	tar TID	NA					
RESPONSE without data for message	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	status	letter(2), mbox(2), msgseg(4)	NA					
RESPONSE with data	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	ttyp e	status	tar TID	dword 0 -> dword n					
DOORBELL	ackl D	rsv = 0	prio	tt	ftyp e	des tID	src ID	rsv = 0			src TID	Info- msb	l n f o - l s b	NA		
MESSAGE	ackl D	rsv = 0	prio	tt	ftyp e	dest ID	src ID	msglen(4), ssize(4), letter(2), mbox(2), msgseg(4)			dword 0 -> dword n					

The large transport field packet formats extends the destination and source IDs to 16-bits each.

15.6.4 RapidIO control symbol summary

Table 15-322 summarizes the 1x/4x LP-Serial control symbols and their format. Refer to the *RapidIO Interconnect Specification, Revision 1.2, Part IV: Physical Layer 1x/4x LP-Serial Specifications, Chapter 4, PCS and PMA Layers*, for 8B/10B data and special (/ PD/, /SC/, idle, sync, skip, align) characters. The 32-bit LP-Serial control symbol is comprised of the eight bit special character and the 24-bit control symbol format.

Table 15-322. 1x/4x LP-serial control symbol format

Bits						Description
24						
stype0	param0	param1	stype1	cmd	CRC	
3	5	5	3	3	5	
000	pkt_ackID	buf_stat	-		crc	Packet accepted
001	pkt_ackID	buf_stat	-		crc	Packet retry
010	pkt_ackID	cause	-		crc	Packet not accepted cause: 00001: Received unexpected ackID on packet 00010: Received a control symbol with bad CRC 00011: Non-maintenance packet reception is stopped 00100: Received packet with bad CRC 00101: Received invalid character or a valid but illegal character 11111: General error
100	ackID_stat	buf_stat	-		crc	Status ackID_stat: 00000: Expecting ackID 0 00001: Expecting ackID 1 00010: Expecting ackID 2 00011: Expecting ackID 3 00100: Expecting ackID 4 00101: Expecting ackID 5 00110: Expecting ackID 6 00111: Expecting ackID 7

Table continues on the next page...

Table 15-322. 1x/4x LP-serial control symbol format (continued)

Bits						Description
24						
stype0	param0	param1	stype1	cmd	CRC	
3	5	5	3	3	5	
110	ackID_stat	port_stat	-		crc	Link-response ackID_stat: 00000: Expecting ackID 0 00001: Expecting ackID 1 00010: Expecting ackID 2 00011: Expecting ackID 3 00100: Expecting ackID 4 00101: Expecting ackID 5 00110: Expecting ackID 6 00111: Expecting ackID 7 port_stat: 00010: Error; unrecoverable 00100: Retry stopped 00101: Error stopped 10000: OK
-			000	000	crc	Start of packet
-			001	000	crc	Stomp
-			010	000	crc	End of packet
-			011	000	crc	Restart from retry
-			100	cmd	crc	Link request cmd: 011: Reset the receiving device 100: Return input port status; functions as a restart-from-error control symbol under error conditions
-			101	000	crc	Multicast-event
-			111	000	crc	NOP (ignore)

15.6.5 Accessing configuration registers using rapidIO packets

There are two suggested methods by which RapidIO transactions can target RapidIO configuration register space in local memory.

The first method is to use RapidIO NREAD and NWRITE_R requests to access the registers through a window defined by the local configuration space base address command and status register (LCSBA1CSR). See [Local configuration space base address 1 command and status register \(SRIO_LCSBA1CSR\)](#), for more information. Subject to the following conditions, inbound requests whose RapidIO address hits the window defined by LCSBA1CSR are translated to the local address range indicated by the CCSRBAR:

- Access to the registers are restricted to 32-bit requests.
- Only NREAD and NWRITE_R requests are supported. If an NWRITE, SWRITE, or ATOMIC request hits the LCSBA1CSR window, an illegal transaction decode error is generated and logged.
- If external configuration accesses are disabled (LLCR[ECRAB] = 1), any configuration access through the LCSBA1CSR window is denied. In this case, writes are ignored and a 32-bit data payload of all zeros is returned for a non-maintenance configuration read.

The second method is based on a RapidIO MAINT requests. This method allows an external device limited access to local RapidIO configuration register space only. Any maintenance access beyond the first 64 Kbytes (the lower 64 Kbytes contain the RapidIO architecture registers; the upper 64 Kbytes contain the RapidIO implementation registers) of RapidIO configuration register space, is denied. A 32-bit data payload of all zeros is returned if for a read response.

A third method, though not suggested, would use an inbound ATMU window to translate RapidIO NREAD and NWRITE_R requests to configuration accesses. This method does not support the configuration access protection features offered by the LCSBACSR window and RapidIO MAINT requests.

15.6.5.1 Guidelines

The RapidIO endpoint limits requests to configuration register space to 32-bit data accesses.

If the order of completion is important, inbound configuration accesses should be assumed incomplete until an appropriate response has been received. It is suggested that there be only one outstanding configuration request at a time to ensure that requests are completed in the order they were intended.

15.6.5.2 Outbound maintenance accesses

Outbound OCeaN NREAD_R or NWRITE requests can be translated to a RapidIO maintenance request if the OCeaN address falls within the bounds of an outbound ATMU window that is setup for generating a maintenance request.

The ATMU window specifies the configuration offset, hop count, source and destination ID, critical request flow, and priority for the outbound RapidIO packet.

15.6.6 RapidIO outbound ATMU

15.6.6.1 Valid hits to multiple ATMU windows

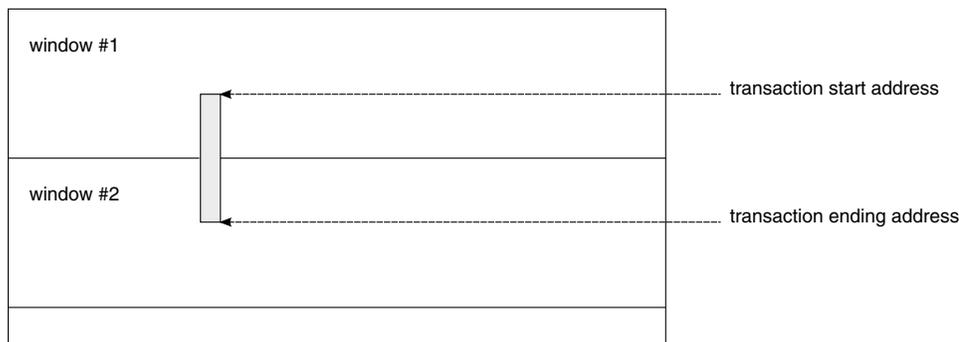
If a request hits multiple ATMU windows, window 1 has the highest priority of the nine outbound ATMU windows (windows 1-8, default).

Window 2 is given the next higher priority and is followed by windows 3 through 8. The default window has the lowest priority.

If a request hits (base address match) multiple ATMU windows and the transaction's end address is contained within the boundary of each hit window and does not extend into another ATMU window, the translation window is the highest priority window that is hit.

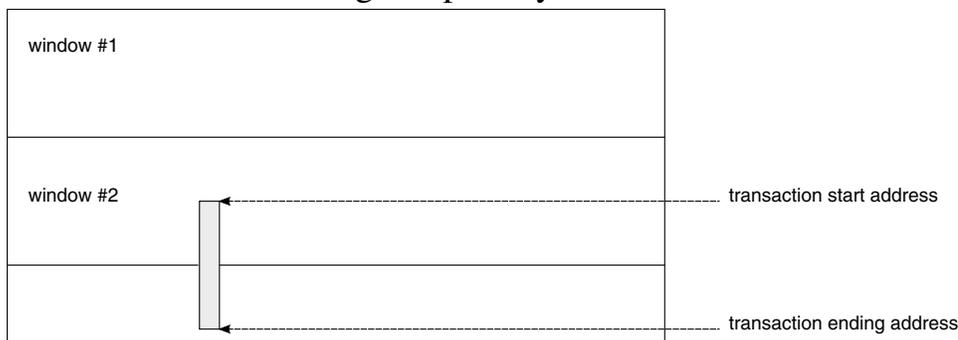
If a lower priority window is programmed to lie entirely within a higher priority window, then it is possible for a transaction to cross window boundaries. Although not a practical programming application, the RapidIO endpoint handles this as follows:

1. If a request hits (base address match) an ATMU window (1-8) and the transaction's end address extends into another ATMU window with lower priority but is still contained within the boundary of the hit window, the translation window is the hit window.



2. If a request hits (base address match) multiple ATMU windows (1-8) and the transaction's end address extends beyond the boundary of a lower priority hit window

but is still contained within the boundary of a higher priority hit window, the translation window is the highest priority window that is hit.

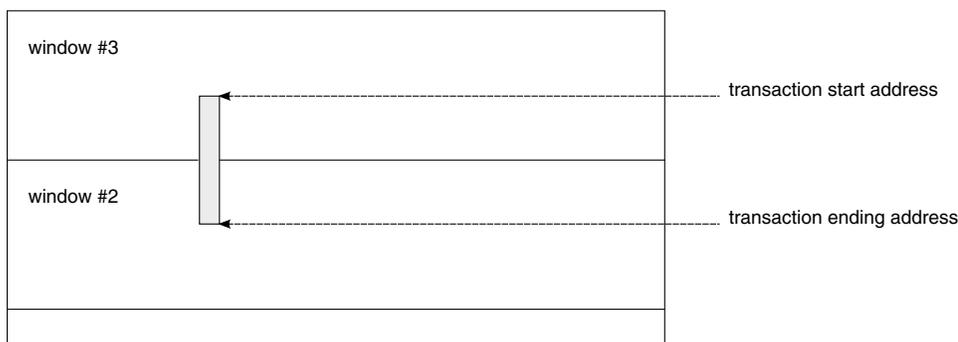


15.6.6.2 Outbound window boundary crossing errors

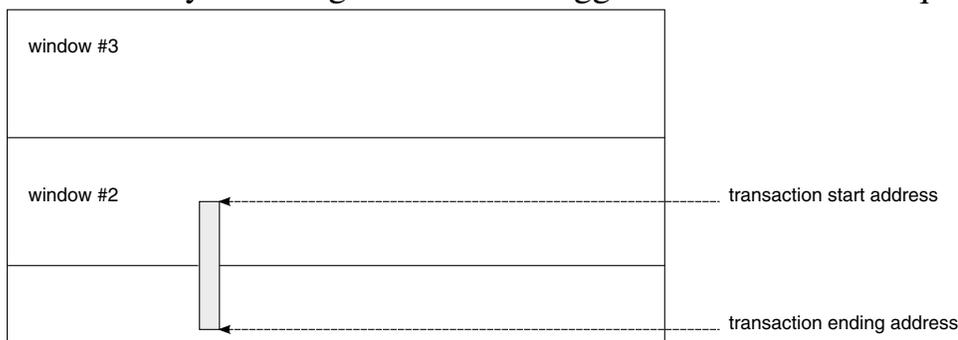
If a higher priority window is programmed to lie entirely within a lower priority window, then it is possible for a transaction to cross window boundaries.

The RapidIO endpoint handles this as follows:

1. If a request hits (base address match) an ATMU window (1-8, default) and the transaction's end address extends into another ATMU window with higher priority, an ATMU crossed boundary error is generated and logged. The outbound request is discarded.

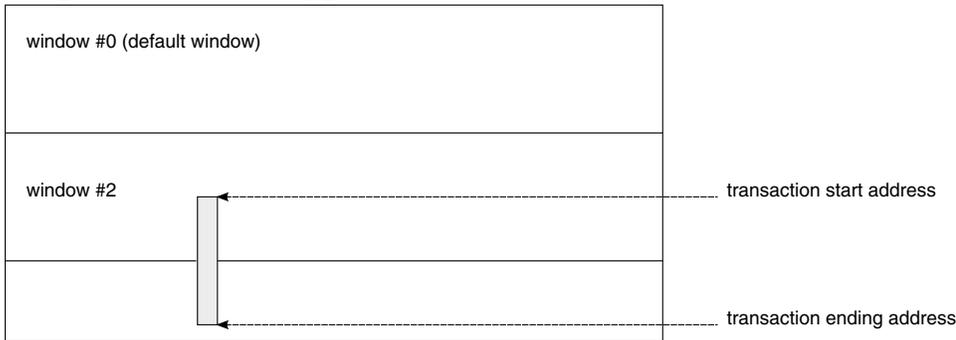


2. If a request hits multiple ATMU windows (1-8, default) and transaction's end address extends beyond the boundary of a higher priority hit window, an outbound ATMU crossed boundary error is generated and logged. The outbound request is discarded.



Other window boundary crossing errors are:

1. If a request hits (base address match) an ATMU window (1-8) and the transaction's end address exceeds the size of the window, an outbound ATMU crossed boundary error is generated and logged. The outbound request is discarded.



Boundary crossing errors (outbound ATMU boundary crossing, segment boundary crossing, and subsegment boundary crossing) are logged in the LTLEDCSR[OACB] configuration register field.

If a request misses all ATMU windows (1-8) and the transaction's end address exceeds the maximum size of the default window, an outbound ATMU crossed boundary error is not generated. The outbound request is forwarded to the RapidIO target device.

15.6.7 RapidIO inbound ATMU

15.6.7.1 Hits to multiple ATMU windows

If a request hits multiple ATMU windows, window 1 has the highest priority of the five inbound ATMU windows (windows 1-4, default).

Window 2 is given the next higher priority and is followed by windows 3 and 4. The default window has the lowest priority.

If a request hits (base address match) multiple ATMU windows and the transaction's end address is contained within the boundary of each hit window and does not extend into another ATMU window, the translation window is the highest priority window that is hit.

If a lower priority window is programmed to lie entirely within a higher priority window, then it is possible for a transaction to cross window boundaries. Although not a practical programming application, the RapidIO endpoint handles this as follows:

1. If a request hits (base address match) an ATMU window (1-4) and the transaction's end address extends into another ATMU window with lower priority but is still

contained within the boundary of the hit window, the translation window is the hit window.

2. If a request hits (base address match) multiple ATMU windows (1-4) and transaction's end address extends beyond the boundary of a lower priority hit window but is still contained within the boundary of a higher priority hit window, the translation window is the highest priority window that is hit.

15.6.7.2 Inbound window boundary crossing errors

If a higher priority window is programmed to lie entirely within a lower priority window, then it is possible for a transaction to cross window boundaries.

The RapidIO endpoint handles this as follows:

1. If a request hits (base address match) an ATMU window (1-4, default) and the transaction's end address extends into another ATMU window with higher priority, an ATMU crossed boundary error is generated and logged.
2. If a request hits multiple ATMU windows (1-4, default) and transaction's end address extends beyond the boundary of a higher priority hit window, an inbound ATMU crossed boundary error is generated and logged.

Other window boundary crossing errors are:

1. If a request hits (base address match) an ATMU window (1-4) and the transaction's end address exceeds the size of the window, an inbound ATMU crossed boundary error is generated and logged.
2. If a NREAD/NWRITE_R/NWRITE/SWRITE request hits (base address match) an ATMU window (1-4, default) and the transaction's end address extends into the region defined as the local configuration space window, an inbound ATMU crossed boundary error is generated and logged.

A RapidIO error response is generated for RapidIO requests that require a response. RapidIO requests that do not require a response are dropped.

Inbound ATMU boundary crossing errors are logged in the LTLEDCSR[IACB] configuration register.

If a request misses all ATMU windows (1-4) and the transaction's end address exceeds the maximum size of the default window, an inbound ATMU crossed boundary error is not generated.

15.6.8 Generating link-request/reset-device

In LP-Serial mode the link partner cannot be reliably reset using the link-request/reset-device control symbols since the input port receiver cannot be disabled independent of the output port driver. The input port driver needs to be disabled to prevent non idle control symbols from being transmitted between the four link-request/reset-device control symbols. For example, if a packet was received on the inbound side after one of the four link-request/reset-device control symbols was sent outbound, the required sequence of four link-request/reset-device symbols would be interrupted with the packet acknowledgement (packet accept, packet retry or packet not accept).

15.6.9 Outbound drain mode

- The RapidIO port is placed into Drain mode when one of the following occurs:
 - $PmPCR[OB DEN]$ is set
 - the Failed Threshold has been encountered and the $PmCCSR[SPF]$ and $PmCCSR[DPE]$ are both set
 - the packet time-to-live counter expires causing $PmPCR[OB DEN]$ to be set
- When the RapidIO port is placed into Drain mode, the RapidIO port discards all packets in the outgoing data stream. Since the data stream is invalid, the RapidIO port also puts its Outbound port back to normal state. Any received acknowledgements and link-responses are considered invalid during this period (since the RapidIO port has cleared out all acknowledgement history).
- The RapidIO port's outbound and outstanding ackID shows that all outstanding packets at the time Drain mode was entered were accepted, whether they truly were accepted or not. If the outbound ackID is not acceptable, then software should change it prior to taking the RapidIO port out of the Drain mode. Also, if the link-partner needs to be put back into inbound OK state, then software should send a link-request/input-status. The recommended sequence for recovering from Drain mode is given in [Software assisted error recovery register support](#).
- $PmPCR[OB DEN]$ also causes any queued up packet acknowledgements to be discarded if the port is uninitialized (the RapidIO port lets them be transferred if the port is initialized). Drain mode due to failed threshold does not cause any packet acknowledgements to be dropped.
- If a discarded packet in the outgoing data stream requires a logical response, a packet response time-out occurs as long as the packet response timer is enabled ($PRTOCCSR$ is non 0).

15.6.10 Input port disable mode

- The RapidIO port is placed into Input Port Disable mode when *PmCCSR*[PD] is set.
- When the RapidIO port is placed in Input Port Disable mode, the RapidIO port discards all incoming data stream (obviously). Since the incoming stream is invalid, the RapidIO port also puts its inbound port back to normal state.
- When the RapidIO port is placed in input port disable mode, the RapidIO port also:
 - ends any packet capture that was in progress.
 - clears the link-request/reset-device count.
- The RapidIO port's inbound ackID shows that all packets successfully received by the RapidIO port at the time input port disable mode was entered were accepted. If output port disable mode was entered at the same time, some packet acknowledgements may be discarded by the RapidIO port. If the inbound ackID is not acceptable, then software should change it prior to taking the RapidIO port out of Input Port Disable mode.

15.6.11 Software assisted error recovery register support

- *PmLMREQCSR* is only supported for recovery from drain mode, including hot-swap support. Consistent with this statement, software should only write this register when the port is in Drain mode.
- The proper sequence for recovering from Drain mode is:
 - a. Software ensures that link activity is stopped. This should include:
 1. software polls for Port OK bit to be set
 2. software waits longer than the link time-out value
 - b. Software generates a link-request/input-status to obtain the link-partner's inbound ackID value.
 - c. Software changes the RapidIO port's outbound ackID to this value (if necessary).
 - d. If the link-partner was hot-inserted, software changes the RapidIO port's inbound ackID to zero. (Note that software needs to know when the link-partner was hot-inserted.)
 - e. NOTE: If software can guarantee that the link-partner does not attempt to forward any packets to this RapidIO port, then software may write the outbound ackID of the link-partner to match the inbound ackID of this RapidIO port. However, if the link-partner attempts to forward another packet while this write is still outstanding, and the ackIDs already happened to line-up, then the write actually causes the ackIDs to not match, and the link can not be recovered.
 - f. Software should cause the link-partner to send a link-request/input-status just to ensure that the RapidIO port's inbound port is in Normal operation.
 - g. Software clears the Failed Encountered bit or the Output Buffer Drain Enable bit; whichever one caused the Drain mode (thus clearing Drain mode).

- Software is responsible for timing software generated link-requests. If the response valid bit is not set in some reasonable period of time, the software should write another request in the register.
- When software writes *PmLMREQCSR*, software should make sure to successfully read *PmLMRESPCSR* set, otherwise, software may read a "stale" *ackID* status/link status later.
- Note that when the RapidIO port's outbound *ackID* is written by software using *PmLASCSR*, the inbound *ackID* is also written. Care must be taken to ensure that the inbound *ackID* is not written to an incorrect value. For example, *PmCCSR[PL]* could be set to prevent the inbound *ackID* from changing before *PmLASCSR* is written.

15.6.12 Hot-swap support

The two basic hot-insertion approaches described in the RapidIO Error Management Extensions are supported although the second approach is only partially supported.

Method one has a host bringing a field replacement unit (FRU) into the system. Method two has the FRU bringing itself into the system. Note that this RapidIO port is most likely the device being hot-inserted/extracted since typically this device would be connected to a switch but it could be the link partner of a device being hot-inserted/extracted.

15.6.12.1 Method 1

One possible sequence for when the host brings the FRU into the system is given.

This RapidIO port can be either the device being hot-inserted/extracted or the link partner of the device being hot-inserted/extracted. The goal of this sequence is to bring the FRU into the system cleanly without generating errors.

15.6.12.1.1 Extraction-method 1

- The host determines that the FRU has failed by getting a port response time-out when polling the port OK bit (PO bit in port *m* error and status CSR) of the FRU. Note that there are other ways to determine that the FRU has failed. This is one possible method to detect FRU failure.
- The host must perform the following steps to the FRU link partner before hot insertion of the FRU occurs.
 - set the port lockout bit (PL = 1 in the port *m* control CSR) preventing packet reception and transmission

- prevent any new packets from arriving to the FRU link partner RapidIO port by all other RapidIO devices and discard all pending packets destined for the FRU so that...
 - congestion to this RapidIO port does not occur and cause other system problems
 - the FRU is not immediately flooded with old packets after insertion causing other errors like unsolicited responses
- note that this RapidIO endpoint has a specific bit to enable the discard of pending packets (OBDEN in *PmPCR*). Also note that setting OBDEN in *PmPCR* forces the output state from error or retry to normal. In a switch, the time-to-live feature may be used to discard packets.
- force the input state to normal. In this RapidIO endpoint, this occurs by disabling the input port receiver.
- leave the input port receivers and output port drivers enabled so that initialization can complete when the FRU is inserted
- clear all errors pertaining to the extracted RapidIO port in the FRU link partner and any other RapidIO port that was affected by the FRU failing (port response time-outs)
- if RapidIO endpoint, clear OBDEN in *PmPCR* for normal operation
- set the outbound and inbound ackID to 0x00 in the port *m* local ackID status CSR so that the ackID is correct and packets can be transmitted and received when the FRU is inserted
- The host indicates that the FRU should be removed.
- The FRU is removed from the system.

15.6.12.1.2 Insertion-method 1

- The FRU is inserted into the system
- Link initialization occurs and initialization complete status is indicated (port OK bit in the port *merror* and status CSR) in both RapidIO ports
- The host determines that this link partner has been inserted by periodically polling the initialization complete status in the FRU link partner (PO bit = 1 in the port *merror* and status command and status register)
- The host clears the output and input port enable bits and clears the port lockout bit in the port *mcontrol* CSR in the FRU link partner allowing only maintenance transactions
- The host sets the master enabled and discovered bits in the FRU (M = 1 and D = 1 in the port general control CSR). Note that the master enable bit does not prevent the RapidIO endpoint from transmitting packets.
- The host completes configuration of the FRU
- The host sets the output and input port enable bits in the port *m* control CSR in the FRU link partner allowing transmission and reception of all packet types

- The host re-enables packets to be sent to this RapidIO port by other RapidIO devices
- System operation resumes

15.6.12.2 Method 2 with RapidIO port hot-swapped

One possible sequence for when the FRU brings itself into the system is given.

This RapidIO Port can be either the device being hot-inserted/extracted or the link partner of the device being hot-inserted/extracted. The goal of this sequence is to bring the FRU into the system cleanly without generating errors.

15.6.12.2.1 Extraction-method 2

- The FRU fails.
- New packets are prevented from arriving to the FRU link partner by all other RapidIO devices and all pending packets destined for the FRU are discarded so that the following occur:
 - Congestion to this RapidIO port does not occur and cause other system problems.
 - The FRU is not immediately flooded with old packets after insertion causing other errors such as unsolicited responses.
- This FRU link partner continues to have its drivers enabled.
- The FRU link partner continues to allow the transmission and reception of all packet types since its output and input port enable bits are set and the port lockout bit is cleared in the port *m* control CSR.
- The FRU is removed from the system.

15.6.12.2.2 Insertion-method 2

- The FRU is inserted.
- Link initialization occurs, initialization complete status is indicated in both RapidIO ports (PO bit = 1 in the port *m* error and status command and status register).
- After initialization is complete, both devices set the port OK bit = 1 in the port *m* control CSR.
- The FRU sets its port lockout bit in the port *m* control CSR.
- The FRU generates a link-request/input-status to its link partner using the port *m* link maintenance request register. The FRU determines the link partner's inbound ackID by reading the port *m* link maintenance response register. The FRU then sets its outbound ackID in the port *m* local ackID status CSR.
- The FRU only enables maintenance transactions by clearing its output and input port enable bits in the port *m* control CSR and by clearing its port lockout bit in the port *m* control CSR.

- The FRU generates a maintenance write to its link partner's port m local ackID status CSR to set the link partner's outbound ackID value to 0. Upon receipt of the maintenance write, the link partner sets its outbound ackID value and generates the maintenance response using the new value. Note that if all packets intended for the link partner have not been discarded, or if any new packets intended for the link partner arrive, this step may fail (the RapidIO endpoint has no way to protect against this).
- The FRU completes configuration.
- The FRU enables all packets to be transmitted and received by setting its output and input port enable bits in the port m control CSR.
- System operation resumes.

15.6.13 Software re-training

A user can issue a software controlled re-training of the link in LP-Serial mode in order to update certain configurations by following the steps below.

- Software on the host must ensure that all RapidIO transactions have completed and link activity has stopped.
- Set the $P_nCCSR[PD]$ bit on the host to disable the port and force the initialization state machines to reset.
- Set the $P_nPCR[OBDEN]$ bit on the host to enable the discarding of any pending packets. (There should be none if proper steps were taken to ensure there was no link activity.)
- Clear the $P_nPCR[OBDEN]$ bit on the host.
- Configure new operating width (via $P_nCCSR[PWO]$) or any other new configurations.
- Clear the $P_nCCSR[PD]$ bit on the host to re-enable the drivers.
- Poll the $P_nESCSR[PO]$ bit on the host until it is set, which indicates the link has attained port and link initialization.
- Poll the $P_nESCSR[PO]$ bit on the agent until it is set, which indicates the link has attained port and link initialization.
- Poll the $P_nESCSR[OES]$ bit on the agent until it is clear, which indicates the link has completed the error recovery sequence initiated from the port disable.
- Poll the $P_nESCSR[OES]$ bit on the host until it is clear, which indicates the link has completed the error recovery sequence initiated from the port disable.
- Clear the agent's error and status registers.
- Clear the host's error and status registers.
- Begin normal packet transfer.

15.6.14 Errors and error handling

This section describes how the logical and physical layers detect and react to RapidIO errors.

The action of the core on notification of any of these errors is described minimally here. See *RapidIO Interconnect Specification, Revision 1.2 Part VII (Error Management Extensions Specifications)* for more details on specific errors described below.

15.6.14.1 RapidIO error description

RapidIO errors are classified under three categories: recoverable errors, notification errors, and fatal errors.

Recoverable errors are non-fatal transmission errors (such as corrupt packet or control symbols, and general protocol errors) that RapidIO supports hardware detection of and a recovery mechanism for, as described in the *RapidIO Interconnect Specification, Revision 1.2*. In these cases, the appropriate bit is set in the port *merror* detect CSR. Only the packet containing the first detected recoverable error that is enabled for error capture (by port *merror* enable CSR) is captured in the port *merror* capture CSRs. No interrupt is generated or actions required for a recoverable error. Recoverable errors are detected in the physical layer only.

Notification errors are non-recoverable non-fatal errors detected by RapidIO (such as degraded threshold, port-write received, and all logical/transport layer (LTL) errors captured). Because they are non-recoverable (and in some cases have caused a packet to be dropped), notification by interrupt is available. However, because they are non-fatal, response to the interrupt is not crucial to port performance; that is, the port is still functional. When a notification error is detected, the appropriate bit is set in the error-specific register, an interrupt is generated, and in some cases, the error is captured. In all cases, the RapidIO port continues operating. Notification errors are detected in both the physical and logical layer.

NOTE

To prevent processor stalls, the *PmLOPTTL*CR registers must be initialized to a non-zero value. See [Port 1 Logical Outbound Packet time-to-live configuration register \(SRIO_P1LOPTTL](#)CR).

The RapidIO controller detects two fatal errors: exceeded failed threshold and exceeded consecutive retry threshold. In these cases, the port has failed because its recoverable error rate has exceeded a predefined failed threshold or because it has received too many

packet retries in a row. In the first case, the controller sets the output failed-encountered bit in the port *merror* and status CSR; the RapidIO output hardware may or may not stop (based on stop-on-port-failed-encounter-enable and drop-packet-enable bits). In the second case, the controller sets the retry counter threshold trigger exceeded bit in the port *mimplementation error* CSR; the RapidIO hardware continues to operate. In both cases, an interrupt is generated, and while the port continues operating at least partially, a system-level fix (such as reset) is recommended to clean up the controller's internal queues and resume normal operation. Fatal errors are detected in the physical layer only.

15.6.14.2 Physical layer RapidIO errors detected

Table 15-323 lists all the RapidIO link errors detected by the RapidIO endpoint physical layer and the actions taken by the RapidIO endpoint. The error enable column lists the control bits that may disable the error checking associated with a particular error (if blank, error checking cannot be disabled). The cause field column indicates what cause field is used with the associated packet-not-accept control symbol for input error recovery. The EME error enable/detect column indicates which bit of the *PmERECSR* allows the error to increment the error rate counter and lock the port *merror* capture registers, and likewise which bit of the *PmEDCSR* is set when the error has been detected.

Table 15-324 lists the RapidIO endpoint behavior after exceeding certain preset limits (degraded threshold, failed threshold, retry threshold).

Table 15-323. Physical RapidIO errors detected

Level	Error	Error Enable	RapidIO Endpoint Action	Cause Field	EME Error Type	EME Error Enable / Detect
Recoverable Errors						

Table continues on the next page...

Table 15-323. Physical RapidIO errors detected (continued)

Level	Error	Error Enable	RapidIO Endpoint Action	Cause Field	EME Error Type	EME Error Enable / Detect
1a	Received character had a disparity error. (serial)	-	Enter Input Error Stopped. Enter Output Error Stopped.	5: Received invalid/illegal character	Delineation Error	DE
1a	Received an invalid character, or valid but illegal character (serial)	-	Enter Input Error Stopped. Enter Output Error Stopped.	5: Received invalid/illegal character		
1b	The four control character bits associated with the received symbol do not make sense (not 0000, 1000, 1111) (serial)	-	Enter Input Error Stopped. Enter Output Error Stopped.	5: Received invalid/illegal character		
1b	Control symbol does not begin with an /SC/ or /PD/ control character. (serial)	-	Enter Input Error Stopped. Enter Output Error Stopped.	5: Received invalid/illegal character		
1c	Received packet with embedded idles. (serial)	-	Enter Input Error Stopped.	5: Received invalid/illegal character		
1d	Received a control symbol with a bad CRC	PmPCR[CC] enables detect.	Enter Input Error Stopped. Enter Output Error Stopped.	2: Received a control symbol with bad CRC	Received corrupt control symbol	CCS
1d	Missing start: Packet data received without previous SOP control symbol	-	Enter Input Error Stopped.	7/31: General error	Protocol Error (unexpected packet/control symbol received)	PE
1e	Received packet that is < 64 bits	-	Enter Input Error Stopped.	7/31: General error		
1e	Received an EOP control symbol when there is no packet being received.	-	Enter Input Error Stopped.	7/31: General error		
1e	Received a stomp control symbol when there is no packet being received.	-	Enter Input Error Stopped.	7/31: General error		
2a	Received a Restart-from-retry control symbol when in the "OK" state	-	Enter Input Error Stopped	7/31: General error	Protocol Error (unexpected packet/control symbol received)	PE

Table continues on the next page...

Table 15-323. Physical RapidIO errors detected (continued)

Level	Error	Error Enable	RapidIO Endpoint Action	Cause Field	EME Error Type	EME Error Enable / Detect
2a	Received packet with a bad CRC value.	PmPCR[CP] enables detect.	Enter Input Error Stopped.	4: bad CRC on packet.	Received packet with bad CRC	CRC
2a	Received packet which exceeds the maximum allowed size by the RapidIO spec.	-	Enter Input Error Stopped.	7/31: General error	Received packet exceeds 276 Bytes	EM
2b	Received packet with unexpected ackID value (out-of-sequence ACKID)	-	Enter Input Error Stopped.	1: Received unexpected ACKID on packet	Received packet with unexpected ackID	UA
2c	Received a non-maintenance packet when non-maintenance packet reception is stopped	Non-maint. packet reception is stopped when 'Input Port Enable' = 0.	Enter Input Error Stopped.	3: Non-maintenance packet reception is stopped	Not Captured	-
2d	Any packet received while Port Lockout bit is set	All packet reception is stopped when Port Lockout bit is set.	Enter Input Error Stopped.	3: Non-maintenance packet reception is stopped	Not Captured	-
-	Received a Link request control symbol before servicing previous link request.	Not detected.				
2a	Received packet-not-accepted ACK control symbol.	-	Enter Output Error Stopped.	-	Received packet-not-accepted symbol	PNA
2b	Received an ACK (accepted, or retry) control symbol when there are no outstanding packets	-	Enter Output Error Stopped.	-	Unsolicited ACK symbol	UCS
2b	Received packet ACK (accepted) for a packet whose transmission has not finished	-	Enter Output Error Stopped.	-		
2b	Received a Link response control symbol when no outstanding request.	-	Enter Output Error Stopped.	-		

Table continues on the next page...

Table 15-323. Physical RapidIO errors detected (continued)

Level	Error	Error Enable	RapidIO Endpoint Action	Cause Field	EME Error Type	EME Error Enable / Detect
2c	Received an ACK (accepted or retry) control symbol with an unexpected ACKID.	-	Enter Output Error Stopped.	-	Received ack. control symbol with unexpected ackID	AUA
2c	Link_response received with an ackID that is not outstanding	-	Re-enter Output Error Stopped.	-	Non-outstanding ackID	NOA
2d	An ACK control symbol is not received within the specified time-out interval.	PLTOCCSR[TV] > 0 enables detect.	Enter Output Error Stopped.	-	Link time-out	LTO
2d	A Link response is not received within the specified time-out interval	PLTOCCSR[TV] > 0 enables detect.	(re-) Enter Output Error Stopped.	-		

Table 15-324. Physical RapidIO threshold response

Error	Error Enable	RapidIO Endpoint Action	EME Error Type	Error Detect	Interrupt clear ¹
Notification Errors					
Error Rate Counter has met or exceeded the Degraded Threshold.	PmERTCSR[ERDTT] > 0 & any bit in PmERCSR enables detect and interrupt generation.	Generate Interrupt. Continue to operate normally.	Degraded Threshold	PmESCSR[ODE]	Write 1 to PmESCSR[ODE]
Fatal Errors					
Consecutive Retry Counter has met or exceeded the Retry Counter Threshold Trigger	PRETCR[RET] > 0 enables detect and interrupt generation	Generate Interrupt. Port is in priority order.	Consecutive Retry Threshold	PmIECSR[RETE]	Write 1 to PmIECSR[RETE]
Error Rate Counter has met or exceeded the Failed Threshold.	PmERTCSR[ERFTT] > 0 & any bit in PmERCSR enables detect and interrupt generation.	Generate Interrupt. Port behavior depends on PmCCSR[SPF] and PmCCSR[DPE] -- port can continue transmitting packets or can stop sending output packets, keeping or dropping them.	Failed Threshold	PmESCSR[OFE]	Write 1 to PmESCSR[OFE].

1. Information given here is minimal for clearing the interrupt. More detailed steps should be taken to find the cause of the interrupt.

15.6.14.3 Logical layer errors and error handling

This section describes how the logical layer detects and reacts to RapidIO errors.

The action of the core on notification of any of these errors is described minimally here. Reference *RapidIO Interconnect Specification, Revision 1.2* Part VII (Error Management Extensions Specifications).

15.6.14.3.1 Logical layer RapidIO errors detected

[Table 15-325](#) through [Table 15-338](#) lists all the errors detected by the RapidIO endpoint logical layer and the actions taken by the RapidIO endpoint. Note that when the RapidIO endpoint action includes sending an error response to either OCN or RapidIO, an error response is only sent if the original transaction was a request that required a response. Otherwise, no error response is sent. When dealing with multiple errors, discard of packet has higher priority than error response.

For misaligned transactions, the error management extension registers are updated with each child.

All packet field positions are assumed to be in the mode (small or large transport) configured. For example, when configured for small transport mode and a large transport mode NREAD packet is received, the transaction type field bit positions checked correspond to a small transport type NREAD packet.

Table 15-325. Hardware errors for NRead transaction

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Priority of Read transaction is	Yes if LTLEECR[ITD] is set	LTLEDCSR[ITD]	No	Using the incoming RapidIO packet, for Small Transport type packet, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDIDCCSR[DIDMSB] gets 0's, LTLDIDCCSR[DID] gets packet bits 16-23, LTLDIDCCSR[SIDMSB] gets 0's, LTLDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets r. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDIDCCSR[DIDMSB] gets 16-23, LTLDIDCCSR[DID] LTL gets packet bits 24-31, LTLDIDCCSR[SIDMSB] gets bits 32-39, LTLDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped.
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCSR[TSE]	No	Same as first entry	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when pass_through is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCSR[TSE]	No	Same as first entry	RapidIO packet is dropped
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (pass_through or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCSR[ITTE]	Yes	Same as first entry	OCN error response is generated to self

Table continues on the next page...

Table 15-325. Hardware errors for NRead transaction (continued)

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
SourceID Not Checked for error.	-	-	-	-	-
TransactionType Received RapidIO packet with reserved TType for this ftype	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	OCN error response is generated to self
RdSize Not Checked for error.	-	-	-	-	-
SrcTID Not Checked for error.	-	-	-	-	-
Address:WdPtr:Xambs Read request hits overlapping ATMU windows Refer to Inbound window boundary crossing errors	Yes if LTLEECS R[IACB] is set	LTLEDCS R[IACB]	Yes	Same as first entry	OCN error response is generated to self
Address:WdPtr:Xambs Request hits a protected ATMU window	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	OCN error response is generated to self
Address:WdPtr:Xambs Beginning address matches LCSBA1CSR with non 32 bit read request. Performed only when ttype == 4'b0100	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	OCN error response is generated to self
Header Size Header size is not 12 Bytes for small Transport packet or not 16 Bytes for Large Transport packet. Large Transport packet has 14 valid bytes and two bytes of padding of 0's. Padding of 0's is not checked.	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	OCN error response is generated to self
PayloadSize Not Applicable	-	-	-	-	-

Table 15-326. Hardware errors for maintenance read/write req transaction

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Priority of maintenance read or write request transaction is 3	Yes if LTLEECR[ITD] is set	LTLEDCR[ITD]	No	Using the incoming RapidIO packet, for Small Transport type packets, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDIDCCSR[DIDMSB] gets 0's, LTLDIDCCSR[DID] gets packet bits 16-23, LTLDIDCCSR[SIDMSB] gets 0's, LTLDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDIDCCSR[DIDMSB] gets 16-23, LTLDIDCCSR[DID] LTL gets packet bits 24-31, LTLDIDCCSR[SIDMSB] gets bits 32-39, LTLDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped
TransportType Received reserved TT	Yes	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	Yes	Same as first entry	OCN error response is generated to self
SourceID Not Checked for error.	-	-	-	-	-

Table continues on the next page...

Table 15-326. Hardware errors for maintenance read/write req transaction (continued)

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
TransactionType Reserved Transaction Type for this ftype	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	RapidIO packet is dropped
RdSize Read/Write request size is not for 4 bytes	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	OCN error response is generated to self
SrcTID Not Checked for error.	-	-	-	-	-
HopCount Not Checked for error.	-	-	-	-	-
Config Offset Not Checked for error.	-	-	-	-	-
Header Size Maintenance Read request - Header size is not 12 Bytes for small Transport packet or not 16 Bytes for Large Transport packet. Maintenance Write request - total header size is not 12 Bytes for Small Transport packet or not 16 Bytes for Large Transport packet. Padding of 0's in last two bytes of Large Transport packet is not checked.	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	OCN error response is generated to self
PayloadSize Write request with payload not equal to 8 bytes. Read request with payload not 0 bytes	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	OCN error response is generated to self

Table 15-327. Hardware errors for atomic (inc, dec, set, or clr) read transaction

Error	Interrupt Generated	Status Bit Set if corresponding bit is enabled	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Priority of read transaction is 3	Yes if LTLEECR[ITD] is set	LTLEDCR[ITD]	No	Using the incoming RapidIO packet, for Small Transport type packets, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLIDCCSR[DIDMSB] gets 0's, LTLIDCCSR[DID] gets packet bits 16-23, LTLIDCCSR[SIDMSB] gets 0's, LTLIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLIDCCSR[DIDMSB] gets 16-23, LTLIDCCSR[DID] gets packet bits 24-31, LTLIDCCSR[SIDMSB] gets bits 32-39, LTLIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	Yes	Same as first entry	OCN error response is generated to self

Table continues on the next page...

Table 15-327. Hardware errors for atomic (inc, dec, set, or clr) read transaction (continued)

Error	Interrupt Generated	Status Bit Set if corresponding bit is enabled	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
SourceID Not Checked for error.	-	-	-	-	-
TransactionType Non-Atomic ttype is tested with Nread	-	-	-	-	-
TransactionType Received Atomic Increment request with DOCAR[AI] disabled. Received Atomic decrement request with DOCAR[AD] disabled. Received Atomic Set request with DOCAR[AS] disabled. Received Atomic Clear request with DOCAR[AC] disabled.	Yes if LTLEECS R[UT] is set	LTLEDCS R[UT]	Yes	Same as second entry	Error response is generated to self
RdSize Not unsupported RdSize request is not for contiguous one, two or four bytes	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	Error response is generated to self
SrcTID Not Checked for error.	-	-	-	-	-
Address:WdPtr:Xambs Not unsupported Request hits a protected ATMU window or the LCSBA1CSR Refer to Outbound window boundary crossing errors	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	Error response is generated to self
Address:WdPtr:Xambs Read request hits overlapping ATMU windows Refer to Outbound window boundary crossing errors	Yes if LTLEECS R[IACB] is set	LTLEDCS R[IACB]	Yes	Same as first entry	Error response is generated to self
Header Size Not unsupported Header size is not 12 Bytes for small Transport packet and not 16 Bytes for Large Transport packet Padding of 0's in last two bytes of Large Transport packet is not checked	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	Yes	Same as first entry	Error response is generated to self

Table 15-328. Hardware errors for NWrite, NWrite_r, and unsupported atomic test-and-swap transactions

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Transaction priority is 3	Yes if LTLEECR[ITD] is set	LTLEDCR[ITD]	No	Using the incoming RapidIO packet, for Small Transport type packets, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLIDCCSR[DIDMSB] gets 0's, LTLIDCCSR[DID] gets packet bits 16-23, LTLIDCCSR[SIDMSB] gets 0's, LTLIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLIDCCSR[DIDMSB] gets 16-23, LTLIDCCSR[DID] gets packet bits 24-31, LTLIDCCSR[SIDMSB] gets bits 32-39, LTLIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped.
TransportType Received reserved TT	Yes	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	Yes for Nwrite_r. No for Nwrite.	Same as first entry	OCN error response is generated to self for Nwrite_r. RapidIO packet is dropped for Nwrite
SourceID Not applicable	-	-	-	-	-

Table continues on the next page...

Table 15-328. Hardware errors for NWrite, NWrite_r, and unsupported atomic test-and-swap transactions (continued)

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
TransactionType Received RapidIO packet for Atomic test-and-swap transaction	Yes if LTLEECRSR[UT] is set	LTLEDCRSR[UT]	Yes	Same as first entry	OCN error response is generated to self
TransactionType Received RapidIO packet with reserved TType for this ftype Packet is treated as Nwrite Transaction	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped
WrSize Not unsupported transaction WrSize request is for one of reserved sizes	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	Yes for Nwrite_r. No for Nwrite.	Same as first entry	OCN error response is generated to self for Nwrite_r. RapidIO packet is dropped for Nwrite
Address:WdPtr:Xambs Not unsupported transaction Nwrite request hits LCSBA1CSR	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No for Nwrite.	Same as first entry	RapidIO packet is dropped for Nwrite
Address:WdPtr:Xambs Not unsupported transaction Request hits a protected ATMU window	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	Yes for Nwrite_r. No for Nwrite.	Same as first entry	OCN error response is generated to self for Nwrite_r. RapidIO packet is dropped for Nwrite
Address:WdPtr:Xambs Write request hits overlapping ATMU windows Refer to Outbound window boundary crossing errors	Yes if LTLEECRSR[IACB] is set	LTLEDCRSR[IACB]	Yes for Nwrite_r. No for Nwrite.	Same as first entry	OCN error response is generated to self for Nwrite_r. RapidIO packet is dropped for Nwrite
SrcTID Not Checked for error.	-	-	-	-	-

Table continues on the next page...

Table 15-328. Hardware errors for NWrite, NWrite_r, and unsupported atomic test-and-swap transactions (continued)

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Address:WdPtr:Xambs Nwrite_r address matches LCSBA1CSR with non 32 bit read request. Performed only for Nwrite_r packet	Yes if LTLEECSR[ITD] is set	LTLEDCSR[ITD]	Yes for Nwrite_r.	Same as first entry	OCN error response is generated to self
Header Size Not unsupported transaction Header size is less than 12 bytes for small Transport packet or less than 16 bytes for Large Transport packet - that is, no payload present. Large Transport packet has 14 valid bytes and 2 bytes of padding of 0s. Padding of 0s is not checked.	Yes if LTLEECSR[ITD] is set	LTLEDCSR[ITD]	Yes for Nwrite_r. No for Nwrite.	Same as first entry	OCN error response is generated to self for Nwrite_r. RapidIO packet is dropped for Nwrite
PayloadSize Not unsupported transaction Payload is greater than that indicated by {wdptr:wrsiz} field, payload is not double word aligned or does not have any payload	Yes if LTLEECSR[ITD] is set	LTLEDCSR[ITD]	Yes for Nwrite_r. No for Nwrite.	Same as first entry	OCN error response is generated to self for Nwrite_r. RapidIO packet is dropped for Nwrite

Table 15-329. Hardware errors for SWrite transactions

Error	Interrupt Generated	Status Bit Set	Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Swrite transaction priority is 3	Yes if LTLEECSR[ITD] is set	LTLEDCSR[ITD]	No	Same as third entry	RapidIO packet is dropped

Table continues on the next page...

Table 15-329. Hardware errors for SWrite transactions (continued)

Error	Interrupt Generated	Status Bit Set	Error Response Generated	Logical/Transport Layer Capture Register	Comments
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Using the incoming RapidIO packet, for Small Transport type packets, LTLACCSR[XA] gets packet bits 62-63, LTLACCSR[A] gets packet bits 32-60, LTLDDCCSR[DIDMSB] gets 0's, LTLDDCCSR[DID] gets packet bits 16-23, LTLDDCCSR[SIDMSB] gets 0's, LTLDDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets. LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLTLTLDDCCSR[DIDMSB] gets 16-23, LTLDDCCSR[DID] gets packet bits 24-31, LTLDDCCSR[SIDMSB] gets bits 32-39, LTLDDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as third entry	RapidIO packet is dropped
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	No	Same as third entry	RapidIO packet is dropped
SourceID Not Checked for error.	-	-	-	-	-
Address:WdPtr:Xambs Swrite request hits overlapping ATMU windows. Refer to Outbound window boundary crossing errors	Yes if LTLEECR[IACB] is set	LTLEDCR[IACB]	No	Same as third entry	RapidIO packet is dropped

Table continues on the next page...

Table 15-329. Hardware errors for SWrite transactions (continued)

Error	Interrupt Generated	Status Bit Set	Error Response Generated	Logical/Transport Layer Capture Register	Comments
Address:WdPtr:Xambs Request hits a protected ATMU window or the LCSBA1CSR	Yes if LTLEECR[ITD] is set	LTLEDCSR[ITD]	No	Same as third entry	RapidIO packet is dropped
PayloadSize Payload size is not in DWs, has exceeded 256 bytes or has no payload.	Yes if LTLEECR[ITD] is set	LTLEDCSR[ITD]	No	Same as third entry	RapidIO packet is dropped

Table 15-330. Hardware errors for maintenance response transactions

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Not UR Response priority is not higher than RapidIO maintenance request priority	Yes if LTLEECR[ITD] is set	LTLEDCSR[ITD]	No	Using the incoming RapidIO packet, for Small Transport type packets, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLIDCCSR[DIDMSB] gets 0's, LTLIDCCSR[DID] gets packet bits 16-23, LTLIDCCSR[SIDMSB] gets 0's, LTLIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLIDCCSR[DIDMSB] gets 16-23, LTLIDCCSR[DID] gets packet bits 24-31, LTLIDCCSR[SIDMSB] gets bits 32-39, LTLIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped and ignored
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCSR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored

Table continues on the next page...

Table 15-330. Hardware errors for maintenance response transactions (continued)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECRSR[TSE] is set	LTLEDCSR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECRSR[ITTE] is set	LTLEDCSR[ITTE]	No	Same as first entry	RapidIO packet is dropped and ignored
SourceID Does not match the request's DestID	Yes if LTLEECRSR[UR] is set	LTLEDCSR[UR]	No	Same as first entry	RapidIO packet is dropped and ignored
TransactionType Not UR Received RapidIO packet with reserved TType for this ftype	Yes if LTLEECRSR[ITD] is set	LTLEDCSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
Not UR Maintenance read/write response does not correspond to an outstanding valid message read/write request.	Yes if LTLEECRSR[ITD] is set	LTLEDCSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
HopCount Not Checked for error.	-	-	-	-	-
Status Not UR Is not "Done" or "Error" Not "Done" status for "read_response" transaction type with payload "Error" status with payload.	Yes if LTLEECRSR[ITD] is set	LTLEDCSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored

Table continues on the next page...

Table 15-330. Hardware errors for maintenance response transactions (continued)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Status Not UR Error Response	Yes if LTLEECRSR[IER] is set	LTLEDCRSR[IER]	Yes	Same as first entry except error capture is done from original request	OCN error response is generated to requestor.
TargetTID No outstanding transaction for this TargetTID	Yes if LTLEECRSR[UR] is set	LTLEDCRSR[UR]	No	Same as first entry	RapidIO packet is dropped and ignored
Header Size Not UR Maintenance Read response - total payload size with done status is not greater than 4 Bytes. Maintenance Write response - total header size is less than 12 Bytes for Small Transport packet or is less than 16 Bytes for Large Transport packet. Padding of 0's for Small or Large Transport packets is not verified.	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
PayloadSize Not UR Maintenance write response has payload. Maintenance read response with done status and payload not matching valid request size, request size for the response is invalid or payload size is not dword aligned.	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
Packet response time-out Response is not received by configured time	Yes if LTLEECRSR[PRT] is set	LTLEDCRSR[PRT]	Yes	Same as first entry except error capture is done from original request.	OCN error response is generated to requestor.

Table 15-331. Hardware errors for IO/GSM response transactions (not maintenance)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
<p>Priority Not UR</p> <p>Response priority is not higher than RapidIO request priority</p>	Yes if LTLEECR[ITD] is set	LTLEDCR[ITD]	No	<p>Using the incoming RapidIO packet, for Small Transport type packets, LTLACCSR[XA] gets packet bits 78-79 (if available), LTLACCSR[A] gets packet bits 48-76 (if available), LTLDDIDCCSR[DIDMSB] gets 0's, LTLDDIDCCSR[DID] gets packet bits 16-23, LTLDDIDCCSR[SIDMSB] gets 0's, LTLDDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's</p> <p>For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95 (if available), LTLACCSR[A] gets packet bits 64-92 (if available), LTLTLTLDDIDCCSR[DIDMSB] gets 16-23, LTLDDIDCCSR[DID] gets packet bits 24-31, LTLDDIDCCSR[SIDMSB] gets bits 32-39, LTLDDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's</p>	RapidIO packet is dropped and ignored
<p>TransportType</p> <p>Received reserved TT for this ftype</p>	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored
<p>Received TT which is not enabled.</p> <p>- Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.</p>	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored
<p>DestID</p> <p>DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled.</p> <p>Error valid when (passthrough or accept_all) is false</p>	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	No	Same as first entry	RapidIO packet is dropped and ignored

Table continues on the next page...

**Table 15-331. Hardware errors for IO/GSM response transactions (not maintenance)
(continued)**

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
SourceID Does not match the request's DestID	Yes if LTLEECRSR[UR] is set	LTLEDCRSR[UR]	No	Same as first entry	RapidIO packet is dropped and ignored
TransactionType Not UR Received RapidIO packet with reserved TType	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
Not UR IO read response does not correspond to an outstanding valid IO/GSM read request. IO write response does not correspond to an outstanding valid IO/GSM write request.	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
Status Not UR IO transaction - Is not "Done" or "Error" GSM transaction IO_Read_Home Is not "Done - Data-Only", "Done - Done-Intervention", "Done", "Retry" or "Error". Flush_w_Data response is not "Done", "Retry" or "Error" Transaction type of "Response_with_data" and status is not done	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored.
Status Not UR GSM Error response	Yes if LTLEECRSR[GER] is set	LTLEDCRSR[GER]	Yes if data is not received for this request.	Same as first entry except error capture is done from original request packet.	OCN error response is generated to requestor if data is not forwarded to it. Else the RapidIO packet is dropped.

Table continues on the next page...

**Table 15-331. Hardware errors for IO/GSM response transactions (not maintenance)
(continued)**

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Status Not UR IO Error Response	Yes if LTLEECRSR[IER] is set	LTLEDCRSR[IER]	Yes	Same as first entry except error capture is done from original request packet.	OCN error response is generated to requestor.
TargetTID No outstanding transaction for this TargetTID	Yes if LTLEECRSR[UR] is set	LTLEDCRSR[UR]	No	Same as first entry	RapidIO packet is dropped and ignored.
Packet Size Not UR (All non-maintenance and non-message) Write response - Header size is not 8 Bytes for Small Transport packet or not 12 Bytes for Large Transport packet GSM - "Done" response packet size to "Flush" is not 8 Bytes for Small Transport packet or not 12 Bytes for Large Transport packet. "Done-Intervention" is not 8 Bytes for Small Transport and 12 Bytes for Large Transport field. Two byte padding of 0's in Large Transport field packet is not checked.	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored.
Payload Size Not UR IO - Read Response - total payload is not of the size requested. "Done" or "Done-Data_Only" response to IO_Read_Home with incorrect payload size. Response with transaction type "response_with_no_data" has payload	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored.

Table continues on the next page...

**Table 15-331. Hardware errors for IO/GSM response transactions (not maintenance)
(continued)**

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Retry Not UR GSM request has had one more than configured number of retries for non misaligned request. The misaligned GSM request has had one to four (cumulative for the corresponding child requests) more than configured number of retries.	Yes if LTLEECRSR[RETE] is set	LTLEDCRSR[RETE]	Yes	Same as first entry except error capture is done from original request.	OCN error response is generated to requestor.
Packet response time-out Response is not received by configured time for packets requiring RapidIO response. "GSM - IO_Read_Home" - Done_With_Data is not received in configured time when Done_Intervention is received for non misaligned request or last child of misaligned request. Done response is not received in configured time for non misaligned request or last child of misaligned request. EME capture occurs for each child packet response time-out.	Yes if LTLEECRSR[PRT] is set	LTLEDCRSR[PRT]	Yes	Same as first entry except error capture is done from original request.	OCN error response is generated to requestor.

Table continues on the next page...

Table 15-331. Hardware errors for IO/GSM response transactions (not maintenance) (continued)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
<p>Packet response time-out</p> <p>Response is not received by configured time for packets requiring RapidIO response.</p> <p>"GSM - IO_Read_Home" - Done_Intervention is not received in configured time when Done_With_Data is received. This is true for both non misaligned or misaligned requests.</p>	Yes if LTLEECRSR[PRT] is set	LTLEDCRSR[PRT]	No	Same as first entry except error capture is done from original request.	An OCN done response is generated when the Done_With_Data is received for non misaligned requests or the last child of a misaligned request. Therefore, an error response cannot be sent when the packet response time-out occurs.
<p>GSM - IO_Read_Home</p> <p>Not UR</p> <p>Done response, Retry response, or Error response is after Done_Intervention response or Data_only is received.</p>	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored.
<p>Not UR</p> <p>Response for OCN packets not requiring response, but converted to "response" type packet by ATMU receives Error response. For example, NWrite converted to NWrite_r received "Error" response</p>	Yes if LTLEECRSR[IER]/[GER]/[RETE] is set	LTLEDCRSR[IER]/[GER]/[RETE]	No	Same as first entry except error capture is done from original request.	RapidIO packet is dropped and ignored.
<p>Response for OCN packets not requiring response, but converted to "response" type packet by ATMU is not received by configured time.</p>	Yes if LTLEECRSR[PRT] is set	LTLEDCRSR[PRT]	No	Same as first entry except error capture is done from original request.	No error response is generated.

Table 15-332. Hardware errors for DMA message response transactions

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Not UR Response priority is not higher than RapidIO request priority	Yes if LTLEECR[ITD] is set	LTLEDCSR[ITD]	No	Using the incoming RapidIO packet, for Small Transport type packets, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDDIDCCSR[DIDMSB] gets 0's, LTLDDIDCCSR[DID] gets packet bits 16-23, LTLDDIDCCSR[SIDMSB] gets 0's, LTLDDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets packet bits 40-47 For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDDIDCCSR[DIDMSB] gets 16-23, LTLDDIDCCSR[DID] gets packet bits 24-31, LTLDDIDCCSR[SIDMSB] gets bits 32-39, LTLDDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets packet bits 56-63	RapidIO packet is dropped and ignored
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCSR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored
Received TT which is not enabled. Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCSR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestId does not match either Alternate DeviceID or DeviceId if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCSR[ITTE]	No	Same as first entry	RapidIO packet is dropped and ignored
SourceID Does not match the request's DestID	Yes if LTLEECR[UR] is set	LTLEDCSR[UR]	No	Same as first entry	RapidIO packet is dropped and ignored

Table continues on the next page...

Table 15-332. Hardware errors for DMA message response transactions (continued)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
TransactionType Not checked. To be a message response TType has to be 0x1	-	-	-	-	-
Status Not UR DMA Message Error response	Yes if LTLEECS R[DMAME R] is set	LTLEDCS R[DMAME R]	Yes	Same as first entry except error capture is done from original request	OCN error response is generated to requestor.
Status Not UR Received status of reserved type	Yes if LTLEECS R[ITD] is set	Yes if LTLEDCS R[ITD] is set	No	Same as first entry	RapidIO packet is dropped and ignored
No outstanding transaction for this letter, mailbox and message segment	Yes if LTLEECS R[UR] is set	LTLEDCS R[UR]	No	Same as first entry	RapidIO packet is dropped and ignored.
Packet Size Not UR (All non-maintenance and non-message) DMA response - Header size is not 8 Bytes for Small Transport packet or not 12 Bytes for Large Transport packet Two byte padding of 0's in Large Transport field packet is not checked.	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored.
Payload Size Not UR Payload size is not zero	Yes if LTLEECS R[ITD] is set	LTLEDCS R[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored.
Retry Not UR DMA request - has had more than configured number of retries	Yes if LTLEECS R[RETE] is set	LTLEDCS R[RETE]	Yes	Same as first entry except error capture is done from original request	OCN error response is generated to requestor.

Table continues on the next page...

Table 15-332. Hardware errors for DMA message response transactions (continued)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Packet response time-out Response is not received by configured time.	Yes if LTLEECR[PRT] is set	LTLEDCSR[PRT]	Yes	Same as first entry except error capture is done from original request	OCN error response is generated to requestor.

Table 15-333. Hardware errors for message request transactions

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Not Applicable	-	-	-	-	-
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCSR[TSE]	No	Using the original request RapidIO packet, for small Transport type, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDIDCCSR[DIDMSB] gets 0's, LTLDIDCCSR[DID] gets packet bits 16-23, LTLDIDCCSR[SIDMSB] gets 0's, LTLDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets packet bits 40-47. For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDIDCCSR[DIDMSB] gets 16-23, LTLDIDCCSR[DID] LTL gets packet bits 24-31, LTLDIDCCSR[SIDMSB] gets bits 32-39, LTLDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets packet bits 56-63.	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCSR[TSE]	No	Same as third entry	RapidIO packet is dropped

Table continues on the next page...

Table 15-333. Hardware errors for message request transactions (continued)

Error	Interrupt Generated	Status Bit Set	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
<p>DestID</p> <p>DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled.</p> <p>Error valid when (passthrough or accept_all) is false</p>	Yes if LTLEECRSR[ITTE] is set	LTLEDCRSR[ITTE]	Yes if priority is not 3. Else packet is dropped	Same as third entry	OCN error response is sent to self if request priority is not 3. Else packet is dropped.
<p>SourceID</p> <p>Not Checked for error.</p>	-	-	-	-	-
<p>MsgLen, Ssize, Ltr, Mbox, MsgSeg</p> <p>Not Checked for error.</p>	-	-	-	-	-
<p>PayloadSize</p> <p>Message payload size is larger than the specified ssize, or is of size 0 when seg_len == msg_len.</p> <p>Or</p> <p>Message payload size is not equal to specified ssize when seg_len != msg_len.</p>	Yes if LTLEECRSR[MFE] is set	LTLEDCRSR[MFE]	Yes if priority is not 3. Else packet is dropped	Same as third entry.	OCN error response is sent to self if request priority is not 3. Else packet is dropped.
<p>Reserved ssize field</p>	Yes if LTLEECRSR[MFE] is set	LTLEDCRSR[MFE]	Yes if priority is not 3. Else packet is dropped	Same as third entry.	OCN error response is sent to self if request priority is not 3. Else packet is dropped.
<p>Other</p> <p>Received Message request with SOCAR[M] disabled</p>	Yes if LTLEECRSR[UT] is set	LTLEDCRSR[UT]	Yes if priority is not 3. Else packet is dropped	Same as third entry	OCN error response is sent to self if request priority is not 3. Else packet is dropped.

Table 15-334. Hardware errors for message response transactions

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Not Checked for error.	-	-	-	Using the original request RapidIO packet, for small Transport type, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDIDCCSR[DIDMSB] gets 0's, LTLDIDCCSR[DID] gets packet bits 16-23, LTLDIDCCSR[SIDMSB] gets 0's, LTLDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets packet bits 40-47. For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDIDCCSR[DIDMSB] gets 16-23, LTLDIDCCSR[DID] LTL gets packet bits 24-31, LTLDIDCCSR[SIDMSB] gets bits 32-39, LTLDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets packet bits 56-63.	-
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLLEDCSR[TSE]	No	Same as first entry except capture registers are loaded from the response RapidIO packet	RapidIO packet is dropped and ignored
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLLEDCSR[TSE]	No	Same as first entry except capture registers are loaded from the response RapidIO packet	RapidIO packet is dropped and ignored
DestID (All non-maintenance) DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLLEDCSR[ITTE]	No	Same as first entry except capture registers are loaded from the response RapidIO packet	RapidIO packet is dropped and ignored
SourceID Not Checked for error.	-	-	-	-	-

Table continues on the next page...

Table 15-334. Hardware errors for message response transactions (continued)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Status Not Checked for error.	-	-	-	-	-
Other Received Message response with SOCAR[M] disabled.	Yes if LTLEECR[UR] is set	LTLEDCR[UR]	No	Same as first entry except capture registers are loaded from the response RapidIO packet	RapidIO packet is dropped and ignored

Table 15-335. Hardware errors for doorbell request transaction

Error	Interrupt Generated	Status Bit Set	Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Not Applicable	-	-	-	-	-
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Using the original request RapidIO packet, for small Transport type, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDDCCSR[DIDMSB] gets 0's, LTLDDCCSR[DID] gets packet bits 16-23, LTLDDCCSR[SIDMSB] gets 0's, LTLDDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDDCCSR[DIDMSB] gets 16-23, LTLDDCCSR[DID] LTL gets packet bits 24-31, LTLDDCCSR[SIDMSB] gets bits 32-39, LTLDDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as third entry	RapidIO packet is dropped

Table continues on the next page...

Table 15-335. Hardware errors for doorbell request transaction (continued)

Error	Interrupt Generated	Status Bit Set	Error Response Generated	Logical/Transport Layer Capture Register	Comments
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECRSR[ITTE] is set	LTLEDCSR[ITTE]	Yes if priority is not 3. Else packet is dropped	Same as third entry	OCN error response is sent to self if request priority is not 3. Else packet is dropped.
SourceID Not Checked for error.	-	-	-	-	-
SrcTID Not Checked for error.	-	-	-	-	-
Other Received Doorbell request with DOCAR[D] disabled.	Yes if LTLEECRSR[UT] is set	LTLEDCSR[UT]	Yes if priority is not 3. Else packet is dropped	Same as third entry	OCN error response is sent to self if request priority is not 3. Else packet is dropped.

Table 15-336. Hardware errors for doorbell response transactions

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Not UR or UT Response priority is not higher than RapidIO request priority	Yes if LTLEECR[ITD] is set	LTLEDCR[ITD]	No	Using the incoming RapidIO packet, for small Transport type, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDIDCCSR[DIDMSB] gets 0's, LTLDIDCCSR[DID] gets packet bits 16-23, LTLDIDCCSR[SIDMSB] gets 0's, LTLDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's For Large Transport type packets r. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDIDCCSR[DIDMSB] gets 16-23, LTLDIDCCSR[DID] LTL gets packet bits 24-31, LTLDIDCCSR[SIDMSB] gets bits 32-39, LTLDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped and ignored
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored
Received TT Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped and ignored
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	No	Same as first entry	RapidIO packet is dropped and ignored

Table continues on the next page...

Table 15-336. Hardware errors for doorbell response transactions (continued)

Error	Interrupt Generated	Status Bit Set	OCN Error Response Generated	Logical/Transport Layer Capture Register	Comments
SourceID Does not match the request's DestID	Yes if LTLEECRSR[UR] is set	LTLEDCRSR[UR]	No	Same as first entry	RapidIO packet is dropped and ignored
Status Not UR or UT Not one of Done/Error/Retry	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
TransactionType Not UR or UT Anything other than Done_No_Data	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored
TargetTID No outstanding transaction for this TargetTID	Yes if LTLEECRSR[UR] is set	LTLEDCRSR[UR]	No	Same as first entry	RapidIO packet is dropped and ignored.
Packet Size Not UR or UT Header size is not 8 Bytes for Small Transport packet or not 12 Bytes for Large Transport packet Note: Two byte padding of 0's in Large Transport field packet is not checked.	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as first entry	RapidIO packet is dropped and ignored.
Packet response timeout Response is not received by configured time	Yes if LTLEECRSR[PRT] is set	LTLEDCRSR[PRT]	Yes	Same as first entry except capture registers are loaded from original request RapidIO packet.	OCN doorbell PRT response is generated to requestor.

Table 15-337. Hardware errors for PortWrite transaction

Error	Interrupt Generated	Status Bit Set	Error Response Generated	Logical/Transport Layer Capture Register	Comments
Priority Not Applicable	-	-	-	-	-
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Using the original request RapidIO packet, for small Transport type, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDIDCCSR[DIDMSB] gets 0's, LTLDIDCCSR[DID] gets packet bits 16-23, LTLDIDCCSR[SIDMSB] gets 0's, LTLDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDIDCCSR[DIDMSB] gets 16-23, LTLDIDCCSR[DID] LTL gets packet bits 24-31, LTLDIDCCSR[SIDMSB] gets bits 32-39, LTLDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as third entry	RapidIO packet is dropped
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	No	Same as third entry	RapidIO packet is dropped
SourceID Not Checked for error.	-	-	-	-	-
TransactionType Not Checked for error.	-	-	-	-	-

Table continues on the next page...

Table 15-337. Hardware errors for PortWrite transaction (continued)

Error	Interrupt Generated	Status Bit Set	Error Response Generated	Logical/Transport Layer Capture Register	Comments
WrSize Not unsupported transaction Is one of reserved sizes or less than 4 bytes	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as third entry	RapidIO packet is dropped
SrcTID Not Checked for error.	-	-	-	-	-
HopCount Not Checked for error.	-	-	-	-	-
ConfigOffset Not Checked for error.	-	-	-	-	-
PayloadSize Not unsupported transaction An incorrect port-write wr_size encoding (not 4, 8, 16, 24, 32, 40, 48, 56 or 64 bytes). Payload size is greater than the value defined by wr_size. Payload size is not dword aligned when the wr_size is not 4 bytes.	Yes if LTLEECRSR[ITD] is set	LTLEDCRSR[ITD]	No	Same as third entry	RapidIO packet is dropped
Other Received PortWrite transaction with DOCAR[PW] disabled.	Yes if LTLEECRSR[UT] is set	LTLEDCRSR[UT]	No	Same as third entry	RapidIO packet is dropped

Table 15-338. Hardware errors for reserved Ftype

Error	Interrupt Generated	Status Bit Set if corresponding bit is enabled	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Ftype Ftype is not IO Read, IO Write, SWrite, Maintenance request, Maintenance Response, Response (Ftype 13), Doorbell or Message class and it is not a passthrough transaction. (passthrough is not enabled or accept_all is enabled or transaction is addressed to this port)	Yes if LTLEECR[UT] is set	LTLEDCR[UT]	No	Using the original request RapidIO packet, for small Transport type, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLIDCCSR[DIDMSB] gets 0's, LTLIDCCSR[DID] gets packet bits 16-23, LTLIDCCSR[SIDMSB] gets 0's, LTLIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's Large Transport type packets. LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLIDCCSR[DIDMSB] gets 16-23, LTLIDCCSR[DID] LTL gets packet bits 24-31, LTLIDCCSR[SIDMSB] gets bits 32-39, LTLIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's	RapidIO packet is dropped
TransportType Received reserved TT	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
Received TT which is not enabled. - Error valid when passthrough is disabled and accept_all is disabled Or when accept_all is enabled.	Yes if LTLEECR[TSE] is set	LTLEDCR[TSE]	No	Same as first entry	RapidIO packet is dropped
DestID DestID does not match this port's DeviceID if Alternate DeviceID is disabled or DestID does not match either Alternate DeviceID or DeviceID if Alternate DeviceID is enabled. Error valid when (passthrough or accept_all) is false	Yes if LTLEECR[ITTE] is set	LTLEDCR[ITTE]	No	Same as first entry	RapidIO packet is dropped

Table continues on the next page...

Table 15-338. Hardware errors for reserved Ftype (continued)

Error	Interrupt Generated	Status Bit Set if corresponding bit is enabled	RapidIO Error Response Generated	Logical/Transport Layer Capture Register	Comments
Address:WdPtr:Xambs Request hits overlapping ATMU windows. Refer to Outbound window boundary crossing errors Packet checked as non Swrite packet	Yes if LTLEECR[IACB] is set	LTLEDCR[IACB]	No	Same as first entry	RapidIO packet is dropped
Address:WdPtr:Xambs Not unsupported transaction Request hits a protected ATMU window or the LCSBA1CSR Packet checked as non Swrite packet	Yes if LTLEECR[ITD] is set	LTLEDCR[ITD]	No	Same as first entry	RapidIO packet is dropped

Table 15-339. Hardware errors for outbound transaction crossed ATMU boundary

Error	Interrupt Generated	Status Bit Set if corresponding bit is enabled		Logical/Transport Layer Capture Register	Comments
<p>OCN Address and payload size</p> <p>OCN address range for Outbound RapidIO transaction hits multiple ATMU windows.</p> <p>Refer to Outbound window boundary crossing errors</p>	<p>Yes if LTLEECR[OACB] is set and/or LTLEECR[PRT] is set</p>	<p>LTLEDCSR[OACB], LTLEDCSR[PRT]</p>		<p>Using the original request RapidIO packet send out by OB, for small Transport type, LTLACCSR[XA] gets packet bits 78-79, LTLACCSR[A] gets packet bits 48-76, LTLDIDCCSR[DIDMSB] gets 0's, LTLDIDCCSR[DID] gets packet bits 16-23, LTLDIDCCSR[SIDMSB] gets 0's, LTLDIDCCSR[SID] gets packet bits 24-31, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 32-35, LTLCCCSR[MI] gets 0's</p> <p>For Large Transport type packets, LTLACCSR[XA] gets packet bits 94-95, LTLACCSR[A] gets packet bits 64-92, LTLTLTLDIDCCSR[DIDMSB] gets 16-23, LTLDIDCCSR[DID] LTL gets packet bits 24-31, LTLDIDCCSR[SIDMSB] gets bits 32-39, LTLDIDCCSR[SID] gets bits 40-47, LTLCCCSR[FT] gets packet bits 12-15, LTLCCCSR[TT] gets packet bits 48-51, LTLCCCSR[MI] gets 0's</p>	<p>-</p>

memory while software processes them sequentially. The depths of the queues in the producer and consumer are configurable by software. A multicast function allows single segment messages to be sent to multiple consumers.

The message unit is compliant with the message passing logical specification contained in the *RapidIO Interconnect Specification, Revision 1.2* (but assumes the msgseg field is redefined for more mailboxes when single segment messages are being sent). The most common use of the message passing model is in systems where a processing element is only allowed to access memory that is local to itself, and communication between processing elements is achieved through message passing and communication is address independent.

Each inbound message controller has a dedicated interrupt that can be used to notify software that a configured number of messages have been received. Similarly, each outbound message controller has a dedicated interrupt that can be used to notify software that there are no more messages for the outbound mailbox controller to process.

The message controller is managed through a set of run-time registers.

15.7.2 Message unit features

- Support for two outbound message controllers with the following features
 - chaining and direct modes
 - extended mailboxes (XMBOX) for single segment messages
 - multicast up to 32 RapidIO destinations for single segment messages
 - transmitting to any mailbox and extended mailbox for a single segment message (letter 3 is reserved for an alternative message source like the DMA but can be used if the DMA does not generate messages and the RapidIO endpoint is configured appropriately)
 - transmitting to any mailbox for multi segment message (letter 3 is reserved for an alternative message source like the DMA, but can be used if the DMA does not generate messages and the RapidIO endpoint is configured appropriately)
 - segment size up to 256 bytes
 - up to 16 segment messages with a total payload of up to 4 Kbytes
 - one entire message (up to a 16-message segments) can be transmitted before receiving any response
 - one entire single segment message to all multicast destinations (up to 32) can be transmitted before receiving any response
 - all message segment transfers for a message transaction must complete before the next message transaction begins

- pipelined transmission of a full message in each message controller but all responses must be received before the next message can be transmitted
- in chaining mode the next descriptor can be fetched before the current message completes (descriptor prefetching)
- Support for two inbound message controllers with the following features
 - reception of any mailbox and letter for a single or multi segment message
 - segment size up to 256 bytes
 - up to 16 segment messages with a total payload of up to 4 Kbytes
 - full inbound line rate performance
 - back-to-back message reception of the same or lower priority
 - out-of-order message segment reception
 - concurrent inbound message controller operation

15.7.3 Outbound modes of operation

- Direct mode: Software is expected to program all the necessary registers for sending an outbound message.
- Chaining mode: A descriptor that describes the message information is fetched from local memory before a message is sent.
- Multicast mode: A single segment message (256 bytes or less) is sent to multiple destinations. This mode is supported in direct or chaining mode.

The following sections describe the structure and operation of the outbound message controller and inbound message controller hardware in the data message controller.

15.7.4 Outbound message controller operation

The outbound message controller is responsible for sending messages stored in local memory.

The outbound message controller supports three modes of operation, direct, chaining, and multicast mode. In direct mode, software programs the necessary registers to point to the beginning of the message in memory. In chaining mode, software programs the necessary registers to point to the beginning of the first valid descriptor in memory. The descriptor provides all the necessary registers to start the message transfer. In multicast mode, a single segment message can be sent to multiple destinations. Multicast mode is supported in direct or chaining mode.

Each outbound message controller uses a unique letter number. For example, if there are two outbound message controllers, message controller 0 uses letter 0 and message controller 1 uses letter 1. The mailbox number is programmable through the destination port register, $OMnDPR$. Each message controller can only operate on one message at a time.

15.7.4.1 Direct mode operation

In direct mode ($OMnMR[MUTM]$ is set in the outbound message mode register) the outbound message controller does not read descriptors from memory, but instead uses the current parameters programmed in the outbound message controller registers to start the transfer.

In direct mode, software is responsible for initializing all the parameters in all the necessary registers to start the message transmission. The message transfer is started when the outbound message controller start bit, $OMnMR[MUS]$, in the outbound message mode register transitions from a 0 to 1 and the outbound message controller is not already busy. If the outbound message controller is already busy, $OMnMR[MUS]$ transitioning from 0 to 1 is ignored. Software is expected to program all the appropriate registers before setting $OMnMR[MUS]$.

There are many ways in which software can interact with the message controller. One sequence of events to start and complete a transfer in direct mode is as follows:

- Poll the status register message unit busy bit, $OMnSR[MUB]$, to make sure the outbound message controller is not busy with a previously initiated message.
- Clear the status bits ($OMnSR[MER]$, $OMnSR[RETE]$, $OMnSR[PRT]$, $OMnSR[TE]$, $OMnSR[QOI]$, $OMnSR[QFI]$, $OMnSR[EOMI]$, and $OMnSR[QEI]$)
- Initialize the source address ($EOMnSAR$, $OMnSAR$), destination port ($OMnDPR$), destination attributes ($OMnDATR$), retry error threshold ($OMnRETCR$), and double-word count ($OMnDCR$) registers. If multicast mode is enabled ($OMnDATR[MM]$) initialize the multicast group and list ($OMnMGR$, $OMnMLR$).
- Initialize the outbound message mode register message unit transfer mode bit, $OMnMR[MUTM] = 1$, to indicate direct mode. Other control parameters must also be initialized in the mode register.
- Clear, then set the mode register message unit start bit, $OMnMR[MUS]$, to start the message transfer.
- $OMnSR[MUB]$ bit is set by the outbound message controller to indicate the message transfer is in progress.
- The outbound message controller reads a message segment from local memory using the source address register ($EOMnSAR$, $OMnSAR$).

- If a message has multiple segments, the outbound message controller reads the other message segments from local memory.
- After the message read to local memory completes, the message is sent.
- The $OMnSR[MUB]$ is cleared by the outbound message controller after the message operation completes. A non multicast message transfer completes after all message segments complete. A multicast message transfer completes after all message segments complete for each destination. A message segment completes when one of the following occurs:
 - done response received
 - error response
 - a packet response time-out received
 - retry error threshold exceeded
 - an internal error occurs during the local memory access
- After the outbound message operation completes, the outbound message interrupt is generated if the end of message outbound message interrupt event is enabled ($OMnDATR[EOMIE]$).

15.7.4.1.1 Direct mode interrupts

The outbound message controller interrupt can be generated for one reason in direct mode.

- End-Of-Message - generates an interrupt after the completion of a message if this interrupt event is enabled ($OMnDATR[EOMIE]$ is a 1). $OMnSR[EOMI]$ provides the event that caused this interrupt. The interrupt is held until the $OMnSR[EOMI]$ bit is cleared by writing a 1.

The error/port-write interrupt can be generated for the following reasons in direct mode.

- message error response - an interrupt is generated after a message error response is received and this interrupt event is enabled ($OMnMR[EIE]$)
- packet response time-out - an interrupt is generated after a packet response time-out occurs and this interrupt event is enabled ($OMnMR[EIE]$)
- retry error threshold exceeded - an interrupt is generated after a retry threshold exceeded error occurs and this interrupt event is enabled ($OMnMR[EIE]$)
- transaction error - an interrupt is generated after an OCN error response is received and this interrupt event is enabled ($OMnMR[EIE]$)

15.7.4.1.2 Direct mode message error response errors

When a message error response is received by the message controller the following occurs.

- The message controller sets the message error response status bit (OM_nSR[MER])
- If OM_nMR[EIE] is set, the interrupt SRIO error/port-write is generated.
- After the message operation completes (indicated by OM_nSR[MUB]) the message controller stops

15.7.4.1.3 Direct mode packet response time-out errors

When a packet response time-out occurs for a message segment the following occurs.

- The message controller sets the packet response time-out status bit (OM_nSR[PRT])
- If OM_nMR[EIE] is set, the interrupt SRIO error/port-write is generated.
- After the message operation completes (indicated by OM_nSR[MUB]) the message controller stops

15.7.4.1.4 Direct mode retry error threshold exceeded errors

When a retry error threshold exceeded error occurs for a message segment the following occurs.

- The message controller sets the retry threshold exceed status bit (OM_nSR[RETE])
- If OM_nMR[EIE] is set, the interrupt SRIO error/port-write is generated.
- After the message operation completes (indicated by OM_nSR[MUB]) the message controller stops

15.7.4.1.5 Direct mode transaction errors

When an internal error occurs during a local memory read by the message controller the following occurs.

- The message controller sets the transaction error bit (OM_nSR[TE])
- Message segments that have an internal error are not sent because the message data is not available
- Memory reads that already were generated before the internal error occurred are also not transferred
- Additional memory reads for the same message operation are generated but not transferred
- All subsequent message segments for the same message operation are not transferred. This includes retried message segments.
- After the message operation completes (indicated by OM_nSR[MUB]) the message controller stops
- If OM_nMR[EIE] is set, the interrupt SRIO error/port-write is generated

15.7.4.1.6 Direct mode error handling

When an error occurs and the SRIO error/port-write interrupt is generated, the following occurs.

- Software determines the cause of the interrupt and processes the error
- Software verifies the message controller has stopped operation by polling OM n SR[MUB]
- Software disables the message controller by clearing OM n MR[MUS]
- Software clears the error by writing a 1 to the corresponding status bit (OM n SR[MER], OM n SR[PRT], OM n SR[RETE], or OM n SR[TE])

When an error occurs and the SRIO error/port-write interrupt is not enabled, the following occurs.

- Software determines that an error has occurred by polling the status bits (OM n SR[MER], OM n SR[PRT], OM n SR[RETE], or OM n SR[TE])
- Software verifies the message controller has stopped operation by polling OM n SR[MUB]
- Software disables the message controller by clearing OM n MR[MUS]
- Software clears the error by writing a 1 to the corresponding status bit (OM n SR[MER], OM n SR[PRT], OM n SR[RETE], or OM n SR[TE])

15.7.4.1.7 Disabling and enabling the message controller

Once the message controller is started, it cannot be stopped.

15.7.4.1.8 Direct mode hardware error handling

The following list possible error conditions and what occurs when they are encountered.

The error checking level indicates the order in which errors are checked. Multiple errors can be checked at an error checking level. Once an error is detected no additional error checking beyond the current level is performed. The first error detected in the processing pipeline updates the error management extensions registers.

These error condition checks are provided by the messaging unit. These checks are in addition to the error condition checks provided by the RapidIO port given in [Errors and error handling](#).

Table 15-341. Outbound message direct mode hardware errors

Transaction	Error	Error Checking Level	Interrupt Generated	Status Bit Set	Message Segment Sent	Logical/Transport Layer Capture Register	Comments
Message Request	An internal error occurred during a read of the message segment from local memory	1	SRIO error/port-write if OMnMR[EE] set	Transaction error in the outbound message status register (OMnSR[TE]). Message Failed in the Mailbox CSR (MCSR[FA]).	No	-	Message controller stops after the current message operation completes. The descriptor dequeue pointer is not incremented in chaining mode.
Message Request	An internal error occurred for an earlier message segment local memory read. An internal error for a subsequent message segment local memory read for the same message may or may not occur.	2	No	None	Yes	-	Message controller stops after the current message operation completes.
Undefined Packet	Reserved ftype encoding ¹	3	SRIO error/port-write if LTLEECR[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UT]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Message Response	Reserved tt encoding ¹	3	SRIO error/port-write if LTLEECR[TSE] set	Illegal transaction target in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITTE] Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[TSE]	Yes	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-341. Outbound message direct mode hardware errors (continued)

Transaction	Error	Error Checking Level	Interrupt Generated	Status Bit Set	Message Segment Sent	Logical/Transport Layer Capture Register	Comments
Message Response	Large transport size when operating in small transport size or small transport size when operating in large transport size ¹	3	SRIO error/port-write if LTLEECR[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[TSE]	Yes	Updated with the packet ²	Packet is ignored and discarded. An error or illegal transaction target error response is not generated.
Message Response	Illegal Destination ID ¹	3	SRIO error/port-write if LTLEECR[ITTE] set	Illegal transaction target in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITTE]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Message Response	type (transaction field) is not message response ¹	3	SRIO error/port-write if LTLEECR[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Message Response	Message response received and no outbound mailboxes are supported ¹	3	SRIO error/port-write if LTLEECR[UR] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UR]	No	Updated with the packet ²	Packet is ignored and discarded.
Message Response	reserved response status (not done, retry, or error)	4a	SRIO error/port-write if LTLEECR[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Message Response	message response packet size is incorrect	4a	SRIO error/port-write if LTLEECR[ITD] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Message Response	Incorrect Source ID	4b	SRIO error/port-write if LTLEECR[UR] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-341. Outbound message direct mode hardware errors (continued)

Transaction	Error	Error Checking Level	Interrupt Generated	Status Bit Set	Message Segment Sent	Logical/Transport Layer Capture Register	Comments
Message Response	letter, mbox and msgseg not outstanding or letter, mbox and xmbx not outstanding	4b	SRIO error/port-write if LTLEECSR[UR] set	Unsolicited response in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UR]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Message Response	RapidIO priority is less than or equal to message request	4c	SRIO error/port-write if LTLEECSR[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Message Response	error response	5	SRIO error/port-write if LTLEECSR[MER] set. SRIO error/port-write if OMnMR[EE].	Message error response in the Logical/Transport Layer Error Detect CSR LTLEDCSR[MER]. OMnSR[MER] bit set if in direct mode or if in chaining mode.	Yes	Updated with the corresponding message request packet ²	Message segment transfer complete. The descriptor dequeue pointer is not incremented in chaining mode.
Message Response	number of retries exceeds limit	5	SRIO error/port-write if LTLEECSR[RETE] set. SRIO error/port-write if OMnMR[EE].	Retry error threshold exceeded in the Logical/Transport Layer Error Detect CSR LTLEDCSR[RETE]. OMnSR[RETE] bit set if in direct mode or if in chaining mode.	Yes	Updated with the corresponding message request packet ²	Message segment transfer complete. The descriptor dequeue pointer is not incremented in chaining mode.
Message Response	packet response time-out	unrelated	SRIO error/port-write if LTLEECSR[PRT] set. SRIO error/port-write if OMnMR[EE].	Packet response time-out in the Logical/Transport Layer Error Detect CSR LTLEDCSR[PRT]. OMnSR[PRT] bit set if in direct mode or if in chaining mode.	Yes	Updated with the corresponding message request packet ² . The LTLIDCCSR[SIDMSB] and LTLIDCCSR[SID] field is 0.	Message segment transfer complete. The descriptor dequeue pointer is not incremented in chaining mode.

1. These error types are actually detected in the RapidIO port, not in the message controller
2. If operating in small transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 78-79), LTLACCSR[A] gets the address (packet bits 48-76), LTLIDCCSR[MDID] gets 0, LTLIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 16-23), LTLIDCCSR[MSID] gets 0, LTLIDCCSR[SID] gets the least significant byte of the source ID (packet bits 24-31), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 32-35), LTLCCCSR[MI] gets the msg info (packet bits 40-47). If operating in large transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 94-95),

LTLACCSR[A] gets the address (packet bits 64-92), LTLDIDCCSR[MDID] gets the most significant byte of the destination ID (packet bits 16-23), LTLDIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 24-31), LTLDIDCCSR[MSID] gets the most significant byte of the source ID (packet bits 32-39), LTLDIDCCSR[SID] gets the least significant byte of the source ID (packet bits 40-47), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 48-51) if the message request packet is captured or 0 if the message response packet is captured, LTLCCCSR[MI] gets the msg info (packet bits 56-63).

15.7.4.1.9 Direct mode programming errors

The following is the list of programming errors that result in undefined or undesired hardware operation.

Table 15-342. Outbound message direct mode programming errors

Error	Interrupt Generated	Status Bit Set	Comments
Double-word count greater than 256 bytes when multicast mode selected	No	None	Undefined operation results
Double-word count set to a reserved value	No	None	Undefined operation results
Transaction flow level set to 3	No	None	Undefined operation results
Target interface set to an invalid RapidIO port	No	None	Undefined operation results
Source address for message read is invalid	No	No	Local memory captures the transaction and generates an interrupt.
Address for error enqueue address pointer is invalid	No	No	Local memory captures the transaction and generates an interrupt.
Register values changed during operation	No	No	Undefined operation results

15.7.4.2 Chaining mode

In chaining mode, $OM_nMR[MUTM] = 0$ in the outbound message mode register, message descriptors are built in local memory in a circular queue.

There are several options available to the programmer in chaining mode. Software can build one or more descriptors in memory before initializing the outbound message controller registers, or it can initialize the outbound message controller registers and then build the descriptors. The enqueue pointer ($OM_nDQEPAR$, $EOM_nDQEPAR$) is maintained by software. The outbound message controller dequeues descriptors, processes them and increments the dequeue pointer ($OM_nDQDPAR$, $EOM_nDQDPAR$) to point to the next descriptor in the queue.

Figure 15-321 depicts a sample structure of the outbound portion of the message controller, and a descriptor queue, with each valid descriptor queue entry pointing to a valid message. In this example, the descriptor queue has eight entries, four of which are currently valid. The local processor enqueues descriptors and the outbound message controller dequeues the descriptors.

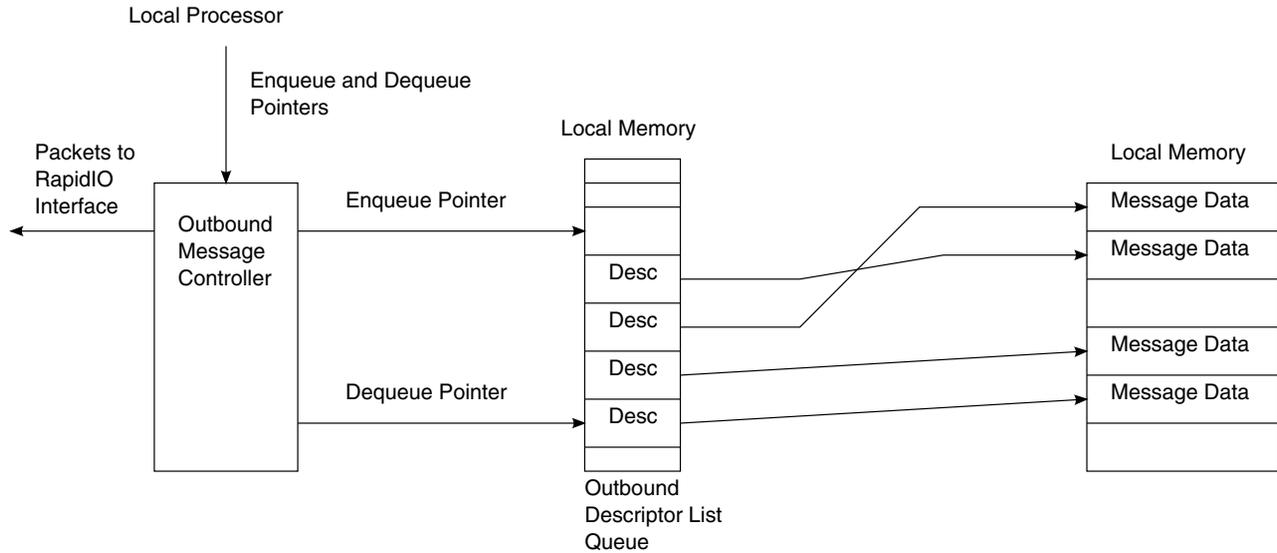


Figure 15-321. Outbound frame queue structure

15.7.4.2.1 Message controller initialization

There are many ways in which software can interact with the message controller.

One method to initialize the message controller is as follows:

- Poll the status register message unit busy bit, $OM_nSR[MUB]$, to make sure the outbound message controller is not busy with a previously initiated message.
- Clear the message unit start bit ($OM_nMR[MUS]$).
- Initialize the descriptor queue dequeue pointer address registers ($OM_nDQDPAR$, $EOM_nDQDPAR$) and the descriptor queue enqueue pointer address registers ($OM_nDQEPAR$, $EOM_nDQEPAR$). These need to be initialized to the same value for proper operation.
- These pairs of registers must also be queue size aligned (that is, the queue must be aligned on a boundary equal to number of queue entries x 32 bytes (size of each queue descriptor)). For example, if there are 16 entries in the queue, the queue must be 512-byte aligned.

The number of queue entries is set in $OM_nMR[CIRQ_SIZ]$; see [Outbound message n mode register \(SRIO_OM_nMR\)](#).

- Initialize the retry error threshold in the outbound message retry error threshold configuration register (OM n RETCCR).
- Clear OM n MR[MUTM] for chaining mode.
- If using single segment multicast mode, set OM n DATR[MM].
- Configure the other control parameters in the mode register (OM n MR).
- Clear OM n SR[MER, PRT, RETE, TE, QOI, QFI, EOMI and QEI]. If OM n SR[MER, PRT, RETE, TE or QOI] are not cleared, the message controller can not start a new message operation. Incorrect status will be indicated if the other status bits are not cleared.
- Set the message unit start (OM n MR[MUS]). This enables the outbound message controller and causes the descriptor dequeue pointer (OM n DQDPAR, EOM n DQDPAR) to be saved off as the base address of the descriptor queue.

15.7.4.2.2 Chaining mode operation

The method to start and complete transfers by adding descriptors after initializing the message unit is as follows:

- Create one or more descriptors in local memory starting at the address pointed to by the descriptor queue enqueue pointer address register (OM n DQEPAR, EOM n DQEPAR)
- Either increment the enqueue pointer address registers (OM n DQEPAR, EOM n DQEPAR) by setting OM n MR[MUI] for each descriptor entry added or directly change the enqueue pointer address register (OM n DQEPAR). If OM n MR[MUI] is set by software, the message controller clears this bit after successfully incrementing the enqueue pointer.
- When the descriptor queue is not empty, the message controller reads the descriptor from local memory using the address pointed to by the dequeue pointer (OM n DQDPAR, EOM n DQDPAR) and sets the busy bit (OM n SR[MUB]).
- If another descriptor is available, the message controller reads the next descriptor from local memory using the address pointed to by the dequeue pointer (OM n DQDPAR, EOM n DQDPAR). The message controller can not prefetch more than one descriptor.
- The message controller sets OM n SR[MUB] to indicate that the message transfer is in progress. OM n SR[MUB] remains set until the descriptor queue is empty or a transaction error occurs.
- Additional descriptors can be created and enqueued by software while the message controller is busy (OM n SR[MUB]). Software can continue adding descriptors as long as the descriptor queue is not full. If software is adding descriptors using the OM n MR[MUI] bit, overflowing the queue can be prevented by polling the queue full bit (OM n SR[QF]) before creating and enqueueing the next descriptor.

- After the descriptor memory read completes, the corresponding message segment is read from local memory.
- If a message has multiple segments, the outbound message controller reads the other message segments from local memory.
- After the message read to local memory completes, the message is sent.
- If multicast is enabled, all of the indicated targets are sent the same message.
- A non multicast message transfer completes after all message segments complete. A multicast message transfer completes after all message segments complete for each destination. A message segment completes when one of the following occurs:
 - done response received
 - error response received
 - a packet response time-out occurred
 - retry error threshold exceeded
 - an internal error occurred during the descriptor (all message segments complete) or message read of local memory
- When processing for the current descriptor completes, if another descriptor is available the above steps are repeated.
- If a RapidIO error response is received, the message error response bit is set ($OM_nSR[MER]$) and outbound message controller operation stops after all message segments complete. If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.
- If a packet response time-out occurs, the packet response time-out bit is set ($OM_nSR[PRT]$). If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.
- If the retry error threshold value is exceeded for a specific segment, the retry error threshold exceeded bit is set ($OM_nSR[RETE]$) and outbound message controller operation stops after all message segments complete. If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.
- If an internal error occurs while reading local memory, the transaction error bit is set ($OM_nSR[TE]$) and outbound message controller operation stops after all message segments complete. If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.
- The above process continues until the descriptor queue is empty (dequeue pointer equals the enqueue pointer).
- The message unit clears $OM_nSR[MUB]$ after completing the processing of the last descriptor or a transaction error occurs.
- If an error occurs the message unit must be disabled, re-initialized and re-enabled before another message can be sent.

15.7.4.2.3 Changing descriptor queues in chaining mode

When software wants to switch to another descriptor queue in local memory, it must wait for the processing of the current queue to complete as indicated by the busy bit ($OMnSR[MUB]$).

Software then disables the message controller by clearing $OMnMR[MUS]$, changes the enqueue and dequeue descriptor pointers ($EOMnDQEPAR$, $OMnDQEPAR$ and $EOMnDQDPAR$, $OMnDQDPAR$) and re-enables the message unit by setting $OMnMR[MUS]$.

15.7.4.2.4 Preventing queue overflow in chaining mode

Software must guarantee that descriptors are not added to an already full queue.

When the increment bit is used ($OMnMR[MUI]$) software can poll the queue full bit ($OMnSR[QF]$) before enqueueing another descriptor. When software sets the enqueue pointer directly, software is responsible for not overflowing the descriptor queue.

15.7.4.2.5 Switching between direct and chaining modes

The message unit architecture allows switching from direct mode to chaining mode and vice-versa once all the required parameters have been initialized in the appropriate registers and when the message unit is not busy with a current transfer as indicated by $OMnSR[MUB]$ being cleared.

When switching from direct mode to chaining mode, if $OMnMR[MUS]$ is cleared and set, the message unit is re-initialized in chaining mode, and the outbound message descriptor queue dequeue pointer address is saved off as the new base address of the circular queue in memory. When switching from chaining mode to direct mode, $OMnMR[MUS]$ must also be cleared and set.

15.7.4.2.6 Chaining mode descriptor format

The descriptor contains information for the message unit controller to transfer data.

Software must ensure that each descriptor is aligned on a 32-byte boundary and that the descriptor queue is on a queue boundary (in other words, on a boundary equal to number of queue entries x 32 bytes (size of each queue descriptor)). For each descriptor in the queue, the message unit controller starts a new message operation with the control parameters specified by the descriptor.

Table 15-343. Outbound message unit descriptor summary

Descriptor Field	Description
Source Extended Address	Contains the source address of the message operation for local addresses of greater than 32 bits. After the message controller reads the descriptor from memory, this field is loaded into the source extended address register.
Source Address	Contains the source address of the message operation. After the message controller reads the descriptor from memory, this field is loaded into the source address register.
Destination Port	Contains the destination port of the message operation. After the message controller reads the descriptor from memory, this field is loaded into the destination port register.
Destination Attributes	Contains transaction attributes of the message operation. After the message controller reads the descriptor from memory, this field is loaded into the destination attributes register.
Multicast Group	Contains the logical multicast group. Groups are defined a list of 32 numerically consecutive destinations by deviceID.
Multicast List	Contains a bit vector list of consecutive destinations by deviceID.
Double-word Count	Contains the number of doublewords for the message operation. After the message controller reads the descriptor from memory, this field is loaded into the double-word count register.
Reserved	-

Figure 15-322 depicts the queue dequeue pointer and an associated descriptor. The descriptor is only valid if the enqueue and dequeue pointers are not equal.

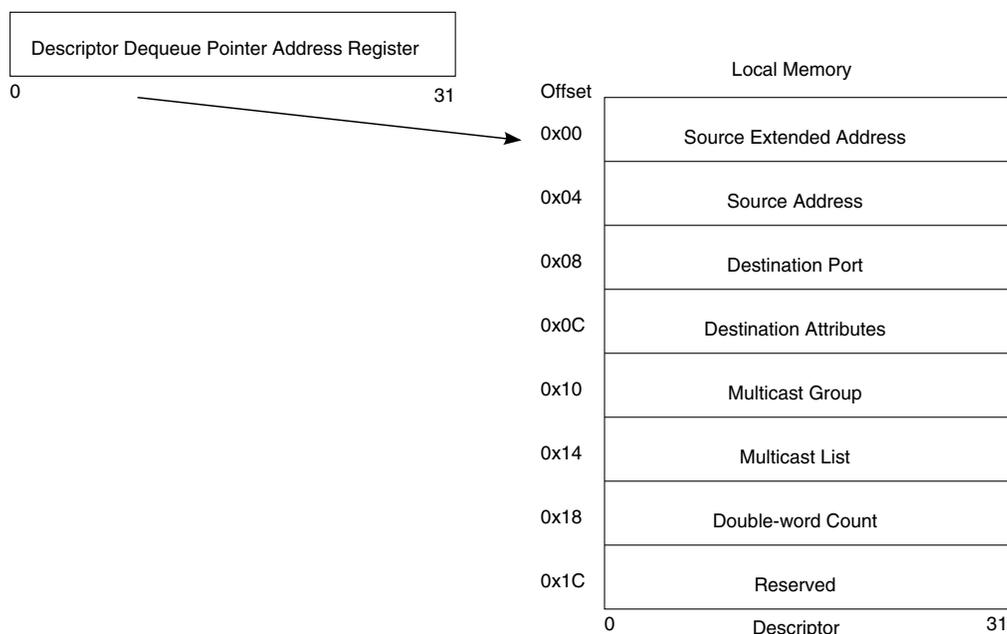


Figure 15-322. Descriptor dequeue pointer and descriptor

15.7.4.2.7 Chaining mode controller interrupts

The outbound message interrupt can be generated for several reasons.

- Queue empty-an interrupt is generated to the interrupt controller if the queue goes empty and the interrupt event is enabled ($OMnMR[QEIE]$ is a 1). The event that caused the outbound message interrupt is indicated by $OMnSR[QEI]$. The interrupt is held until the queue is no longer empty and the $OMnSR[QEI]$ bit is cleared by writing a 1.
- Queue full-an interrupt is generated to the interrupt controller if the queue is full and the interrupt event is enabled ($OMnMR[QFIE]$ is a 1). The event that caused the outbound message interrupt is indicated by $OMnSR[QFI]$. The interrupt is held until the queue is no longer full and the $OMnSR[QFI]$ bit is cleared by writing a 1.
- Queue overflow-an interrupt is generated to the interrupt controller if the queue is full, the increment bit is set ($OMnMR[MUI]$) and the interrupt event is enabled ($OMnMR[QOIE]$ is a 1). The event that caused the outbound message interrupt is indicated by $OMnSR[QOI]$. The message unit must be re-initialized. The interrupt is held until the $OMnSR[QOI]$ bit has been cleared by writing a 1. This interrupt is also generated if the enqueue pointer is directly written and causes an overflow.
- End-of-message-an interrupt is generated after the completion of the message if this interrupt event is enabled ($OMnDATR[EOMIE]$ is a 1). The event that caused this interrupt is indicated by $OMnSR[EOMI]$. The interrupt is held until the $OMnSR[EOMI]$ bit is cleared by writing a 1. This allows an interrupt to be generated after a particular descriptor has been processed.

The error/port-write interrupt can be generated for the following reasons in direct mode.

- Message error response-an interrupt is generated after a message error response is received and this interrupt event is enabled ($OMnMR[EIE]$)
- Packet response time-out-an interrupt is generated after a packet response time-out occurs and this interrupt event is enabled ($OMnMR[EIE]$)
- Retry error threshold exceeded-an interrupt is generated after a retry threshold exceeded error occurs and this interrupt event is enabled ($OMnMR[EIE]$)
- Transaction error-an interrupt is generated after a message error response is received and this interrupt event is enabled ($OMnMR[EIE]$)

15.7.4.2.8 Chaining mode message error response errors

When a message error response is received by the message controller the following occurs.

- The message controller sets the message error response status bit ($OMnSR[MER]$)

- If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.
- After the message operation completes (indicated by $OM_nSR[MUB]$) the message controller stops

15.7.4.2.9 Chaining mode packet response time-out errors

When a packet response time-out occurs for a message segment the following occurs.

- The message controller sets the packet response time-out status bit ($OM_nSR[PRT]$)
- If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.
- After the message operation completes (indicated by $OM_nSR[MUB]$) the message controller stops

15.7.4.2.10 Chaining mode retry error threshold exceeded errors

When a retry error threshold exceeded error occurs for a message segment the following occurs.

- The message controller sets the retry threshold exceed status bit ($OM_nSR[RETE]$)
- If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.
- After the message operation completes (indicated by $OM_nSR[MUB]$) the message controller stops

15.7.4.2.11 Chaining mode transaction errors

When an internal error occurs during a local memory read by the message controller the following occurs.

- The message controller sets the transaction error bit ($OM_nSR[TE]$)
- If the internal error occurs during the descriptor memory read by the message controller, no message segment memory reads are generated and no message segments are sent
- Message segments that have an internal error are not sent because the message data is not available
- Memory reads that already were generated before the internal error occurred are also not transferred.
- Additional memory reads for the same message operation are generated but not transferred.
- All subsequent message segments for the same message operation are not transferred. This includes retried message segments.
- After the message operation completes (indicated by $OM_nSR[MUB]$) the message controller stops
- If $OM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.

15.7.4.2.12 Chaining mode error handling

When an error occurs and the SRIO error/port-write interrupt is generated, the following occurs.

- Software determines the cause of the interrupt and processes the error
- Software verifies the message controller has stopped operation by polling OM_nSR[MUB]
- Software disables the message controller by clearing OM_nMR[MUS]
- Software clears the error by writing a 1 to the corresponding status bit (OM_nSR[MER], OM_nSR[PRT], OM_nSR[RETE], or OM_nSR[TE])

When an error occurs and the SRIO error/port-write interrupt is not enabled, the following occurs.

- Software determines that an error has occurred by polling the status bits (OM_nSR[MER], OM_nSR[PRT], OM_nSR[RETE], or OM_nSR[TE])
- Software verifies the message controller has stopped operation by polling OM_nSR[MUB]
- Software disables the message controller by clearing OM_nMR[MUS]
- Software clears the error by writing a 1 to the corresponding status bit (OM_nSR[MER], OM_nSR[PRT], OM_nSR[RETE], or OM_nSR[TE])

15.7.4.2.13 Chaining mode hardware error handling

The following list additional error conditions beyond the errors listed for direct mode.

All the errors listed in the direct mode section can also occur. The error checking level indicates the order in which errors are checked. Multiple errors can be checked at an error checking level. Once an error is detected no additional error checking beyond the current level is performed. The first error detected in the processing pipeline updates the error management extensions registers.

These error condition checks are provided by the messaging unit. These checks are in addition to the error condition checks provided by the RapidIO port given in [Errors and error handling](#).

Table 15-344. Outbound message chaining mode hardware errors

Transaction	Error	Error Checking Level	Interrupt Generated	Status Bit Set	Message Segment Sent	Logical/Transport Layer Capture Register	Comments
Message Request	An internal error occurred during a read of the descriptor from local memory	0	SRIO error/port-write if OM _n MR[EI] set	Transaction error in the outbound message status register (OM _n SR[TE]). Message Failed in the Mailbox CSR (MCSR[FA]).	No	-	Message controller stops. Note that the descriptor dequeue pointer is not incremented.

15.7.4.2.14 Chaining mode programming errors

The following is the list of programming errors that result in undefined or undesired hardware operation.

These errors are in addition to the programming errors listed for direct mode.

Table 15-345. Outbound message chaining mode programming errors

Error	Interrupt Generated	Status Bit Set	Comments
Enqueued descriptor address is invalid	No	No	Local memory captures the transaction and generate an interrupt.
Address for descriptor enqueue address pointer is invalid	No	No	Local memory captures the transaction and generate an interrupt.
Descriptor queue enqueue and dequeue pointers are not initialized to the same value	No	No	Undefined operation results
Descriptor queue size set to a reserved value	No	No	Undefined operation results
Address of descriptor enqueue pointer set to a value outside of queue	No	No	Undefined operation results
Enqueueing of descriptors causes descriptor queue overflow	Outbound message interrupt enable set (OM _n MR[QOIE])	Queue overflow (OM _n SR[QOI])	Message controller stops.
Queue misaligned	No	No	May result in duplicate messages being sent

15.7.4.3 Message controller arbitration

Service control defines the order in which each message controller sends messages from its message queues when there are more than one outbound message units.

There are two options:

- Fixed priority-the lowest numbered message unit has the highest priority. Message unit 0 has the highest priority.
- Rotating priority- in this mode the order in which the message units take turns sending messages is round robin (message units 0, 1, 2, 3, 0, and so on). The number of messages a message unit can send is from 1 to 64.

OM n MR[SCNTL] configures the message units for these modes and defines operation in direct or chaining mode.

In fixed priority mode, all message segments for message unit 0 must complete before another lower priority message unit (message units 1, 2, 3, and so on) can start. However, in rotating priority mode, another message unit can start processing a message as soon as all message segments have been transmitted. Also, if in fixed priority mode and a message unit other than message unit 0 is processing a message, message unit 0 can start processing a message as soon as all of the message segments have been transmitted.

15.7.5 Inbound message controller operation

The inbound message controller is responsible for receiving messages and placing them in a circular frame queue in local memory.

The inbound message controller can receive segments of a message in any order. The address of where to write the message is computed as: Base address + (msgseg x ssize in double-words). Unlike the outbound message controller where the enqueue pointer is controlled by software and the dequeue pointer is controlled by hardware, the inbound message controller controls the enqueue pointer and the software controls the dequeue pointer.

[Figure 15-323](#) depicts a sample structure of the inbound message frames and the frame pointers. In this example, the frame queue has eight entries, three of which are currently valid. The inbound controller controls the enqueue pointer and the software controls the dequeue pointer. After a configured number of messages have been received, an interrupt is generated to the processor. After processing a received message, the local processor can either write the inbound message mode register mailbox increment bit (IM n MR[MI]) causing the dequeue pointer to point to the next message frame in the queue or wait until all the received messages have been processed and write the dequeue pointer.

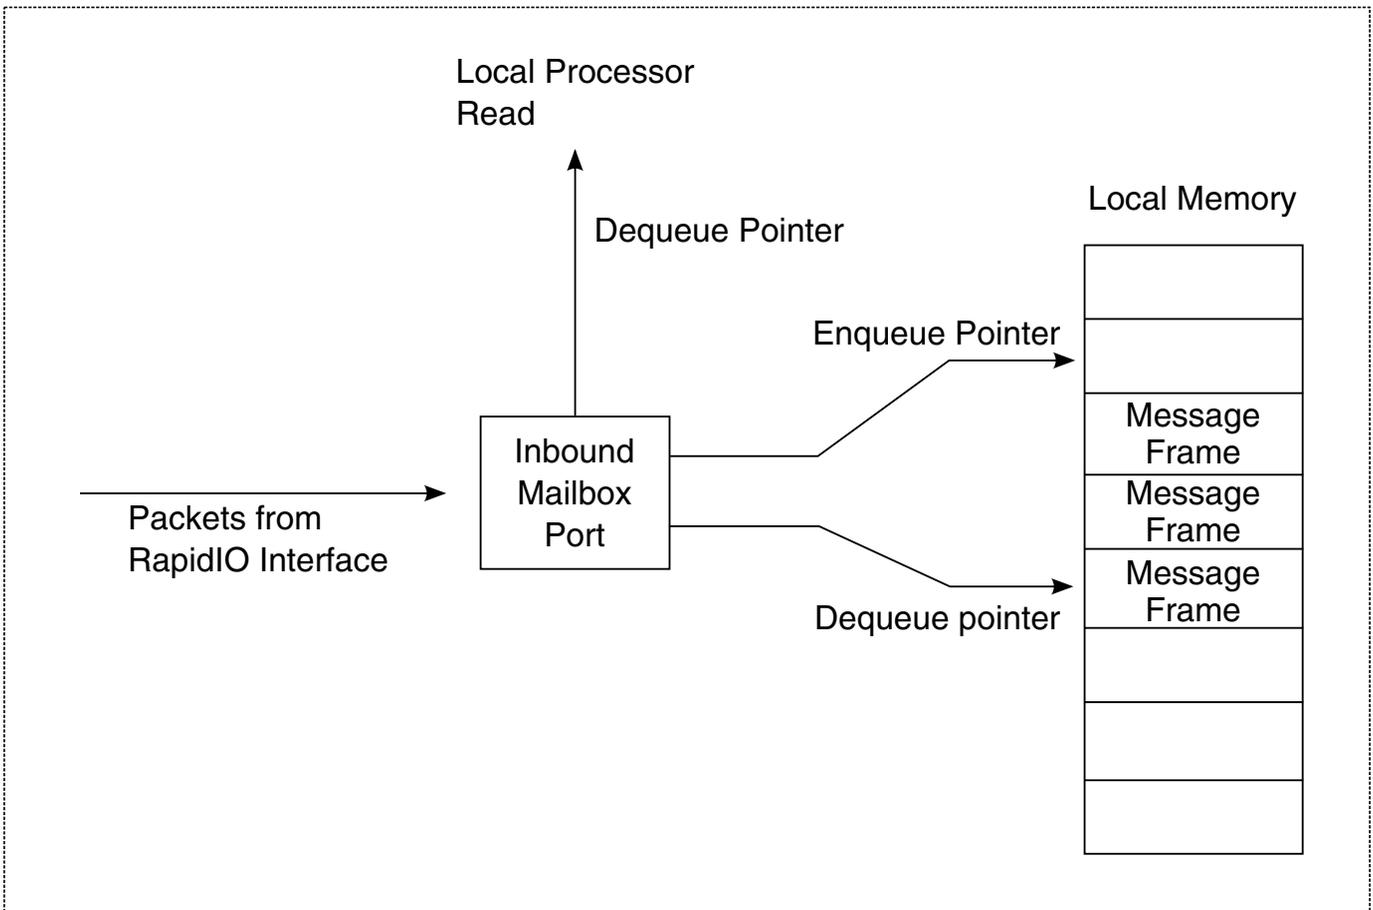


Figure 15-323. Inbound message structure

15.7.5.1 Inbound message controller initialization

The sequence of events to initialize the message controller is as follows:

- Initialize the frame queue dequeue pointer address registers ($IM_nFQDPAR$, $EIM_nFQDPAR$) and the frame queue enqueue pointer address registers ($EIM_nFQEPAR$, $IM_nFQEPAR$). These need to be initialized to the same value for proper operation.

These also must be queue size aligned (in other words, the queue to which they point must be aligned on a boundary equal to number of queue entries x frame size). For example, if there are eight entries in the queue and the frame size is 128 bytes, the queue must be 1024-byte aligned.

- Clear the status register (IM_nSR).
- Set the mailbox enable bit ($IM_nMR[ME]$) along with the other control parameters (frame queue size, message-in-queue threshold, frame size, snoop enable, and the various interrupt enables) in the inbound message mode register (IM_nMR).

15.7.5.2 Inbound controller operation

There are many ways in which software can interact with the message controller. One method is as follows:

- The inbound message controller receives a message segment request from the RapidIO port. If the inbound message controller is enabled ($IMnMR[ME] = 1$), the inbound message controller has received all the segments for the previous message, and the frame queue is not full then the message segment is accepted.
- The inbound message controller computes the address for each segment of the message (up to 16 segments per message) using the value of the inbound message frame queue enqueue pointer address registers and the segment number.
- The inbound message controller writes each segment to the circular queue in local memory at the computed address.
- Once the entire message is received and the write of all message segments have completed, the enqueue pointer is incremented to point to the next message frame in local memory by the inbound message controller. A message operation completes after all message segments for the message complete. A message segment completes when one of the following occurs:
 - the memory write completes (either successfully or an internal error occurred)
 - a message request time-out occurs (all message segments that have not yet been received are now complete)
- The message segments for a message can immediately be followed by the message segments for another message if certain rules are followed. If a message segment arrives for a new message before all of the previous message memory writes complete for the previous message and the RapidIO priority (the prio field value in the message packet) of the message is equal to or lower than the RapidIO priority of all of the previous messages, the message is processed by the message controller and a memory write is generated to the appropriate frame queue entry. If the RapidIO priority of the message is higher than any of the previous message memory writes that have not completed, the message controller generates a retry. Also, if a message segment arrives before all of the previous message memory writes for the same message complete and the new message segment is higher priority than the previous message segments, then the message controller generates a retry. The $IMnSR[MB]$ is cleared by the message controller when all message operations complete.
- An inbound message interrupt is generated to the local processor if the number of messages in the queue is greater than or equal to the configured message-in-queue threshold ($IMnMR[MIQ_THRESH]$) and this event is enabled to generate the interrupt ($IMnMR[MIQIE]$).

- Software determines that the message-in-queue event caused the interrupt by detecting that the message-in-queue interrupt bit is set when reading the inbound message status register (IM_nSR[MIQI]).
- Software processes the frame queue entry pointed to by the frame dequeue pointer address registers (IM_nFQDPAR, EIM_nFQDPAR).
- Software increments the dequeue pointer address registers (IM_nFQDPAR, EIM_nFQDPAR) by setting the message increment bit (IM_nMR[MI]). Software determines if there are any more messages to process by reading the queue empty bit (IM_nSR[QE]). If the queue is not empty, the previous two steps are repeated.
- Optionally, software reads the enqueue pointer address registers (IM_nFQEPAR, EIM_nFQEPAR) and processes all the received messages. After the message processing is complete the dequeue pointer address registers (IM_nFQDPAR, EIM_nFQDPAR) are written.
- Software clears the message-in-queue interrupt bit (IM_nSR[MIQI]) by writing a 1 to the IM_nSR[MIQI] bit.

15.7.5.3 Message steering

Messages are forwarded to the inbound message controllers as follows:

- Messages directed to mailbox 0, regardless of letter number, are forwarded to message controller 0.
- Messages directed to mailbox 1, 2 or 3, regardless of letter number, are forwarded to message controller 1.

A message controller which has received part of a multi-segment message can not process another message until all segments for that message have arrived. Subsequent messages will be retried if targeted for the same message controller. A maximum of two multi-segment messages can be reassembled at one time if handled by different message controllers.

15.7.5.4 Inbound message controller retry response conditions

The conditions to generate a logical layer retry (response retry) are:

- Local memory circular queue is full and a message is received.
- The inbound message controller has already received at least one segment of a multi-segment message but has not received all message segments (determined by a different RapidIO source ID, RapidIO destination ID or mailbox)
- Message received with a higher priority than all of the previous messages that are being written to memory but have not completed.

Note that if all inbound messages are the same RapidIO priority, the third condition for generating a retry can not occur.

15.7.5.5 Inbound message controller interrupts

The inbound message controller generates the inbound message interrupt for several different events.

Each event can be individually enabled.

- Message-in-queue-an interrupt is generated whenever
 - The circular queue has accumulated the specified number of messages and this interrupt event is enabled ($IM_nMR[MIQIE]$). The event that caused this interrupt is indicated by $IM_nSR[MIQI]$. The interrupt is held until the dequeue pointer and enqueue pointer indicate that the specified number of messages is not in the frame queue and the $IM_nSR[MIQI]$ bit has been cleared by writing a 1.
 - The circular queue has one or message in it, the specified number of message has not accumulated, a message has not been dequeued for the maximum interrupt report interval and this interrupt event is enabled ($IM_nSR[MIQIE]$). The event that caused this interrupt is indicated by $IM_nSR[MIQI]$. The interrupt is held until the $IM_nSR[MIQI]$ bit has been cleared by writing a 1.
- Queue full-an interrupt is generated whenever the circular queue becomes full and this interrupt event is enabled ($IM_nSR[QFIE]$). The event that caused this interrupt is indicated by $IM_nSR[QFI]$. The interrupt is held until the queue is not full and the $IM_nSR[QFI]$ bit has been cleared by writing a 1.

The error/port-write interrupt can be generated for the following reasons.

- Message request time-out-an interrupt is generated after a message request time-out occurs and this interrupt event is enabled ($IM_nMR[EIE]$)
- Transaction error- an interrupt is generated after a OCN error response is received and this interrupt event is enabled ($IM_nMR[EIE]$)

15.7.5.5.1 Message request time-out errors

The message request time-out counter starts after the first valid segment of a multisegment message is received and the time-out counter is enabled.

When a message request time-out occurs the following occurs.

- The message controller sets the message request time-out status bit ($IM_nSR[MRT]$)
- All message segments that have not yet been received are considered complete

- If $IM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated
- Once all message segments complete (indicated by $IM_nSR[MB]$), the message controller stops

15.7.5.5.2 Inbound message controller transaction errors

When an internal error occurs during a local memory write by the message controller the following occurs.

- The message controller sets the transaction error bit ($IM_nSR[TE]$) and enters the error state
- The message controller returns an error response
- Memory writes that already were generated before the internal error occurred also return an error response
- After the message operation completes (indicated by $IM_nSR[MB]$) the message controller stops
- If $IM_nMR[EIE]$ is set, the interrupt SRIO error/port-write is generated.

15.7.5.5.3 Inbound message controller error handling

When an error occurs and the SRIO error/port-write interrupt is generated, the following occurs.

- Software determines the cause of the interrupt and processes the error
- Software verifies the message controller has stopped operation by polling $IM_nSR[MB]$
- Software clears the error by writing a 1 to the corresponding status bit ($IM_nSR[MRT]$, and/or $IM_nSR[TE]$)
- The message unit must be disabled, re-initialized and re-enabled before another message can be received.

When an error occurs and the SRIO error/port-write interrupt is not enabled, the following occurs.

- Software determines that an error has occurred by polling the status bits ($IM_nSR[MRT]$ and/or $IM_nSR[TE]$)
- Software verifies the message controller has stopped operation by polling $IM_nSR[MB]$
- Software disables the message controller by clearing $IM_nMR[ME]$
- Software clears the error by writing a 1 to the corresponding status bit ($IM_nSR[MRT]$ and/or $IM_nSR[TE]$)

15.7.5.5.4 Inbound message controller hardware error handling

The following lists possible error conditions and what occurs when they are encountered.

The error checking level indicates the order in which errors are checked. Multiple errors can be checked at an error checking level. Once an error is detected no additional error checking beyond the current level is performed. Note that messages are processed in a pipeline fashion. The first error detected in the processing pipeline updates the error management extensions registers.

These error condition checks are provided by the messaging unit. These check are in addition to the error condition checks provided by the RapidIO port given in [Table 15-346](#).

Table 15-346. Inbound message hardware errors

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/Transport Layer Capture Register	Comments
A reserved ftype ¹	1	SRIO error/port-write if LTLEECS R[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UT]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
Reserved tt encoding ¹	1	SRIO error/port-write if LTLEECS R[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[TSE].	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
Large transport size when operating in small transport size or small transport size when operating in large transport size ¹	1	SRIO error/port-write if LTLEECS R[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[TSE]	No	No Response	Updated with the packet ²	Packet is ignored and discarded. An error or illegal transaction target error response is not generated.
Illegal Destination ID ¹	1	SRIO error/port-write if LTLEECS R[ITTE] set	Illegal transaction target in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITTE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-346. Inbound message hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/Transport Layer Capture Register	Comments
An incorrect message packet size (payload is not than the specified ssize except for the last segment. The last segment's payload can be less than or equal to the segment size but not 0. Note that payload sizes are dword multiples.) ¹	1	SRIO error/port-write if LTLEECS R[MFE] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
Reserved ssize field ¹	1	SRIO error/port-write if LTLEECS R[MFE] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
Message received and no mailboxes are supported as indicated by DOCAR[M] ¹	1	SRIO error/port-write if LTLEECS R[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UT]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
Message packet size larger than the configured frame size and message controller is enabled	2	SRIO error/port-write if LTLEECS R[MFE] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
An inbound message packet with a RapidIO priority of 3	2	SRIO error/port-write if LTLEECS R[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-346. Inbound message hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/ Transport Layer Capture Register	Comments
Not an error. The RapidIO priority is not consistent for all message segments of a message	2	-	-	Yes	Done or Retry Response	-	Retry response occurs if the higher priority message segment is received while the memory write for a corresponding lower priority message segment is outstanding
Message segment number is larger than the number of message segments in the message	2	SRIO error/port-write if LTLEECR[MFE] set	Message format error in the Logical/ Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
Duplicate message segment is received (Note that all segments of a multi segment message must be received before the next message begins).	2	SRIO error/port-write if LTLEECR[MFE] set	Message format error in the Logical/ Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
msglen (number of segments in a message) is not consistent in all segments of a multi segment message	2	SRIO error/port-write if LTLEECR[MFE] set	Message format error in the Logical/ Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
ssize (segment size) is not consistent in all segments of a multi segment message	2	SRIO error/port-write if LTLEECR[MFE] set	Message format error in the Logical/ Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-346. Inbound message hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/Transport Layer Capture Register	Comments
Message received for an unsupported mailbox but at least one mailbox is supported	2	SRIO error/port-write if LTLEECS R[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UT]	No	Error	-	Packet is ignored and discarded. This error only applies if a mailbox is not supported. This error is not currently supported since all mailboxes are supported.
Message Controller Disabled and Message Received	2	No	None	No	Error	-	Packet is ignored and discarded.
Message Controller Enabled but in the Error State and Message Received	2	No	None	No	Error	-	Packet is ignored and discarded.
Internal error occurred during the write of the frame queue entry to memory	3	SRIO error/port-write if IMnMR[EIE] set	Transaction error in the Message status register (IMnSR[TE]). Message Failed in the Message CSR (MCSR[FA])	No	Error	-	Message controller stops after the current message operation completes. The enqueue pointer is not incremented.

Table continues on the next page...

Table 15-346. Inbound message hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/Transport Layer Capture Register	Comments
An internal error occurred for an earlier posted frame queue entry memory write and a subsequent frame queue entry memory write is posted before the internal error is detected. An internal error may or may not occur during the subsequent frame queue entry memory write. The frame queue could be for the same message or a different message.	4	No	None	Yes	Error	-	-
Message request to request time-out for multi segment messages	unrelated	SRIO error/port-write if LTLEECSR[MRT] set. SRIO error/port-write if OMnMR[E]E].	Message request time-out in the Logical/Transport Layer Error Detect CSR LTLEDCSR[MRT]. IMnSR[MRT] bit set.	No	No	Updated with the previous message request packet except that the message segment field (bits 4 to 7 of LTLCCCSR[MI]) is updated with the lowest message segment number that has not yet been received. ²	All message segments received before the time-out update memory. The enqueue pointer is not incremented. The message operation completes.

1. These error types are actually detected in the RapidIO port, not in the message controller.
2. If operating in small transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 78-79), LTLACCSR[A] gets the address (packet bits 48-76), LTLDIDCCSR[MDID] gets 0, LTLDIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 16-23), LTLDIDCCSR[MSID] gets 0, LTLDIDCCSR[SID] gets the least significant byte of the source ID (packet bits 24-31), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 32-35), LTLCCCSR[MI] gets the msg info (packet bits 40-47). If operating in large transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 94-95), LTLACCSR[A] gets the address (packet bits 64-92), LTLDIDCCSR[MDID] gets the most significant byte of the destination ID (packet bits 16-23), LTLDIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 24-31), LTLDIDCCSR[MSID] gets the most significant byte of the source ID (packet bits 32-39), LTLDIDCCSR[SID] gets the least

significant byte of the source ID (packet bits 40-47), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 48-51), LTLCCCSR[MI] gets the msg info (packet bits 56-63).

15.7.5.5.5 Programming errors

The following is a partial list of programming errors that result in undefined or undesired hardware operation.

Table 15-347. Inbound message programming errors

Error	Interrupt Generated	Status Bit Set	Comments
Reserved value of the message in queue threshold (IMnMR[MIQ_THRESH]) or reserved value of the circular frame queue size (IMnMR[CIRQ_SIZE])	No	No	Undefined operation results
The message in-queue threshold is equal to the frame queue size	No	No	Message in queue interrupt occurs when queue is full
The message in-queue threshold is greater than the frame queue size	No	No	Message in queue interrupt never occurs
Frame queue entry written to non-existent memory	No	No	Memory controller causes the interrupt and updates capture registers. IMnSR[TE] is set due to the internal error.
Message enqueue and dequeue pointers are not initialized to the same value	No	No	Undefined operation results
The dequeue frame pointer register is set incorrectly.	No	No	Undefined operation results
Queue misaligned	No	No	May cause unpredictable behavior

15.7.5.5.6 Disabling and enabling the inbound message controller

When the message controller is disabled by clearing IMnMR[ME] the following occurs.

- Queue full clears (IMnSR[QF])
- Message-in-queue clears (IMnSR[MIQ])
- Queue empty is set (IMnSR[QE])

Once the message controller is disabled, an error response is generated for all new message packets. If the message controller is disabled before all of the message segments for a multisegment message are received, a message request time-out must occur and all pending frame queue writes must complete before message busy clears (IMnSR[MB]).

For proper message controller operation, before the message controller is re-enabled, the following register writes must be performed: clear (IM n SR[MB]) and (IM n MR[ME]), and initialize the frame queue dequeue pointer address registers (IM n FQDPAR, EIM n FQDPAR) and the frame queue enqueue pointer address registers (IM n FQEPAR, EIM n FQEPAR) to the same value.

15.7.6 Message unit message passing logical specification registers

The mailbox command and status register (MCSR) provides the status for the four inbound and the four outbound controllers.

These read-only status bits indicate the state of each of the message controllers.

- Available (MCSR[A])-Indicates that the inbound message controller is enabled (IM n MR[ME]), the inbound message controller is not in the internal error state (IM n SR[TE] = 0) and the inbound message controller did not detect a message request time-out (IM n SR[MRT] = 0)
- Full (MCSR[FU])-This bit reflects the inbound message controller queue full status
- Empty (MCSR[EM])-This bit reflects the state of the outbound message controller message empty status
- Busy (MCSR[B])-This bit reflects the state of the inbound message controller busy bit IM n SR[MB]
- Failed (MCSR[FA])-This bit is set if any of the following bits are set:
 - The inbound message controller transaction error status bit IM n SR[TE]
 - The inbound message controller message request time-out status bit IM n SR[MRT]
 - The outbound message controller transaction error status bit OM n SR[TE] is set
 - The outbound message controller packet response time-out bit OM n SR[PRT] is set
 - The outbound message controller message error response received status bit OM n SR[MER] is set
 - The outbound message controller retry error threshold exceeded status bit OM n SR[RETE] is set
- Error (MCSR[ERR])-This bit is always 0

15.8 RapidIO doorbell and port-write unit

This section describes the operation of the doorbell and port-write controllers, which are part of the RapidIO message unit.

The doorbell and port-write controllers are compliant with the message passing logical specification contained in the *RapidIO Interconnect Specification, Revision 1.2*. The doorbell controller generates and receives doorbells. The port-write controller receives but does not generate port-writes.

The doorbell and port-write controllers are controlled through a set of run-time registers.

15.8.1 Doorbell and port-write unit features

- Support for one outbound doorbell controller
- Support for one inbound doorbell controller
 - The doorbell controller can sustain back-to-back inbound doorbells
- Support for one inbound port-write controller with the following features
 - Up to 64 bytes of payload
 - Only one inbound port-write queue entry

15.8.2 Doorbell controller

RapidIO supports a doorbell type that contains no data payload.

The RapidIO architecture references outbound and inbound doorbell controllers. This doorbell unit supports both inbound and outbound doorbells. Inbound doorbells are handled by the doorbell controller similar to how the message controller handles inbound data messages. The doorbell controller receives the doorbells and places it in a circular queue located in local memory. Additional doorbells can be received and forwarded to memory before the previous doorbell memory write completes as long as the RapidIO priority is the same or lower than the previous doorbell. The doorbell is retried if the RapidIO priority is higher than the previous doorbell. Outbound doorbells are generated through a memory-mapped write port rather than a queue. An outbound doorbell must complete before another outbound doorbell can be generated.

[Figure 15-324](#) depicts an example of the structure of the inbound doorbell queue and pointers. The doorbell queue has eight entries, three of which are currently valid. The doorbell controller enqueues doorbells and the local processor dequeues doorbells.

The doorbell entry size is fixed at 64 bits because doorbell packets only pass a small amount of information, making the enqueue and dequeue pointers double-word addresses.

The outbound doorbell controller has a dedicated interrupt that can be used to notify software that a doorbell has been sent. Each inbound doorbell controller has a dedicated interrupt that can be used for notification of incoming doorbells.

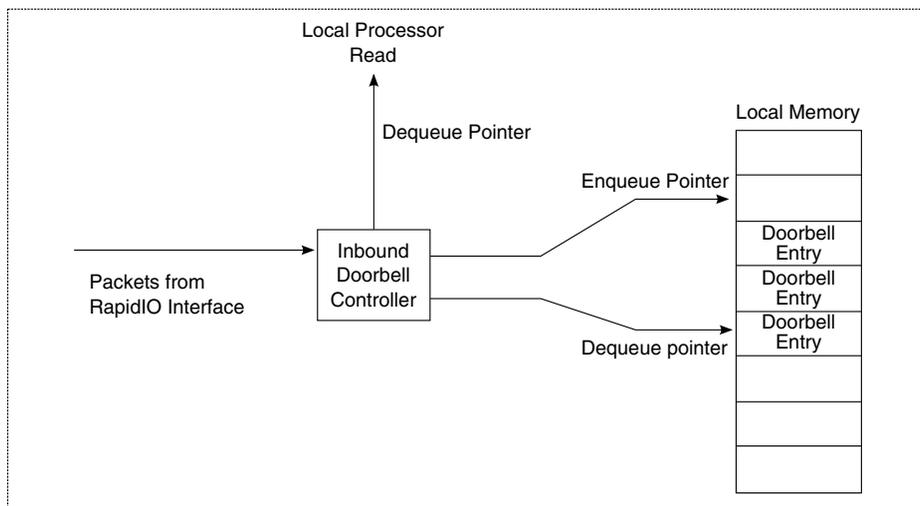


Figure 15-324. Inbound doorbell queue and pointer structure

15.8.2.1 Outbound doorbell controller

The outbound doorbell controller is used to generate doorbells.

Software is responsible for initializing all the parameters in all the necessary registers to start the doorbell transmission. The doorbell transfer is started when the doorbell start bit, `OD MR[DUS]`, in the outbound doorbell mode register transitions from a 0 to 1 and the doorbell controller is not already busy. Software is expected to program all the appropriate registers before setting `ODMR[DUS]`.

There are many ways in which software can interact with the doorbell controller. One method to generate a doorbell is as follows:

- Poll the status register doorbell unit busy bit, `ODSR[DUB]`, to make sure the outbox is not busy with a previously initiated doorbell.
- Clear the status bits (`ODSR[MER]`, `ODSR[RETE]`, `ODSR[PRT]`, and `ODSR[EODI]`)
- Initialize the destination port (`OD DPR`), destination attributes (`OD DATR`) and retry error threshold (`OD RETCR`) registers
- Clear, then set the doorbell start bit, `ODMR[DUS]`, to start the doorbell transfer.
- `ODSR[DUB]` is set when `ODMR[DUS]` transitions from a 0 to a 1 to indicate the doorbell transfer is in progress.
- The outbound doorbell controller sends the doorbell
- The `ODSR[DUB]` is cleared by the outbound doorbell controller after the doorbell operation completes. A doorbell completes when one of the following occurs:
 - done response received
 - error response received

- a packet response time-out occurred
- the retry limit was exceeded
- The outbound doorbell interrupt is generated if the end of doorbell outbound doorbell interrupt event is enabled (ODDATR[EODIE]).

15.8.2.1.1 Outbound doorbell controller interrupts

The "SRIO outbound doorbell" controller interrupt is generated after the completion of a doorbell (done, error, packet response time-out or retry limit exceeded) if this interrupt event is enabled (ODDATR[EODIE] is a 1).

The event that caused this interrupt is indicated by OD SR[EODI]. The interrupt is held until the OD SR[EODI] bit has been cleared by writing a 1.

The "SRIO error/port-write" interrupt can be generated for the following reasons:

- RapidIO error response. An interrupt is generated after a RapidIO error response is received and this interrupt event is enabled (LTLEECSSR[MER])
- Packet response time-out. An interrupt is generated after a packet response time-out occurs and this interrupt event is enabled (LTLEECSSR[PRT])
- Retry error threshold exceeded. An interrupt is generated after a retry threshold exceeded error occurs and this interrupt event is enabled (LTLEECSSR[RETE])

15.8.2.1.2 Error response errors

When a RapidIO error response is received by the doorbell controller the following occurs:

- The doorbell controller sets the message error response status bits (OD SR[MER] and LTLEDCSSR[MER])
- If LTLEECSSR[MER] is set, the interrupt "SRIO error/port-write" is generated.
- After the doorbell operation completes (indicated by OD SR[DUB]) the doorbell controller stops.

15.8.2.1.3 Outbound doorbell controller packet response time-out errors

When a packet response time-out occurs for a doorbell the following occurs:

- The doorbell controller sets the packet response time-out status bits (OD SR[PRT] and LTLEDCSSR[PRT])
- If LTLEECSSR[PRT] is set, the interrupt "SRIO error/port-write" is generated and EPWISR[PINT] is set.
- After the doorbell operation completes (indicated by OD SR[DUB]) the doorbell controller stops.

15.8.2.1.4 Outbound doorbell controller retry error threshold exceeded errors

When a retry error threshold exceeded error occurs for a doorbell the following occurs:

- The doorbell controller sets the retry threshold exceed status bits (OD SR[RETE] and LTLEDCSR[RETE])
- If LTLEDCSR[RETE] is set, the interrupt "SRIO error/port-write" is generated.
- After the doorbell operation completes (indicated by OD SR[DUB]) the doorbell controller stops.

15.8.2.1.5 Error handling

When an error occurs and the "SRIO error/port-write" interrupt is generated, the following occurs:

- Software determines the cause of the interrupt and processes the error
 - LTLEDCSR and ODSR capture the error condition for outbound doorbells
 - EPWISR[PINT] is set for PRT error since it is detected by the SRIO controller
 - Note that LTLEDCSR is a capture once register, so ODSR should be examined to make sure an outbound doorbell error did not occur immediately after another captured error
- Software verifies the doorbell controller has stopped operation by polling OD SR[DUB]
- Software disables the doorbell controller by clearing OD MR[DUS]
- Software clears the error by writing a 1 to the corresponding outbound doorbell status bits (OD SR[PRT], OD SR[PRT], and/or OD SR[RETE]) as well as LTLEDCSR)

When an error occurs and the "SRIO error/port-write" interrupt is not enabled, the following occurs:

- Software determines that an error has occurred by polling the status bits (OD SR[MER], OD SR[PRT], and/or OD SR[RETE])
- Software verifies the doorbell controller has stopped operation by polling OD SR[DUB]
- Software disables the doorbell controller by clearing OD MR[DUS]
- Software clears the error by writing a 1 to the corresponding status bits (OD SR[MER], OD SR[PRT], and/or OD SR[RETE])

15.8.2.1.6 Disabling and enabling the doorbell controller-outbound

Once the doorbell controller is started, it cannot be stopped.

15.8.2.1.7 Outbound doorbell controller hardware error handling

The following is a list possible error conditions and what occurs when they are encountered.

The error checking level indicates the order in which errors are checked. Multiple errors can be checked at an error checking level. Once an error is detected, no additional error checking beyond the current level is performed. Note that outbound doorbell responses are processed in a pipeline fashion. The first error detected in the processing pipeline updates the error management extensions registers.

These error condition checks are provided by the messaging unit. These checks are in addition to the error condition checks provided by the RapidIO port given in [Errors and error handling](#).

Table 15-348. Outbound doorbell hardware errors

Transaction	Error	Error Checking Level	Interrupt Generated	Status Bit Set	Doorbell Sent	Logical/Transport Layer Capture Register	Comments
Undefined Packet	Reserved ftype encoding ¹	1	SRIO error/port-write if LTLEEC[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDC[UT]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	Reserved tt encoding ¹	1	SRIO error/port-write if LTLEEC[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDC[TSE].	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	Large transport size when operating in small transport size or small transport size when operating in large transport size ¹	1	SRIO error/port-write if LTLEEC[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDC[TSE].	Yes	Updated with the packet ²	Packet is ignored and discarded. An error or illegal transaction target error response is not generated.
Doorbell Response	Illegal Destination ID ¹	1	SRIO error/port-write if LTLEEC[ITTE] set	Illegal transaction target in the Logical/Transport Layer Error Detect CSR LTLEDC[ITTE]	Yes	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-348. Outbound doorbell hardware errors (continued)

Transaction	Error	Error Checking Level	Interrupt Generated	Status Bit Set	Doorbell Sent	Logical/Transport Layer Capture Register	Comments
Doorbell Response	doorbell not outstanding ¹	1	SRIO error/port-write if LTLEEC[SR] set	Unsolicited response in the Logical/Transport Layer Error Detect CSR LTLEDC[SR]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	transaction type (transaction field) is not doorbell response ¹	1	SRIO error/port-write if LTLEEC[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDC[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	RapidIO priority is less than or equal to outbound request ¹	2	SRIO error/port-write if LTLEEC[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDC[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	Incorrect Source ID ¹	2	SRIO error/port-write if LTLEEC[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDC[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	reserved response status ¹	2	SRIO error/port-write if LTLEEC[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDC[ITD]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	doorbell response packet size is incorrect ¹	2	SRIO error/port-write if LTLEEC[MFE] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDC[MFE]	Yes	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Response	error response	3	SRIO error/port-write if LTLEEC[MER] set.	Message error response in the Logical/Transport Layer Error Detect CSR LTLEDC[MER]. ODSR[MER] bit set	Yes	Updated with the corresponding doorbell request packet ²	doorbell transfer complete

Table continues on the next page...

Table 15-348. Outbound doorbell hardware errors (continued)

Transaction	Error	Error Checking Level	Interrupt Generated	Status Bit Set	Doorbell Sent	Logical/Transport Layer Capture Register	Comments
Doorbell Response	number of retries exceeds limit	3	SRIO error/port-write if LTLEECRSR[RETE] set.	Retry limit exceeded in the Logical/Transport Layer Error Detect CSR LTLEDCRSR[RETE] ODSR[RETE] bit set.	Yes	Updated with the corresponding doorbell request packet ²	doorbell transfer complete
Doorbell Response	packet response time-out ¹	unrelated	SRIO error/port-write if LTLEECRSR[PRT] set.	Packet response time-out in the logical/transport layer error detect CSR LTLEDCRSR[PRT] in RapidIO endpoint. ODSR[PRT] bit set.	Yes	Updated with the doorbell request packet in the RapidIO endpoint ²	doorbell transfer complete. Note that the RapidIO endpoint sends special priority 3 pkt indicating doorbell time-out.

1. These error types are actually detected in the RapidIO port, not in the doorbell controller
2. If operating in small transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 78-79), LTLACCSR[A] gets the address (packet bits 48-76), LTLIDCCSR[MDID] gets 0, LTLIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 16-23), LTLIDCCSR[MSID] gets 0, LTLIDCCSR[SID] gets the least significant byte of the source ID (packet bits 24-31), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 32-35), LTLCCCSR[MI] gets 0. If operating in large transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 94-95), LTLACCSR[A] gets the address (packet bits 64-92), LTLIDCCSR[MDID] gets the most significant byte of the destination ID (packet bits 16-23), LTLIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 24-31), LTLIDCCSR[MSID] gets the most significant byte of the source ID (packet bits 32-39), LTLIDCCSR[SID] gets the least significant byte of the source ID (packet bits 40-47), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 48-51), LTLCCCSR[MI] gets 0.

15.8.2.1.8 Outbound doorbell controller programming errors

The following is the list of programming errors that result in undefined or undesired hardware operation.

Table 15-349. Outbound doorbell programming errors

Error	Interrupt Generated	Status Bit Set	Comments
Transaction flow level set to reserved (2'b11)	No	None	Unit hangs. If the transaction flow level is then changed to a value other than reserved, the doorbell operation starts using this new transaction flow level.
Target interface set to an invalid RapidIO port	No	None	Undefined operation results
Register values changed during operation	No	No	Undefined operation results

15.8.2.2 Inbound doorbell controller

The inbound doorbell controller is responsible for receiving doorbells and placing them in a circular doorbell queue in local memory.

The inbound controller controls the enqueue pointer and the software controls the dequeue pointer. After a configured number of doorbells have been received, an interrupt is generated to the processor. After processing a received doorbell, the local processor can either write the doorbell mode register increment bit (ID MR[DI]) causing the dequeue pointer to point to the next doorbell in the queue or wait until all the received doorbells have been processed and write the dequeue pointer.

15.8.2.2.1 Inbound doorbell controller initialization

There are many ways in which software can interact with the doorbell controller.

One method to initialize the doorbell controller is as follows:

- Initialize the doorbell queue dequeue pointer address registers (DQDPAR, EDQDPAR) and the doorbell queue enqueue pointer address registers (DQEPAR, EDQEPAR). These need to be initialized to the same value for proper operation. These also must be queue size aligned (in other words, the queue to which they point must be aligned on a boundary equal to number of queue entries x 8 bytes). For example, if there are eight entries in the queue, the queue must be 64-byte aligned.
- Clear the status register (ID SR).
- Set the doorbell enable bit (ID MR[DE]) along with the other control parameters (doorbell queue size, doorbell-in-queue threshold, snoop enable, and the various interrupt enables) in the doorbell mode register (ID MR).

15.8.2.2.2 Inbound doorbell controller operation

There are many ways in which software can interact with the doorbell controller.

One method is as follows:

- The doorbell controller receives a doorbell. If the inbound doorbell controller is enabled (ID MR[DE] = 1) and the doorbell queue is not full, then the doorbell is accepted.
- The 16-bit information field along with the RapidIO source ID and RapidIO destination ID are stored in local memory by the doorbell controller using the value of the doorbell queue enqueue pointer address registers (DQEPAR, EDQEPAR).
- Once the memory write completes the enqueue pointer is incremented to point to the next doorbell queue entry in local memory.

- If another doorbell arrives before all of the previous doorbell memory writes complete and the RapidIO priority of the doorbell is equal to or lower than the RapidIO priority of all of the previous doorbells, the doorbell is processed by the doorbell controller and a memory write is generated to the appropriate doorbell queue entry. If the RapidIO priority of the doorbell is higher than any of the previous doorbell memory writes that have not completed, the doorbell controller generates a retry.
- An inbound doorbell interrupt is generated to the local processor because the number of doorbells in the queue is greater than or equal to the configured doorbell-in-queue threshold (ID MR[DIQ_THRESH]) and this event is enabled to generate the interrupt (ID MR[DIQIE]).
- Software determines that the doorbell-in-queue event caused the interrupt by detecting that the doorbell-in-queue interrupt bit is set when reading the doorbell status register (ID SR[DIQI]).
- Software processes the doorbell queue entry pointed to by the doorbell dequeue pointer address registers (DQDPAR, EDQDPAR).
- Software increments the dequeue pointer address registers (DQDPAR, EDQDPAR) by setting the doorbell increment bit (ID MR[DI]).
- Software determines if there are any more doorbells to process by reading the queue empty bit (ID SR[QE]). If the queue is not empty, the previous two steps are repeated.
- Software clears the doorbell-in-queue interrupt bit (ID SR[DIQI]) by writing a 1 to the ID SR[DIQI] bit.

15.8.2.2.3 Inbound doorbell queue entry format

This section defines the format of the doorbell information written to memory by the doorbell controller.

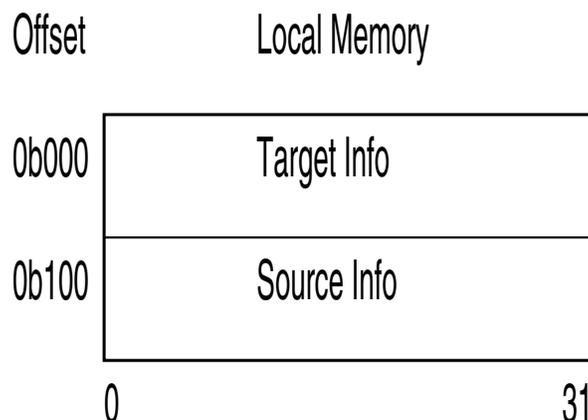
Each doorbell entry in the queue has 2 offsets of 32 bits, one for target information and one for source information. The target information is stored because a RapidIO port can be configured to accept packets from any destination. In addition, when there are multiple RapidIO ports on one device each port may be configured with a different destination ID.

Table 15-350. Inbound doorbell target info definition

Bits	Name	Description
0-15	-	Reserved
16-23	ETID	Extended target ID when in large transport mode, reserved when in small transport mode
24-31	TID	Target ID field from the received doorbell packet

Table 15-351. Source info definition

Bits	Name	Description
0-7	ESID	Extended Source ID when in large transport mode, reserved when in small transport mode
8-15	SID	Source ID field from the received doorbell packet
16-23	INFO MSB	Most significant byte of the info field from the received doorbell packet
24-31	INFO LSB	Least significant byte of the info field from the received doorbell packet


Figure 15-325. Doorbell entry format

15.8.2.2.4 Inbound doorbell controller retry response conditions

There are two conditions in which a doorbell is retried at the logical layer (response retry).

- Doorbell received and there are no entries available in the doorbell queue.
- Doorbell received with a higher priority than all of the previous doorbells that are being written to memory but have not completed.

If all inbound doorbells are the same RapidIO priority, the second condition for generating a retry can not occur.

15.8.2.2.5 Doorbell controller interrupts

There is one doorbell controller interrupt per inbound doorbell controller.

The following events can generate this interrupt.

- Doorbell-in-queue-this event generates an interrupt under the following conditions:
 - the circular queue has accumulated the configured number of doorbells specified by the doorbell-in-queue threshold (ID MR[DIQ_THRESH]) and this interrupt event is enabled (ID MR[DIQIE]). The event that caused this interrupt is indicated by ID SR[DIQI]. The interrupt is held until the dequeue pointer and

- enqueue pointer indicate that the specified number of doorbells is not in the doorbell queue and the ID SR[DIQI] bit has been cleared by writing a 1.
- the circular queue has one or more doorbells in it, the specified number of doorbells has not accumulated, a doorbell has not been dequeued for the maximum interrupt report interval and this interrupt event is enabled (ID MR[DIQIE]). The event that caused this interrupt is indicated by ID SR[DIQI]. The interrupt is held until either ID MR[DI] has been set or DQDPAR[DQDPA] has been written followed by clearing ID SR[DIQI].
- Queue full-an interrupt is generated each time the circular queue becomes full and this interrupt event is enabled (ID SR[QFIE]). The event that caused this interrupt is indicated by ID SR[QFI]. The interrupt is held until the queue is not full and the ID SR[QFI] bit has been cleared by writing a 1.

The error/port-write interrupt can be generated for the following reasons.

- Transaction error-an interrupt is generated after a OCN error response is received and this interrupt event is enabled (ID MR[EIE])

15.8.2.2.6 Doorbell controller transaction errors

When an internal error occurs during a local memory write by the doorbell controller the following occurs.

- The doorbell controller sets the transaction error bit (ID SR[TE]) and enters the error state
- The doorbell controller returns an error response
- If ID MR[EIE] is set, the interrupt SRIO error/port-write is generated.

15.8.2.2.7 Doorbell controller error handling

When an error occurs and the SRIO error/port-write interrupt is generated, the following occurs.

- Software determines the cause of the interrupt and processes the error
- Software verifies the doorbell controller has stopped operation by polling ID SR[DB]
- Software clears the error by writing a 1 to the corresponding status bit (ID SR[TE])
- The doorbell unit must be disabled, re-initialized and re-enabled before another doorbell can be received.

15.8.2.2.8 Inbound doorbell controller hardware error handling

The following list possible error conditions and what occurs when they are encountered. The error checking level indicates the order in which errors are checked.

Multiple errors can be checked at an error checking level. Once an error is detected no additional error checking beyond the current level is performed. Note that doorbells are processed in a pipeline fashion. The first error detected in the processing pipeline updates the error management extensions registers.

These error condition checks are provided by the messaging unit. These check are in addition to the error condition checks provided by the RapidIO port given in [Errors and error handling](#).

Table 15-352. Inbound doorbell hardware errors

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/Transport Layer Capture Register	Comments
Reserved ftype ¹	1	SRIO error/port-write if LTLEECR[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCR[UT]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
Reserved tt encoding ¹	1	SRIO error/port-write if LTLEECR[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCR[TSE].	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
Large transport size when operating in small transport size or small transport size when operating in large transport size ¹	1	SRIO error/port-write if LTLEECR[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCR[TSE].	No	No Response	Updated with the packet ²	Packet is ignored and discarded. An error or illegal transaction target error response is not generated.
Illegal Destination ID ¹	1	SRIO error/port-write if LTLEECR[ITTE] set	Illegal transaction target in the Logical/Transport Layer Error Detect CSR LTLEDCR[ITTE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
Inbound doorbell received and inbound doorbells are not supported as indicated by DOCAR[D] ¹	1	SRIO error/port-write if LTLEECR[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCR[UT]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
An inbound doorbell packet with a RapidIO priority of 3	2	SRIO error/port-write if LTLEECR[ITD] set	Illegal transaction decode in the Logical/Transport Layer Error Detect CSR LTLEDCR[ITD]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-352. Inbound doorbell hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/Transport Layer Capture Register	Comments
An incorrect doorbell packet size (not one datum in small transport mode or not two datums in large transport mode)	2	SRIO error/port-write if LTLEECsr[MFE] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[MFE]	No	Error	Updated with the packet ²	Packet is ignored and discarded.
Doorbell Controller Disabled and Doorbell Received	3	No	None	No	Error	-	Packet is ignored and discarded.
Doorbell Controller Enabled but in the Error State and Doorbell Received	3	No	None	No	Error	-	Packet is ignored and discarded.
Internal error occurred during the write of the doorbell queue entry to memory	4	SRIO error/port-write if OMMR[EIE] set	Transaction error in the Doorbell status register (ISSR[TE]). Doorbell Failed in the Port-write and Doorbell CSR (PWDCSR[FA])	No	Error	-	Doorbell controller stops after the current doorbell operation completes. The enqueue pointer is not incremented.

Table continues on the next page...

Table 15-352. Inbound doorbell hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue Entry Written in Local Memory	Response Status	Logical/Transport Layer Capture Register	Comments
An internal error occurred for an earlier posted write of a doorbell queue entry to memory and a subsequent write of a doorbell queue entry to memory is posted before the internal error is detected. An internal error may or may not occur during the subsequent write of a doorbell queue entry to memory.	5	No	None	Yes	Error	-	-

1. These error types are actually detected in the RapidIO port, not in the doorbell controller.
2. If operating in small transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 78-79), LTLACCSR[A] gets the address (packet bits 48-76), LTLDIDCCSR[MDID] gets 0, LTLDIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 16-23), LTLDIDCCSR[MSID] gets 0, LTLDIDCCSR[SID] gets the least significant byte of the source ID (packet bits 24-31), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 32-35), LTLCCCSR[MI] gets 0. If operating in large transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 94-95), LTLACCSR[A] gets the address (packet bits 64-92), LTLDIDCCSR[MDID] gets the most significant byte of the destination ID (packet bits 16-23), LTLDIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 24-31), LTLDIDCCSR[MSID] gets the most significant byte of the source ID (packet bits 32-39), LTLDIDCCSR[SID] gets the least significant byte of the source ID (packet bits 40-47), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 48-51), LTLCCCSR[MI] gets 0.

15.8.2.2.9 Inbound doorbell controller programming errors

The following is the list of programming errors that result in undefined or undesired hardware operation.

Table 15-353. Inbound doorbell programming errors

Error	Interrupt Generated	Status Bit Set	Comments
Reserved value of the doorbell in queue threshold (ID MR[DIQ_THTRES]) or reserved value of the circular doorbell queue size (ID MR[CIRQ_SIZE])	No	No	Undefined operation results
The doorbell in-queue threshold is equal to the doorbell queue size	No	No	Doorbell in queue interrupt occurs when queue is full
The doorbell in-queue threshold is greater than the doorbell queue size	No	No	Doorbell in queue interrupt never occurs
Doorbell queue entry written to non-existent memory	No	No	Memory controller causes the interrupt and update capture registers
Doorbell enqueue and dequeue pointers are not initialized to the same value	No	No	Undefined operation results
The dequeue pointer register is set incorrectly.	No	No	Undefined operation results

15.8.2.2.10 Disabling and enabling the doorbell controller-inbound

When the doorbell controller is disabled by clearing ID MR[DE] the following occurs.

- Queue full clears (ID SR[QF]).
- Doorbell-in-queue clears (ID SR[DIQ]).
- Queue empty is set (ID SR[QE]).
- Doorbell busy clears (ID SR[DUB]) after all pending doorbell queue entry writes to local memory complete.

Before the doorbell controller is re-enabled the doorbell busy bit must be cleared (ID SR[DB]) and the doorbell queue dequeue pointer address registers (DQDPAR, EDQDPAR) and the doorbell queue enqueue pointer address registers (DQEPAR, EDQEPAR) must be initialized to the same value for proper doorbell controller operation.

15.8.2.3 Doorbell message passing logical specification registers

The port-write and doorbell command and status register (PWDCSR) includes several doorbell controller status bits.

These read-only status bits indicate the state of doorbell controller 0.

- Available (PWDCSR[A]). Indicates that the inbound doorbell controller is enabled (ID MR[DE]) and the doorbell controller is not in the internal error state (ID SR[TE] = 0).
- Full (PWDCSR[FU]). This bit reflects the inbound doorbell controller queue full status bit (ID SR[QF]).
- Empty (PWDCSR[EM]). This bit reflects the inverted state of the outbound doorbell busy bit (OD SR[DUB] = 0).
- Busy (PWDCSR[B]). This bit reflects the state of the inbound doorbell controller busy bit ID SR[DUB].
- Failed (PWDCSR[FA]). This bit reflects the state of the transaction error status bit ID SR[TE].
- Error (PWDCSR[ERR]). This bit is always a 0.

15.8.3 Port-write controller

The implementation of the port-write controller is very similar to the inbound message and doorbell controllers except only one queue entry is supported.

The port write is intended as an error reporting mechanism from an end point-less device to a control processor or other system host.

[Figure 15-326](#) depicts an example of the structure of the inbound queue and pointer. The port-write queue only contains one entry with a fixed size of 64 bytes and aligned to a cache line boundary.

The port-write controller uses the error/port-write interrupt (SRIO error/port-write) for notification of incoming port-writes.

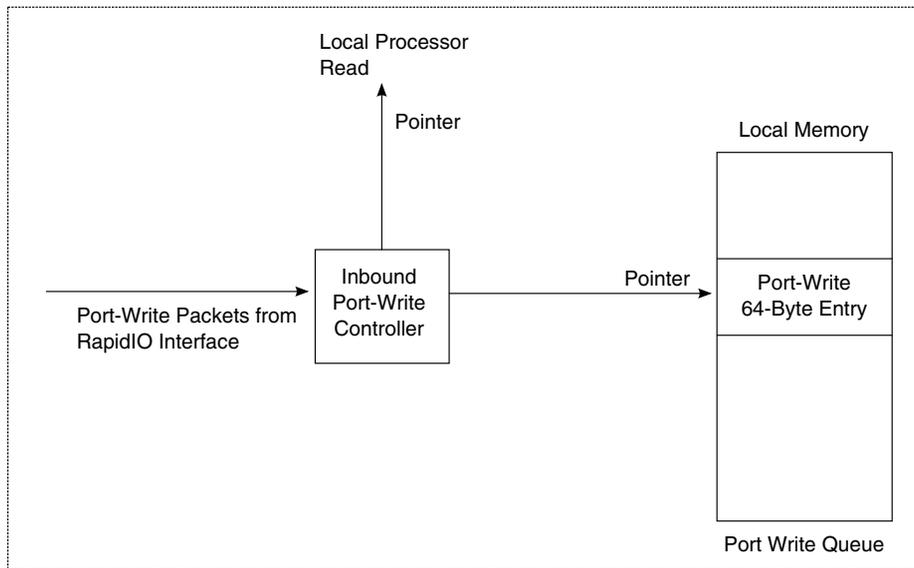


Figure 15-326. Inbound port-write structure

15.8.3.1 Port-write controller initialization

There are many ways in which software can interact with the port-write controller.

One method to initialize the port-write controller is as follows:

- Initialize the port-write queue base address registers (IPWQBAR, EIPWQBAR).
- Clear the status register (IPWSR).
- Set the port-write enable bit (IPWMR[PWE]) along with the other control parameters (snoop enable and the interrupt enable) in the inbound port-write mode register (IPWMR).

15.8.3.2 Port-write controller operation

There are several ways in which software can interact with the port-write controller.

One method is as follows:

- The port-write controller receives a port-write from the RapidIO port. If the inbound port-write controller is enabled (IPW MR[PWE] = 1) and the port-write queue is not full, then the port-write is accepted.
- 64 bytes of payload is stored by the port-write controller in local memory using the value of the port-write queue base address registers (IPW QBAR, EIPW QBAR). Valid payload sizes include 4, 8, 16, 24, 32, 40, 48, 56 or 64 bytes. Note that 64 bytes are always written to memory. If the actual payload size is less than 64 bytes, the non payload data written is undefined.

- If the queue full interrupt enable bit is set (IPW MR[QFIE]) after the memory write completes the error/port-write interrupt is generated by the port-write controller.
- An inbound error/port-write interrupt is generated to the local processor because a port-write was received and this event is enabled to generate the interrupt (IPW MR[QFIE]). Note that the RMU actually generates the SRIO error/port-write output and this is combined with the error interrupt to generate the error/port-write interrupt.
- Software determines that the queue full event caused the interrupt by detecting that the queue full interrupt bit is set when reading the inbound port-write status register (IPW SR[QFI]). Note there are many events that can generate this interrupt. Software must read several registers to determine that the interrupt was generated due to a port-write.
- Software processes the port-write queue entry pointed to by the port-write base address registers (IPW QBAR, EIPW QBAR).
- Software sets the clear queue bit (IPW MR[CQ]) re-enabling the hardware to receive another port-write.
- Software clears the queue full interrupt bit (IPW SR[QFI]) by setting the IPW SR[QFI].

15.8.3.3 Port-write controller interrupt

The error/port-write interrupt is used by the port-write controller.

This interrupt is used to notify the processor that some type of error event has occurred in a RapidIO port, message controller, doorbell controller or port-write controller. There are many events that can generate this interrupt. For example, the error management extensions use this interrupt to notify that error events have occurred. In the port-write controller the following event can generate this interrupt.

- Queue full-an interrupt is generated when this interrupt event is enabled (IPW MR[QFIE]) and a port-write is received and has been written to memory. The event that caused this interrupt is indicated by IPW SR[QFI]. The interrupt is held until the queue is no longer full and the IPW SR[QFI] bit has been cleared by writing a 1.
- Transaction error-an interrupt is generated after a OCN error response is received and this interrupt event is enabled (IPW MR[EIE]).

15.8.3.4 Discarding port-writes

While the queue full bit is set or if a port-write is currently being written to memory but has not completed all received port-writes are discarded.

When a port-write is discarded for one of these reasons the controller sets the port write discarded bit (IPW SR[PWD]). Note that the port-write busy bit (IPW SR[PWB]) indicates that a port-write is currently being written to memory but has not completed.

15.8.3.5 Transaction errors

When an internal error occurs during a local memory write by the port-write controller the following occurs:

- The port-write controller sets the transaction error bit (IPW SR[TE]) and enters the error state.
- If IPW MR[EIE] is set, the interrupt SRIO error/port-write is generated.

15.8.3.5.1 Port-write controller error handling

When an error occurs and the SRIO error/port-write interrupt is generated, the following occurs:

- Software determines the cause of the interrupt and processes the error.
- Software verifies the port-write controller has stopped operation by polling IPW SR[PWB].
- Software disables the port-write controller by clearing IPW MR[PWE].
- Software clears the error by writing a 1 to the corresponding status bit (IPW SR[TE]).
- The port-write unit must be disabled, re-initialized and re-enabled before another maintenance port-write can be received.

When an error occurs and the SRIO error/port-write interrupt is not enabled, the following occurs.

- Software determines that an error has occurred by polling the status bits (IPW SR[TE]).
- Software verifies the port-write controller has stopped operation by polling IPW SR[PWB].
- Software disables the port-write controller by clearing IPW MR[PWE].
- Software clears the error by writing a 1 to the corresponding status bit (IPW SR[TE]).

15.8.3.6 Port-write controller hardware error handling

The following list shows the possible error conditions and what occurs when they are encountered.

The error checking level indicates the order in which errors are checked. Multiple errors can be checked at an error checking level. Once an error is detected no additional error checking beyond the current level is performed. Note that port-writes are processed in a pipeline fashion. The first error detected in the processing pipeline updates the error management extensions registers.

These error condition checks are provided by the messaging unit.

These check are in addition to the error condition checks provided by the RapidIO port given in [Errors and error handling](#).

Table 15-354. Inbound port-write hardware errors

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue entry written in local memory	Response status	Logical/Transport Layer Capture Register	Comments
reserved ftype ¹	1	SRIO error/port-write if LTLEECSR[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UT]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
Reserved tt encoding ¹	1	SRIO error/port-write if LTLEECSR[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[TSE].	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
Large transport size when operating in small transport size or small transport size when operating in large transport size ¹	1	SRIO error/port-write if LTLEECSR[TSE] set	Transport size error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[TSE].	No	No Response	Updated with the packet ²	Packet is ignored and discarded. An error or illegal transaction target error response is not generated.

Table continues on the next page...

Table 15-354. Inbound port-write hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue entry written in local memory	Response status	Logical/Transport Layer Capture Register	Comments
Illegal Destination ID ¹	1	SRIO error/port-write if LTLEECSR[ITTE] set	Illegal transaction target in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITTE]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
An incorrect wr_size encoding (not 4, 8, 16, 24, 32, 40, 48, 56 or 64 bytes), payload size greater than the size defined by wr_size, or not dword aligned when the size is not 4 bytes. ¹	1	SRIO error/port-write if LTLEECSR[ITD] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
wr_size value reserved ¹	1	SRIO error/port-write if LTLEECSR[ITD] set	Message Format error in the Logical/Transport Layer Error Detect CSR LTLEDCSR[ITD]	No	No Response	Updated with the packet ²	Packet is ignored and discarded.
Inbound maintenance port-write received and inbound maintenance port-writes are not supported as indicated by DOCAR[PW] ¹	1	SRIO error/port-write if LTLEECSR[UT] set	Unsupported transaction in the Logical/Transport Layer Error Detect CSR LTLEDCSR[UT]	No	Error	Updated with the packet ²	Packet is ignored and discarded.

Table continues on the next page...

Table 15-354. Inbound port-write hardware errors (continued)

Error	Error Checking Level	Interrupt Generated	Status Bit Set	Queue entry written in local memory	Response status	Logical/Transport Layer Capture Register	Comments
Not an error - An inbound port-write packet with a RapidIO priority of 3	2	-	-	Yes	No Response	Inbound port write considered priority 2 by the inbound port write controller since response from memory is required at priority 3.	-
Port-write Controller Disabled and Port-write Received	2	No	None	No	No Response	-	Packet is ignored and discarded.
Port-write Controller Enabled but in the Error State and Port-write Received	2	No	None	No	No Response	-	Packet is ignored and discarded.
Internal error occurred during the write of the port-write queue entry to memory	3	SRIO error/port-write if IPWMR[EIE] set	Transaction error in the Port-write status register (IPWSR[TE]). Port-write Failed in the Port-write and Doorbell CSR (PWDCSR[PFA])	No	No Response	-	Port-write controller stops after the current port-write operation completes.

1. These error types are actually detected in the RapidIO port, not in the port-write controller
2. If operating in small transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 78-79), LTLACCSR[A] gets the address (packet bits 48-76), LTLDIDCCSR[MDID] gets 0, LTLDIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 16-23), LTLDIDCCSR[MSID] gets 0, LTLDIDCCSR[SID] gets the least significant byte of the source ID (packet bits 24-31), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 32-35), LTLCCCSR[MI] gets 0. If operating in large transport size configuration using the packet LTLACCSR[XA] gets the extended address (packet bits 94-95), LTLACCSR[A] gets the address (packet bits 64-92), LTLDIDCCSR[MDID] gets the most significant byte of the destination ID (packet bits 16-23), LTLDIDCCSR[DID] gets the least significant byte of the destination ID (packet bits 24-31), LTLDIDCCSR[MSID] gets the most significant byte of the source ID (packet bits 32-39), LTLDIDCCSR[SID] gets the least significant byte of the source ID (packet bits 40-47), LTLCCCSR[FT] gets the ftype (packet bits 12-15), LTLCCCSR[TT] gets the ttype (packet bits 48-51), LTLCCCSR[MI] gets 0.

15.8.3.6.1 Port-write controller programming errors

The following is the list of programming errors.

Table 15-355. Inbound port-write programming errors

Error	Interrupt Generated	Status Bit Set	Comments
Port-write queue entry written to non-existent memory	No	No	When a write to memory occurs, the memory controller causes its own interrupt and update its own capture registers. An internal error response will be returned. When the port-write controller receives the error response it will set the transaction error bit (IPWSR[TE]) and be in the error state.

15.8.3.7 Disabling and enabling the port-write controller

When the port-write controller is disabled by clearing IPW MR[PWE] the following occurs:

- Queue full clears (IPW SR[QF]).
- Port-write busy clears (IPW SR[PWB]) after a pending port-write queue entry write completes.

Before the port-write controller is re-enabled (IPW MR[PWE]) the port-write busy bit must be clear (IPW SR[PWB]).

15.8.3.8 Port-write controller message passing logical specification registers

The port-write and doorbell command and status register (PWDCSR) includes several port-write controller status bits.

These read-only status bits only indicate the state of the port-write controller.

- Available (PWDCSR[PA]). Indicates that the port-write controller is enabled (IPW MR[PWE]), the only port-write queue entry is available to be written (IPW SR[QF]) = 0 and the port-write controller is not in the internal error state (IPW SR[TE] = 0).
- Full (PWDCSR[PFU]). This bit reflects the state of the queue full status bit (IPW SR[QF]).
- Empty (PWDCSR[PEM]). This bit always a 1 since a port-write cannot be generated.
- Busy (PWDCSR[PB]). This bit reflects the state of the busy bit IPW SR[PWB].

- Failed (PWDCSR[PFA]). This bit reflects the state of the transaction error status bit IPW SR[TE].
- Error (PWDCSR[PE]). This bit is always a 0.



Chapter 16

PCI Express Interface Controller

The PCI Express controller provides the mechanism to communicate with PCI Express devices.

16.1 Introduction

The PCI Express controller provides the mechanism to communicate with PCI Express devices.

The PCI Express interface is compatible with the *PCI Express™ Base Specification, Revision 1.0a* (available at <http://www.pcisig.org>). It is beyond the scope of this manual to document the intricacies of the PCI Express protocol. This chapter describes the PCI Express controller of this device and provides a basic description of the PCI Express protocol. The specific emphasis is directed at how the device implements the PCI Express specification. Designers of systems incorporating PCI Express devices should refer to the specification for a thorough description of PCI Express.

NOTE

Much of the available PCI Express literature refers to a 16-bit quantity as a WORD and a 32-bit quantity as a DWORD. Note that this is inconsistent with the terminology in the rest of this manual where the terms 'word' and 'double word' refer to a 32-bit and 64-bit quantity, respectively. Where necessary to avoid confusion, the precise number of bits or bytes is specified.

The PCI Express controller connects the internal platform to a 2.5-GHz serial interface. The P2020 offers one, two, or three PCI Express interfaces with up to x4 link width.

Notes are included to indicate variations for multiple instantiations.

As both an initiator and a target device, the PCI Express interface is capable of high-bandwidth data transfer and is designed to support next generation I/O devices. Upon coming out of reset, the PCI Express interface performs link width negotiation and exchanges flow control credits with its link partner. Once link autonegotiation is successful, the controller is in operation.

Internally, the design contains queues to keep track of inbound and outbound transactions. There is control logic that handles buffer management, bus protocol, transaction spawning and tag generation. In addition, there are memory blocks used to store inbound and outbound data.

The PCI Express controller can be configured to operate as either a PCI Express root complex (RC) or an endpoint (EP) device. An RC device connects the host CPU/memory subsystem to I/O devices while an EP device typically denotes a peripheral or I/O device. In RC mode, a PCI Express type 1 configuration header is used; in EP mode, a PCI Express type 0 configuration header is used.

This figure shows a high-level block diagram of the PCI Express controller.

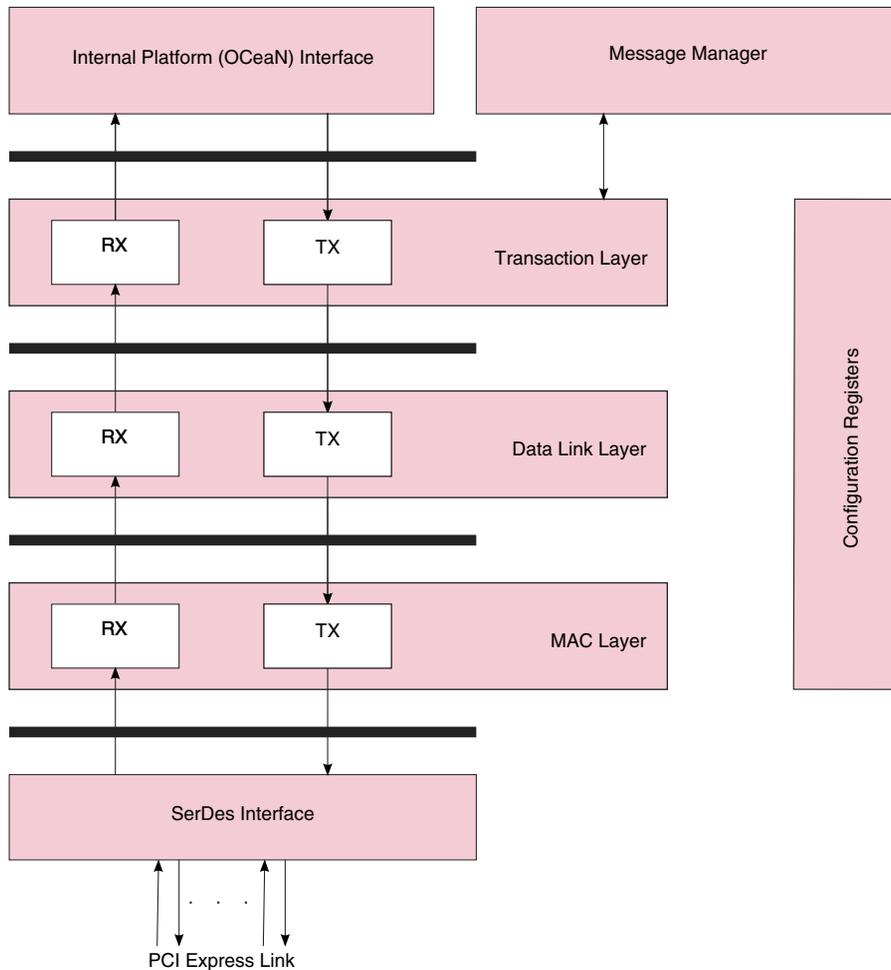


Figure 16-1. PCI Express Controller Block Diagram

As an initiator, the PCI Express controller supports memory read and write operations with a maximum transaction size of 256 bytes. In addition, configuration and I/O transactions are supported if the PCI Express controller is in RC mode. As a target interface, the PCI Express controller accepts read and write operations to local memory space. When configured as an EP device, the PCI Express controller accepts configuration transactions to the internal PCI Express configuration registers. Message generation and acceptance are supported in both RC and EP modes. Locked transactions and inbound I/O transactions are not supported.

16.1.1 Outbound transactions

Outbound internal platform transactions to PCI Express are first mapped to a translation window to determine what PCI Express transactions are to be issued.

A transaction from the internal platform can become a PCI Express Memory, I/O, Message, or Configuration transaction depending on the window attributes.

A transaction may be broken up into smaller sized transactions depending on the original request size, transaction type, and either the PCI Express device control register [MAX_PAYLOAD_SIZE] field for write requests or the PCI Express device control register [MAX_READ_SIZE] field for read requests. The controller performs PCI Express ordering rule checking to determine which transaction is to be sent on the PCI Express link.

In general, transactions are serviced in the order that they are received from the internal platform. Only when there is a stalled condition does the controller apply PCI Express ordering rules to outstanding transactions. For posted write transactions, once all data has been received from the internal platform, the data is forwarded to the PCI Express link and the transaction is considered as done. For non-posted write transactions, the controller waits for the completion packets to return before considering the transaction finished. For non-posted read transactions, the controller waits for all completion packets to return and then forwards all data back to the internal platform before terminating the transaction.

Note that after reset or when recovering from a link down condition, external transactions should not be attempted until the link has successfully trained. Software can poll the LTSSM state status register (PEX_LTSSM_STAT) to check the status of link training before issuing external requests.

16.1.2 Inbound transactions

Inbound PCI Express transactions to internal platform are first mapped to a translation window to determine what internal platform transactions are to be issued.

A transaction may be broken up into smaller sized transactions when sending to the internal platform depending on the original request size, byte enables and starting/ending addresses. The controller performs PCI Express ordering rule checking to determine what transaction is to be sent next to the internal platform.

In general, transactions are serviced in the order that they are received from the PCI Express link. Only when there is a stalled condition does the controller apply PCI Express ordering to outstanding transactions. For posted write transactions, once all data has been received from the PCI Express link, the data is forwarded to the internal platform and the transaction is considered as done. For non-posted read transactions, the controller forwards internal platform data back to the PCI Express link.

Note that the controller splits transactions at the crossing of every 256-byte-aligned boundary when sending data back to the PCI Express link.

16.2 PCI Express features summary

The following is a list of features supported by the PCI Express controller:

- Compatible with the *PCI Express Base Specification, Revision 1.0a*
- Supports root complex (RC) and endpoint (EP) configurations
- 32- and 64-bit PCI Express address support
- 36-bit internal platform address support
- x4, x2, or x1 link support
- Supports accesses to all PCI Express memory and I/O address spaces (requestor only)
- Supports posting of processor-to-PCI Express and PCI Express-to-memory writes
- Supports strong and relaxed transaction ordering rules
- Enforces outbound PCI Express ordering rules and inbound internal platform priority
- PCI Express configuration registers (type 0 in EP mode, type 1 in RC mode)
- Baseline and advanced error reporting support
- One virtual channel (VC0)
- 256-byte maximum payload size (MAX_PAYLOAD_SIZE)
- Supports three inbound general-purpose translation windows and one configuration window
- Supports four outbound translation windows and one default window
- Supports eight non-posted and four posted PCI Express transactions
- Supports up to six internal platform reads and eight internal platform writes. (The maximum number of outstanding transactions at any given time is eight.)
- Credit-based flow control management
- Supports PCI Express messages and interrupts
- Accepts up to 256-byte transactions from the internal platform

16.3 PCI Express modes of operation

Several parameters that affect the PCI Express controller's modes of operation are determined at power-on reset (POR) by reset configuration signals as described in the following table.

Table 16-1. POR parameters for PCI Express controller

Reset configuration signals	Description
Host/Agent Configuration	Selects between root complex (RC) and endpoint (EP) modes
I/O Port Selection	Selects the width of the PCI Express link

16.3.1 Root complex/endpoint modes

The PCI Express controller can function as either a root complex (RC) or an endpoint (EP) on the PCI Express link.

The host/agent configuration input signals `cfg_host_agt[0:2]` determine the RC/EP mode.

16.3.2 Link width

The I/O port selection configuration input signals `cfg_IO_ports[0:3]`, `TSEC1_TXD[3:1]`, and `TSEC2_TX_ER` determine the initial link width. (See [I/O port selection](#).)

See [Minimum frequency requirements](#), for proper selection of CCB clock frequency with regard to PCI Express link width selection.

Note that the link width also depends on the configuration of the serial RapidIO width.

16.4 PCI Express signal descriptions

This topic describes the signals of the PCI Express controller.

PCI Express defines the connection between two devices as a link, which can be composed of a single or multiple lanes. Each lane consists of a differential pair for transmitting (`SD_TXn` and `SD_TXn_B`) and a differential pair for receiving (`SD_RXn` and `SD_RXn_B`) with an embedded data clock.

Although the generic PCI Express controller described here accommodates up to a single x8 link, there are three PCI Express controllers instantiated on the P2020. See [I/O port selection](#) for specific pin muxing details.

This table contains detailed descriptions of the external PCI Express interface signals.

Table 16-2. PCI Express interface signals-Detailed signal descriptions

Signal	I/O	Description	
SD_RX[3:0]	I	Receive data. The receive data signals carry PCI Express packet information. The mapping of SerDes lane and PCI Express controller depends on the POR configuration as described I/O port selection .	
		State Meaning	Asserted/Negated-PCI Express interface.
		Timing	Assertion/Negation-As described in the <i>PCI Express Base Specification, Revision 1.0a</i> .
SD_RX_B[3:0]	I	Receive data, inverted. SD_RX_B[3:0] are the inverted forms of the receive data signals (SD_RX[3:0]).	
		State Meaning	Asserted/Negated-Represents the inverse of data being received from the PCI Express interface.
		Timing	Assertion/Negation-As described in the <i>PCI Express Base Specification, Revision 1.0a</i> .
SD_TX[3:0]	O	Transmit data. The transmit data signals carry PCI Express packet information. The mapping of SerDes lane and PCI Express controller depends on the POR configuration as described in I/O port selection .	
		State Meaning	Asserted/Negated-Represents data being transmitted to the PCI Express interface.
		Timing	Assertion/Negation-As described in the <i>PCI Express Base Specification, Revision 1.0a</i> .
SD_TX_B[3:0]	O	Transmit data, inverted. SD_TX_B[3:0] are the inverted form of the transmit data signals SD_TX[3:0] .	
		State Meaning	Asserted/Negated-Represents the inverse of data being transmitted to the PCI Express interface.
		Timing	Assertion/Negation-As described in the <i>PCI Express Base Specification, Revision 1.0a</i> .

16.5 Memory map/register overview

The PCI Express interface supports the following register types:

- Memory-mapped registers-these registers control PCI Express address translation, PCI error management, and PCI Express configuration register access. These registers are described in the sections below.
- PCI Express configuration registers contained within the PCI Express configuration space-these registers are specified by the PCI Express specification for every PCI Express device. These registers are described in [PCI Express configuration space access](#) and its subsections.

16.6 PCI Express memory-mapped registers

The PCI Express memory mapped registers are accessed by reading and writing to an address comprised of the base address (specified in the CCSRBAR on the local side or the PEXCSRBAR on the PCI Express side) plus the block base address, plus the offset of the specific register to be accessed. Note that all memory-mapped registers (except the PCI Express configuration data register, PEX_CONFIG_DATA) must only be accessed as 32-bit quantities.

Also note that although the table explicitly lists only the registers for the PCI Express controller 1, the register map for PCI Express controllers 2 and 3 are the same except for the block base address. Memory-mapped registers for PCI Express controller 1 begin at block base address 0x0_A000, controller 2 registers begin at 0x0_9000, and controller 3 registers begin at 0x0_8000.

PEX memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
8000	PCI Express configuration address register (PEX3_PEX_CONFIG_ADDR)	32	R/W	0000_0000h	16.6.1/1331
8004	PCI Express configuration data register (PEX3_PEX_CONFIG_DATA)	32	R/W	0000_0000h	16.6.2/1332
800C	PCI Express outbound completion timeout register (PEX3_PEX_OTB_CPL_TOR)	32	R/W	0010_FFFFh	16.6.3/1333
8010	PCI Express configuration retry timeout register (PEX3_PEX_CONF_RTY_TOR)	32	R/W	0400_FFFFh	16.6.4/1334
8014	PCI Express configuration register (PEX3_PEX_CONFIG)	32	R/W	0000_0000h	16.6.5/1335
8020	PCI Express PME & message detect register (PEX3_PEX_PME_MES_DR)	32	w1c	0000_0000h	16.6.6/1336
8024	PCI Express PME & message disable register (PEX3_PEX_PME_MES_DISR)	32	R/W	0000_0000h	16.6.7/1339
8028	PCI Express PME & message interrupt enable register (PEX3_PEX_PME_MES_IER)	32	R/W	0000_0000h	16.6.8/1341
802C	PCI Express power management command register (PEX3_PEX_PMCR)	32	R/W	0000_0000h	16.6.9/1343
8BF8	IP block revision register 1 (PEX3_PEX_IP_BLK_REV1)	32	R	0208_0101h	16.6.10/1344
8BFC	IP block revision register 2 (PEX3_PEX_IP_BLK_REV2)	32	R	0000_0000h	16.6.11/1344
8C00	PCI Express outbound translation address register n (PEX3_PEXOTAR0)	32	R/W	0000_0000h	16.6.12/1345

Table continues on the next page...

PEX memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
8C04	PCI Express outbound translation extended address register n (PEX3_PEXOTEAR0)	32	R/W	0000_0000h	16.6.13/ 1345
8C10	PCI Express outbound window attributes register n (PEX3_PEXOWAR0)	32	R/W	8004_4023h	16.6.14/ 1347
8C20	PCI Express outbound translation address register n (PEX3_PEXOTAR1)	32	R/W	0000_0000h	16.6.12/ 1345
8C24	PCI Express outbound translation extended address register n (PEX3_PEXOTEAR1)	32	R/W	0000_0000h	16.6.13/ 1345
8C28	PCI Express outbound window base address register n (PEX3_PEXOWBAR1)	32	R/W	0000_0000h	16.6.15/ 1349
8C30	PCI Express outbound window attributes register n (PEX3_PEXOWAR1)	32	R/W	0004_4023h	16.6.16/ 1350
8C40	PCI Express outbound translation address register n (PEX3_PEXOTAR2)	32	R/W	0000_0000h	16.6.12/ 1345
8C44	PCI Express outbound translation extended address register n (PEX3_PEXOTEAR2)	32	R/W	0000_0000h	16.6.13/ 1345
8C48	PCI Express outbound window base address register n (PEX3_PEXOWBAR2)	32	R/W	0000_0000h	16.6.15/ 1349
8C50	PCI Express outbound window attributes register n (PEX3_PEXOWAR2)	32	R/W	0004_4023h	16.6.16/ 1350
8C60	PCI Express outbound translation address register n (PEX3_PEXOTAR3)	32	R/W	0000_0000h	16.6.12/ 1345
8C64	PCI Express outbound translation extended address register n (PEX3_PEXOTEAR3)	32	R/W	0000_0000h	16.6.13/ 1345
8C68	PCI Express outbound window base address register n (PEX3_PEXOWBAR3)	32	R/W	0000_0000h	16.6.15/ 1349
8C70	PCI Express outbound window attributes register 3 (PEX3_PEXOWAR3)	32	R/W	0000_0000h	16.6.17/ 1353
8C80	PCI Express outbound translation address register n (PEX3_PEXOTAR4)	32	R/W	0000_0000h	16.6.12/ 1345
8C84	PCI Express outbound translation extended address register n (PEX3_PEXOTEAR4)	32	R/W	0000_0000h	16.6.13/ 1345
8C88	PCI Express outbound window base address register n (PEX3_PEXOWBAR4)	32	R/W	0000_0000h	16.6.15/ 1349
8C90	PCI Express outbound window attributes register 4 (PEX3_PEXOWAR4)	32	R/W	0004_4023h	16.6.18/ 1355
8DA0	PCI Express inbound translation address register n (PEX3_PEXITAR3)	32	R/W	0000_0000h	16.6.19/ 1357
8DA8	PCI Express inbound window base address register n (PEX3_PEXIWBAR3)	32	R/W	0000_0000h	16.6.20/ 1358
8DAC	PCI Express inbound window base extended address register n (PEX3_PEXIWBEAR3)	32	R/W	0000_0000h	16.6.21/ 1359
8DB0	PCI Express inbound window attributes register n (PEX3_PEXIWAR3)	32	R/W	20F4_4023h	16.6.22/ 1360

Table continues on the next page...

PEX memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
8DC0	PCI Express inbound translation address register n (PEX3_PEXITAR2)	32	R/W	0000_0000h	16.6.19/1357
8DC8	PCI Express inbound window base address register n (PEX3_PEXIWBAR2)	32	R/W	0000_0000h	16.6.20/1358
8DCC	PCI Express inbound window base extended address register n (PEX3_PEXIWBEAR2)	32	R/W	0000_0000h	16.6.21/1359
8DD0	PCI Express inbound window attributes register n (PEX3_PEXIWAR2)	32	R/W	20F4_4023h	16.6.22/1360
8DE0	PCI Express inbound translation address register n (PEX3_PEXITAR1)	32	R/W	0000_0000h	16.6.19/1357
8DE8	PCI Express inbound window base address register n (PEX3_PEXIWBAR1)	32	R/W	0000_0000h	16.6.20/1358
8DF0	PCI Express inbound window attributes register n (PEX3_PEXIWAR1)	32	R/W	20F4_4023h	16.6.22/1360
8E00	PCI Express error detect register (PEX3_PEX_ERR_DR)	32	w1c	0000_0000h	16.6.23/1363
8E08	PCI Express error interrupt enable register (PEX3_PEX_ERR_EN)	32	R/W	0000_0000h	16.6.24/1366
8E10	PCI Express error disable register (PEX3_PEX_ERR_DISR)	32	R/W	0000_0000h	16.6.25/1369
8E20	PCI Express error capture status register (PEX3_PEX_ERR_CAP_STAT)	32	R/W	0000_0000h	16.6.26/1371
8E28	PCI Express error capture register n (PEX3_PEX_ERR_CAP_R0)	32	R/W	0000_0000h	16.6.27/1372
8E2C	PCI Express error capture register n (PEX3_PEX_ERR_CAP_R1)	32	R/W	0000_0000h	16.6.27/1372
8E30	PCI Express error capture register n (PEX3_PEX_ERR_CAP_R2)	32	R/W	0000_0000h	16.6.27/1372
8E34	PCI Express error capture register n (PEX3_PEX_ERR_CAP_R3)	32	R/W	0000_0000h	16.6.27/1372
9000	PCI Express configuration address register (PEX2_PEX_CONFIG_ADDR)	32	R/W	0000_0000h	16.6.1/1331
9004	PCI Express configuration data register (PEX2_PEX_CONFIG_DATA)	32	R/W	0000_0000h	16.6.2/1332
900C	PCI Express outbound completion timeout register (PEX2_PEX_OTB_CPL_TOR)	32	R/W	0010_FFFFh	16.6.3/1333
9010	PCI Express configuration retry timeout register (PEX2_PEX_CONF_RTY_TOR)	32	R/W	0400_FFFFh	16.6.4/1334
9014	PCI Express configuration register (PEX2_PEX_CONFIG)	32	R/W	0000_0000h	16.6.5/1335
9020	PCI Express PME & message detect register (PEX2_PEX_PME_MES_DR)	32	w1c	0000_0000h	16.6.6/1336
9024	PCI Express PME & message disable register (PEX2_PEX_PME_MES_DISR)	32	R/W	0000_0000h	16.6.7/1339

Table continues on the next page...

PEX memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
9028	PCI Express PME & message interrupt enable register (PEX2_PEX_PME_MES_IER)	32	R/W	0000_0000h	16.6.8/1341
902C	PCI Express power management command register (PEX2_PEX_PMCR)	32	R/W	0000_0000h	16.6.9/1343
9BF8	IP block revision register 1 (PEX2_PEX_IP_BLK_REV1)	32	R	0208_0101h	16.6.10/1344
9BFC	IP block revision register 2 (PEX2_PEX_IP_BLK_REV2)	32	R	0000_0000h	16.6.11/1344
9C00	PCI Express outbound translation address register n (PEX2_PEXOTAR0)	32	R/W	0000_0000h	16.6.12/1345
9C04	PCI Express outbound translation extended address register n (PEX2_PEXOTEAR0)	32	R/W	0000_0000h	16.6.13/1345
9C10	PCI Express outbound window attributes register n (PEX2_PEXOWAR0)	32	R/W	8004_4023h	16.6.14/1347
9C20	PCI Express outbound translation address register n (PEX2_PEXOTAR1)	32	R/W	0000_0000h	16.6.12/1345
9C24	PCI Express outbound translation extended address register n (PEX2_PEXOTEAR1)	32	R/W	0000_0000h	16.6.13/1345
9C28	PCI Express outbound window base address register n (PEX2_PEXOWBAR1)	32	R/W	0000_0000h	16.6.15/1349
9C30	PCI Express outbound window attributes register n (PEX2_PEXOWAR1)	32	R/W	0004_4023h	16.6.16/1350
9C40	PCI Express outbound translation address register n (PEX2_PEXOTAR2)	32	R/W	0000_0000h	16.6.12/1345
9C44	PCI Express outbound translation extended address register n (PEX2_PEXOTEAR2)	32	R/W	0000_0000h	16.6.13/1345
9C48	PCI Express outbound window base address register n (PEX2_PEXOWBAR2)	32	R/W	0000_0000h	16.6.15/1349
9C50	PCI Express outbound window attributes register n (PEX2_PEXOWAR2)	32	R/W	0004_4023h	16.6.16/1350
9C60	PCI Express outbound translation address register n (PEX2_PEXOTAR3)	32	R/W	0000_0000h	16.6.12/1345
9C64	PCI Express outbound translation extended address register n (PEX2_PEXOTEAR3)	32	R/W	0000_0000h	16.6.13/1345
9C68	PCI Express outbound window base address register n (PEX2_PEXOWBAR3)	32	R/W	0000_0000h	16.6.15/1349
9C70	PCI Express outbound window attributes register 3 (PEX2_PEXOWAR3)	32	R/W	0000_0000h	16.6.17/1353
9C80	PCI Express outbound translation address register n (PEX2_PEXOTAR4)	32	R/W	0000_0000h	16.6.12/1345
9C84	PCI Express outbound translation extended address register n (PEX2_PEXOTEAR4)	32	R/W	0000_0000h	16.6.13/1345
9C88	PCI Express outbound window base address register n (PEX2_PEXOWBAR4)	32	R/W	0000_0000h	16.6.15/1349

Table continues on the next page...

PEX memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
9C90	PCI Express outbound window attributes register 4 (PEX2_PEXOWAR4)	32	R/W	0004_4023h	16.6.18/ 1355
9DA0	PCI Express inbound translation address register n (PEX2_PEXITAR3)	32	R/W	0000_0000h	16.6.19/ 1357
9DA8	PCI Express inbound window base address register n (PEX2_PEXIWBAR3)	32	R/W	0000_0000h	16.6.20/ 1358
9DAC	PCI Express inbound window base extended address register n (PEX2_PEXIWBEAR3)	32	R/W	0000_0000h	16.6.21/ 1359
9DB0	PCI Express inbound window attributes register n (PEX2_PEXIWAR3)	32	R/W	20F4_4023h	16.6.22/ 1360
9DC0	PCI Express inbound translation address register n (PEX2_PEXITAR2)	32	R/W	0000_0000h	16.6.19/ 1357
9DC8	PCI Express inbound window base address register n (PEX2_PEXIWBAR2)	32	R/W	0000_0000h	16.6.20/ 1358
9DCC	PCI Express inbound window base extended address register n (PEX2_PEXIWBEAR2)	32	R/W	0000_0000h	16.6.21/ 1359
9DD0	PCI Express inbound window attributes register n (PEX2_PEXIWAR2)	32	R/W	20F4_4023h	16.6.22/ 1360
9DE0	PCI Express inbound translation address register n (PEX2_PEXITAR1)	32	R/W	0000_0000h	16.6.19/ 1357
9DE8	PCI Express inbound window base address register n (PEX2_PEXIWBAR1)	32	R/W	0000_0000h	16.6.20/ 1358
9DF0	PCI Express inbound window attributes register n (PEX2_PEXIWAR1)	32	R/W	20F4_4023h	16.6.22/ 1360
9E00	PCI Express error detect register (PEX2_PEX_ERR_DR)	32	w1c	0000_0000h	16.6.23/ 1363
9E08	PCI Express error interrupt enable register (PEX2_PEX_ERR_EN)	32	R/W	0000_0000h	16.6.24/ 1366
9E10	PCI Express error disable register (PEX2_PEX_ERR_DISR)	32	R/W	0000_0000h	16.6.25/ 1369
9E20	PCI Express error capture status register (PEX2_PEX_ERR_CAP_STAT)	32	R/W	0000_0000h	16.6.26/ 1371
9E28	PCI Express error capture register n (PEX2_PEX_ERR_CAP_R0)	32	R/W	0000_0000h	16.6.27/ 1372
9E2C	PCI Express error capture register n (PEX2_PEX_ERR_CAP_R1)	32	R/W	0000_0000h	16.6.27/ 1372
9E30	PCI Express error capture register n (PEX2_PEX_ERR_CAP_R2)	32	R/W	0000_0000h	16.6.27/ 1372
9E34	PCI Express error capture register n (PEX2_PEX_ERR_CAP_R3)	32	R/W	0000_0000h	16.6.27/ 1372
A000	PCI Express configuration address register (PEX1_PEX_CONFIG_ADDR)	32	R/W	0000_0000h	16.6.1/1331
A004	PCI Express configuration data register (PEX1_PEX_CONFIG_DATA)	32	R/W	0000_0000h	16.6.2/1332

Table continues on the next page...

PEX memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
A00C	PCI Express outbound completion timeout register (PEX1_PEX_OTB_CPL_TOR)	32	R/W	0010_FFFFh	16.6.3/1333
A010	PCI Express configuration retry timeout register (PEX1_PEX_CONF_RTY_TOR)	32	R/W	0400_FFFFh	16.6.4/1334
A014	PCI Express configuration register (PEX1_PEX_CONFIG)	32	R/W	0000_0000h	16.6.5/1335
A020	PCI Express PME & message detect register (PEX1_PEX_PME_MES_DR)	32	w1c	0000_0000h	16.6.6/1336
A024	PCI Express PME & message disable register (PEX1_PEX_PME_MES_DISR)	32	R/W	0000_0000h	16.6.7/1339
A028	PCI Express PME & message interrupt enable register (PEX1_PEX_PME_MES_IER)	32	R/W	0000_0000h	16.6.8/1341
A02C	PCI Express power management command register (PEX1_PEX_PMCR)	32	R/W	0000_0000h	16.6.9/1343
ABF8	IP block revision register 1 (PEX1_PEX_IP_BLK_REV1)	32	R	0208_0101h	16.6.10/1344
ABFC	IP block revision register 2 (PEX1_PEX_IP_BLK_REV2)	32	R	0000_0000h	16.6.11/1344
AC00	PCI Express outbound translation address register n (PEX1_PEXOTAR0)	32	R/W	0000_0000h	16.6.12/1345
AC04	PCI Express outbound translation extended address register n (PEX1_PEXOTEAR0)	32	R/W	0000_0000h	16.6.13/1345
AC10	PCI Express outbound window attributes register n (PEX1_PEXOWAR0)	32	R/W	8004_4023h	16.6.14/1347
AC20	PCI Express outbound translation address register n (PEX1_PEXOTAR1)	32	R/W	0000_0000h	16.6.12/1345
AC24	PCI Express outbound translation extended address register n (PEX1_PEXOTEAR1)	32	R/W	0000_0000h	16.6.13/1345
AC28	PCI Express outbound window base address register n (PEX1_PEXOWBAR1)	32	R/W	0000_0000h	16.6.15/1349
AC30	PCI Express outbound window attributes register n (PEX1_PEXOWAR1)	32	R/W	0004_4023h	16.6.16/1350
AC40	PCI Express outbound translation address register n (PEX1_PEXOTAR2)	32	R/W	0000_0000h	16.6.12/1345
AC44	PCI Express outbound translation extended address register n (PEX1_PEXOTEAR2)	32	R/W	0000_0000h	16.6.13/1345
AC48	PCI Express outbound window base address register n (PEX1_PEXOWBAR2)	32	R/W	0000_0000h	16.6.15/1349
AC50	PCI Express outbound window attributes register n (PEX1_PEXOWAR2)	32	R/W	0004_4023h	16.6.16/1350
AC60	PCI Express outbound translation address register n (PEX1_PEXOTAR3)	32	R/W	0000_0000h	16.6.12/1345
AC64	PCI Express outbound translation extended address register n (PEX1_PEXOTEAR3)	32	R/W	0000_0000h	16.6.13/1345

Table continues on the next page...

PEX memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
AC68	PCI Express outbound window base address register n (PEX1_PEXOWBAR3)	32	R/W	0000_0000h	16.6.15/ 1349
AC70	PCI Express outbound window attributes register 3 (PEX1_PEXOWAR3)	32	R/W	0000_0000h	16.6.17/ 1353
AC80	PCI Express outbound translation address register n (PEX1_PEXOTAR4)	32	R/W	0000_0000h	16.6.12/ 1345
AC84	PCI Express outbound translation extended address register n (PEX1_PEXOTEAR4)	32	R/W	0000_0000h	16.6.13/ 1345
AC88	PCI Express outbound window base address register n (PEX1_PEXOWBAR4)	32	R/W	0000_0000h	16.6.15/ 1349
AC90	PCI Express outbound window attributes register 4 (PEX1_PEXOWAR4)	32	R/W	0004_4023h	16.6.18/ 1355
ADA0	PCI Express inbound translation address register n (PEX1_PEXITAR3)	32	R/W	0000_0000h	16.6.19/ 1357
ADA8	PCI Express inbound window base address register n (PEX1_PEXIWBAR3)	32	R/W	0000_0000h	16.6.20/ 1358
ADAC	PCI Express inbound window base extended address register n (PEX1_PEXIWBEAR3)	32	R/W	0000_0000h	16.6.21/ 1359
ADB0	PCI Express inbound window attributes register n (PEX1_PEXIWAR3)	32	R/W	20F4_4023h	16.6.22/ 1360
ADC0	PCI Express inbound translation address register n (PEX1_PEXITAR2)	32	R/W	0000_0000h	16.6.19/ 1357
ADC8	PCI Express inbound window base address register n (PEX1_PEXIWBAR2)	32	R/W	0000_0000h	16.6.20/ 1358
ADCC	PCI Express inbound window base extended address register n (PEX1_PEXIWBEAR2)	32	R/W	0000_0000h	16.6.21/ 1359
ADD0	PCI Express inbound window attributes register n (PEX1_PEXIWAR2)	32	R/W	20F4_4023h	16.6.22/ 1360
ADE0	PCI Express inbound translation address register n (PEX1_PEXITAR1)	32	R/W	0000_0000h	16.6.19/ 1357
ADE8	PCI Express inbound window base address register n (PEX1_PEXIWBAR1)	32	R/W	0000_0000h	16.6.20/ 1358
ADF0	PCI Express inbound window attributes register n (PEX1_PEXIWAR1)	32	R/W	20F4_4023h	16.6.22/ 1360
AE00	PCI Express error detect register (PEX1_PEX_ERR_DR)	32	w1c	0000_0000h	16.6.23/ 1363
AE08	PCI Express error interrupt enable register (PEX1_PEX_ERR_EN)	32	R/W	0000_0000h	16.6.24/ 1366
AE10	PCI Express error disable register (PEX1_PEX_ERR_DISR)	32	R/W	0000_0000h	16.6.25/ 1369
AE20	PCI Express error capture status register (PEX1_PEX_ERR_CAP_STAT)	32	R/W	0000_0000h	16.6.26/ 1371
AE28	PCI Express error capture register n (PEX1_PEX_ERR_CAP_R0)	32	R/W	0000_0000h	16.6.27/ 1372

Table continues on the next page...

PEX memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
AE2C	PCI Express error capture register n (PEX1_PEX_ERR_CAP_R1)	32	R/W	0000_0000h	16.6.27/ 1372
AE30	PCI Express error capture register n (PEX1_PEX_ERR_CAP_R2)	32	R/W	0000_0000h	16.6.27/ 1372
AE34	PCI Express error capture register n (PEX1_PEX_ERR_CAP_R3)	32	R/W	0000_0000h	16.6.27/ 1372

16.6.1 PCI Express configuration address register (PEX_x_PEX_CONFIG_ADDR)

The PCI Express configuration address register contains address information for accesses to PCI Express internal and external configuration registers.

Both root complex (RC) and endpoint (EP) configuration headers contain 4096 bytes of address space. To access a register within the header, both the extended register number and the register number fields are concatenated to form the 4-byte aligned address of the register. That is, the register address is extended to register number 0b00.

Address: Base address + 0h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	EN	Reserved			EXTREGN				BUSN							
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	DEVN				FUNCN				REGN							Reserved
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PEX_x_PEX_CONFIG_ADDR field descriptions

Field	Description
0 EN	Enable. This bit allows a PCI Express configuration access when PEX_CONFIG_DATA is accessed. If this bit is cleared, writing to PEX_CONFIG_DATA has no effect and reading PEX_CONFIG_DATA returns unknown data.
1–3 -	This field is reserved. Reserved
4–7 EXTREGN	Extended register number. This field allows access to extended PCI Express configuration space (that is, the registers in the offset range from 0x100 to 0xFFFF).
8–15 BUSN	Bus number. PCI bus number to access
16–20 DEVN	Device number. Device number to access on specified bus

Table continues on the next page...

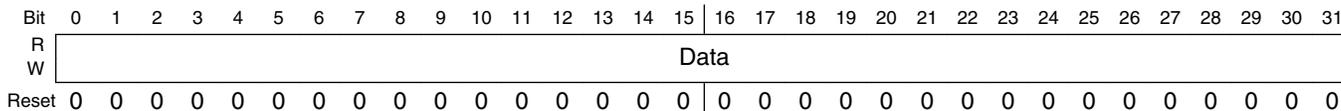
PEXx_PEX_CONFIG_ADDR field descriptions (continued)

Field	Description
21–23 FUNCN	Function number. Function to access within specified device
24–29 REGN	Register number. 32-bit register to access within specified device
30–31 -	This field is reserved. Reserved

16.6.2 PCI Express configuration data register (PEXx_PEX_CONFIG_DATA)

The PCI Express configuration data register is a 32-bit port for internal and external configuration access. Note that accesses of 1, 2, or 4 bytes to the PCI Express configuration data register are allowed. Also note that accesses to the little-endian PCI Express configuration space must be properly formatted. See [Byte order for configuration transactions](#) for more information.

Address: Base address + 4h offset



PEXx_PEX_CONFIG_DATA field descriptions

Field	Description
0–31 Data	A read or write to this register starts a PCI Express configuration cycle if the PEX_CONFIG_ADDR enable bit is set (PEX_CONFIG_ADDR[EN] = 1).

16.6.3 PCI Express outbound completion timeout register (PEX_x_PEX_OTB_CPL_TOR)

The PCI Express outbound completion timeout register contains the maximum wait time for a response to come back as a result of an outbound non-posted request before a timeout condition occurs.

Address: Base address + Ch offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	TC															
W	TC															
Reset	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

PEX_x_PEX_OTB_CPL_TOR field descriptions

Field	Description
0 TD	Timeout disable. This bit controls the enabling/disabling of the timeout function. 0 Enable completion timeout 1 Disable completion timeout
1–7 -	This field is reserved. Reserved
8–31 TC	Timeout counter. This is the value that is used to load the response counter of the completion timeout. One TC unit is 8x the PCI Express controller clock period; that is, one TC unit is 20 ns at 400 MHz, and 30 ns at 266.66 MHz. The following are examples of timeout periods based on different TC settings: 0x00_0000 Reserved 0x10_FFFF 22.28 ms at 400 MHz controller clock; 33.34 ms at 266.66 MHz controller clock 0xFF_FFFF 335.54 ms at 400 MHz controller clock; 503.31 ms at 266.66 MHz controller clock

16.6.4 PCI Express configuration retry timeout register (PEX_x_PEX_CONF_RTY_TOR)

The PCI Express configuration retry timeout register contains the maximum time period during which retries of configuration transactions which resulted in a CRS response occur.

Address: Base address + 10h offset

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R																	
W	RD	Reserved			TC												
Reset	0	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	TC																
W	TC																
Reset	1	1	1	1	1	1	1	1		1	1	1	1	1	1	1	1

PEX_x_PEX_CONF_RTY_TOR field descriptions

Field	Description
0 RD	<p>Retry disable. This bit disables the retry of a configuration transaction that receives a CRS status response packet.</p> <p>0 Enable retry of a configuration transaction in response to receiving a CRS status response until the timeout counter (defined by the PEX_CONF_RTY_TOR[TC] field) has expired.</p> <p>1 Disable retry of a configuration transaction regardless of receiving a CRS status response.</p>
1–3 -	This field is reserved. Reserved
4–31 TC	<p>Timeout counter. This is the value that is used to load the CRS response counter.</p> <p>One TC unit is 8x the PCI Express controller clock period; that is, one TC unit is 20 ns at 400 MHz and 30 ns at 266.66 MHz.</p> <p>Timeout period based on different TC settings:</p> <p>0x000_0000 Reserved</p> <p>0x400_FFFF 1.34 s at 400 MHz controller clock, 2.02 s at 266.66 MHz controller clock</p> <p>0xFFF_FFFF 5.37 s at 400 MHz controller clock, 8.05 s at 266.66 MHz controller clock</p>

16.6.5 PCI Express configuration register (PEXx_PEX_CONFIG)

The PCI Express configuration register contains various control switches for the controller.

Address: Base address + 14h offset

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved																	
W	Reserved																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved												SAC	Reserved		SP	SCC	
W	Reserved												SAC	Reserved		SP	SCC	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	

PEXx_PEX_CONFIG field descriptions

Field	Description
0–26 -	This field is reserved. Reserved
27 SAC	Sense ASPM Control. This bit controls the default value of ASPM of PEX Link Control Register's bit 0. See PCI Express Link Control Register (Link_Control_Register) for more information.
28–29 -	This field is reserved. Reserved
30 SP	Slot Present. This bit controls the default value of the PCI Express capabilities register [slot] bit. See PCI Express Capabilities Register (Capabilities_Register) for more information.
31 SCC	Slot Clock Configuration. This bit controls the default value of the PCI Express link status register [SCC] bit. See PCI Express Link Status Register (Link_Status_Register) for more information.

16.6.6 PCI Express PME & message detect register (PEXx_PEX_PME_MES_DR)

The PCI Express PME and message detect register logs inbound messages and PME events that are detected by the PCI Express controller. This register is a write-1-to-clear type register.

Address: Base address + 20h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	PTO	Reserved	ENL23	EXL23	Reserved	HRD	LDD	Reserved	AION	AIB	AIOF	PION	PIB	PIOF	ABP	
W	w1c		w1c	w1c		w1c	w1c		w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PEXx_PEX_PME_MES_DR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16 PTO	PME turn off. This bit indicates the detection of a PME_Turn_Off message. This bit is only valid in EP mode. 1 A PME_Turn_Off_message is detected 0 No PME_Turn_Off message detected
17 -	This field is reserved. Reserved. Note that during normal operation, this bit may be set (falsely). The bit may be ignored and cleared (w1c) without consequence.
18 ENL23	Entered L2/L3 ready state. This bit indicates that the PCI Express controller has entered L2/L3 state. This is only valid in RC mode. 1 L2/L3 ready state has been entered 0 L2/L3 ready state has not been entered
19 EXL23	Exit L2/L3 ready state. This bit indicates that the PCI Express controller has exited the L2/L3 state. This is only valid in RC mode. 1 Exit L2/L3 state has been detected 0 Exit L2/L3 state not detected
20 -	This field is reserved. Reserved. Note that during normal operation, this bit may be set (falsely). The bit may be ignored and cleared (w1c) without consequence.
21 HRD	Hot reset detected. This bit indicates that the PCI Express controller has detected a hot reset condition on the link. The controller is reset and cleans up all outstanding transactions. Link retraining takes place once hot reset state is exited. This is valid only in EP mode. 1 Hot reset request has been detected 0 Hot reset request not detected
22 LDD	Link down detected. This bit indicates that a link down condition has been detected. The controller is reset and then cleans up all outstanding transactions. Link retraining takes place once the controller has cleaned itself up. Note that for EP, this bit and HRD are typically set when a hot reset event is detected. 1 Link down has been detected 0 Link down not detected
23–24 -	This field is reserved. Reserved
25 AION	Attention indicator on. This bit indicates the detection of an Attention_Indicator_On message. This bit is only valid in EP mode. 1 Attention indicator on message is detected 0 No attention indicator on message detected
26 AIB	Attention indicator blink. This bit indicates the detection of an Attention_Indicator_Blink message. This bit is only valid in EP mode. 1 Attention indicator blink message is detected 0 No attention indicator blink message detected
27 AIOF	Attention indicator off. This bit indicates the detection of an Attention_Indicator_Off message. This bit is only valid in EP mode.

Table continues on the next page...

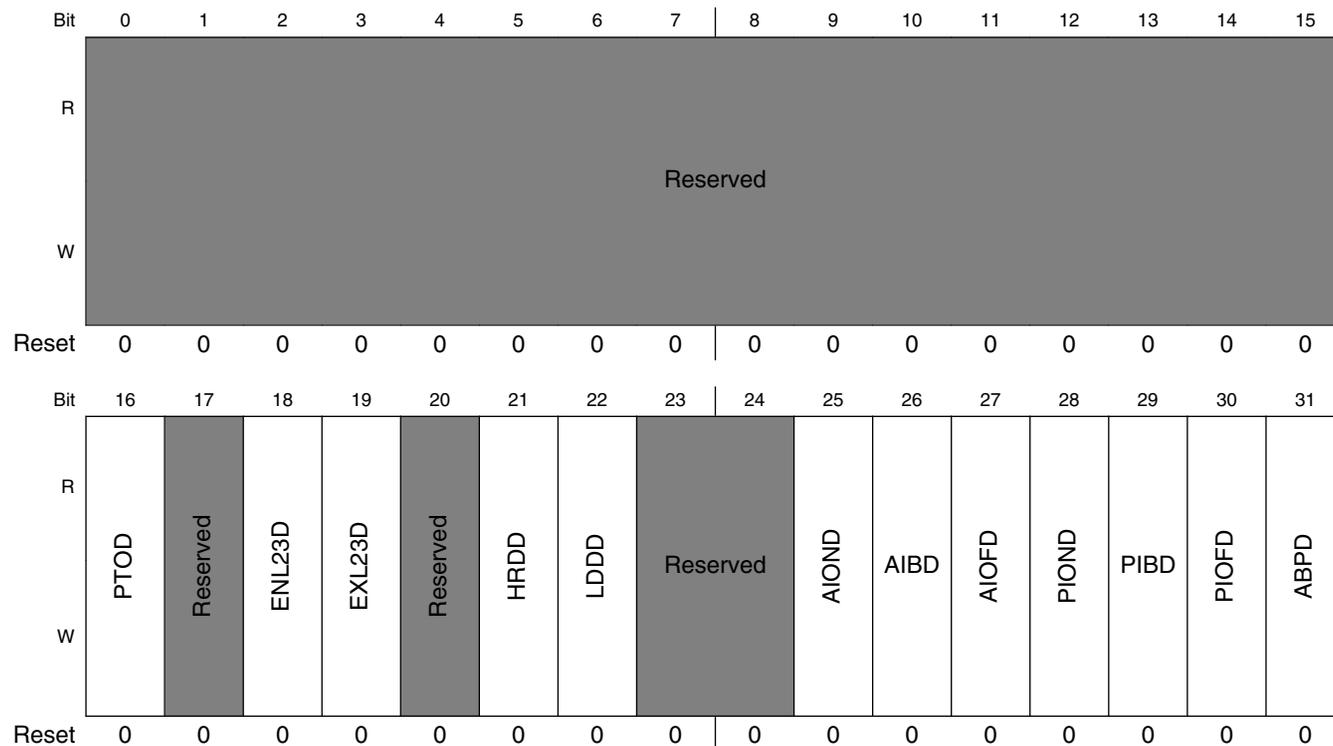
PEXx_PEX_PME_MES_DR field descriptions (continued)

Field	Description
	1 Attention indicator off message is detected 0 No attention indicator off message detected
28 PION	Power indicator on. This bit indicates the detection of a Power_Indicator_On message. This bit is only valid in EP mode. 1 Power indicator on message is detected 0 No power indicator on message detected
29 PIB	Power indicator blink. This bit indicates the detection of an Power_Indicator_Blink message. This bit is only valid in EP mode. 1 Power indicator blink message is detected 0 No power indicator blink message detected
30 PIOF	Power indicator off. This bit indicates the detection of an Power_Indicator_Off message. This bit is only valid in EP mode. 1 Power indicator off message is detected 0 No power indicator off message detected
31 ABP	Attention button pressed. This bit indicates the detection of an Attention_Button_Pressed message. This bit is only valid in RC mode. 1 Attention button press message is detected 0 No attention button press message detected

16.6.7 PCI Express PME & message disable register (PEX_x_PEX_PME_MES_DISR)

The PCI Express PME and message disable register when set, prevents the detection of the corresponding bits in the PCI Express PME and message detect register.

Address: Base address + 24h offset



PEX_x_PEX_PME_MES_DISR field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16 PTOD	PME turn off disable. When set, this bit disables the setting of PEX_PME_MES_DR[PTO] bit. 1 Disable PME_Turn_Off_message detection 0 Enable PME_Turn_Off message detection
17 -	This field is reserved. Reserved
18 ENL23D	Entered_L2/L3 ready disable. When set, this bit disables the setting of PEX_PME_MES_DR[ENL23] bit. 1 Disable Entered_L2/L3 ready state detection 0 Enable Entered_L2/L3 ready state detection

Table continues on the next page...

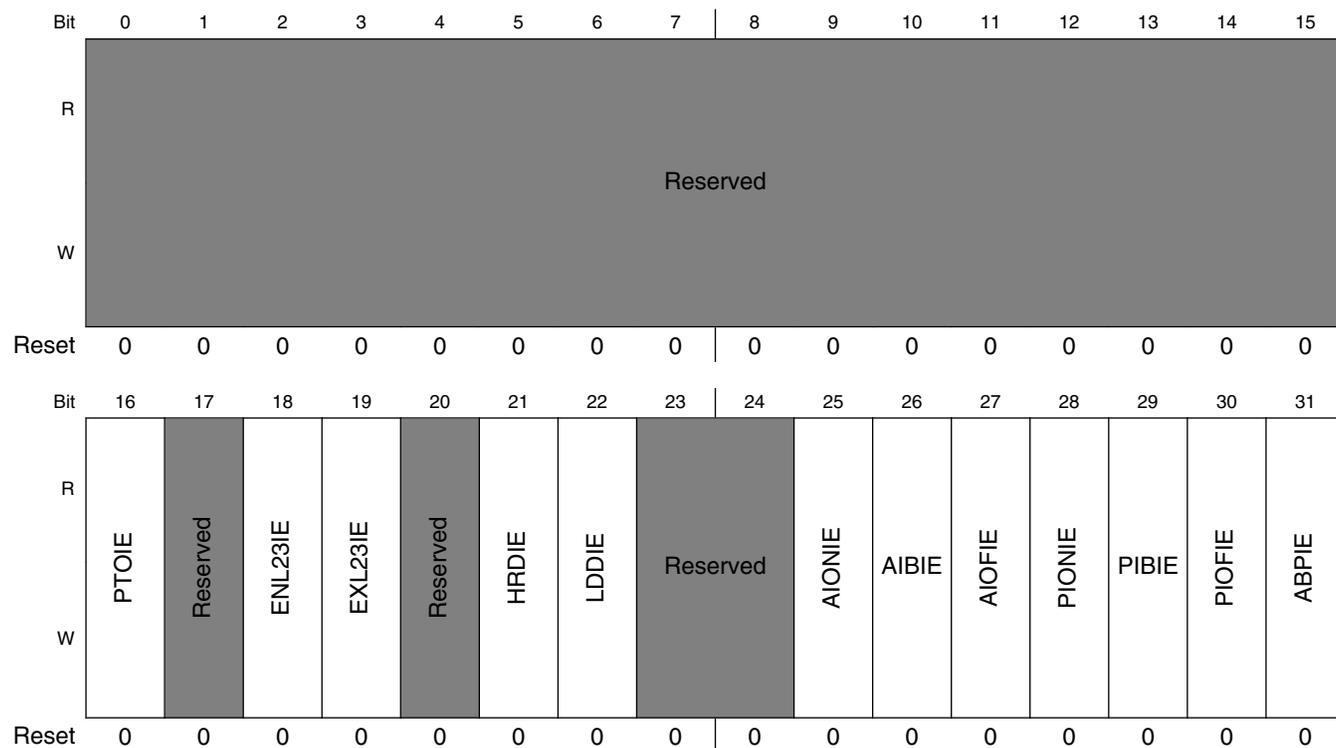
PEXx_PEX_PME_MES_DISR field descriptions (continued)

Field	Description
19 EXL23D	Exited_L2/L3 ready disable. When set, this bit disables the setting of PEX_PME_MES_DR[EXL23] bit. 1 Disable Exited_L2/L3 ready state detection 0 Enable Exited_L2/L3 ready state detection
20 -	This field is reserved. Reserved
21 HRDD	Hot reset detected disable. When set, this bit disables the setting of PEX_PME_MES_DR[HRD] bit. 1 Disable hot reset state detection 0 Enable hot reset state detection
22 LDDD	Link down detected disable. When set, this bit disables the setting of PEX_PME_MES_DR[LDD] bit. 1 Disable link down state detection 0 Enable link down state detection
23–24 -	This field is reserved. Reserved
25 AIOND	Attention indicator on disable. When set, this bit disables the setting of PEX_PME_MES_DR[AION] bit. 1 Disable attention indicator on message detection 0 Enable attention indicator on message detection
26 AIBD	Attention indicator blink disable. When set, this bit disables the setting of PEX_PME_MES_DR[AIB] bit. 1 Disable attention indicator blink message detection 0 Enable attention indicator blink message detection
27 AIOFD	Attention indicator off disable. When set, this bit disables the setting of PEX_PME_MES_DR[AIOF] bit. 1 Disable attention indicator off message detection 0 Enable attention indicator off message detection
28 PIOND	Power indicator on disable. When set, this bit disables the setting of PEX_PME_MES_DR[PION] bit. 1 Disable power indicator on message detection 0 Enable power indicator on message detection
29 PIBD	Power indicator blink disable. When set, this bit disables the setting of PEX_PME_MES_DR[PIB] bit. 1 Disable power indicator blink message detection 0 Enable power indicator blink message detection
30 PIOFD	Power indicator off disable. When set, this bit disables the setting of PEX_PME_MES_DR[PIOF] bit. 1 Disable power indicator off message detection 0 Enable power indicator off message detection
31 ABPD	Attention button pressed disable. When set, this bit disables the setting of PEX_PME_MES_DR[ABP] bit. 1 Disable attention button press message detection 0 Enable attention button press message detection

16.6.8 PCI Express PME & message interrupt enable register (PEXx_PEX_PME_MES_IER)

The PCI Express PME and message interrupt enable register allows for the detection of a message or a PME event to generate an interrupt, provided that the corresponding bit in the PCI Express PME and message detect register is set.

Address: Base address + 28h offset



PEXx_PEX_PME_MES_IER field descriptions

Field	Description
0–15 -	This field is reserved. Reserved
16 PTOIE	PME turn off interrupt enable. When set and PEX_PME_MES_DR[PTO]=1 generates an interrupt. 1 Enable PME_Turn_Off_message interrupt generation 0 Disable PME_Turn_Off message interrupt generation
17 -	This field is reserved. Reserved
18 ENL23IE	Entered L2/L3 ready interrupt enable. When set and PEX_PME_MES_DR[ENL23]=1 generates an interrupt.

Table continues on the next page...

PEXx_PEX_PME_MES_IER field descriptions (continued)

Field	Description
	1 Enable Entered_L2/L3 ready state interrupt generation 0 Disable Entered_L2/L3 ready state interrupt generation
19 EXL23IE	Exited L2/L3 ready interrupt enable. When set and PEX_PME_MES_DR[EXL23]=1 generates an interrupt. 1 Enable Exited_L2/L3 ready state interrupt generation 0 Disable Exited_L2/L3 ready state interrupt generation
20 -	This field is reserved. Reserved
21 HRDIE	Hot reset detected interrupt enable. When set and PEX_PME_MES_DR[HRD]=1 generates an interrupt. 1 Enable hot reset state interrupt generation 0 Disable hot reset state interrupt generation
22 LDDIE	Link down detected interrupt enable. When set and PEX_PME_MES_DR[LDD]=1 generates an interrupt. 1 Enable link down state interrupt generation 0 Disable link down state interrupt generation
23–24 -	This field is reserved. Reserved
25 AIONIE	Attention indicator on interrupt enable. When set and PEX_PME_MES_DR[AION]=1 generates an interrupt. 1 Enable attention indicator on message interrupt generation 0 Disable attention indicator on message interrupt generation
26 AIBIE	Attention indicator blink interrupt enable. When set and PEX_PME_MES_DR[AIB]=1 generates an interrupt. 1 Enable attention indicator blink message interrupt generation 0 Disable attention indicator blink message interrupt generation
27 AIOFIE	Attention indicator off interrupt enable. When set and PEX_PME_MES_DR[AIOF]=1 generates an interrupt. 1 Enable attention indicator off message interrupt generation 0 Disable attention indicator off message interrupt generation
28 PIONIE	Power indicator on interrupt enable. When set and PEX_PME_MES_DR[PION]=1 generates an interrupt. 1 Enable power indicator on message interrupt generation 0 Disable power indicator on message interrupt generation
29 PIBIE	Power indicator blink interrupt enable. When set and PEX_PME_MES_DR[PIB]=1 generates an interrupt. 1 Enable power indicator blink message interrupt generation 0 Disable power indicator blink message interrupt generation
30 PIOFIE	Power indicator off interrupt enable. When set and PEX_PME_MES_DR[PIOF]=1 generates an interrupt. 1 Enable power indicator off message interrupt generation 0 Disable power indicator off message interrupt generation
31 ABPIE	Attention button pressed interrupt enable. When set and PEX_PME_MES_DR[ABP]=1 generates an interrupt.

Table continues on the next page...

PEXx_PEX_PME_MES_IER field descriptions (continued)

Field	Description
1	Enable attention button press message interrupt generation
0	Disable attention button press message interrupt generation

16.6.9 PCI Express power management command register (PEXx_PEX_PMCR)

The PCI Express power management command register provides software a mechanism to allow the PCI Express controller to get back to L0 link state.

Address: Base address + 2Ch offset

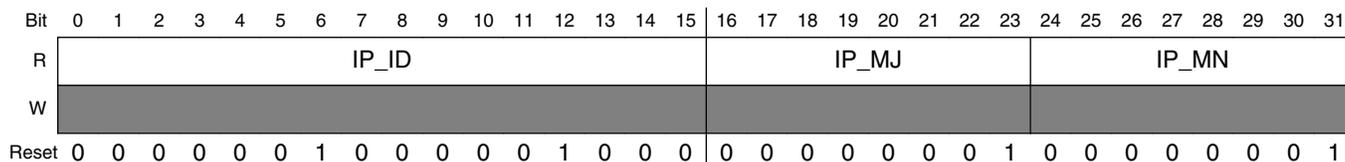
Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved													SPMES	EXL2S	PTOMR
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PEXx_PEX_PMCR field descriptions

Field	Description
0–28 -	This field is reserved. Reserved
29 SPMES	Set PME status. This sets the PME status bit and if PME is enabled (see PCI Express Power Management Status and Control Register (Power_Management_Status_and_Control_Register) , for more information) it transmits a PM_PME message upstream. This bit should not be used when in RC mode. This bit is self-clearing.
30 EXL2S	Exit L2 state. When set exits the link state out of L2/L3 ready state in order to send new requests. The request is only made when entered_L2/L3 ready state is active. This bit is self-clearing. When the link has exited L2/L3 ready state, the status bit Exit_L2/L3 ready state is set. This bit should not be used when in EP mode.
31 PTOMR	PME_Turn_Off message request. When set broadcasts a PME turn_off message. This bit should not be used when in EP mode. This bit is self-clearing

16.6.10 IP block revision register 1 (PEX_x_PEX_IP_BLK_REV1)

Address: Base address + BF8h offset

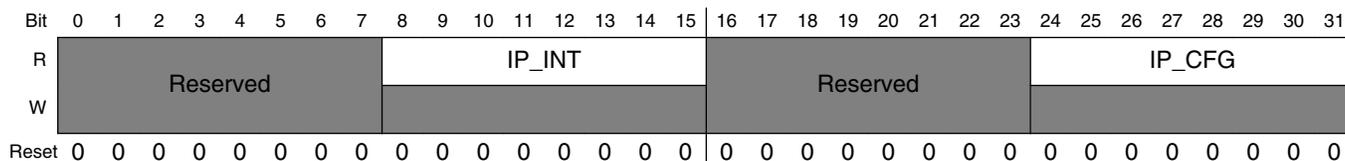


PEX_x_PEX_IP_BLK_REV1 field descriptions

Field	Description
0–15 IP_ID	Block ID
16–23 IP_MJ	Block Major Revision
24–31 IP_MN	Block Minor Revision

16.6.11 IP block revision register 2 (PEX_x_PEX_IP_BLK_REV2)

Address: Base address + BFC_h offset



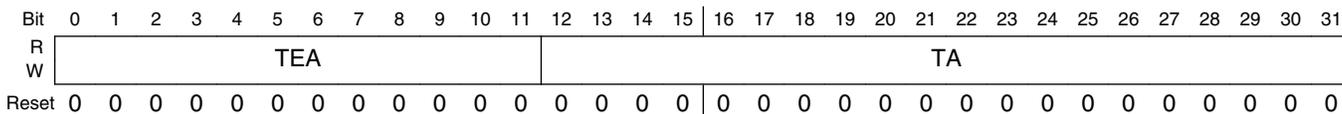
PEX_x_PEX_IP_BLK_REV2 field descriptions

Field	Description
0–7 -	This field is reserved. Reserved
8–15 IP_INT	Block integration option
16–23 -	This field is reserved. Reserved
24–31 IP_CFG	Block configuration option

16.6.12 PCI Express outbound translation address register n (PEXx_PEXOTARn)

The PCI Express outbound translation address registers select the starting addresses in the system address space for window hits within the PCI Express outbound address translation windows. The new translated address is created by concatenating the transaction offset to this translation address.

Address: Base address + C00h offset + (32d × i), where i=0d to 4d



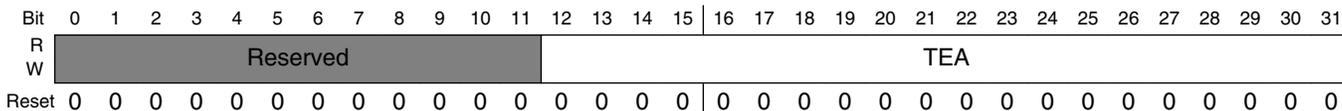
PEXx_PEXOTARn field descriptions

Field	Description
0–11 TEA	Translation extended address. System address which indicates the starting point of the outbound translated address. The translation address must be aligned based on the size field. Corresponds to PCI Express address bits [43:32] (bit 32 is the lsb).
12–31 TA	Translation address. System address which indicates the starting point of the outbound translated address. The translation address must be aligned based on the size field. This corresponds to PCI Express address bits [31:12].

16.6.13 PCI Express outbound translation extended address register n (PEXx_PEXOTEARN)

The PCI Express outbound translation extended address registers contain the most-significant bits of a 64 bit translation address.

Address: Base address + C04h offset + (32d × i), where i=0d to 4d



PEXx_PEXOTEARN field descriptions

Field	Description
0–11 -	This field is reserved. Reserved

Table continues on the next page...

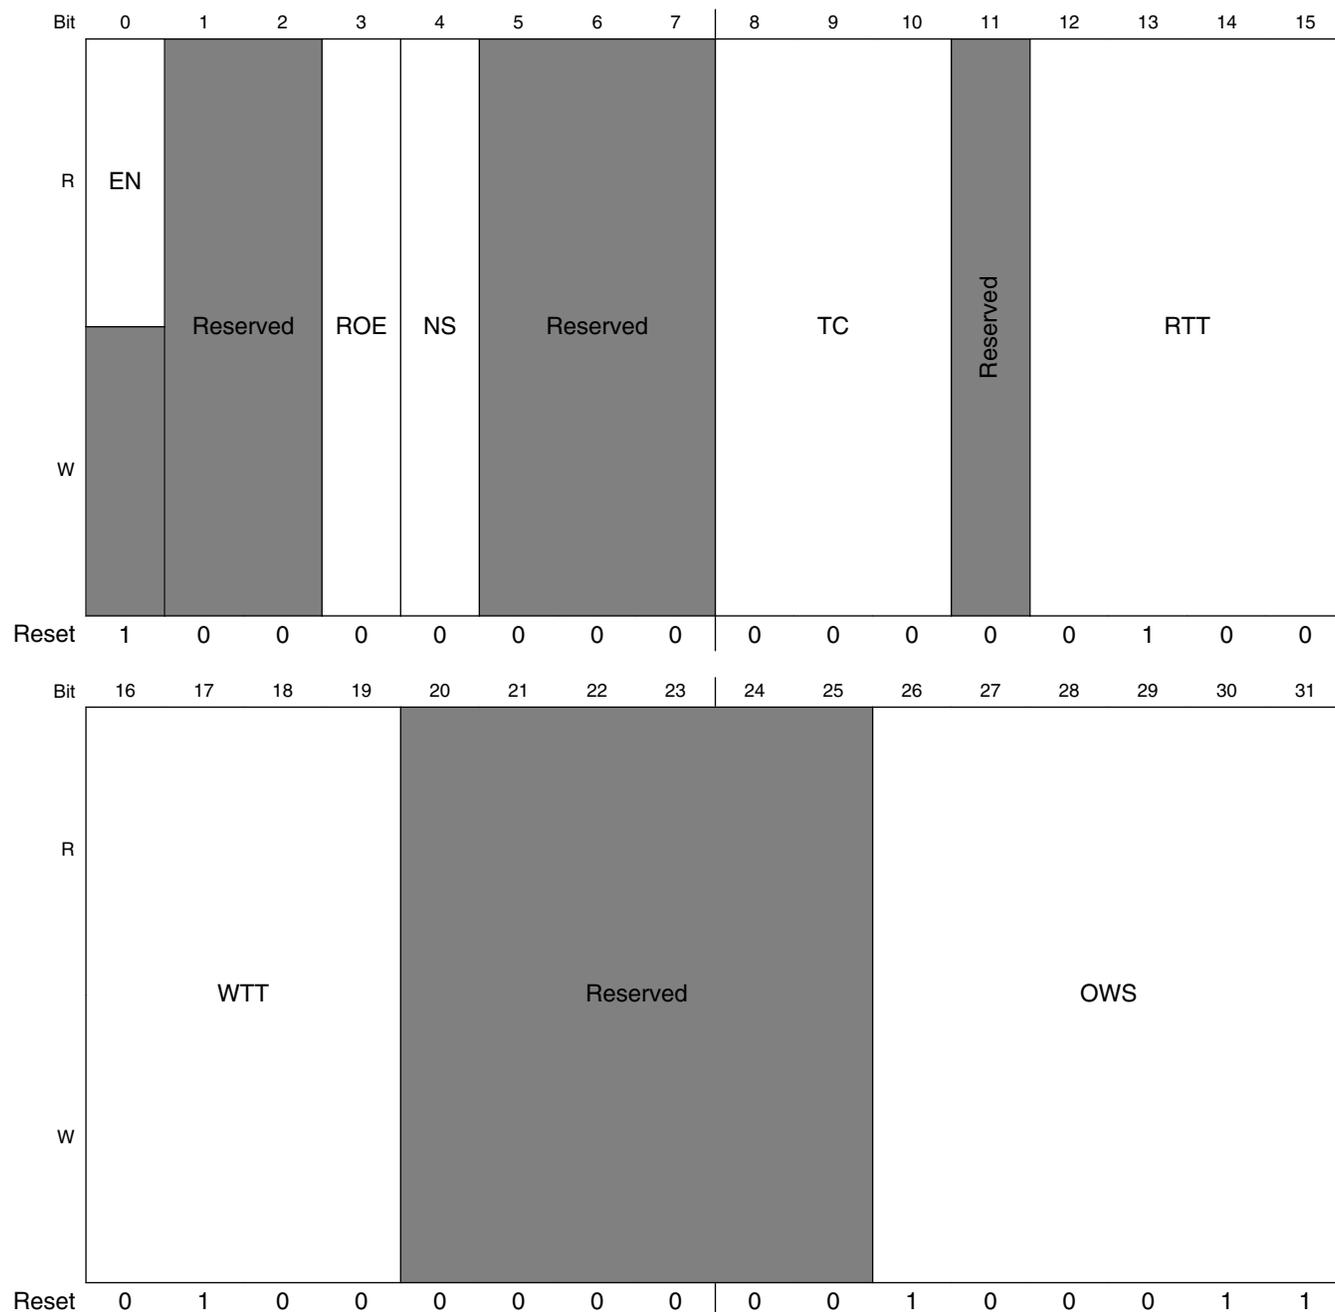
PEXx_PEXOTEAR_n field descriptions (continued)

Field	Description
12–31 TEA	Translation extended address. System address which indicates the starting point of the outbound translated address. The translation address must be aligned based on the size field. Corresponds to PCI Express address bits [63:44].

16.6.14 PCI Express outbound window attributes register n (PEXx_PEXOWAR0)

The PCI Express outbound window attributes registers define the window sizes to translate and other attributes for the translations. 64 Gbytes is the largest window size allowed.

Address: Base address + C10h offset



PEXx_PEXOWAR0 field descriptions

Field	Description
0 EN	<p>Enable. This bit enables this address translation window. For the default window, this bit is read-only and always hardwired to 1.</p> <p>0 Disable outbound translation window 1 Enable outbound translation window</p>
1–2 -	This field is reserved. Reserved
3 ROE	<p>Relaxed ordering enable. When this bit and the RO bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the relaxed ordering bit for the packet is enabled. This bit only applies to memory transactions.</p> <p>0 Default ordering 1 Relaxed ordering</p>
4 NS	<p>No snoop enable. When this bit and the NSE bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the no snoop bit for the packet is enabled. This bit only applies to memory transactions.</p> <p>0 Snoopable 1 No snoop</p>
5–7 -	This field is reserved. Reserved
8–10 TC	<p>Traffic class. This field indicates the traffic class of the outbound packet. This field only applies to memory transaction. All other transaction types should set the TC field to 0.</p> <p>NOTE: Traffic class settings are passed through to the PCI Express link, but no specific actions are taken in the device based on traffic class.</p> <p>000 TC0 001 TC1 010 TC2 011 TC3 100 TC4 101 TC5 110 TC6 111 TC7</p>
11 -	This field is reserved. Reserved
12–15 RTT	<p>Read transaction type. Read transaction type to run on the PCI Express link. Settings not shown are reserved.</p> <p>0010 Configuration read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. 0100 Memory read 1000 IO read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary.</p>
16–19 WTT	<p>Write transaction type. Write transaction type to run on the PCI Express link. Settings not shown are reserved.</p> <p>0010 Configuration write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI</p>

Table continues on the next page...

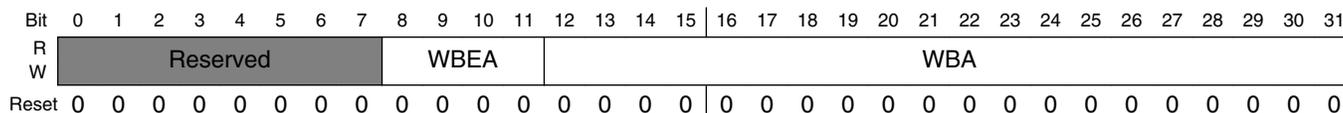
PEXx_PEXOWAR0 field descriptions (continued)

Field	Description
	express port must not translate to outbound configuration write transactions on another PCI Express port. 0100 Memory write 0101 Message write. Only support 4-byte size access on a 4-byte address boundary. 1000 IO Write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI express port must not translate to outbound I/O write transactions on another PCI Express port.
20–25 -	This field is reserved. Reserved
26–31 OWS	Outbound window size. Outbound translation window size N which is the encoded $2^{(N+1)}$ -byte window size. The smallest window size is 4 Kbytes. Note that for the default window (window 0), the outbound window size may be programmed less than the 64-Gbyte maximum. However, accesses that miss all other windows and hit outside the default window is aliased to the default window. Patterns not shown are reserved. 001011 4-Kbyte window size 001100 8-Kbyte window size ... 011111 4-Gbyte window size 100000 8-Gbyte window size 100001 16-Gbyte window size 100010 32-Gbyte window size 100011 64-Gbyte window size

16.6.15 PCI Express outbound window base address register n (PEXx_PEXOWBARn)

The PCI Express outbound window base address registers select the base address for the windows which are translated to the external address space. Addresses for outbound transactions are compared to these windows. If such a transaction does not fall within one of these spaces the transaction is forwarded through a default register set.

Address: Base address + C28h offset + (32d × i), where i=0d to 3d



PEXx_PEXOWBARn field descriptions

Field	Description
0–7 -	This field is reserved. Reserved

Table continues on the next page...

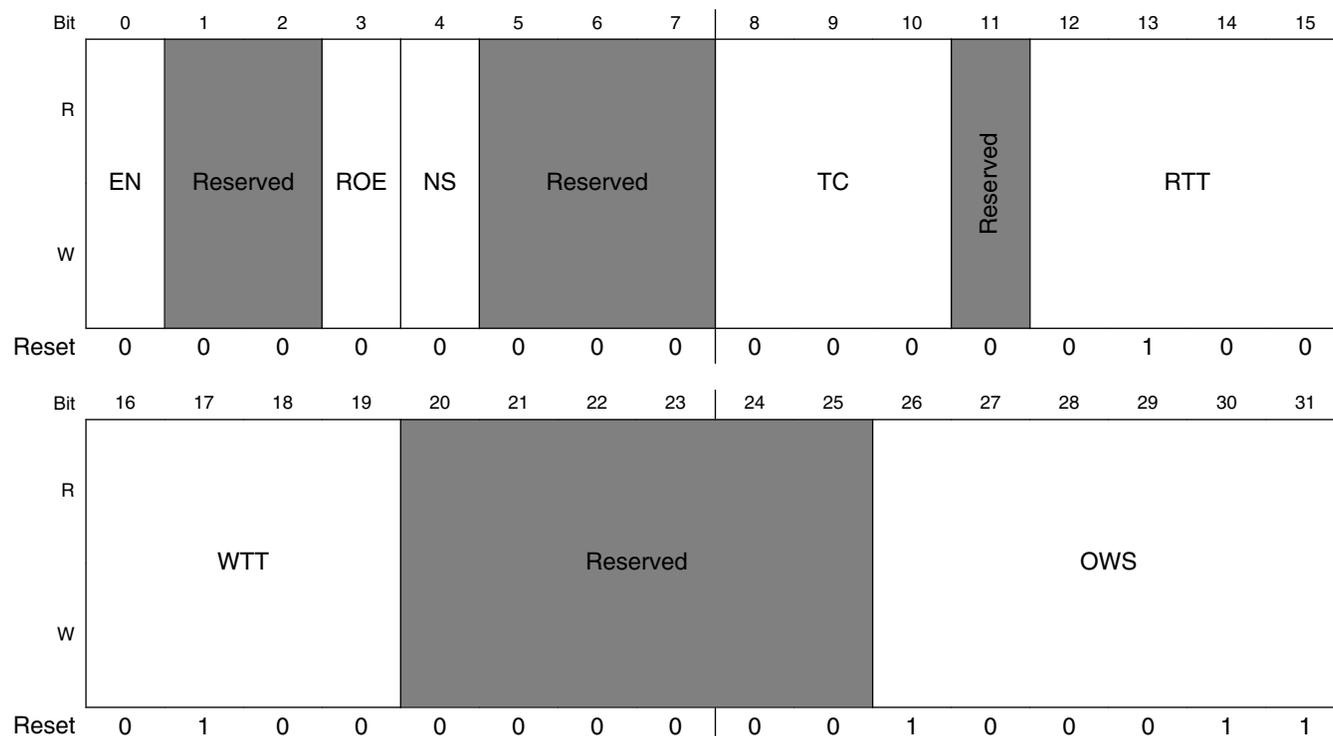
PEXx_PEXOWBARn field descriptions (continued)

Field	Description
8–11 WBEA	Window base extended address. Source address which is the starting point for the outbound translation window. The window must be aligned based on the size selected in the window size bits. Correspond to internal platform address bits [0:3] . (where 0 is the msb of the internal platform address)
12–31 WBA	Window base address. Source address which is the starting point for the outbound translation window. The window must be aligned based on the size selected in the window size bits. This corresponds to internal platform address bits [4:23] .

16.6.16 PCI Express outbound window attributes register n (PEXx_PEXOWARn)

The PCI Express outbound window attributes registers define the window sizes to translate and other attributes for the translations. 64 Gbytes is the largest window size allowed.

Address: Base address + C30h offset + (32d × i), where i=0d to 1d



PEXx_PEXOWARn field descriptions

Field	Description
0 EN	Enable. This bit enables this address translation window.

Table continues on the next page...

PEXx_PEXOWARn field descriptions (continued)

Field	Description
	0 Disable outbound translation window 1 Enable outbound translation window
1–2 -	This field is reserved. Reserved
3 ROE	Relaxed ordering enable. When this bit and the RO bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the relaxed ordering bit for the packet is enabled. This bit only applies to memory transactions. 0 Default ordering 1 Relaxed ordering
4 NS	No snoop enable. When this bit and the NSE bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the no snoop bit for the packet is enabled. This bit only applies to memory transactions. 0 Snoopable 1 No snoop
5–7 -	This field is reserved. Reserved
8–10 TC	Traffic class. This field indicates the traffic class of the outbound packet. This field only applies to memory transaction. All other transaction types should set the TC field to 0. NOTE: Traffic class settings are passed through to the PCI Express link, but no specific actions are taken in the device based on traffic class. 000 TC0 001 TC1 010 TC2 011 TC3 100 TC4 101 TC5 110 TC6 111 TC7
11 -	This field is reserved. Reserved
12–15 RTT	Read transaction type. Read transaction type to run on the PCI Express link. Settings not shown are reserved. 0010 Configuration read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. 0100 Memory read 1000 IO read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary.
16–19 WTT	Write transaction type. Write transaction type to run on the PCI Express link. Settings not shown are reserved. 0010 Configuration write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI express port must not translate to outbound configuration write transactions on another PCI Express port.

Table continues on the next page...

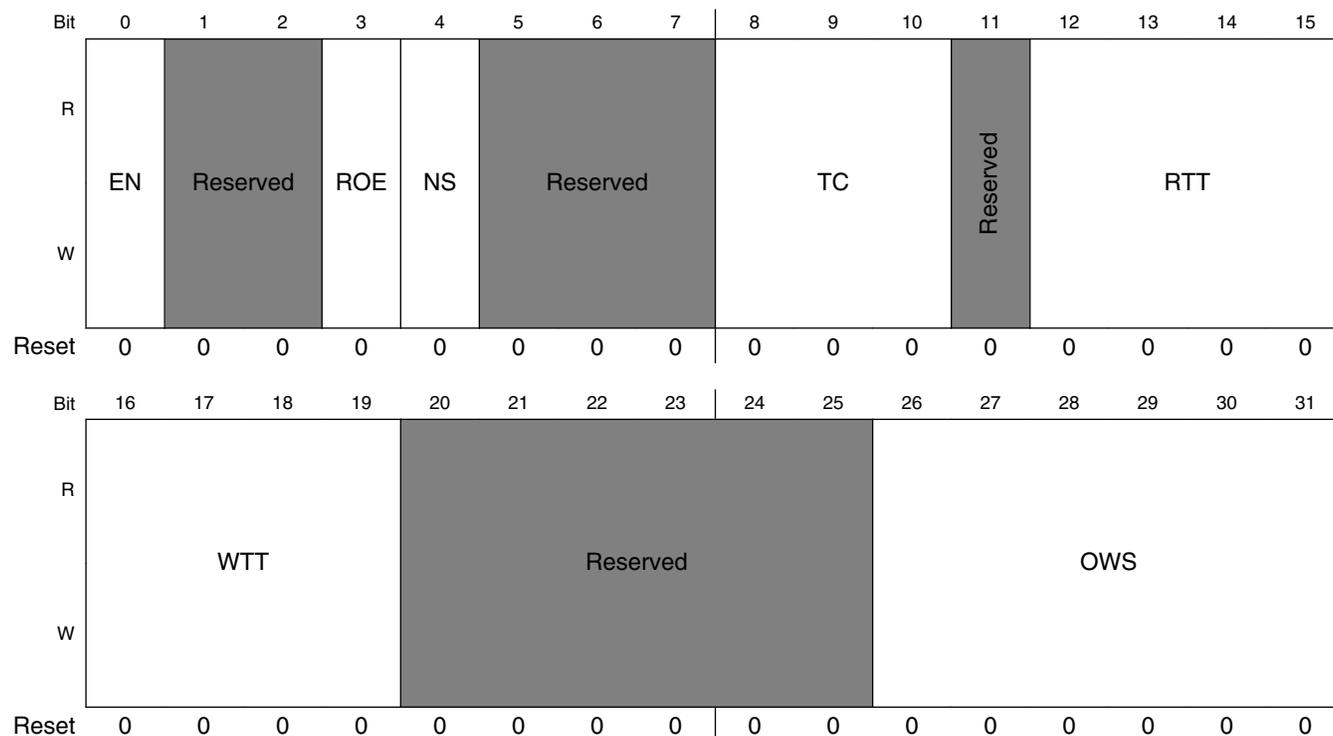
PEXx_PEXOWARn field descriptions (continued)

Field	Description
	0100 Memory write 0101 Message write. Only support 4-byte size access on a 4-byte address boundary. 1000 IO Write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI express port must not translate to outbound I/O write transactions on another PCI Express port.
20–25 -	This field is reserved. Reserved
26–31 OWS	Outbound window size. Outbound translation window size N which is the encoded $2^{(N+1)}$ -byte window size. The smallest window size is 4 Kbytes. Note that for the default window (window 0), the outbound window size may be programmed less than the 64-Gbyte maximum. However, accesses that miss all other windows and hit outside the default window is aliased to the default window. 000000 Reserved 001010 Reserved 001011 4-Kbyte window size 001100 8-Kbyte window size 011111 4-Gbyte window size 100000 8-Gbyte window size 100001 16-Gbyte window size 100010 32-Gbyte window size 100011 64-Gbyte window size 100100 Reserved 111111 Reserved

16.6.17 PCI Express outbound window attributes register 3 (PEXx_PEXOWAR3)

The PCI Express outbound window attributes registers define the window sizes to translate and other attributes for the translations. 64 Gbytes is the largest window size allowed.

Address: Base address + C70h offset



PEXx_PEXOWAR3 field descriptions

Field	Description
0 EN	Enable. This bit enables this address translation window. 0 Disable outbound translation window 1 Enable outbound translation window
1-2 -	This field is reserved. Reserved
3 ROE	Relaxed ordering enable. When this bit and the RO bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the relaxed ordering bit for the packet is enabled. This bit only applies to memory transactions. 0 Default ordering 1 Relaxed ordering

Table continues on the next page...

PEXx_PEXOWAR3 field descriptions (continued)

Field	Description
4 NS	No snoop enable. When this bit and the NSE bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the no snoop bit for the packet is enabled. This bit only applies to memory transactions. 0 Snoopable 1 No snoop
5–7 -	This field is reserved. Reserved
8–10 TC	Traffic class. This field indicates the traffic class of the outbound packet. This field only applies to memory transaction. All other transaction types should set the TC field to 0. NOTE: Traffic class settings are passed through to the PCI Express link, but no specific actions are taken in the device based on traffic class. 000 TC0 001 TC1 010 TC2 011 TC3 100 TC4 101 TC5 110 TC6 111 TC7
11 -	This field is reserved. Reserved
12–15 RTT	Read transaction type. Read transaction type to run on the PCI Express link. Settings not shown are reserved. 0010 Configuration read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. 0100 Memory read 1000 IO read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary.
16–19 WTT	Write transaction type. Write transaction type to run on the PCI Express link. Settings not shown are reserved. 0010 Configuration write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI express port must not translate to outbound configuration write transactions on another PCI Express port. 0100 Memory write 0101 Message write. Only support 4-byte size access on a 4-byte address boundary. 1000 IO Write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI express port must not translate to outbound I/O write transactions on another PCI Express port.
20–25 -	This field is reserved. Reserved
26–31 OWS	Outbound window size. Outbound translation window size N which is the encoded $2^{(N+1)}$ -byte window size. The smallest window size is 4 Kbytes. Note that for the default window (window 0), the outbound window size may be programmed less than the 64-Gbyte maximum. However, accesses that miss all other windows and hit outside the default window is aliased to the default window.

Table continues on the next page...

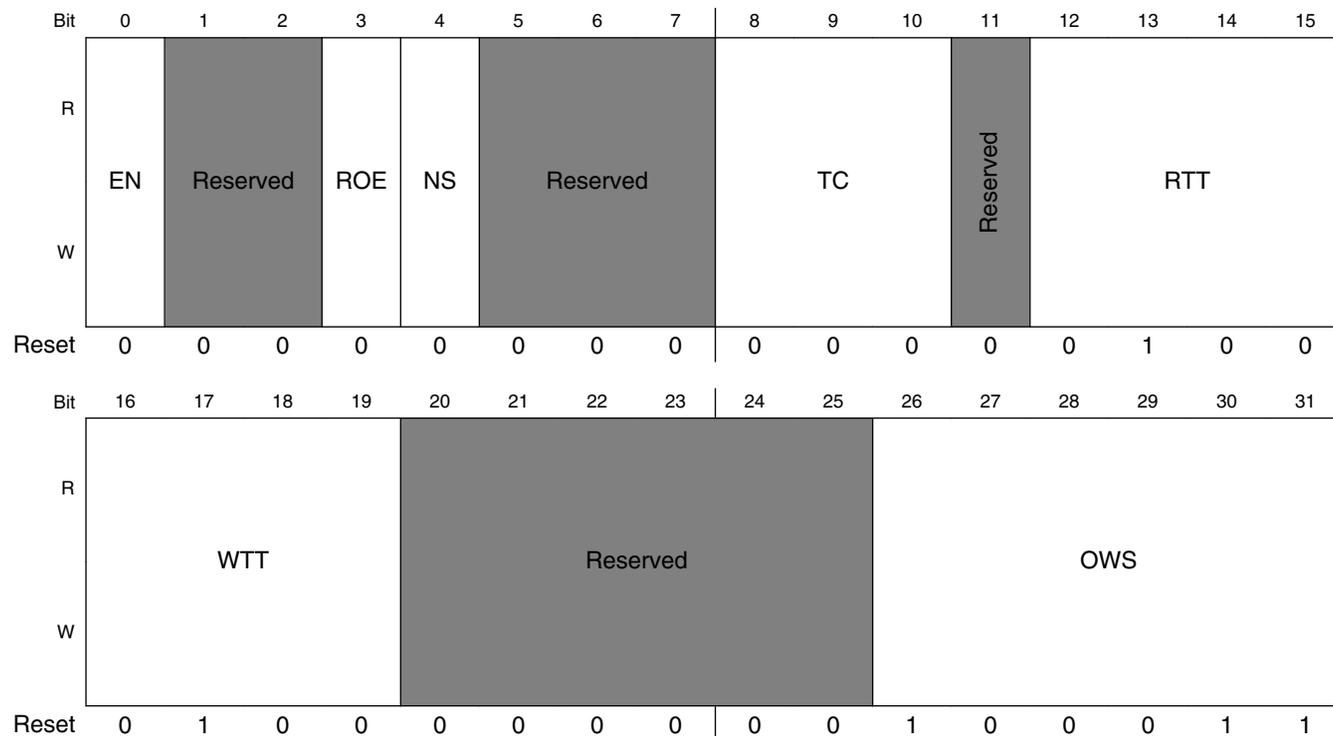
PEX_x_PEXOWAR3 field descriptions (continued)

Field	Description
000000	Reserved
001010	Reserved
001011	4-Kbyte window size
001100	8-Kbyte window size
011111	4-Gbyte window size
100000	8-Gbyte window size
100001	16-Gbyte window size
100010	32-Gbyte window size
100011	64-Gbyte window size
100100	Reserved
111111	Reserved

16.6.18 PCI Express outbound window attributes register 4 (PEX_x_PEXOWAR4)

The PCI Express outbound window attributes registers define the window sizes to translate and other attributes for the translations. 64 Gbytes is the largest window size allowed.

Address: Base address + C90h offset



PEXx_PEXOWAR4 field descriptions

Field	Description
0 EN	Enable. This bit enables this address translation window. 0 Disable outbound translation window 1 Enable outbound translation window
1–2 -	This field is reserved. Reserved
3 ROE	Relaxed ordering enable. When this bit and the RO bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the relaxed ordering bit for the packet is enabled. This bit only applies to memory transactions. 0 Default ordering 1 Relaxed ordering
4 NS	No snoop enable. When this bit and the NSE bit of the PCI Express device control register (described in PCI Express Device Control Register (Device_Control_Register)) is set, the no snoop bit for the packet is enabled. This bit only applies to memory transactions. 0 Snoopable 1 No snoop
5–7 -	This field is reserved. Reserved
8–10 TC	Traffic class. This field indicates the traffic class of the outbound packet. This field only applies to memory transaction. All other transaction types should set the TC field to 0. NOTE: Traffic class settings are passed through to the PCI Express link, but no specific actions are taken in the device based on traffic class. 000 TC0 001 TC1 010 TC2 011 TC3 100 TC4 101 TC5 110 TC6 111 TC7
11 -	This field is reserved. Reserved
12–15 RTT	Read transaction type. Read transaction type to run on the PCI Express link. Settings not shown are reserved. 0010 Configuration read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. 0100 Memory read 1000 IO read. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary.
16–19 WTT	Write transaction type. Write transaction type to run on the PCI Express link. Settings not shown are reserved. 0010 Configuration write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI

Table continues on the next page...

PEXx_PEXOWAR4 field descriptions (continued)

Field	Description
	express port must not translate to outbound configuration write transactions on another PCI Express port. 0100 Memory write 0101 Message write. Only support 4-byte size access on a 4-byte address boundary. 1000 IO Write. Supported only when in RC mode and size of less than or equal to 4 bytes and not crossing 4-byte address boundary. Note that inbound write transactions on one PCI express port must not translate to outbound I/O write transactions on another PCI Express port.
20–25 -	This field is reserved. Reserved
26–31 OWS	Outbound window size. Outbound translation window size N which is the encoded $2^{(N+1)}$ -byte window size. The smallest window size is 4 Kbytes. Note that for the default window (window 0), the outbound window size may be programmed less than the 64-Gbyte maximum. However, accesses that miss all other windows and hit outside the default window is aliased to the default window. 000000 Reserved 001010 Reserved 001011 4-Kbyte window size 001100 8-Kbyte window size 011111 4-Gbyte window size 100000 8-Gbyte window size 100001 16-Gbyte window size 100010 32-Gbyte window size 100011 64-Gbyte window size 100100 Reserved 111111 Reserved

16.6.19 PCI Express inbound translation address register n (PEXx_PEXITARn)

The PCI Express inbound translation address registers contain the translated internal platform address to be used. Note that PEXITAR0 does not exist in the memory-mapped space; it is a fixed translation to the internal configuration (CCSR) space.

Address: Base address + DA0h offset + (32d × i), where i=0d to 2d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								TEA								TA															
W	Reserved								TEA								TA															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PEXx_PEXITARn field descriptions

Field	Description
0–7 -	This field is reserved. Reserved

Table continues on the next page...

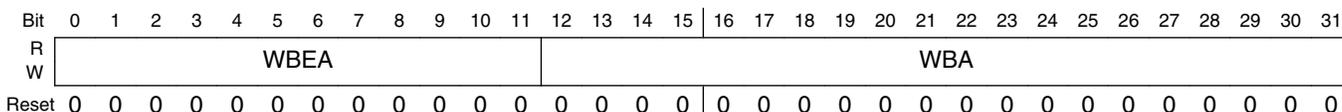
PEXx_PEXITAR_n field descriptions (continued)

Field	Description
8–11 TEA	Translation extended address. Target address which indicates the starting point of the inbound translated address. The translation address must be aligned based on the size field. Corresponds to internal platform address bits [0:3] where bit 0 is the msb of the internal platform address.
12–31 TA	Translation address. Target address which indicates the starting point of the inbound translated address. The translation address must be aligned based on the size field. This corresponds to internal platform address bits [4:23].

16.6.20 PCI Express inbound window base address register n (PEXx_PEXIWBAR_n)

The PCI Express inbound window base address registers select the base address for the windows which are translated to an alternate target address space. In root complex (RC) mode, addresses for inbound transactions are compared to these windows. In RC mode, the BAR for inbound window 0 is located in the PCI Express type 1 configuration header space and the BARs for inbound windows 1-3 (PEXIWBAR[1-3]) are implemented as described in this section. In endpoint (EP) mode, these registers are not implemented in the memory-mapped space. Reading these registers in EP mode returns all zeros and writing to these offsets has no consequences. All base address registers in EP are located in the PCI Express type 0 configuration header space. Note that PEXIWBAR1 only supports 32-bit PCI Express address space.

Address: Base address + DA8h offset + (32d × i), where i=0d to 2d



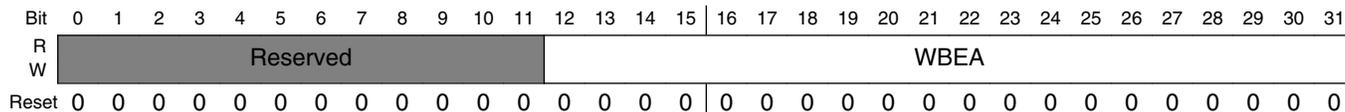
PEXx_PEXIWBAR_n field descriptions

Field	Description
0–11 WBEA	Window base extended address. This field corresponds to PCI Express address bits [43:32]. Note that the extended address is supported for windows 2 and 3 only; for PEXIWBAR1, these bits are reserved and must be 0.
12–31 WBA	Window base address. Source address which is the starting point for the inbound translation window. The window must be aligned based on the size selected in the window size bits. This corresponds to PCI Express address bits [31:12].

16.6.21 PCI Express inbound window base extended address register n (PEXx_PEXIWBEARn)

The PCI Express inbound window base extended address registers contain the most-significant bits of a 64 bit base address.

Address: Base address + DACH offset + (32d × i), where i=0d to 1d



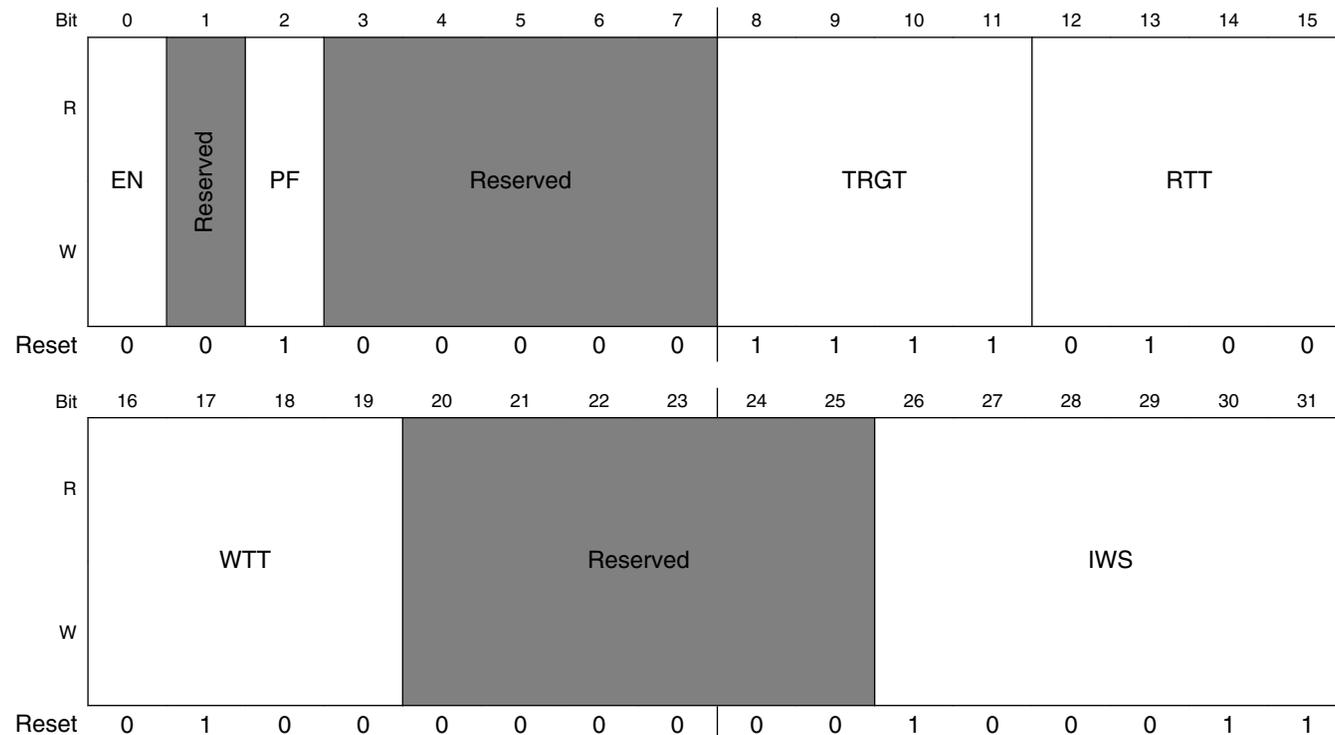
PEXx_PEXIWBEARn field descriptions

Field	Description
0–11 -	This field is reserved. Reserved
12–31 WBEA	Window base extended address. This field corresponds to PCI Express address bits [63:44]

16.6.22 PCI Express inbound window attributes register n (PEXx_PEXIWARn)

The PCI Express inbound window attributes registers define the window sizes to translate and other attributes for the translations. 64 Gbytes is the largest window size allowed.

Address: Base address + DB0h offset + (32d × i), where i=0d to 2d



PEXx_PEXIWARn field descriptions

Field	Description
0 EN	Enable. This bit controls the enabling/disabling of the translation window. 0 Disable inbound window translation 1 Enable inbound window translation
1 -	This field is reserved. Reserved
2 PF	Prefetchable. This bit indicates that the address space is prefetchable. This bit corresponds to the prefetchable bit in the BAR in the PCI Express type 0 header. This bit drives the BAR's prefetchable bit in EP mode. 0 Not prefetchable 1 Prefetchable
3–7 -	This field is reserved. Reserved

Table continues on the next page...

PEXx_PEXIWARn field descriptions (continued)

Field	Description
<p>8–11 TRGT</p>	<p>Target interface. If this field is set to anything other than local memory space, the attributes for the transaction must be assigned in a corresponding outbound window at the target.</p> <p>0000 PCI Express 3-PCI Express controller 3 should not use this encoding 0001 PCI Express 2-PCI Express controller 2 should not use this encoding 0010 PCI Express 1-PCI Express controller 1 should not use this encoding 0011-1011 Reserved 1100 Serial Rapid IO1 1101 Serial Rapid IO2 1110 Reserved 1111 Local memory space</p>
<p>12–15 RTT</p>	<p>Read transaction type. Read transaction type to send to target interface.</p> <p>If the transaction is not going to local memory space</p> <p>0000 Reserved ... 0100 Read 0101 Reserved 0110 Reserved 0111 Reserved 1000 Reserved ... 1111 Reserved</p> <p>If the transaction is going to local memory space</p> <p>0000 Reserved ... 0100 Read, do not snoop local processor 0101 Read, snoop local processor 0110 Reserved 0111 Read, snoop local processor, unlock L2 cache line 1000 Reserved ... 1111 Reserved</p>
<p>16–19 WTT</p>	<p>Write transaction type. Write transaction type to send to target interface.</p> <p>If the transaction is not going to local memory space</p> <p>0000 Reserved ... 0100 Write 0101 Reserved 0110 Reserved 0111 Reserved</p>

Table continues on the next page...

PEXx_PEXIWAR_n field descriptions (continued)

Field	Description
	<p>1000 Reserved</p> <p>...</p> <p>1111 Reserved</p> <p>If the transaction is going to local memory space</p> <p>0000 Reserved</p> <p>...</p> <p>0100 Write, do not snoop local processor</p> <p>0101 Write, snoop local processor</p> <p>0110 Write, snoop local processor, allocate L2 cache line</p> <p>0111 Write, snoop local processor, allocate and lock L2 cache line</p> <p>1000 Reserved</p> <p>...</p> <p>1111 Reserved</p>
<p>20–25 -</p>	<p>This field is reserved. Reserved</p>
<p>26–31 IWS</p>	<p>Inbound window size. Inbound translation window size N which is the encoded $2^{(N+1)}$-bytes window size. The smallest window size is 4 Kbytes. For EP mode, this field directly controls the size of the BARs.</p> <p>000000 Reserved</p> <p>...</p> <p>001010 Reserved</p> <p>001011 4-Kbyte window size</p> <p>001100 8-Kbyte window size</p> <p>...</p> <p>011111 4-Gbyte window size</p> <p>100000 8-Gbyte window size</p> <p>100001 16-Gbyte window size</p> <p>100010 32-Gbyte window size</p> <p>100011 64-Gbyte window size</p> <p>100100 Reserved</p> <p>...</p> <p>111111 Reserved</p>

16.6.23 PCI Express error detect register (PEXx_PEX_ERR_DR)

The PCI Express error detect register contains error status bits that are detected by hardware. The detected error bits are write-1-to-clear type registers. Reading from these registers occurs normally; however, write operations can clear but not set bits. A bit is cleared whenever the register is written, and the data in the corresponding bit location is a 1. For example, to clear bit 6 and not affect any other bits in the register, the value 0b0200_0000 is written to the register. When an error is detected the appropriate error bit is set. Subsequent errors sets the appropriate error bits in the error detection registers, but only the first error for a particular unit have any relevant information captured. The interrupt enable bits are used to allow or block the error reporting to the interrupt mechanism while the disable bits are used to prevent or allow the setting of the status bits.

Address: Base address + E00h offset

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	ME	Reserved							PCT	Reserved	PCAC	PNM	CDNSC	CRSNC	ICCA	IACA
W	w1c	Reserved							w1c	Reserved	w1c	w1c	w1c	w1c	w1c	w1c
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PCI Express memory-mapped registers

Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	CRST	MIS	IOIS	CIS	CIEP	IOIEP	OAC	IOIA	Reserved							
W	w1c	w1c	w1c	w1c	w1c	w1c	w1c	w1c								
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PEXx_PEX_ERR_DR field descriptions

Field	Description
0 ME	<p>Multiple errors. Detecting multiple errors of the same type. An error is considered as multiple error when its detect bit is set and the same error is occurring again. Note that this bit does not track the ordering of when the error occurs.</p> <p>1 Multiple errors were detected. 0 Multiple errors were not detected.</p>
1-7 -	This field is reserved. Reserved
8 PCT	<p>PCI Express completion time-out. A completion time-out condition was detected for a non-posted, outbound PCI Express transaction. An error response is sent back to the requestor. Note that a completion timeout counter only starts when the non-posted request was able to send to the link partner.</p> <p>1 A completion time-out on the PCI Express link was detected. Note that a completion timeout error is a fatal error. If a completion timeout error is detected, the system has become unstable. Hot reset is recommended to restore stability of the system. 0 No completion time-out on the PCI Express link detected.</p>
9 -	This field is reserved. Reserved
10 PCAC	<p>PCI Express completer abort (CA) completion. A completion with CA status was received.</p> <p>1 A completion with CA status was detected. 0 No completion with CA status detected.</p>
11 PNM	<p>PCI Express no map. An inbound transaction that is not mapped to any inbound windows was detected. In RC mode, a posted transaction will be dropped silently and this bit will be set. A nonposted transaction will return a completion without data (Cpl) packet with a UR completion status to the requester and this bit is set. For EP mode, a Cpl packet with a UR completion status is sent back to the requester but does not set this bit.</p>

Table continues on the next page...

PEXx_PEX_ERR_DR field descriptions (continued)

Field	Description
	1 A no-map transaction was detected in RC mode. 0 No no-map transaction detected.
12 CDNSC	Completion with data not successful. A completion with data packet was received with a non successful status (that is, UR, CA or CRS status). 1 Completion with data non successful packet was detected. 0 No completion with data non successful packet detected.
13 CRSNC	CRS non configuration. A completion was detected for a non configuration cycle and with CRS status. See PCI Express configuration space access for more information. 1 CRS non configuration packet was detected. 0 No CRS non configuration packet detected.
14 ICCA	Invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA configuration access. Access to an illegal configuration space from PEX_CONFIG_ADDR/PEX_CONFIG_DATA was detected. 1 Invalid CONFIG_ADDR/PEX_CONFIG_DATA access detected 0 No invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA access detected
15 IACA	Invalid ATMU configuration access. Access to an illegal configuration space from an ATMU window was detected. 1 Invalid ATMU configuration access was detected 0 No invalid ATMU configuration access detected
16 CRST	CRS thresholded. An outbound configuration transaction was retried and thresholded due to a CRS completion status. An error response is sent back to the requestor. See PCI Express configuration retry timeout register (PEX_PEX_CONF_RTY_TOR) , for more information. 1 A CRS threshold condition was detected for an outbound configuration transaction 0 No CRS threshold condition detected
17 MIS	Message invalid size. An outbound message transaction that is greater than 4 bytes or crosses a 4-byte boundary was detected. See Outbound ATMU message generation , for more information. 1 An invalid size outbound message transaction was detected 0 No invalid size outbound message transaction detected
18 IOIS	I/O invalid size. An outbound I/O transaction that is greater than 4 bytes or crosses a 4-byte boundary was detected. 1 An invalid size outbound I/O transaction was detected 0 No invalid size outbound I/O transaction detected
19 CIS	Configuration invalid size. An outbound configuration transaction that is greater than 4 bytes or crosses a 4-byte boundary was detected. 1 An invalid size outbound configuration transaction was detected 0 No invalid size outbound configuration transaction detected
20 CIEP	Configuration invalid EP. An outbound ATMU configuration transaction request was seen when in EP mode. 1 An outbound configuration transaction while in EP was detected 0 No outbound configuration transaction in EP detected
21 IOIEP	I/O invalid EP. An outbound I/O transaction request was seen when in EP mode.

Table continues on the next page...

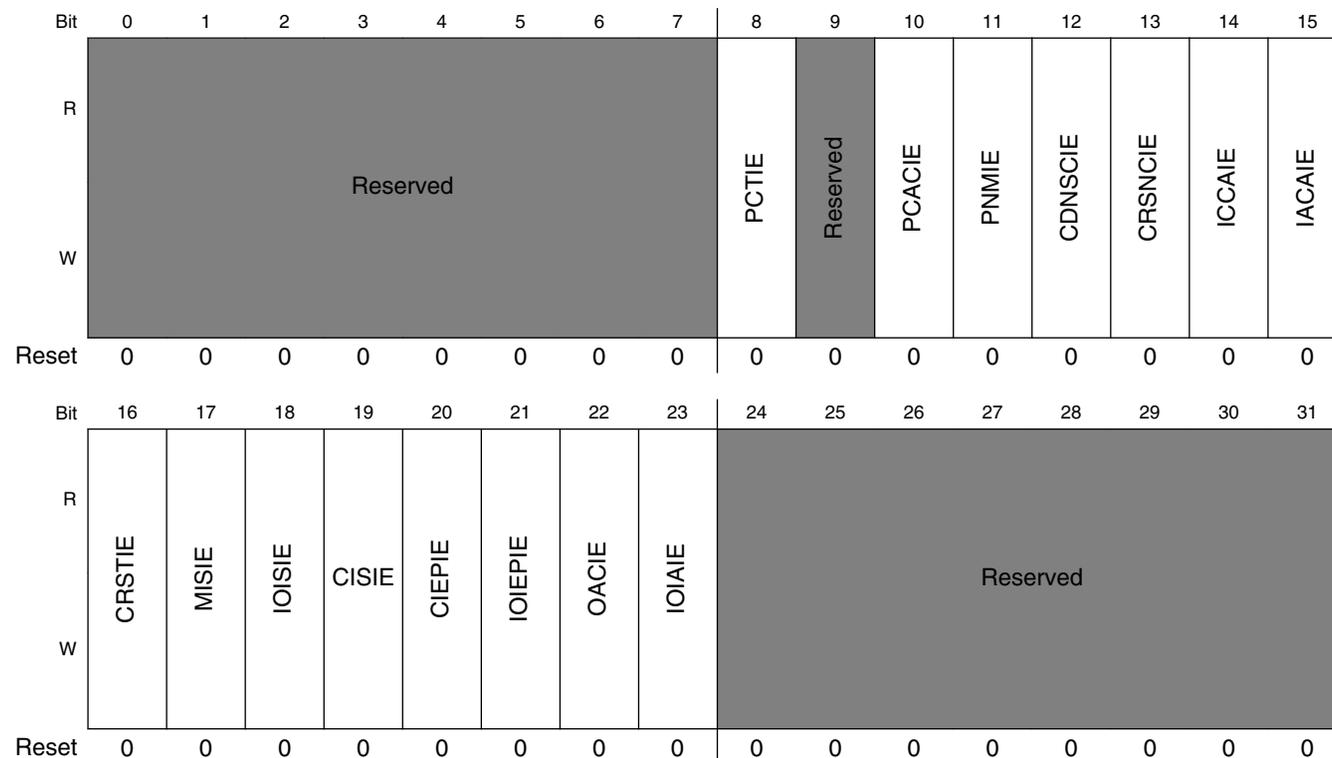
PEXx_PEX_ERR_DR field descriptions (continued)

Field	Description
	1 An outbound I/O transaction while in EP was detected 0 No outbound I/O transaction in EP detected
22 OAC	Outbound ATMU crossing. An outbound crossing ATMU transaction was detected. 1 An outbound transaction that hits in one window and crosses overing it was detected 0 No outbound ATMU crossing condition detected
23 IOIA	I/O invalid address. An outbound I/O transaction with a translated address of greater than 4 Gbytes was detected. 1 A greater than 4-Gbyte I/O address was detected 0 No greater than 4-Gbyte I/O address detected
24–31 -	This field is reserved. Reserved

16.6.24 PCI Express error interrupt enable register (PEXx_PEX_ERR_EN)

The PCI Express error interrupt enable register allows interrupts to be generated when the corresponding PCI Express error detect register bits are set.

Address: Base address + E08h offset



PEXx_PEX_ERR_EN field descriptions

Field	Description
0-7 -	This field is reserved. Reserved
8 PCTIE	PCI Express completion time-out interrupt enable. When set and PEX_ERR_DR[PCT]=1 generates an interrupt. 1 Enable PCI Express completion time-out interrupt generation 0 Disable PCI Express completion time-out interrupt generation
9 -	This field is reserved. Reserved
10 PCACIE	PCI Express CA completion interrupt enable. When set and PEX_ERR_DR[PCAC]=1 generates an interrupt. 1 Enable completion with CA status interrupt generation 0 Disable completion with CA status interrupt generation
11 PNMIE	PCI Express no map interrupt enable. When set and PEX_ERR_DR[PNM]=1 generates an interrupt. 1 Enable no map PCI Express packet interrupt generation 0 Disable no map PCI Express packet interrupt generation
12 CDNSCIE	Completion with data not successful interrupt enable. When this bit is set and PEX_ERR_DR[CDNSC] = 1 generates an interrupt. 1 Enable completion with data non successful interrupt generation 0 Disable completion with data non successful interrupt generation
13 CRSNCIE	CRS non configuration interrupt enable. When this bit is set and PEX_ERR_DR[CRSNC] = 1 generates an interrupt. 1 Enable CRS non configuration interrupt generation 0 Disable CRS non configuration interrupt generation
14 ICCAIE	Invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA configuration access interrupt enable. When set and PEX_ERR_DR[ICCA]=1 generates an interrupt. 1 Enable invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA access interrupt generation 0 Disable invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA access interrupt generation.
15 IACAIE	Invalid ATMU configuration access. When set and PEX_ERR_DR[IACA]=1 generates an interrupt. 1 Enable invalid ATMU configuration access interrupt generation 0 Disable invalid ATMU configuration access interrupt generation
16 CRSTIE	CRS thresholded interrupt enable. When set and PEX_ERR_DR[CRST]=1 generates an interrupt. 1 Enable CRS threshold interrupt generation 0 Disable CRS threshold interrupt generation
17 MISIE	Message invalid size interrupt enable. When set and PEX_ERR_DR[MIS]=1 generates an interrupt. 1 Enable invalid outbound message size interrupt generation 0 Disable invalid outbound message size interrupt generation
18 IOISIE	I/O invalid size interrupt enable. When set and PEX_ERR_DR[IOIS]=1 generates an interrupt. 1 Enable invalid outbound I/O size interrupt generation 0 Disable invalid outbound I/O size interrupt generation

Table continues on the next page...

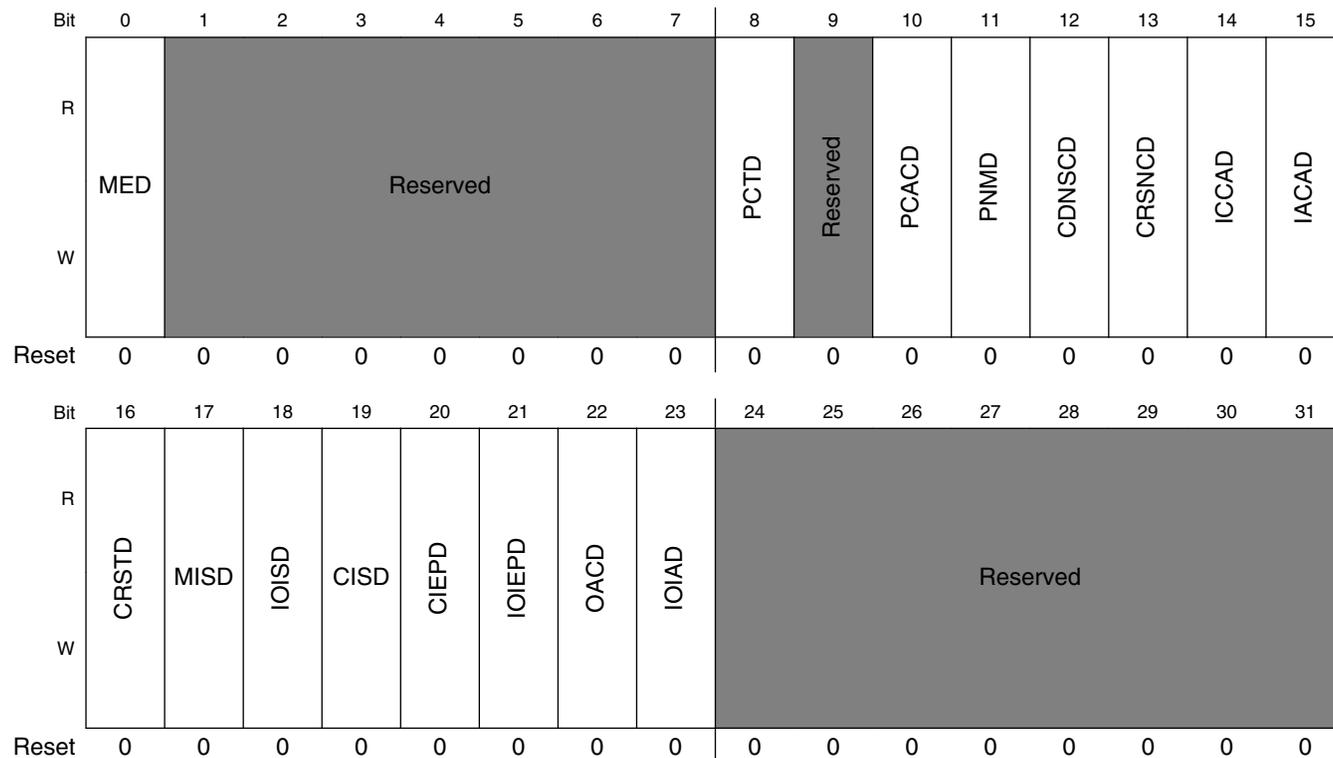
PEXx_PEX_ERR_EN field descriptions (continued)

Field	Description
19 CISIE	Configuration invalid size interrupt enable. When set and PEX_ERR_DR[CIS]=1 generates an interrupt. 1 Enable invalid outbound configuration size interrupt generation 0 Disable invalid outbound configuration size interrupt generation
20 CIEPIE	Configuration invalid EP interrupt enable. When set and PEX_ERR_DR[CIEP]=1 generates an interrupt. 1 Enable outbound configuration transaction while in EP mode interrupt generation 0 Disable outbound configuration transaction in EP mode interrupt generation
21 IOIEPIE	I/O invalid EP interrupt enable. When set and PEX_ERR_DR[IOIEP]=1 generates an interrupt. 1 Enable outbound I/O transaction EP mode interrupt generation 0 Disable outbound I/O transaction EP mode interrupt generation
22 OACIE	Outbound ATMU crossing interrupt enable. When set and PEX_ERR_DR[OAC]=1 generates an interrupt. 1 Enable outbound crossing ATMU interrupt generation 0 Disable outbound crossing ATMU interrupt generation
23 IOIAIE	I/O address invalid enable. When set and PEX_ERR_DR[IOIA]=1 generates an interrupt. 1 Enable greater than 4G I/O address interrupt generation 0 Disable greater than 4G I/O address interrupt generation
24–31 -	This field is reserved. Reserved

16.6.25 PCI Express error disable register (PEX_x_PEX_ERR_DISR)

The PCI Express error disable register controls the setting of the PCI Express error detect register's bits.

Address: Base address + E10h offset



PEX_x_PEX_ERR_DISR field descriptions

Field	Description
0 MED	Multiple errors disable. When set disables the setting of PEX_ERR_DR[ME] bit. 1 Disable multiple errors detection 0 Enable multiple errors detection
1-7 -	This field is reserved. Reserved
8 PCTD	PCI Express completion time-out disable. When set disables the setting of PEX_ERR_DR[PCT] bit. 1 Disable PCI Express completion time-out detection 0 Enable PCI Express completion time-out detection
9 -	This field is reserved. Reserved

Table continues on the next page...

PEXx_PEX_ERR_DISR field descriptions (continued)

Field	Description
10 PCACD	PCI Express CA completion disable. When set disables the setting of PEX_ERR_DR[PCAC] bit. 1 Disable completion with CA status detection 0 Enable completion with CA status detection
11 PNMD	PCI Express no map disable. When set disables the setting of PEX_ERR_DR[PNM] bit. 1 Disable no map PCI Express packet detection 0 Enable no map PCI Express packet detection
12 CDNSCD	Completion with data not successful disable. When set disables the setting of PEX_ERR_DR[CDNSC] bit. 1 Disable completion with data not successful detection 0 Enable completion with data not successful detection
13 CRSNCD	CRS non configuration disable. When set disables the setting of PEX_ERR_DR[CRSNC] bit. 1 Disable CRS non configuration detection 0 Enable CRS non configuration detection
14 ICCAD	Invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA configuration access disable. When set disables the setting of PEX_ERR_DR[ICCA] bit. 1 Disable invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA access detection 0 Enable invalid PEX_CONFIG_ADDR/PEX_CONFIG_DATA access detection
15 IACAD	Invalid ATMU configuration access. When set disables the setting of PEX_ERR_DR[IACA] bit. 1 Disable invalid ATMU configuration access detection 0 Enable invalid ATMU configuration access detection
16 CRSTD	CRS thresholded disable. When set disables the setting of PEX_ERR_DR[CRST] bit. 1 Disable CRS threshold detection 0 Enable CRS threshold detection
17 MISD	Message invalid size disable. When set disables the setting of PEX_ERR_DR[MIS] bit. 1 Disable invalid outbound message size detection 0 Enable invalid outbound message size detection
18 IOISD	I/O invalid size disable. When set disables the setting of PEX_ERR_DR[IOIS] bit. 1 Disable invalid outbound I/O size detection 0 Enable invalid outbound I/O size detection
19 CISD	Configuration invalid size disable. When set disables the setting of PEX_ERR_DR[CIS] bit. 1 Disable invalid outbound configuration size detection 0 Enable invalid outbound configuration size detection
20 CIEPD	Configuration invalid EP disable. When set disables the setting of PEX_ERR_DR[CIEP] bit. 1 Disable outbound configuration transaction EP mode detection 0 Enable outbound configuration transaction EP mode detection
21 IOIEPD	I/O invalid EP disable. When set disables the setting of PEX_ERR_DR[IOEP] bit. 1 Disable outbound I/O transaction EP mode detection 0 Enable outbound I/O transaction EP mode detection

Table continues on the next page...

PEXx_PEX_ERR_DISR field descriptions (continued)

Field	Description
22 OACD	Outbound ATMU crossing disable. When set disables the setting of PEX_ERR_DR[OAC] bit. 1 Disable outbound crossing ATMU detection 0 Enable outbound crossing ATMU detection
23 IOIAD	I/O invalid address disable. When set disables the setting of PEX_ERR_DR[IOIA] bit. 1 Disable greater than 4G I/O address detection 0 Enable greater than 4G I/O address detection
24–31 -	This field is reserved. Reserved

16.6.26 PCI Express error capture status register (PEXx_PEX_ERR_CAP_STAT)

The PCI Express error capture status register allows vital error information to be captured when an error occurs. Note that no further error capturing is performed until the ECV bit is cleared.

Address: Base address + E20h offset

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved																	
W	Reserved																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved									TO	GSID						ECV	
W	Reserved																w1c	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0

PEXx_PEX_ERR_CAP_STAT field descriptions

Field	Description
0–24 -	This field is reserved. Reserved
25 TO	Transaction originator. This field Indicates whether the originator of the transaction is from PEX_CONFIG_ADDR/PEX_CONFIG_DATA. 1 Transaction originated from PEX_CONFIG_ADDR/PEX_CONFIG_DATA. 0 Transaction not originated from PEX_CONFIG_ADDR/PEX_CONFIG_DATA.
26–30 GSID	Global source ID. This field indicates the internal platform global source ID that the error transaction originates. This field only applies to non PEX_CONFIG_ADDR/PEX_CONFIG_DATA transactions.

Table continues on the next page...

PEX_x_PEX_ERR_CAP_STAT field descriptions (continued)

Field	Description
	All settings not shown are reserved.
	00000 PCI Express 3
	00001 PCI Express 2
	00010 PCI Express 1
	00101 USB
	00111 Security
	01010 Boot sequencer
	01011 eSDHC
	01100 Serial RapidIO 1
	01101 Serial RapidIO 2
	10000 Processor 0 (instruction)
	10001 Processor 0 (data)
	10010 Processor 1 (instruction)
	10011 Processor 1 (data)
	10101 DMA 1
	10110 DMA 2
	11000 eTSEC 1
	11001 eTSEC 2
	11010 eTSEC 3
	11100 RapidIO message unit
	11101 RapidIO doorbell unit
	11110 RapidIO port-write unit
31 ECV	Error capture valid. This bit indicates that the capture registers 0-3 contain valid info. This bit when set indicates that the captured registers contain valid capturing information. No new capturing is done unless this bit is cleared by writing a 1 to it.

16.6.27 PCI Express error capture register n (PEX_x_PEX_ERR_CAP_Rn)

PEX_ERR_CAP_Rn allow vital error information to be captured when an error occurs. Different error information is reported depending on whether the error source is from an outbound transaction from an internal source or from an inbound transaction from an external source; the source of the captured error is reflected in PEX_ERR_CAP_STAT[GSID]. Note that after the initial error is captured, no further capturing is performed until the PEX_ERR_CAP_STAT[ECV] bit is clear.

See [Error capture registers](#) for further details.

Address: Base address + E28h offset + (4d × i), where i=0d to 3d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	DATA																															
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

PEX_x_PEX_ERR_CAP_Rn field descriptions

Field	Description
0–31 DATA	Captured Data

16.7 PCI Express configuration-space registers

This table lists the PCI Express configuration-space registers.

Table 16-235. PCI Express configuration-space registers

Address offset (Hex)	Description
000-03F	PCI Compatible configuration header (See PCI compatible configuration headers for more information.)
040-0FF	PCI-compatible device-specific configuration space (See PCI compatible device-specific configuration space for more information.)
100-134	PCI Express Extended Configuration Registers (See PCI Express extended configuration space for more information.)
138-3FF	Reserved
400-6FF	PCI Express Controller Internal CSRs ¹ (See PCI Express controller internal CSRs for more information.)
700-FFF	Reserved

1. Note that the PCI Express Controller Internal CSRs are not accessible by inbound PCI Express configuration transactions. Attempts to access these registers returns all 0s.

16.7.1 PCI compatible configuration headers

The first 64 bytes of the 256-byte, PCI-compatible configuration space consists of a predefined header that every PCI device must support.

Many registers in the predefined header are defined the same for all PCI Express devices.

These registers are shown in this figure.

Type 0 configuration header registers

Reserved				Address Offset (Hex)
Device ID		Vendor ID		00
Status		Command		04
Class Code			Revision ID	08
BIST	Header Type	Latency Timer	Cache Line Size	0C
Header Type (0/1) Specific				10
				14
				18
				1C
				20
				24
				28
				2C
				30
				38
		Interrupt Pin	Interrupt Line	3C

Figure 16-234. PCI Express PCI-compatible configuration header common registers

The remaining registers in the header may have differing layouts depending on the function of the device. There are two header types applicable to PCI Express. Type 0 headers are typically used by endpoints; Type 1 headers are used by root complexes and switches/bridges.

16.8 Type 0 configuration header registers

The figure below shows the type 0 header.

NOTE

The BIST register (0x0F) is optional and reserved on the PCI Express controller.

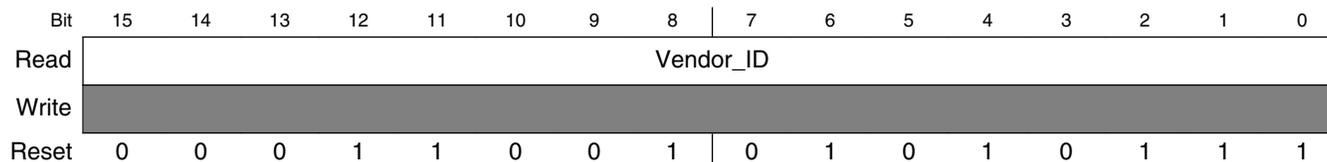
Reserved				Address Offset (Hex)
Device ID		Vendor ID		00
Status		Command		04
Class Code			Revision ID	08
BIST	Header Type	Latency Timer	Cache Line Size	0C
Base Address Registers				10
				14
				18
				1C
				20
				24
				28
Subsystem ID		Subsystem Vendor		2C
Expansion ROM Base Address				30
			Capabilities Pointer	34
				38
MAX_LAT	MIN_GNT	Interrupt Pin	Interrupt Line	3C

Figure 16-235. PCI Express PCI-Compatible Configuration Header—Type 0

16.8.1 PCI Express Vendor ID Register (Vendor_ID_Register)

The vendor ID register is used to identify the manufacturer of the device.

Address: 0h



Vendor_ID_Register field descriptions

Field	Description
15–0 Vendor_ID	0x1957 (Freescale)

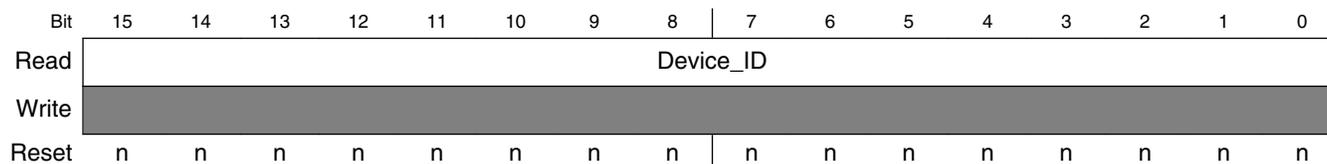
16.8.2 PCI Express Device ID Register (Device_ID_Register)

The device ID register is used to identify the device.

NOTE

See bitfield description table for reset value.

Address: 2h



Device_ID_Register field descriptions

Field	Description
15–0 Device_ID	0x0070 P2020 with security 0x0071 P2020 without security 0x0078 P2010 with security 0x0079 P2010 without security

16.8.3 PCI Express Command Register (Command_Register)

The command register provides control over the ability to generate and respond to PCI Express cycles.

Address: 4h

Bit	15	14	13	12	11	10	9	8
Read	Reserved					Interrupt_Disable	Reserved	SERR
Write	Reserved					Interrupt_Disable	Reserved	SERR
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved	Parity_error_response	Reserved			Bus_master	Memory_space	I_O_space
Write	Reserved	Parity_error_response	Reserved			Bus_master	Memory_space	I_O_space
Reset	0	0	0	0	0	0	0	0

Command_Register field descriptions

Field	Description
15–11 -	This field is reserved. Reserved
10 Interrupt_Disable	Controls the ability to generate INTx interrupt messages. Any INTx emulation interrupts already asserted by this device must be deasserted when this bit is set. 0 Enables INTx interrupt messages 1 Disables INTx interrupt messages
9 -	This field is reserved. Reserved
8 SERR	Controls the reporting of fatal and non-fatal errors detected by the device to the root complex. NOTE: The error control and status bits in the command and status registers control PCI-compatible error reporting. PCI Express advanced error reporting is controlled by the PCI Express device control register described in PCI Express Device Control Register (Device_Control_Register) , and the advance error reporting capability structure described in sections PCI Express Advanced Error Reporting Capability ID Register (Advanced_Error_Reporting_Capability_ID_Register) through PCI Express Error Source ID Register (Error_Source_ID_Register) . 0 Disables reporting 1 Enables reporting
7 -	This field is reserved. Reserved
6 Parity_error_response	Controls whether this PCI Express controller responds to parity errors. NOTE: The error control and status bits in the command and status registers control PCI-compatible error reporting. PCI Express advanced error reporting is controlled by the PCI Express device control register described in PCI Express Device Control Register (Device_Control_Register) , and the advance error reporting capability structure described in sections PCI Express Advanced

Table continues on the next page...

Command_Register field descriptions (continued)

Field	Description
	<p>Error Reporting Capability ID Register (Advanced_Error_Reporting_Capability_ID_Register) through PCI Express Error Source ID Register (Error_Source_ID_Register) .</p> <p>0 Parity errors are ignored and normal operation continues. 1 Parity errors cause the appropriate bit in the PCI Express status register to be set. However, note that errors are reported based on the values set in the PCI Express error enable and detection registers.</p>
5-3 -	This field is reserved. Reserved
2 Bus_master	<p>Indicates whether this PCI Express device is configured as a master.</p> <p>EP mode: Clearing this bit prevents the device from issuing any memory or I/O transactions. Because MSI interrupts are effectively memory writes, clearing this bit also disables the ability of the device to issue MSI interrupts.</p> <p>RC mode: Clearing this bit disables the ability of the device to forward memory transactions upstream. This causes any inbound memory transaction (except those targeting PEXCSRBAR) to be treated as an unsupported request. Note that this description does not apply for inbound memory transactions targeting PEXCSRBAR which are forwarded if the Memory_space bit is set.</p> <p>0 Disables the ability to generate PCI Express accesses 1 Enables this PCI Express controller to behave as a PCI Express bus master</p>
1 Memory_space	<p>Controls whether this PCI Express device (as a target) responds to memory accesses.</p> <p>EP mode: Clearing this bit prevents the device from accepting any memory transaction.</p> <p>RC mode: This bit is ignored for inbound memory transactions except those targeting PEXCSRBAR. Clearing this bit prevents the chip from accepting inbound memory transactions to PEXCSRBAR. This bit does not affect outbound memory transactions.</p> <p>0 This PCI Express device does not respond to PCI Express memory space accesses. 1 This PCI Express device responds to PCI Express memory space accesses.</p>
0 I_O_space	<p>I/O space.</p> <p>EP mode: Clearing this bit prevents the device from accepting any IO transaction. Note that this bit is a don't care in EP mode since the device does not support IO transaction.</p> <p>RC mode: This bit is ignored. It does not affect outbound IO transaction.</p> <p>0 This PCI Express device (as a target) does not respond to PCI Express I/O space accesses. 1 This PCI Express device (as a target) does respond to PCI Express I/O space accesses.</p>

16.8.4 PCI Express Status Register (Status_Register)

The status register is used to record status information for PCI Express related events.

Address: 6h

Bit	15	14	13	12	11	10	9	8
Read	Detected_parity_error	Signaled_system_error	Received_master_abort	Received_target_abort	Signaled_target_abort	Reserved		Master_data_parity_error_detected
Write	w1c	w1c	w1c	w1c	w1c			w1c
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved			Capabilities_List	Interrupt_Status	Reserved		
Write	Reserved					Reserved		
Reset	0	0	0	1	0	0	0	0

Status_Register field descriptions

Field	Description
15 Detected_parity_error	Set whenever a device receives a poisoned TLP regardless of the state of bit 6 in the command register. ³
14 Signaled_system_error	Set whenever a device sends a ERR_FATAL or ERR_NONFATAL message and the SERR enable bit in the command register is set. ²
13 Received_master_abort	Set whenever a requestor receives a completion with unsupported request completion status. ²
12 Received_target_abort	Set whenever a device receives a completion with completer abort completion status. ²
11 Signaled_target_abort	Set whenever a device completes a request using completer abort completion status. ²
10–9 -	This field is reserved. Reserved
8 Master_data_parity_error_detected	Set by the requestor (primary side for Type1 headers) when either the requestor receives a completion marked poisoned or the requestor poisons a write request. Note that the parity error enable bit (bit 6) in the command register must be set for this bit to be set. ²

Table continues on the next page...

Status_Register field descriptions (continued)

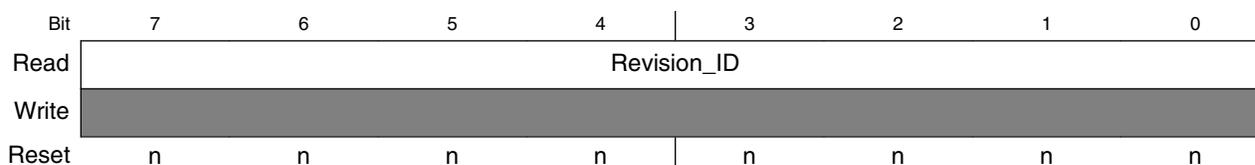
Field	Description
7-5 -	This field is reserved. Reserved
4 Capabilities_List	All PCI Express devices are required to implement the PCI Express capability structure.
3 Interrupt_Status	Set when an INTx interrupt message is pending internally to the device. Note that this bit is associated with INTx messages and not Message Signaled Interrupts.
2-0 -	This field is reserved. Reserved

- The error control and status bits in the command and status registers control PCI-compatible error reporting. PCI Express advanced error reporting is controlled by the PCI Express device control register described in [PCI Express Device Control Register \(Device_Control_Register\)](#) , and the advance error reporting capability structure described in sections [PCI Express Advanced Error Reporting Capability ID Register \(Advanced_Error_Reporting_Capability_ID_Register\)](#) through [PCI Express Error Source ID Register \(Error_Source_ID_Register\)](#) .

16.8.5 PCI Express Revision ID Register (Revision_ID_Register)

The revision ID register is used to identify the revision of the device.

Address: 8h



Revision_ID_Register field descriptions

Field	Description
7-0 Revision_ID	Revision specific.

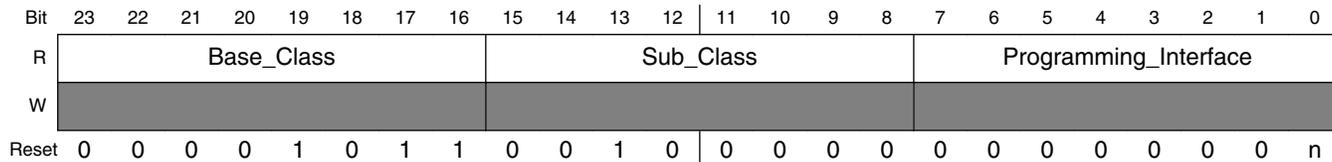
16.8.6 PCI Express Class Code Register (Class_Code_Register)

The class code register is comprised of three single-byte fields-base class (offset 0x0B), sub-class (offset 0x0A), and programming interface (offset 0x09)-that indicate the basic functionality of the function.

NOTE

See bitfield description table for Programming_Interface reset value.

Address: 9h



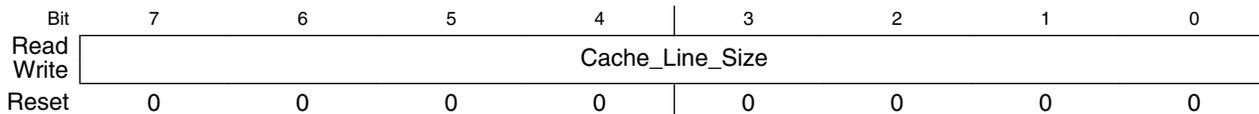
Class_Code_Register field descriptions

Field	Description
23–16 Base_Class	0x0B Processor
15–8 Sub_Class	0x20 PowerPC
7–0 Programming_Interface	0x00 RC mode 0x01 EP mode

16.8.7 PCI Express Cache Line Size Register (Cache_Line_Size_Register)

The cache line size register is provided for legacy compatibility purposes (PCI 2.3); it is not used for PCI Express device functionality.

Address: Ch



Cache_Line_Size_Register field descriptions

Field	Description
7–0 Cache_Line_Size	Represents the cache line size of the processor in terms of 32-bit words (8 32-bit words = 32 bytes). Note that for PCI Express operation this register is ignored.

16.8.8 PCI Express Latency Timer Register (Latency_Timer_Register)

The latency timer register is provided for legacy compatibility purposes (PCI 2.3); it is not used for PCI Express device functionality.

Address: Dh

Bit	7	6	5	4	3	2	1	0
Read	Latency_Timer							
Write								
Reset	0	0	0	0	0	0	0	0

Latency_Timer_Register field descriptions

Field	Description
7-0 Latency_Timer	Note that for PCI Express operation this register is ignored.

16.8.9 PCI Express Header Type Register (Header_Type_Register)

The PCI Express header type register is used to identify the layout of the PCI compatible header.

NOTE

See bitfield description table for reset value.

Address: Eh

Bit	7	6	5	4	3	2	1	0
Read	Multifunction	Header_Layout						
Write								
Reset	0	0	0	0	0	0	0	n

Header_Type_Register field descriptions

Field	Description
7 Multifunction	Identifies whether a device supports multiple functions 0 Single function device 1 Multiple function device

Table continues on the next page...

Header_Type_Register field descriptions (continued)

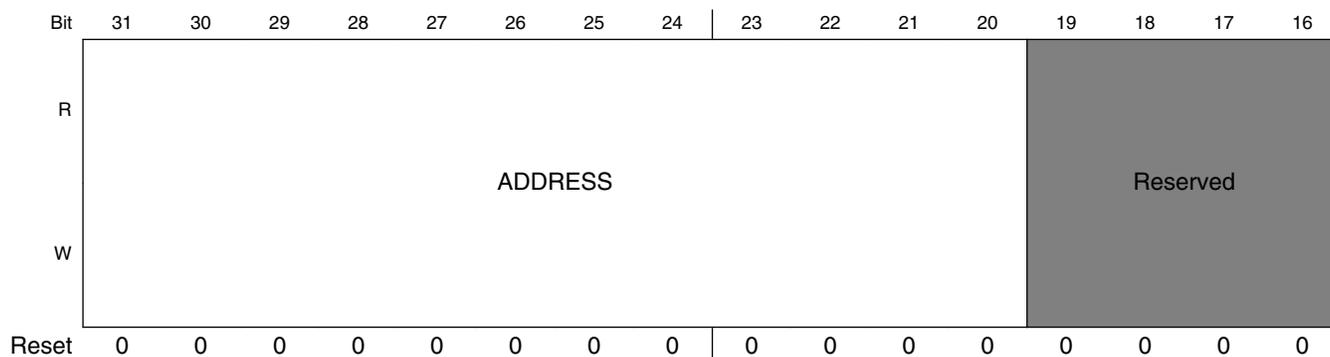
Field	Description
6-0 Header_Layout	All other encodings reserved. 0x00 Endpoint. See Type 0 configuration header registers for type 0 layout. 0x01 Root Complex. See Type 1 configuration header registers for type 1 layout.

16.8.10 PCI Express Base Address Register 0 (PEXCSRBAR)

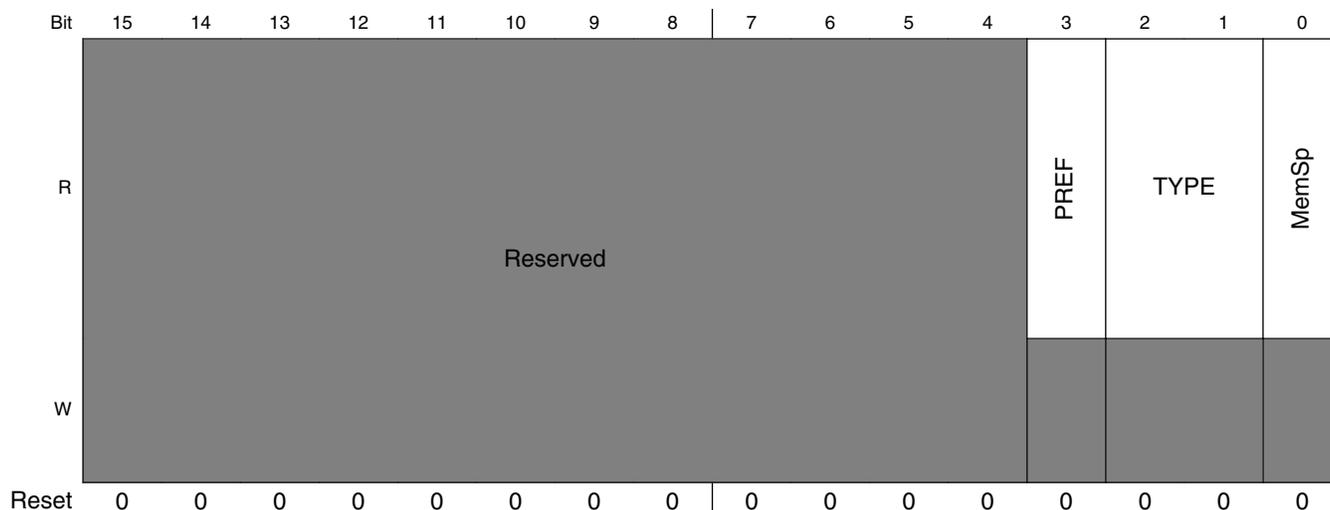
The PCI Express base address registers (BARs) point to the beginning of distinct address ranges which the device should claim. In EP mode, the device supports a configuration space BAR, a 32-bit memory space BAR, and two 64-bit memory space BARs. In RC mode, the device only supports the configuration space BAR in the header; the other memory spaces are defined by the inbound ATMUs. Refer to [PCI Express inbound ATMUs](#) for more information.

Base address register 0, also known as the PCI Express configuration and status register base address register (PEXCSRBAR), at offset 0x10 is a special fixed 1 -Mbyte window that is used for inbound configuration accesses. Note that PEXCSRBAR cannot be updated through the inbound ATMU registers.

Address: 10h



type 0 configuration header registers



PEXCSRBAR field descriptions

Field	Description
31–20 ADDRESS	Indicates the base address that the inbound configuration window occupies. This window is fixed at 1 Mbyte.
19–4 -	This field is reserved. Reserved
3 PREF	Prefetchable
2–1 TYPE	Type. 00 Locate anywhere in 32-bit address space.
0 MemSp	Memory space indicator

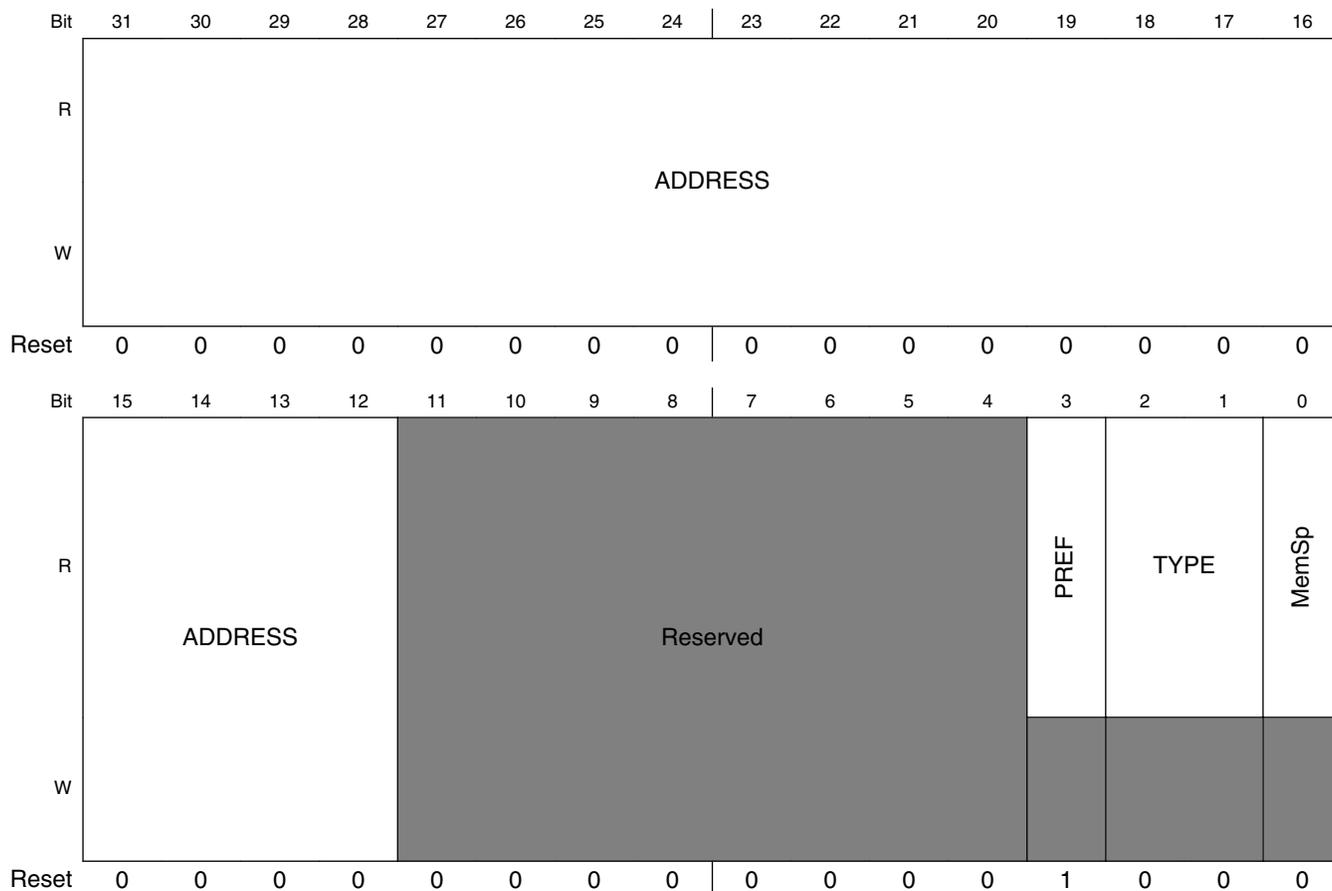
16.8.11 PCI Express Base Address Register 1 (BAR1)

This register is supported only for EP mode.

The PCI Express base address registers (BARs) point to the beginning of distinct address ranges which the device should claim. In EP mode, the device supports a configuration space BAR, a 32-bit memory space BAR, and two 64-bit memory space BARs. Refer to [PCI Express inbound ATMUs](#) for more information.

Base address register 1 at offset 0x14 is used to define the inbound memory window in the 32-bit memory space.

Address: 14h



BAR1 field descriptions

Field	Description
31–12 ADDRESS	Indicates the base address where the inbound memory window begins. The number of upper bits that the device allows to be writable is selected through the inbound window size in the inbound window attributes register (PEXIWAR1).
11–4 -	This field is reserved. Reserved. The device allows a 4 Kbyte window minimum.
3 PREF	Prefetchable. This bit is determined by PEXIWAR1[PF].
2–1 TYPE	Type. 00 Locate anywhere in 32-bit address space.
0 MemSp	Memory space indicator.

16.8.12 PCI Express Base Address Register 2,4 (BAR_n)

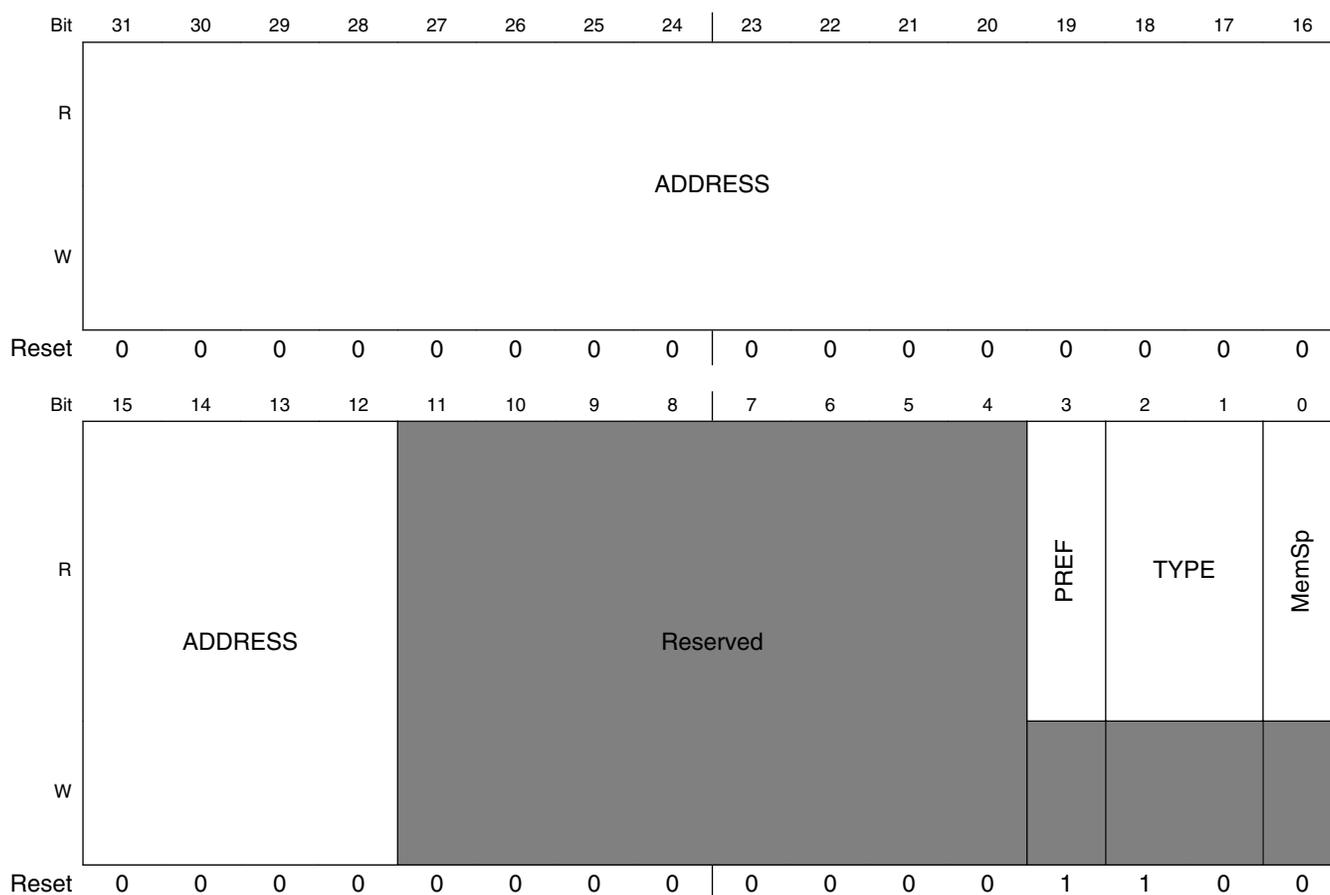
This register is supported only for EP mode.

Type 0 configuration header registers

The PCI Express base address registers (BARs) point to the beginning of distinct address ranges which the device should claim. In EP mode, the device supports a configuration space BAR, a 32-bit memory space BAR, and two 64-bit memory space BARs. In RC mode, the device only supports the configuration space BAR in the header; the other memory spaces are defined by the inbound ATMUs. Refer to [PCI Express inbound ATMUs](#), for more information.

Base address register 2 at offset 0x18 and base address register 4 at offset 0x20 are used to define the lower portion of the 64-bit inbound memory windows.

Address: base + offset + (8d × i), where i=0d to 1d



BAR_n field descriptions

Field	Description
31–12 ADDRESS	Indicates the lower portion of the base address where the inbound memory window begins. The number of bits that the device allows to be writable is selected through the inbound window size in the inbound window attributes registers (PEXIWAR2 for offset 0x18 and PEXIWAR3 for offset 0x20).
11–4 -	This field is reserved. Reserved. The device allows a 4 Kbyte window minimum.
3 PREF	Prefetchable. This bit is determined by PEXIWAR <i>n</i> [2].

Table continues on the next page...

BAR_n field descriptions (continued)

Field	Description
2–1 TYPE	Type. 0b10 Locate anywhere in 64-bit address space.
0 MemSp	Memory space indicator

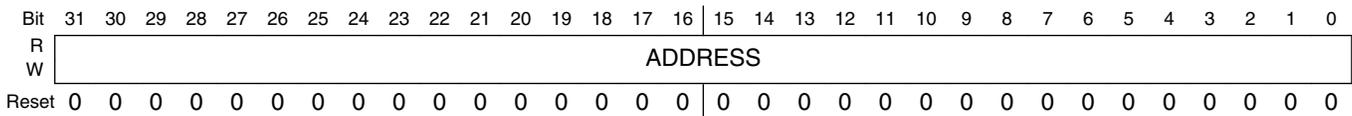
16.8.13 PCI Express Base Address Register 3,5 (BAR_n)

This register is supported only for EP mode.

The PCI Express base address registers (BARs) point to the beginning of distinct address ranges which the device should claim. In EP mode, the device supports a configuration space BAR, a 32-bit memory space BAR, and two 64-bit memory space BARs. In RC mode, the device only supports the configuration space BAR in the header; the other memory spaces are defined by the inbound ATMUs. Refer to [PCI Express inbound ATMUs](#) , for more information.

Base address register 3 at offset 0x1C and base address register 5 at offset 0x24 are used to define the upper portion of the 64-bit inbound memory windows.

Address: base + offset + (8d × i), where i=0d to 1d



BAR_n field descriptions

Field	Description
31–0 ADDRESS	Indicates the upper portion of the base address where the inbound memory window begins. The number of bits that the device allows to be writable is selected through the inbound window size in the inbound window attributes registers (PEXIWAR2 for offset 0x1C and PEXIWAR3 for offset 0x24). If no access to local memory is to be permitted by external requestors, then all bits are programmed.

16.8.14 PCI Express Subsystem Vendor ID Register (Subsystem_Vendor_ID_Register)

This register is supported only for EP mode.

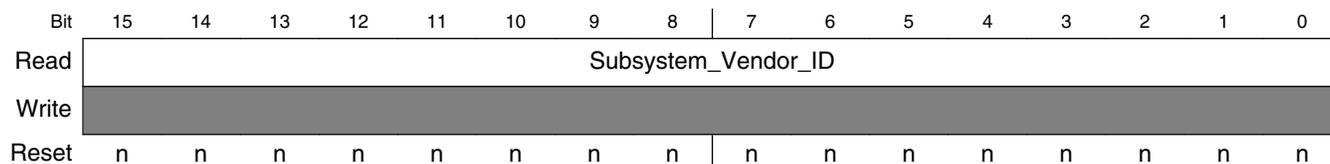
The PCI Express subsystem vendor ID register is used to identify the subsystem.

NOTE

Reset value taken from PEX_SSVID_UPDATE[SS_ID].

Type 0 configuration header registers

Address: 2Ch



Subsystem_Vendor_ID_Register field descriptions

Field	Description
15–0 Subsystem_Vendor_ID	The value for subsystem vendor ID is determined by the PCI Express subsystem vendor ID update register. See PCI Express Subsystem Vendor ID Update Register (Subsystem_Vendor_ID_Update_Register) , for more information.

16.8.15 PCI Express Subsystem ID Register (Subsystem_ID_Register)

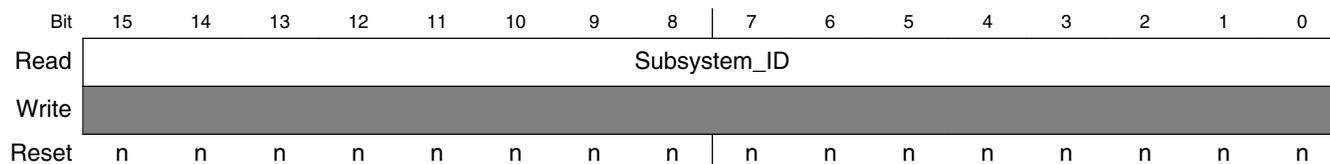
This register is supported only for EP mode.

The PCI Express subsystem ID register is used to identify the subsystem.

NOTE

Reset value taken from PEX_SSVID_UPDATE[SSV_ID].

Address: 2Eh



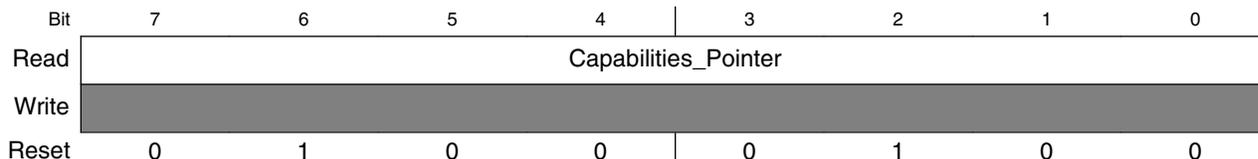
Subsystem_ID_Register field descriptions

Field	Description
15–0 Subsystem_ID	The value for subsystem ID is determined by the PCI Express subsystem vendor ID update register. See PCI Express Subsystem Vendor ID Update Register (Subsystem_Vendor_ID_Update_Register) , for more information.

16.8.16 Capabilities Pointer Register (Capabilities_Pointer_Register)

The capabilities pointer identifies additional functionality supported by the device.

Address: 34h



Capabilities_Pointer_Register field descriptions

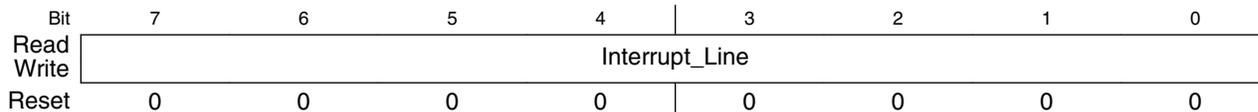
Field	Description
7-0 Capabilities_Pointer	The capabilities pointer provides the offset (0x44) for additional PCI-compatible registers above the common 64-byte header. Refer to PCI compatible device-specific configuration space for more information.

16.8.17 PCI Express Interrupt Line Register (Interrupt_Line_Register)

This register is supported only for EP mode.

The interrupt line register is used by device drivers and OS software to communicate interrupt line routing information. Values in this register are programmed by system software and are system specific.

Address: 3Ch



Interrupt_Line_Register field descriptions

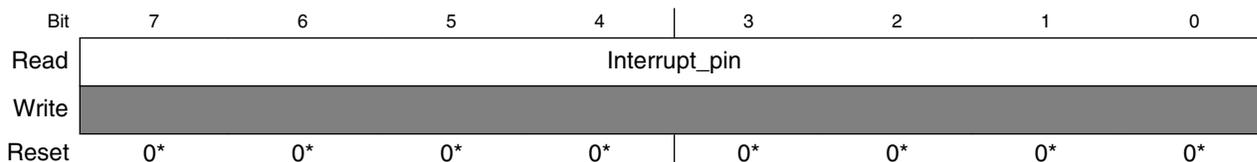
Field	Description
7-0 Interrupt_Line	Used to communicate interrupt line routing information.

16.8.18 PCI Express Interrupt Pin Register (Interrupt_Pin_Register)

The interrupt pin register identifies the legacy interrupt (INTx) messages the device (or function) uses.

This register only applies to EP mode.

Address: 3Dh



* Notes:

- Reset value is 0x00 for RC mode and 0x01 for EP mode.

Interrupt_Pin_Register field descriptions

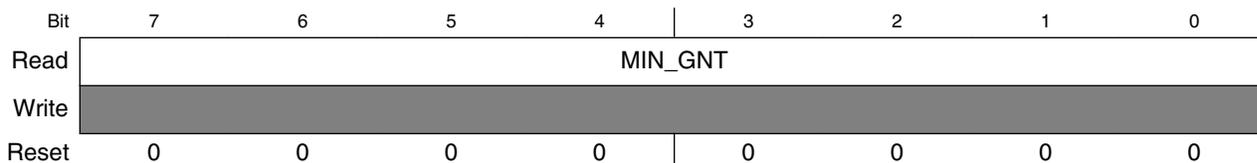
Field	Description
7-0 Interrupt_pin	<p>Legacy INTx message used by this device.</p> <p>All other settings Reserved.</p> <p>0x00 This device does not use legacy interrupt (INTx) messages.</p> <p>0x01 INTA</p> <p>0x02 INTB</p> <p>0x03 INTC</p> <p>0x04 INTD</p>

16.8.19 PCI Express Minimum Grant Register (Minimum_Grant_Register)

This register is supported only for EP mode.

This register does not apply to PCI Express. It is present for legacy purposes.

Address: 3Eh



Minimum_Grant_Register field descriptions

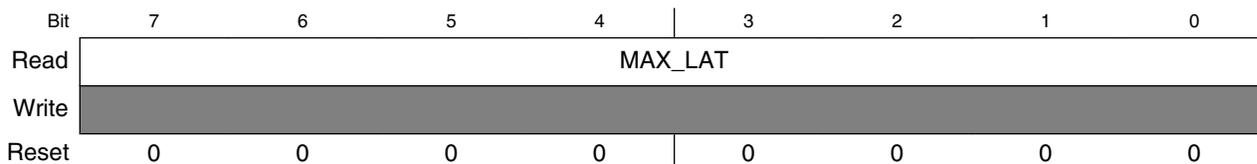
Field	Description
7-0 MIN_GNT	Does not apply for PCI Express.

16.8.20 PCI Express Maximum Latency Register (Maximum_Latency_Register)

This register is supported only for EP mode.

This register does not apply to PCI Express. It is present for legacy purposes.

Address: 3Fh



Maximum_Latency_Register field descriptions

Field	Description
7-0 MAX_LAT	Does not apply for PCI Express.

16.9 Type 1 configuration header registers

The figure below shows the type 1 header.

NOTE

The first 16 bytes of the type 1 header are identical to the type 0 header. See [Type 0 configuration header registers](#) for descriptions of those registers. This section describes the registers that are unique to the type 1 header beginning at offset 0x10.

NOTE

The BIST register (0x0F) is optional and reserved on the PCI Express controller.

Reserved				Address Offset (Hex)
Device ID		Vendor ID		00
Status		Command		04
Class Code			Revision ID	08
BIST	Header Type	Latency Timer	Cache Line Size	0C
Base Address Register 0				10
				14
Second Latency Timer	Subordinate Bus Number	Secondary Bus Number	Primary Bus Number	18
Secondary Status		I/O Limit	I/O Base	1C
Memory Limit		Memory Base		20
Prefetchable Memory Limit		Prefetchable Memory Base		24
Prefetchable Base Upper 32 Bits				28
Prefetchable Limit Upper 32 Bits				2C
I/O Limit upper 16 Bits		I/O Base Upper 16 Bits		30
			Capabilities Pointer	34
Expansion ROM Base Address				38
Bridge Control		Interrupt Pin	Interrupt Line	3C

Figure 16-260. PCI Express PCI-Compatible Configuration Header—Type 1

NOTE

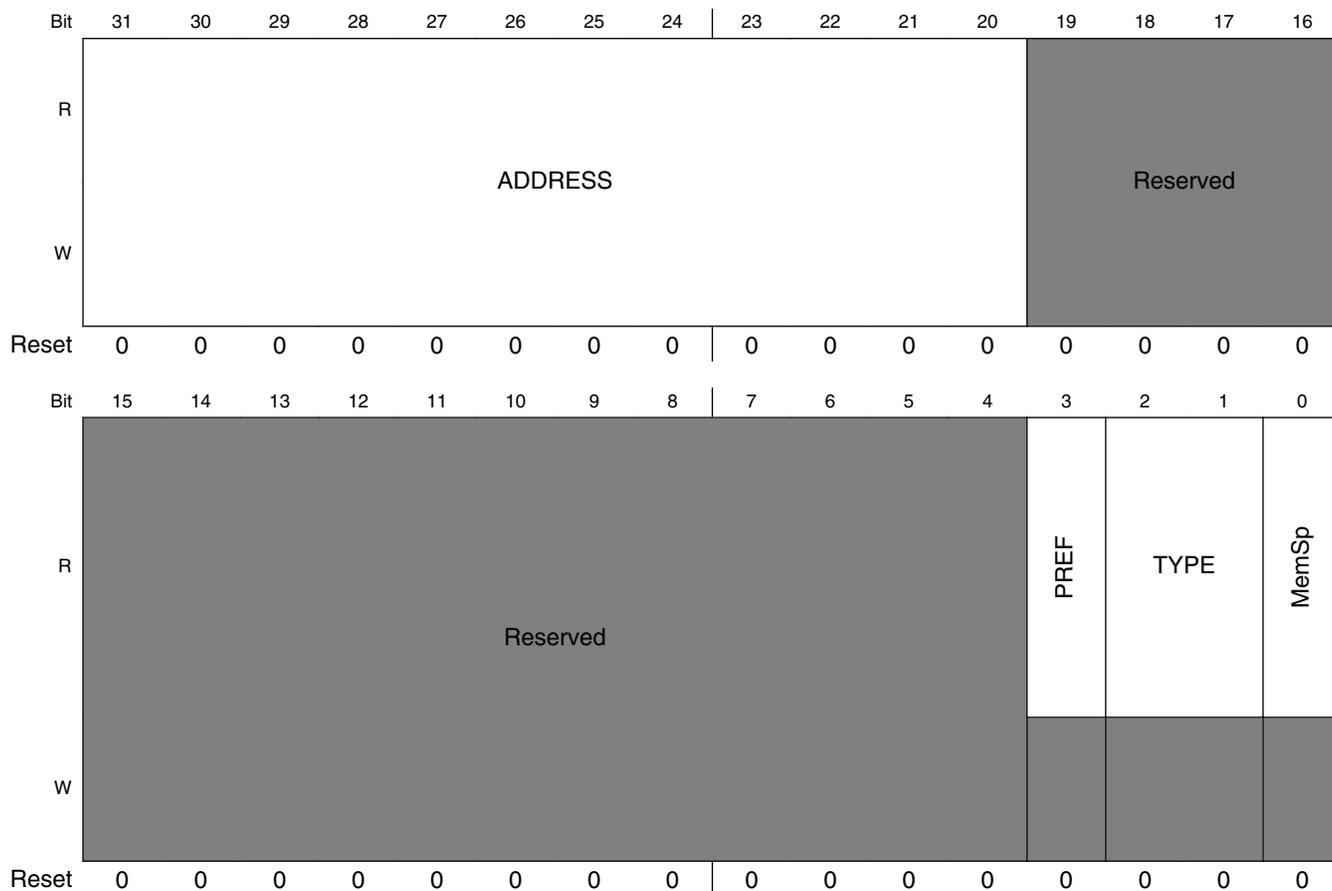
The secondary latency timer register (0x1B) does not apply to PCI Express. It must be read-only and return all zeros when read.

16.9.1 PCI Express Base Address Register 0 (PEXCSRBAR)

Base address register 0, also known as the PCI Express configuration and status register base address register (PEXCSRBAR), at offset 0x10 is a special fixed 1-Mbyte window that is used for inbound configuration accesses.

Note that PEXCSRBAR cannot be updated through the inbound ATMU registers.

Address: 10h



PEXCSRBAR field descriptions

Field	Description
31–20 ADDRESS	Indicates the base address that the inbound configuration window occupies. This window is fixed at 1 Mbyte.
19–4 -	This field is reserved. Reserved
3 PREF	Prefetchable
2–1 TYPE	Type. 00 Locate anywhere in 32-bit address space.
0 MemSp	Memory space indicator

16.9.2 PCI Express Primary Bus Number Register (Primary_Bus_Number_Register)

Address: 18h

Bit	7	6	5	4	3	2	1	0
Read	Primary_Bus_Number							
Write								
Reset	0	0	0	0	0	0	0	0

Primary_Bus_Number_Register field descriptions

Field	Description
7-0 Primary_Bus_Number	Bus that is connected to the upstream interface. Note that this register is programmed during system enumeration; in RC mode this register should remain 0x00.

16.9.3 PCI Express Secondary Bus Number Register (Secondary_Bus_Number_Register)

Address: 19h

Bit	7	6	5	4	3	2	1	0
Read	Secondary_Bus_Number							
Write								
Reset	0	0	0	0	0	0	0	0

Secondary_Bus_Number_Register field descriptions

Field	Description
7-0 Secondary_Bus_Number	Bus that is directly connected to the downstream interface. Note that this register is programmed during system enumeration; in RC mode, this register is typically programmed to 0x01.

16.9.4 PCI Express Subordinate Bus Number Register (Subordinate_Bus_Number_Register)

Address: 1Ah

Bit	7	6	5	4	3	2	1	0
Read	Subordinate_Bus_Number							
Write								
Reset	0	0	0	0	0	0	0	0

Subordinate_Bus_Number_Register field descriptions

Field	Description
7-0 Subordinate_Bus_Number	Highest bus number that is on the downstream interface.

16.9.5 PCI Express I/O Base Register (IO_Base_Register)

Note that this device does not support inbound I/O transactions.

Address: 1Ch

Bit	7	6	5	4	3	2	1	0
Read	IO_Start_Address				Address_Decode_Type			
Write								
Reset	0	0	0	0	0	0	0	0

IO_Base_Register field descriptions

Field	Description
7-4 IO_Start_Address	Specifies bits 15:12 of the I/O space start address. NOTE: I/O Start Address is writeable in root complex, read-only in end point.
3-0 Address_Decode_Type	Specifies the number of I/O address bits. All other settings reserved. 0x00 16-bit I/O address decode 0x01 32-bit I/O address decode

16.9.6 PCI Express I/O Limit Register (IO_Limit_Register)

Note that this device does not support inbound I/O transactions.

Address: 1Dh

Bit	7	6	5	4	3	2	1	0
Read	IO_Limit_Address				Address_Decode_Type			
Write								
Reset	0	0	0	0	0	0	0	0

IO_Limit_Register field descriptions

Field	Description
7-4 IO_Limit_Address	Specifies bits 15:12 of the I/O space ending address
3-0 Address_Decode_Type	Specifies the number of I/O address bits. All other settings reserved. 0x00 16-bit I/O address decode 0x01 32-bit I/O address decode

16.9.7 PCI Express Secondary Status Register (Secondary_Status_Register)

Note that the errors in this register may be masked by corresponding bits in the secondary status interrupt mask register (PEX_SS_INTR_MASK) and that by default all of the errors are masked. See [Secondary Status Interrupt Mask Register \(Secondary_Status_Interrupt_Mask_Register\)](#) for more information.

Address: 1Eh

Bit	15	14	13	12	11	10	9	8
Read	DPE	SSE	RMA	RTA	STA	Reserved		MDPE
Write	w1c	w1c	w1c	w1c	w1c	Reserved		w1c
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved							
Write	Reserved							
Reset	0	0	0	0	0	0	0	0

Secondary_Status_Register field descriptions

Field	Description
15 DPE	Detected parity error. This bit is set whenever the secondary side receives a poisoned TLP regardless of the state of the parity error response bit.
14 SSE	Signaled system error. This bit is set when a device sends a ERR_FATAL or ERR_NONFATAL message, provided the SERR enable bit in the command register is set to enable reporting.
13 RMA	Received master abort. This bit is set when the secondary side receives an unsupported request (UR) completion.
12 RTA	Received target abort. This bit is set when the secondary side receives a completer abort (CA) completion.

Table continues on the next page...

Secondary_Status_Register field descriptions (continued)

Field	Description
11 STA	Signaled target abort. This bit is set when the secondary side issues a CA completion.
10–9 -	This field is reserved. Reserved
8 MDPE	Master data parity error. This bit is set when the parity error response bit is set and the secondary side requestor receives a poisoned completion or poisons a write request. If the parity error response bit is cleared, this bit is never set.
7–0 -	This field is reserved. Reserved

16.9.8 PCI Express Memory Base Register (Memory_Base_Register)

Address: 20h

Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Read	Memory_Base												Reserved			
Write	Memory_Base												Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Memory_Base_Register field descriptions

Field	Description
15–4 Memory_Base	Specifies bits 31:20 of the non-prefetchable memory space start address. Typically used for specifying memory-mapped I/O space. NOTE: Inbound posted transactions hitting into the mem base/limit range are ignored; inbound non-posted transactions hitting into the mem base/limit range results in an unsupported request response.
3–0 -	This field is reserved. Reserved

16.9.9 PCI Express Memory Limit Register (Memory_Limit_Register)

Address: 22h

Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Read	Memory_Limit												Reserved			
Write	Memory_Limit												Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Memory_Limit_Register field descriptions

Field	Description
15–4 Memory_Limit	Specifies bits 31:20 of the non-prefetchable memory space ending address. Typically used for specifying memory-mapped I/O space. NOTE: Inbound posted transactions hitting into the mem base/limit range are ignored; inbound non-posted transactions hitting into the mem base/limit range results in unsupported request response.
3–0 -	This field is reserved. Reserved

16.9.10 PCI Express Prefetchable Memory Base Register (Prefetchable_Memory_Base_Register)

Address: 24h

Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Read	PF_Memory_Base												Address_Decode_Type			
Write	PF_Memory_Base												Address_Decode_Type			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

Prefetchable_Memory_Base_Register field descriptions

Field	Description
15–4 PF_Memory_Base	Specifies bits 31:20 of the prefetchable memory space start address. NOTE: Inbound posted transactions hitting into the mem base/limit range are ignored; inbound non-posted transactions hitting into the mem base/limit range results in unsupported request response.
3–0 Address_Decode_Type	Specifies the number of prefetchable memory address bits. All other settings reserved. 0x00 32-bit memory address decode 0x01 64-bit memory address decode

16.9.11 PCI Express Prefetchable Memory Limit Register (Prefetchable_Memory_Limit_Register)

Address: 26h

Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Read	PF_Memory_Limit												Address_Decode_Type			
Write	PF_Memory_Limit												Address_Decode_Type			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

Prefetchable_Memory_Limit_Register field descriptions

Field	Description
15–4 PF_Memory_Limit	Specifies bits 31:20 of the prefetchable memory space ending address. NOTE: Inbound posted transactions hitting into the mem base/limit range are ignored; inbound non-posted transactions hitting into the mem base/limit range results in unsupported request response.
3–0 Address_Decode_Type	Specifies the number of prefetchable memory address bits. All other settings reserved. 0x00 32-bit memory address decode 0x01 64-bit memory address decode

16.9.12 PCI Express Prefetchable Base Upper 32 Bits Register (Prefetchable_Base_Upper_32_Bits_Register)

Address: 28h

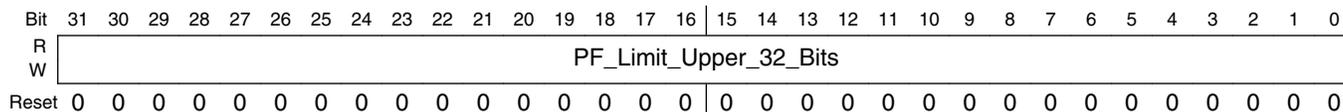


Prefetchable_Base_Upper_32_Bits_Register field descriptions

Field	Description
31–0 PF_Base_Upper_32_Bits	Specifies bits 64:32 of the prefetchable memory space start address when the address decode type field in the prefetchable memory base register is 0x01.

16.9.13 PCI Express Prefetchable Limit Upper 32 Bits Register (Prefetchable_Limit_Upper_32_Bits_Register)

Address: 2Ch



Prefetchable_Limit_Upper_32_Bits_Register field descriptions

Field	Description
31–0 PF_Limit_Upper_32_Bits	Specifies bits 64-32 of the prefetchable memory space ending address when the address decode type field in the prefetchable memory limit register is 0x01.

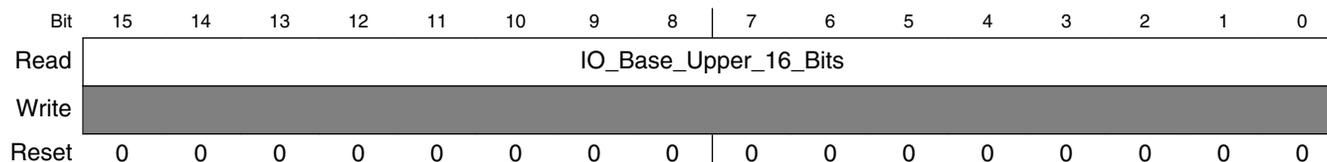
16.9.14 PCI Express I/O Base Upper 16 Bits Register (IO_Base_Upper_16_Bits_Register)

Note that this device does not support inbound I/O transactions.

NOTE

Access is read/write in root complex and read only in endpoint.

Address: 30h



IO_Base_Upper_16_Bits_Register field descriptions

Field	Description
15-0 IO_Base_Upper_16_Bits	Specifies bits 31-16 of the I/O space start address when the address decode type field in the I/O base register is 0x01.

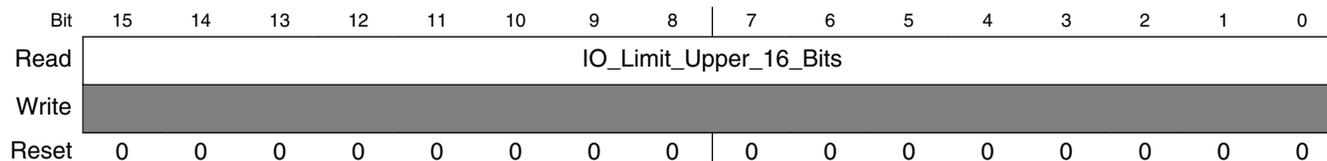
16.9.15 PCI Express I/O Limit Upper 16 Bits Register (IO_Limit_Upper_16_Bits_Register)

Note that this device does not support inbound I/O transactions.

NOTE

Access is read/write in root complex and read only in endpoint.

Address: 32h



IO_Limit_Upper_16_Bits_Register field descriptions

Field	Description
15-0 IO_Limit_Upper_16_Bits	Specifies bits 31-16 of the I/O space ending address when the address decode type field in the I/O limit register is 0x01.

16.9.16 Capabilities Pointer Register (Capabilities_Pointer_Register)

The capabilities pointer identifies additional functionality supported by the device.

Address: 34h

Bit	7	6	5	4	3	2	1	0
Read	Capabilities_Pointer							
Write								
Reset	0	1	0	0	0	1	0	0

Capabilities_Pointer_Register field descriptions

Field	Description
7-0 Capabilities_Pointer	The capabilities pointer provides the offset (0x44) for additional PCI-compatible registers above the common 64-byte header. Refer to PCI compatible device-specific configuration space for more information.

16.9.17 PCI Express Interrupt Line Register (Interrupt_Line_Register)

This register only applies to endpoint mode.

The interrupt line register is used by device drivers and OS software to communicate interrupt line routing information. Values in this register are programmed by system software and are system specific.

Address: 3Ch

Bit	7	6	5	4	3	2	1	0
Read	Interrupt_Line							
Write								
Reset	0	0	0	0	0	0	0	0

Interrupt_Line_Register field descriptions

Field	Description
7-0 Interrupt_Line	Used to communicate interrupt line routing information.

16.9.18 PCI Express Interrupt Pin Register (Interrupt_Pin_Register)

The interrupt pin register identifies the legacy interrupt (INTx) messages the device (or function) uses when configured in EP mode.

This register only applies to EP mode.

Address: 3Dh

Bit	7	6	5	4	3	2	1	0
Read	Interrupt_pin							
Write								
Reset	0*	0*	0*	0*	0*	0*	0*	0*

* Notes:

- Reset value is 0x00 for RC mode and 0x01 for EP mode.

Interrupt_Pin_Register field descriptions

Field	Description
7-0 Interrupt_pin	<p>Legacy INTx message used by this device.</p> <p>All other settings reserved.</p> <p>0x00 This device does not use legacy interrupt (INTx) messages.</p> <p>0x01 INTA</p> <p>0x02 INTB</p> <p>0x03 INTC</p> <p>0x04 INTD</p>

16.9.19 PCI Express Bridge Control Register (Bridge_Control_Register)

Address: 3Eh

Bit	15	14	13	12	11	10	9	8
Read	Reserved							
Write	Reserved							
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved	Scnd_RST	Reserved		VGA_EN	ISA_EN	SERR_EN	PER
Write	Reserved	Scnd_RST	Reserved		VGA_EN	ISA_EN	SERR_EN	PER
Reset	0	0	0	0	0	0	0	0

Bridge_Control_Register field descriptions

Field	Description
15–7 -	This field is reserved. Reserved
6 Scnd_RST	Secondary bus reset
5–4 -	This field is reserved. Reserved
3 VGA_EN	VGA enable
2 ISA_EN	ISA enable
1 SERR_EN	SERR enable. This bit controls the propagation of ERR_COR, ERR_NONFATAL, and ERR_FATAL responses received on the secondary side.
0 PER	Parity error response.

16.10 PCI compatible device-specific configuration space

The PCI compatible device-specific configuration space is a PCI compatible configuration space from 0x40 to 0xFF (just after the 64-byte PCI-compatible configuration header).

Reserved		Address Offset (Hex)	
PCI-Compatible Configuration Header		00	
		3F	
		40	
Power Mgmt Capabilities	Next Pointer (0x4C)	Power Mgmt Capability ID	44
Data		Power Management Status & Control	48
PCI Express Capabilities	Next Pointer (0x70—EP mode) (NULL—RC mode)	PCI Express Capability ID	4C
Device Capabilities		50	
Device Status	Device Control		54
Link Capabilities		58	
Link Status	Link Control		5C
Slot Capabilities (RC mode only)		60	
Slot Status (RC mode only)	Slot Control (RC mode only)		64
	Root Control (RC mode only)		68
Root Status		6C	
MSI Message Control	Next Pointer (NULL)	MSI Message Capability ID	70
MSI Message Address		74	
MSI Upper Message Address		78	
	MSI Message Data		7C
		80	
		FF	

Figure 16-280. PCI Compatible Device-Specific Configuration Space

16.10.1 PCI Express Power Management Capability ID Register (Power_Management_Capability_ID_Register)

Address: 44h

Bit	7	6	5	4	3	2	1	0
Read	Power_Mgmt_Capability_ID							
Write								
Reset	0	0	0	0	0	0	0	1

Power_Management_Capability_ID_Register field descriptions

Field	Description
7-0 Power_Mgmt_Capability_ID	Power Management = 0x01

16.10.2 PCI Express Power Management Capabilities Register (Power_Management_Capabilities_Register)

Address: 46h

Bit	15	14	13	12	11	10	9	8
Read	PME_Support					D2	D1	AUX_Curr
Write								
Reset	n*	1*	1*	1*	1*	1	1	0
Bit	7	6	5	4	3	2	1	0
Read	AUX_Curr	DSI	Reserved		PME_CLK	Version		
Write								
Reset	0	0	0	0	0	0	1	0

* Notes:

- PME_Support field: Reset value is 11111 in RC mode and 01111 in EP mode.

Power_Management_Capabilities_Register field descriptions

Field	Description
15-11 PME_Support	Indicates the power states that this device supports
10 D2	D2 Support
9 D1	D1 Support

Table continues on the next page...

Power_Management_Capabilities_Register field descriptions (continued)

Field	Description
8–6 AUX_Curr	AUX Current
5 DSI	Device Specific Initialization
4 -	This field is reserved. Reserved
3 PME_CLK	Does not apply to PCI Express.
2–0 Version	Setting depends on version of the specification.

16.10.3 PCI Express Power Management Status and Control Register (Power_Management_Status_and_Control_Register)

Address: 48h

Bit	15	14	13	12	11	10	9	8
Read	PME_STAT	Data_Scale		Data_Select				PME_EN
Write	w1c							
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved						Power_State	
Write								
Reset	0	0	0	0	0	0	0	0

Power_Management_Status_and_Control_Register field descriptions

Field	Description
15 PME_STAT	PME Status
14–13 Data_Scale	Obtained directly from the <i>PCI Express Base Specification</i> .
12–9 Data_Select	Obtained directly from the <i>PCI Express Base Specification</i> .
8 PME_EN	PME Enable NOTE: Bitfield access is sticky.
7–2 -	This field is reserved. Reserved
1–0 Power_State	Power state. Indicates the current power state of the function. 0x00 D0

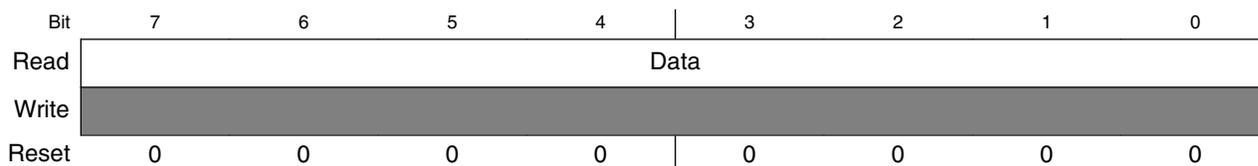
Table continues on the next page...

Power_Management_Status_and_Control_Register field descriptions (continued)

Field	Description
0x01 D1	
0x02 D2	
0x03 D3	

16.10.4 PCI Express Power Management Data Register (Power_Management_Data_Register)

Address: 4Bh

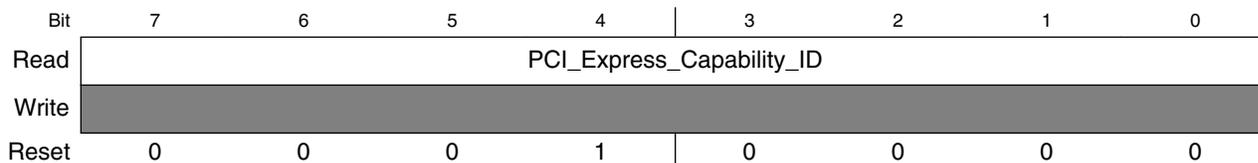


Power_Management_Data_Register field descriptions

Field	Description
7-0 Data	Obtained directly from the <i>PCI Express Base Specification</i> .

16.10.5 PCI Express Capability ID Register (Capability_ID_Register)

Address: 4Ch



Capability_ID_Register field descriptions

Field	Description
7-0 PCI_Express_Capability_ID	PCI Express = 0x10

16.10.6 PCI Express Capabilities Register (Capabilities_Register)

Address: 4Eh

Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Read	Reserved		Interrupt_Message_Number				Slot	Device_Port_Type				Capability_Version				
Write	Reserved															
Reset	0	0	0	0	0	0	0	0	*	*	*	*	0	0	0	1

* Notes:

- Device_Port_Type field: Reset value depends on mode: 0100 (RC mode), 0000 (EP mode)

Capabilities_Register field descriptions

Field	Description
15–14 -	This field is reserved. Reserved
13–9 Interrupt_ Message_ Number	If this function is allocated more than one MSI interrupt number, then this register is required to contain the offset between the base Message Data and the MSI Message that is generated when any of the status bits in either the Slot Status register or the Root Port Status register, of this capability structure, are set.
8 Slot	Slot Implemented (RC mode only)
7–4 Device_Port_ Type	0100 (RC mode) 0000 (EP mode)
3–0 Capability_ Version	Indicates the defined PCI Express capability structure version number. Must be 1h for 1.0, 1.0a, and 1.1 specification.

16.10.7 PCI Express Device Capabilities Register (Device_Capabilities_Register)

Refer to PCI Express Base Specification for register details.

Address: 50h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	
R	Reserved				CSPLS		CSPLV								Reserved		
W	Reserved				Reserved		Reserved								Reserved		
Reset	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	0	0
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
R	Reserved	PIP	AIP	ABP	EP_L1_LAT			EP_L0s_LAT			ET	PHAN_FCT	MAX_PL_SIZE_SUP				
W	Reserved	Reserved	Reserved	Reserved	Reserved			Reserved			Reserved	Reserved	Reserved				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	

Device_Capabilities_Register field descriptions

Field	Description
31–28 -	This field is reserved. Reserved
27–26 CSPLS	Captured Slot Power Limit Scale
25–18 CSPLV	Captured Slot Power Limit Value
17–15 -	This field is reserved. Reserved
14 PIP	Power Indicator Present
13 AIP	Attention Indicator Present
12 ABP	Attention Button Present
11–9 EP_L1_LAT	Endpoint L1 Acceptable Latency
8–6 EP_L0s_LAT	Endpoint L0s Acceptable Latency
5 ET	Extended Tag Field Supported
4–3 PHAN_FCT	Phantom Functions Supported
2–0 MAX_PL_SIZE_SUP	Maximum payload size supported. 001 = 256-bytes

16.10.8 PCI Express Device Control Register (Device_Control_Register)

Address: 54h

Bit	15	14	13	12	11	10	9	8
Read	Reserved	MAX_READ_SIZE			NSE	APE	PFE	ETE
Write								
Reset	0	0	1	0	1	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	MAX_PAYLOAD_SIZE			RO	URR	FER	NFER	CER
Write								
Reset	0	0	0	1	0	0	0	0

Device_Control_Register field descriptions

Field	Description
15 -	This field is reserved. Reserved
14–12 MAX_READ_SIZE	Maximum read request size
11 NSE	No snoop enable
10 APE	AUX power PM enable
9 PFE	Phantom functions enable
8 ETE	Extended tag field enable
7–5 MAX_PAYLOAD_SIZE	Maximum payload size 000 For 128 bytes MAX_PAYLOAD_SIZE 001 For 256 bytes MAX_PAYLOAD_SIZE
4 RO	Relaxed ordering NOTE: Configuration software must clear this bit if relaxed ordering is not desired. 1 Relaxed ordering is enabled.
3 URR	Unsupported request reporting
2 FER	Fatal error reporting
1 NFER	Non-fatal error reporting
0 CER	Correctable error reporting

16.10.9 PCI Express Device Status Register (Device_Status_Register)

Address: 56h

Bit	15	14	13	12	11	10	9	8
Read	Reserved							
Write	Reserved							
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved		TP	APD	URD	FED	NFED	CED
Write	Reserved				w1c	w1c	w1c	w1c
Reset	0	0	0	0	0	0	0	0

Device_Status_Register field descriptions

Field	Description
15–6 -	This field is reserved. Reserved
5 TP	Transactions pending
4 APD	AUX power detected
3 URD	Unsupported request detected
2 FED	Fatal error detected
1 NFED	Non-fatal error detected
0 CED	Correctable error detected

16.10.10 PCI Express Link Capabilities Register (Link_Capabilities_Register)

NOTE

See bitfield descriptions for reset values.

Address: 58h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Port_Number								Reserved						L1_EX_LAT	
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	L1_EX_LAT	L0s_EX_LAT			ASPM		MAX_LINK_W				MAX_LINK_SP					
W																
Reset		1	0	1	0	1	0	0	n	n	n	n	0	0	0	1

Link_Capabilities_Register field descriptions

Field	Description
31–24 Port_Number	Port number for PCI Express link
23–18 -	This field is reserved. Reserved
17–15 L1_EX_LAT	L1 exit latency Note that exit latencies may be influenced by PCI Express reference clock configuration depending upon whether a component uses a common or separate reference clock. 000 Less than 1 μ s 001 1 μ s to less than 2 μ s 010 2 μ s to less than 4 μ s 011 4 μ s to less than 8 μ s 100 8 μ s to less than 16 μ s 101 16 μ s to less than 32 μ s 110 32 μ s - 64 μ s 111 More than 64 μ s
14–12 L0s_EX_LAT	L0s exit latency Note that exit latencies may be influenced by PCI Express reference clock configuration depending upon whether a component uses a common or separate reference clock. 000 Less than 64 ns 001 64 ns to less than 128 ns

Table continues on the next page...

Link_Capabilities_Register field descriptions (continued)

Field	Description
	010 128 ns to less than 256 ns 011 256 ns to less than 512 ns 100 512 ns to less than 1µs 101 1 µs to less than 2 µs 110 2 µs – 4 µs 111 More than 4 µs
11–10 ASPM	Active state power management (ASPM) Support 01 L0s entry supported 11 L0s and L1entry supported
9–4 MAX_LINK_W	Maximum link width. Reset value depends upon the maximum link width of the PCI Express controller is capable of supporting. 000001 x1 000010 x2 000100 x4 001000 x8
3–0 MAX_LINK_SP	Maximum link speed 0001 2.5 GT/s link

16.10.11 PCI Express Link Control Register (Link_Control_Register)

Address: 5Ch

Bit	15	14	13	12	11	10	9	8
Read	Reserved							
Write	Reserved							
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	EXT_SYNC	CCC		LD	RCB	Reserved	ASPM_CTL	
Write	EXT_SYNC	CCC	RL	LD	RCB	Reserved	ASPM_CTL	
Reset	0	0	0	0	0	0	0	0

Link_Control_Register field descriptions

Field	Description
15–8 -	This field is reserved. Reserved
7 EXT_SYNC	Extended synch

Table continues on the next page...

Link_Control_Register field descriptions (continued)

Field	Description
6 CCC	Common clock configuration
5 RL	Retrain link (Reserved for EP devices). In RC mode, setting this bit initiates link retraining by directing the Physical Layer LTSSM to the Recovery state, if the link has already been up; reads of this bit always return 0.
4 LD	Link disable (Reserved for EP devices)
3 RCB	Read completion boundary
2 -	This field is reserved. Reserved
1-0 ASPM_CTL	Active state power management (ASPM) control

16.10.12 PCI Express Link Status Register (Link_Status_Register)

Address: 5Eh

Bit	15	14	13	12	11	10	9	8
Read	Reserved			SCC	LT	Reserved	NEG_LINK_W	
Write	Reserved			Reserved	Reserved	Reserved	Reserved	
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	NEG_LINK_W				LINK_SP			
Write	Reserved				Reserved			
Reset	0	0	0	1	0	0	0	1

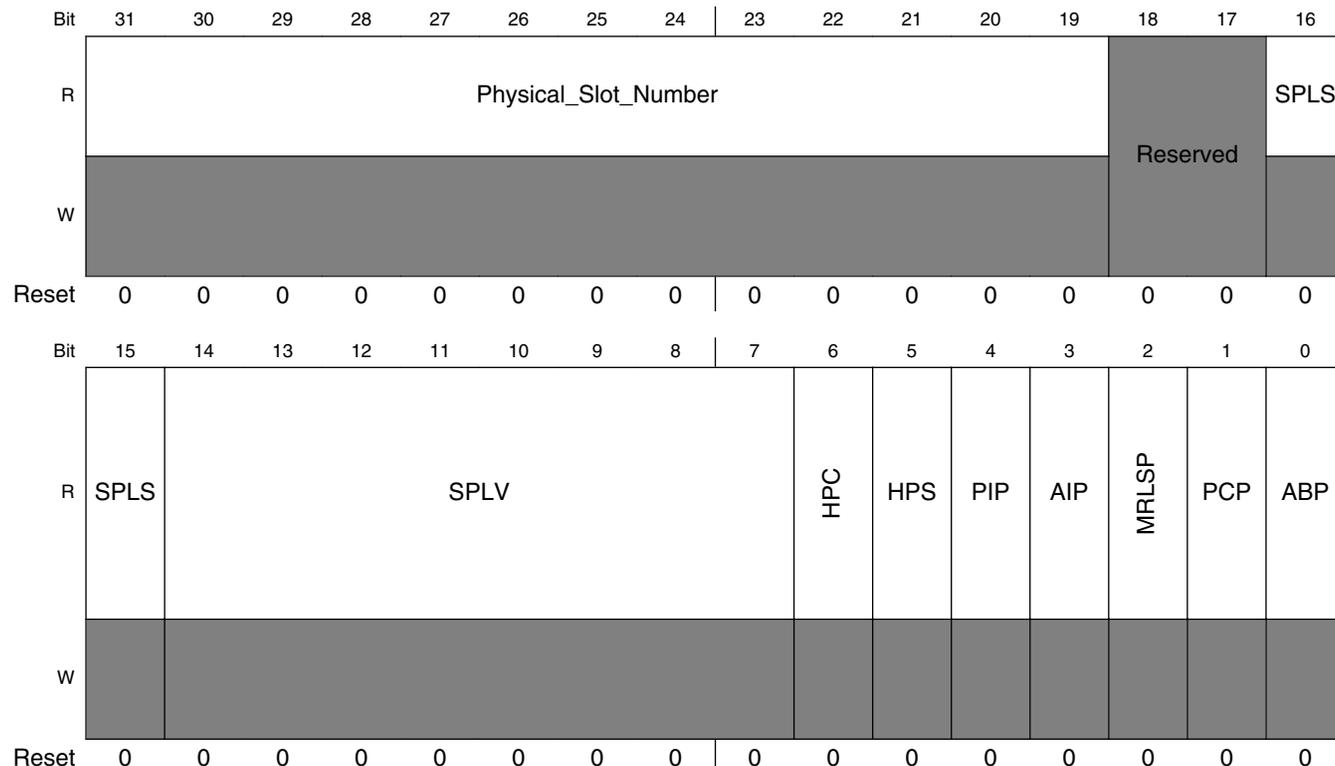
Link_Status_Register field descriptions

Field	Description
15-13 -	This field is reserved. Reserved
12 SCC	Slot clock configuration
11 LT	Link training
10 -	This field is reserved. Reserved.
9-4 NEG_LINK_W	Negotiated link width
3-0 LINK_SP	Link speed.

16.10.13 PCI Express Slot Capabilities Register (Slot_Capabilities_Register)

This register is supported only for RC mode.

Address: 60h



Slot_Capabilities_Register field descriptions

Field	Description
31–19 Physical_Slot_Number	This hardware initialized field indicates the physical slot number attached to this Port. This field must be hardware initialized to a value that assigns a slot number that is globally unique within the chassis. These registers should be initialized to 0 for Ports connected to devices that are either integrated on the system board or integrated within the same silicon as the Switch device or Root Port.
18–17 -	This field is reserved. Reserved
16–15 SPLS	Slot power limit scale.
14–7 SPLV	Slot power limit value.
6 HPC	Hot plug capable.

Table continues on the next page...

Slot_Capabilities_Register field descriptions (continued)

Field	Description
5 HPS	Hot plug surprise.
4 PIP	Power indicator present.
3 AIP	Attention indicator present.
2 MRLSP	MRL sensor present.
1 PCP	Power controller present.
0 ABP	Attention button present.

16.10.14 PCI Express Slot Control Register (Slot_Control_Register)

This register is supported only for RC mode.

Address: 64h

Bit	15	14	13	12	11	10	9	8
Read	Reserved					PCC	PIC	
Write								
Reset	0	0	0	0	0	0	1	1
Bit	7	6	5	4	3	2	1	0
Read	AIC		HPIE	CCIE	PDCE	MRLSCE	PFDE	ABPE
Write								
Reset	1	1	0	0	0	0	0	0

Slot_Control_Register field descriptions

Field	Description
15–11 -	This field is reserved. Reserved
10 PCC	Power controller control.
9–8 PIC	Power indicator control. If a power indicator is implemented, set the power indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. Defined encodings are shown below.

Table continues on the next page...

Slot_Control_Register field descriptions (continued)

Field	Description
	<p>NOTE: The default value of this field must be one of the non-reserved values. If the power indicator present bit in the slot capabilities register is 0, this bit is permitted to be read-only with a value of 00.</p> <p>00 Reserved 01 On 10 Blink 11 Off</p>
7–6 AIC	<p>Attention indicator control.</p> <p>If an attention indicator is implemented, writes to this field set the attention indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined.</p> <p>NOTE: The default value of this field must be one of the non-reserved values. if the attention indicator present bit in the slot capabilities register is 0, this bit is permitted to be readonly with a value of 00.</p> <p>00 Reserved 01 On 10 Blink 11 Off</p>
5 HPIE	Hot plug interrupt enable.
4 CCIE	Command completed interrupt enable.
3 PDCE	Presence detect changed enable.
2 MRLSCE	MRL sensor changed enable.
1 PFDE	Power fault detected enable.
0 ABPE	Attention button pressed enable.

16.10.15 PCI Express Slot Status Register (Slot_Status_Register)

This register is supported only for RC mode.

Address: 66h

Bit	15	14	13	12	11	10	9	8
Read	Reserved							
Write								
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved	PDS	MRLSS	CC	PDC	MRLSC	PFD	ABP
Write		w1c	w1c	w1c	w1c	w1c	w1c	w1c
Reset	0	1	0	0	1	0	0	0

Slot_Status_Register field descriptions

Field	Description
15–7 -	This field is reserved. Reserved
6 PDS	Presence detect state. In the PCI Express specification, this bit indicates the presence of a card in the slot; however, on this device, this bit is set regardless of whether there is a card present or not. 0 Slot empty 1 Card present in slot
5 MRLSS	MRL sensor state. 0 MRL closed 1 MRL open
4 CC	Command completed.
3 PDC	Presence detect changed.
2 MRLSC	MRL sensor changed.
1 PFD	Power fault detected.
0 ABP	Attention button pressed.

16.10.16 PCI Express Root Control Register (Root_Control_Register)

This register is supported only for RC mode.

Address: 68h

Bit	15	14	13	12	11	10	9	8
Read	Reserved							
Write								
Reset	0	0	0	0	0	0	0	0
Bit	7	6	5	4	3	2	1	0
Read	Reserved				PMEIE	SEFEE	SENFEE	SECEE
Write	Reserved							
Reset	0	0	0	0	0	0	0	0

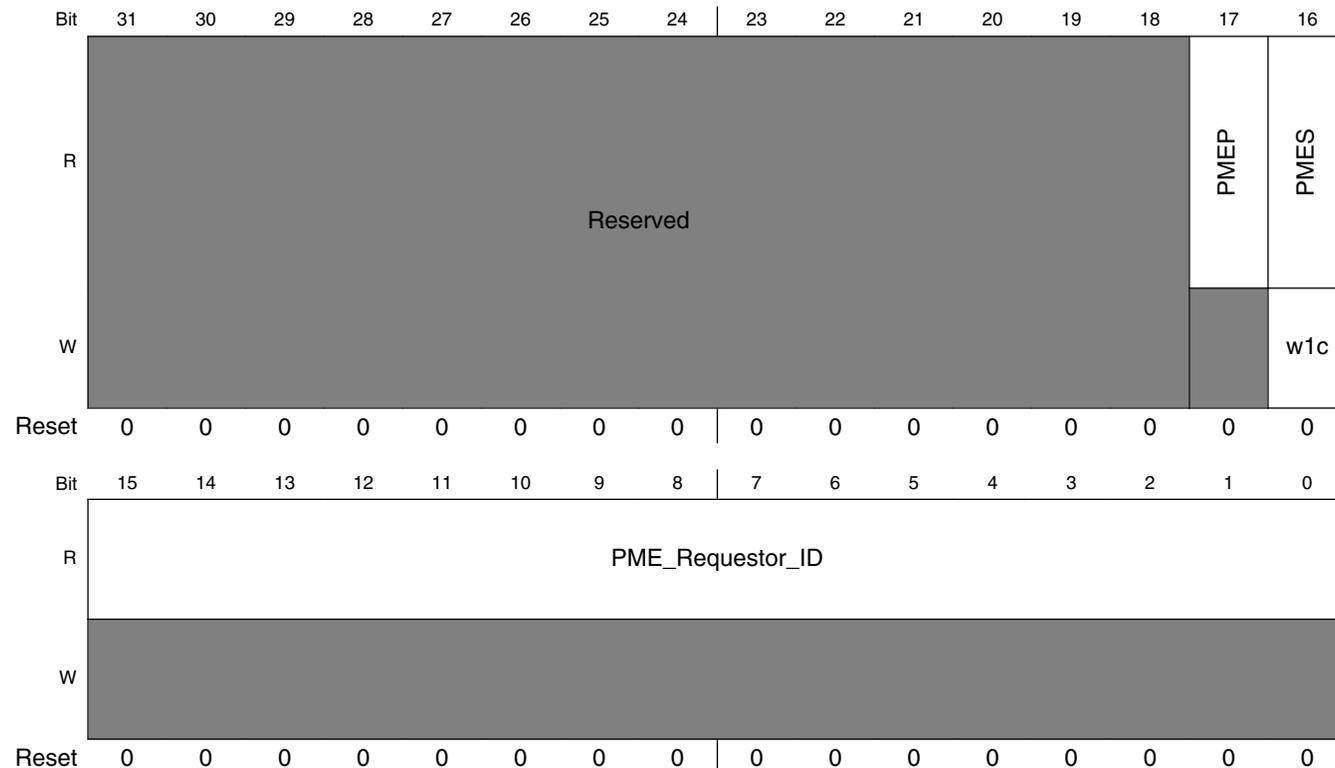
Root_Control_Register field descriptions

Field	Description
15–4 -	This field is reserved. Reserved
3 PMEIE	PME interrupt enable.
2 SEFEE	System error on fatal error enable.
1 SENFEE	System error on non-fatal error enable.
0 SECEE	System error on correctable error enable.

16.10.17 PCI Express Root Status Register (Root_Status_Register)

This register is supported only for RC mode.

Address: 6Ch



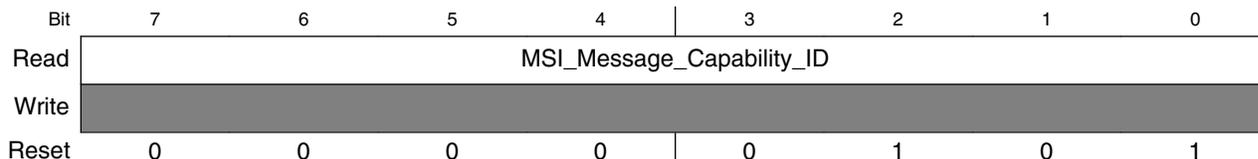
Root_Status_Register field descriptions

Field	Description
31–18 -	This field is reserved. Reserved
17 PMEP	PME pending.
16 PMES	PME status.
15–0 PME_Requestor_ID	PME requestor ID.

16.10.18 PCI Express MSI Message Capability ID Register (MSI_Message_Capability_ID_Register)

This register is supported only for EP mode.

Address: 70h



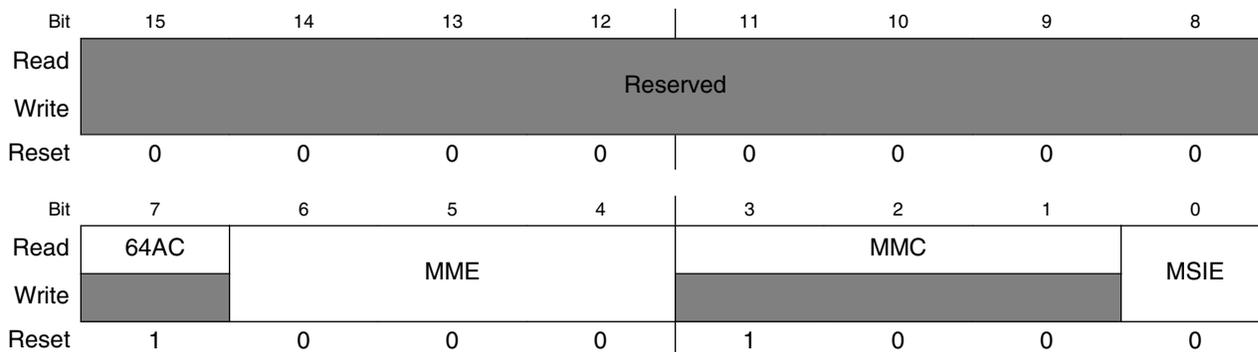
MSI_Message_Capability_ID_Register field descriptions

Field	Description
7-0 MSI_Message_Capability_ID	MSI Message = 0x05

16.10.19 PCI Express MSI Message Control Register (MSI_Message_Control_Register)

This register is supported only for EP mode.

Address: 72h



MSI_Message_Control_Register field descriptions

Field	Description
15-8 -	This field is reserved. Reserved

Table continues on the next page...

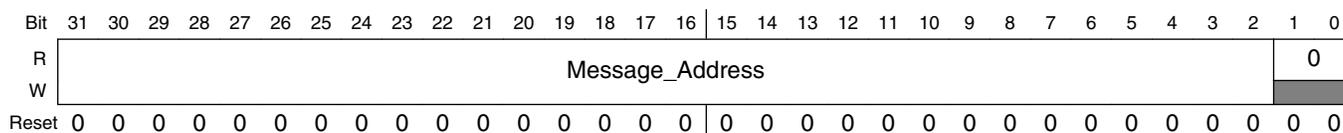
MSI_Message_Control_Register field descriptions (continued)

Field	Description
7 64AC	64-bit address capable.
6–4 MME	Multiple message enable.
3–1 MMC	Multiple message capable.
0 MSIE	MSI enable.

16.10.20 PCI Express MSI Message Address Register (MSI_Message_Address_Register)

This register only applies to endpoint mode.

Address: 74h



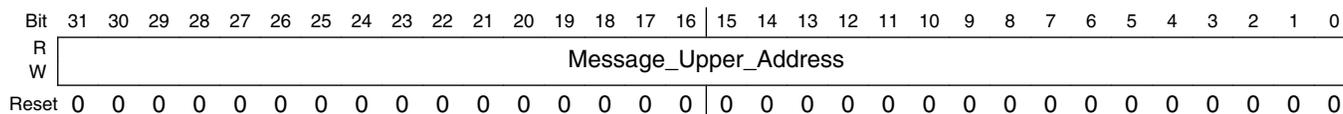
MSI_Message_Address_Register field descriptions

Field	Description
31–2 Message_Address	System-specified message address
1–0 Reserved	This read-only field is reserved and always has the value 0.

16.10.21 PCI Express MSI Message Upper Address Register (MSI_Message_Upper_Address_Register)

This register only applies to endpoint mode.

Address: 78h



MSI_Message_Upper_Address_Register field descriptions

Field	Description
31–0 Message_Upper_Address	System-specified message upper address

16.10.22 PCI Express MSI Message Data Register (MSI_Message_Data_Register)

This register is supported only for EP mode.

Address: 7Ch

Bit	15	14	13	12	11	10	9	8		7	6	5	4	3	2	1	0
Read	Message_Data																
Write																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

MSI_Message_Data_Register field descriptions

Field	Description
15–0 Message_Data	System-specified message.

16.11 PCI Express extended configuration space

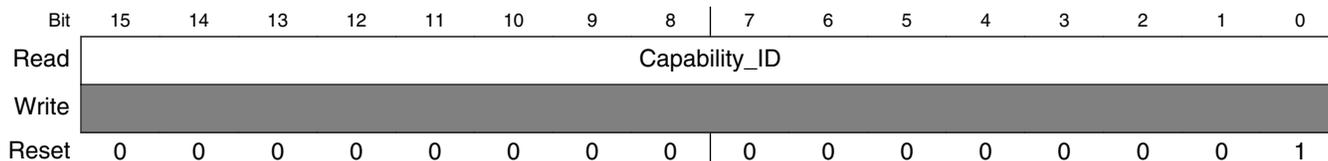
Reserved		Address Offset (Hex)
PCI-Compatible Configuration Header		000 03F
PCI-Compatible Device-Specific Configuration Space		040 0FF
Next Capability Offset (NULL)/Capability Version	Advanced Error Reporting Capability ID	100
Uncorrectable Error Status		104
Uncorrectable Error Mask		108
Uncorrectable Error Severity		10C
Correctable Error Status		110
Correctable Error Mask		114
Advanced Error Capabilities and Control		118
Header Log		11C 120 124 128
Root Error Command (RC mode only)		12C
Root Error Status (RC mode only)		130
Error Source ID	Correctable Error Source ID	134
		138 3FF
PCI Express Controller Internal CSRs*		400 6FF
		700 FFF

*Note that the PCI Express Controller Internal CSRs are not accessible by inbound PCI Express configuration transactions. Attempts to access these registers returns all 0s.

Figure 16-303. PCI Express extended configuration space

16.11.1 PCI Express Advanced Error Reporting Capability ID Register (Advanced_Error_Reporting_Capability_ID_Register)

Address: 100h

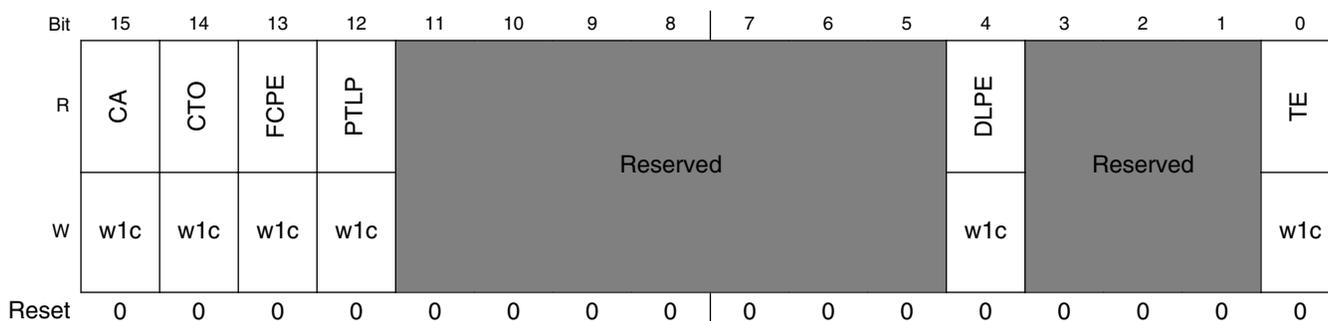
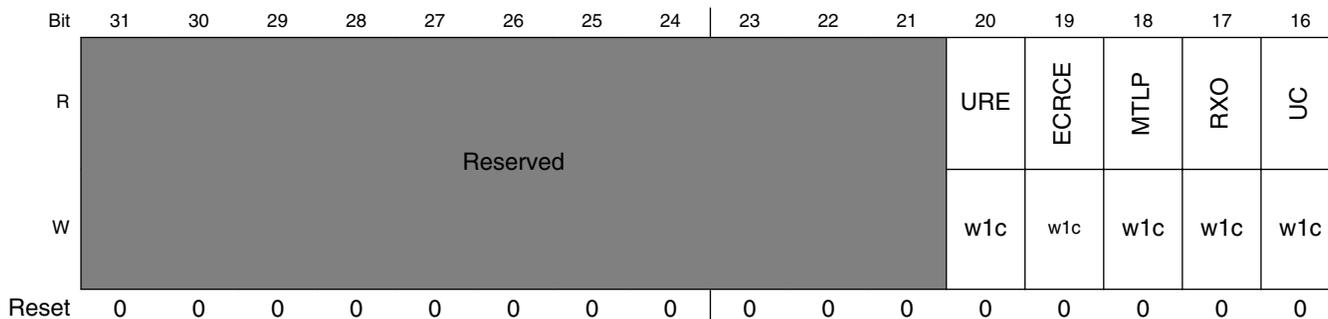


Advanced_Error_Reporting_Capability_ID_Register field descriptions

Field	Description
15-0 Capability_ID	Advanced error reporting capability = 0x0001

16.11.2 PCI Express Uncorrectable Error Status Register (Uncorrectable_Error_Status_Register)

Address: 104h



Uncorrectable_Error_Status_Register field descriptions

Field	Description
31–21 -	This field is reserved. Reserved
20 URE	Unsupported request error status.
19 ECRCE	ECRC error status.
18 MTLP	Malformed TLP status.
17 RXO	Receiver overflow status.
16 UC	Unexpected completion status.
15 CA	Completer abort status.
14 CTO	Completion timeout status. Note that a completion timeout error is a fatal error. If a completion timeout error is detected, the system has become unstable. Hot reset is recommended to restore stability of the system.
13 FCPE	Flow control protocol error status.
12 PTLP	Poisoned TLP status.
11–5 -	This field is reserved. Reserved
4 DLPE	Data link protocol error status.
3–1 -	This field is reserved. Reserved
0 TE	Training error status.

16.11.3 PCI Express Uncorrectable Error Mask Register (Uncorrectable_Error_Mask_Register)

Address: 108h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Reserved											UREM	ECRCCEM	MTLPM	RXOM	UCM
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	CAM	CTOM	FCPEM	PTLPM	Reserved						DLPEM	Reserved			TEM	
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Uncorrectable_Error_Mask_Register field descriptions

Field	Description
31–21 -	This field is reserved. Reserved
20 UREM	Unsupported request error mask
19 ECRCCEM	ECRC error mask
18 MTLPM	Malformed TLP mask
17 RXOM	Receiver overflow mask
16 UCM	Unexpected completion mask
15 CAM	Completer abort mask
14 CTOM	Completion timeout mask
13 FCPEM	Flow control protocol error mask
12 PTLPM	Poisoned TLP mask
11–5 -	This field is reserved. Reserved
4 DLPEM	Data link protocol error mask

Table continues on the next page...

Uncorrectable_Error_Mask_Register field descriptions (continued)

Field	Description
3-1 -	This field is reserved. Reserved
0 TEM	Training error mask

16.11.4 PCI Express Uncorrectable Error Severity Register (Uncorrectable_Error_Severity_Register)

Address: 10Ch

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Reserved											URES	ECRCES	MTLPS	RXOS	UCS
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	CAS	CTOS	FCPES	PTLPS	Reserved				DLPES	Reserved				TES		
W																
Reset	0	0	1	0	0	0	0	0	0	0	0	1	0	0	0	1

Uncorrectable_Error_Severity_Register field descriptions

Field	Description
31-21 -	This field is reserved. Reserved
20 URES	Unsupported request error severity
19 ECRCES	ECRC error severity
18 MTLPS	Malformed TLP severity
17 RXOS	Receiver overflow severity
16 UCS	Unexpected completion severity
15 CAS	Completer abort severity
14 CTOS	Completion timeout severity

Table continues on the next page...

Uncorrectable_Error_Severity_Register field descriptions (continued)

Field	Description
13 FCPES	Flow control protocol error severity
12 PTLPS	Poisoned TLP severity
11–5 -	This field is reserved. Reserved
4 DLPES	Data link protocol error severity
3–1 -	This field is reserved. Reserved
0 TES	Training error severity

16.11.5 PCI Express Correctable Error Status Register (Correctable_Error_Status_Register)

Address: 110h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved		ADVNE	RTTO	Reserved				RNR	BDLLP	BTLP	Reserved				RXE
W			w1c	w1c					w1c	w1c	w1c					w1c
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Correctable_Error_Status_Register field descriptions

Field	Description
31–14 -	This field is reserved. Reserved
13 ADVNF	Advisory Non-Fatal Error Status indicates the occurrence of the advisory error, and the Advisory Non-Fatal Error Mask bit in the Correctable Error Mask register is checked to determine whether to proceed further with logging and signaling
12 RTTO	Replay timer timeout status
11–9 -	This field is reserved. Reserved
8 RNR	REPLAY_NUM Rollover status
7 BDLLP	Bad DLLP status
6 BTLP	Bad TLP status
5–1 -	This field is reserved. Reserved
0 RXE	Receiver error status

16.11.6 PCI Express Correctable Error Mask Register (Correctable_Error_Mask_Register)

Address: 114h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved	Reserved	ADVNF	RTTO	Reserved	Reserved	Reserved	RNR	BDLLP	BTLP	Reserved	Reserved	Reserved	Reserved	Reserved	RXE
W	Reserved	Reserved	ADVNF	RTTO	Reserved	Reserved	Reserved	RNR	BDLLP	BTLP	Reserved	Reserved	Reserved	Reserved	Reserved	RXE
Reset	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0

Correctable_Error_Mask_Register field descriptions

Field	Description
31–14 -	This field is reserved. Reserved

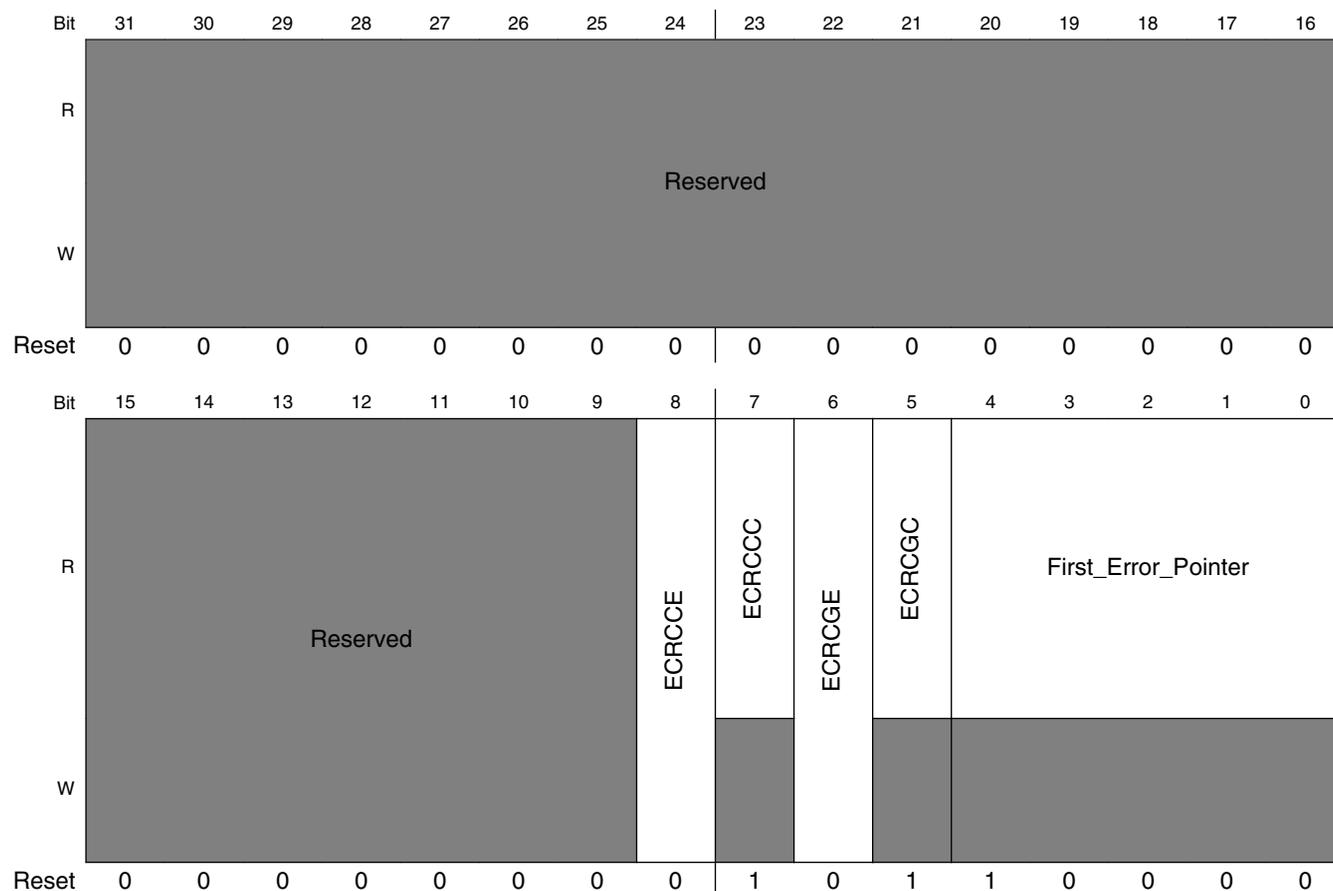
Table continues on the next page...

Correctable_Error_Mask_Register field descriptions (continued)

Field	Description
13 ADVNFEM	Advisory Non-Fatal Error Mask – This bit is Set by default to enable compatibility with software that does not comprehend Role-Based Error Reporting.
12 RTTOM	Replay timer timeout mask
11–9 -	This field is reserved. Reserved
8 RNRM	REPLAY_NUM Rollover mask
7 BDLLPM	Bad DLLP mask
6 BTLPM	Bad TLP mask
5–1 -	This field is reserved. Reserved
0 RXEM	Receiver error mask

16.11.7 PCI Express Advanced Error Capabilities and Control Register (Advanced_Error_Capabilities_and_Control_Register)

Address: 118h



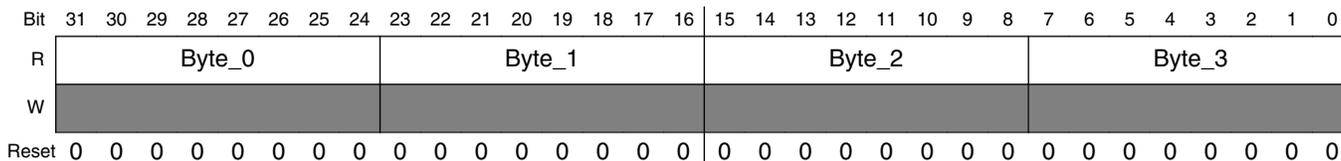
Advanced_Error_Capabilities_and_Control_Register field descriptions

Field	Description
31–9 -	This field is reserved. Reserved
8 ECRCCE	ECRC checking enable
7 ECRCCC	ECRC checking capable
6 ECRCGE	ECRC generation enable
5 ECRCGC	ECRC generation capable
4–0 First_Error_Pointer	The First Error Pointer is a read-only field that identifies the bit position of the first error reported in the uncorrectable error status register (see PCI Express Uncorrectable Error Status Register (Uncorrectable_Error_Status_Register)).

16.11.8 PCI Express Header Log Register 1 (Header_Log_Register_DWORD1)

The first DWORD of the PCI Express header log register is shown in the figure below.

Address: 11Ch



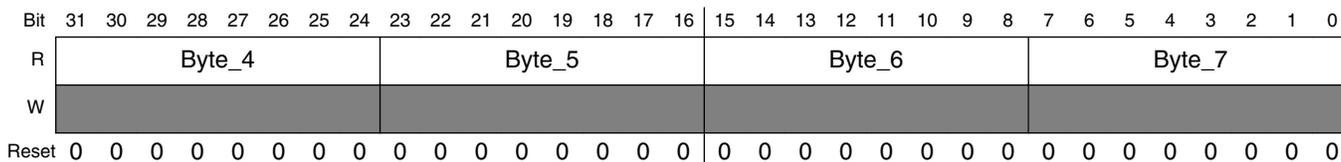
Header_Log_Register_DWORD1 field descriptions

Field	Description
31–24 Byte_0	Byte n of the TLP header associated with the error.
23–16 Byte_1	Byte n of the TLP header associated with the error.
15–8 Byte_2	Byte n of the TLP header associated with the error.
7–0 Byte_3	Byte n of the TLP header associated with the error.

16.11.9 PCI Express Header Log Register 2 (Header_Log_Register_DWORD2)

The second DWORD of the PCI Express header log register is shown in the figure below.

Address: 120h



Header_Log_Register_DWORD2 field descriptions

Field	Description
31–24 Byte_4	Byte n of the TLP header associated with the error.

Table continues on the next page...

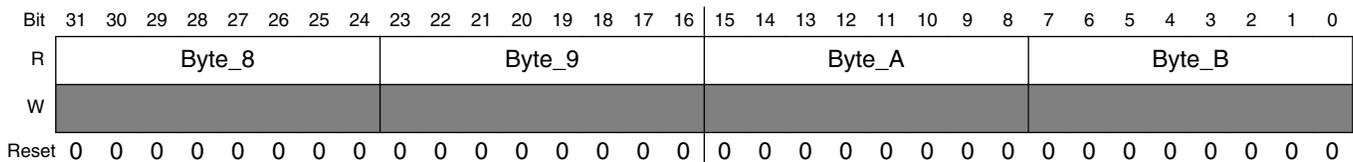
Header_Log_Register_DWORD2 field descriptions (continued)

Field	Description
23–16 Byte_5	Byte n of the TLP header associated with the error.
15–8 Byte_6	Byte n of the TLP header associated with the error.
7–0 Byte_7	Byte n of the TLP header associated with the error.

16.11.10 PCI Express Header Log Register 3 (Header_Log_Register_DWORD3)

The third DWORD of the PCI Express header log register is shown in the figure below.

Address: 124h



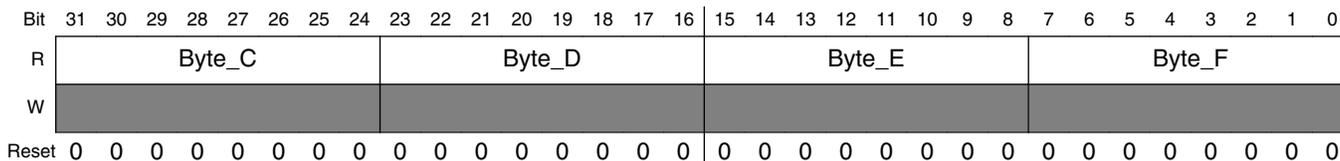
Header_Log_Register_DWORD3 field descriptions

Field	Description
31–24 Byte_8	Byte n of the TLP header associated with the error.
23–16 Byte_9	Byte n of the TLP header associated with the error.
15–8 Byte_A	Byte n of the TLP header associated with the error.
7–0 Byte_B	Byte n of the TLP header associated with the error.

16.11.11 PCI Express Header Log Register 4 (Header_Log_Register_DWORD4)

The fourth DWORD of the PCI Express header log register is shown in the figure below.

Address: 128h



Header_Log_Register_DWORD4 field descriptions

Field	Description
31–24 Byte_C	Byte n of the TLP header associated with the error.
23–16 Byte_D	Byte n of the TLP header associated with the error.
15–8 Byte_E	Byte n of the TLP header associated with the error.
7–0 Byte_F	Byte n of the TLP header associated with the error.

16.11.12 PCI Express Root Error Command Register (Root_Error_Command_Register)

This register is supported only for RC mode.

Address: 12Ch

Bit	31	30	29	28	27	26	25	24		23	22	21	20	19	18	17	16
R	Reserved																
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8		7	6	5	4	3	2	1	0
R	Reserved													FERE	NFERE	CERE	
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

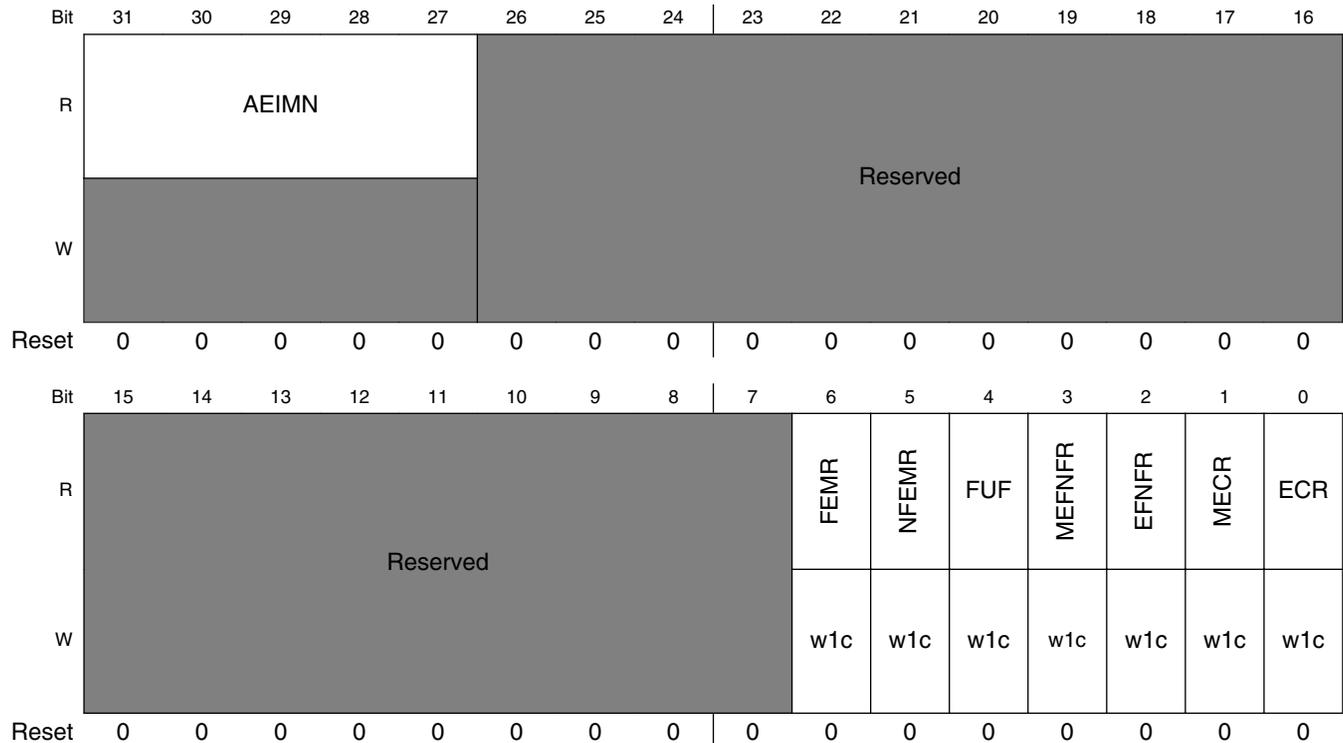
Root_Error_Command_Register field descriptions

Field	Description
31–3 -	This field is reserved. Reserved
2 FERE	Fatal error reporting enable.
1 NFERE	Non-fatal error reporting enable
0 CERE	Correctable error reporting enable

16.11.13 PCI Express Root Error Status Register (Root_Error_Status_Register)

This register is supported only for RC mode.

Address: 130h



Root_Error_Status_Register field descriptions

Field	Description
31–27 AEIMN	Advanced error interrupt message number.
26–7 -	This field is reserved. Reserved
6 FEMR	Fatal error messages received.
5 NFEMR	Non-fatal error messages received.
4 FUF	First uncorrectable fatal.
3 MEFNFR	Multiple ERR_FATAL/NONFATAL received.

Table continues on the next page...

Root_Error_Status_Register field descriptions (continued)

Field	Description
2 EFNFR	ERR_FATAL/NONFATAL received.
1 MECR	Multiple ERR_COR received.
0 ECR	ERR_COR received.

16.11.14 PCI Express Correctable Error Source ID Register (Correctable_Error_Source_ID_Register)

Address: 134h

Bit	15	14	13	12	11	10	9	8		7	6	5	4	3	2	1	0
Read	ERR_COR_Source_ID																
Write																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

Correctable_Error_Source_ID_Register field descriptions

Field	Description
15–0 ERR_COR_Source_ID	Loaded with the Requestor ID indicated in the received ERR_COR Message when the ERR_COR Received register is not already set. Default value of this field is 0.

16.11.15 PCI Express Error Source ID Register (Error_Source_ID_Register)

Address: 136h

Bit	15	14	13	12	11	10	9	8		7	6	5	4	3	2	1	0
Read	Error_Source_ID																
Write																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

Error_Source_ID_Register field descriptions

Field	Description
15–0 Error_Source_ID	ERR_FATAL/NONFATAL source ID

16.11.16 LTSSM State Status Register (LTSSM_State_Status_Register)

The PCI Express link training and status state machine (LTSSM) state status register provides detailed information about link training status. This register is useful for debugging link training failures.

NOTE

The Status Code in PEX_LTSSM_STAT changes while reacting to the link status. Therefore, software may need to read PEX_LTSSM_STAT multiple times to ensure the value has stabilized.

The following table provides the encodings for the status code field of the PEX_LTSSM_STAT register.

Table 16-317. PEX_LTSSM_STAT Status Codes

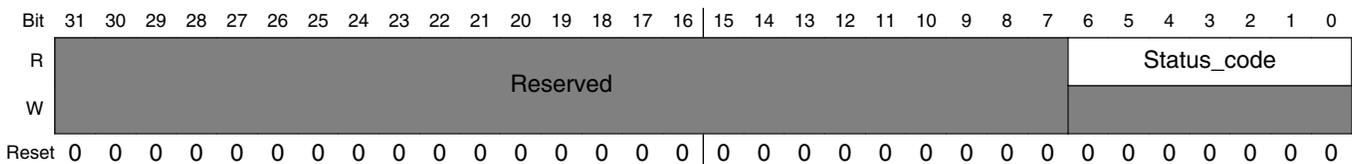
Status Code (Hex)	LTSSM State Description	Status Code (Hex)	LTSSM State Description
00	Detect quiet	27	TX L0s FTS; RX L0s FTS
01	Detect active (0)	28	L0 to L1 (0)
02	Detect active (1)	29	L0 to L1 (1)
03	Detect active (2)	2A	L1 entry
04	Polling active (0)	2B	L1 idle (0)
05	Polling active (1)	2C	L1 idle (1)
06	Polling config (0)	2D	L0 to L2 (0)
07	Polling config (1)	2E	L0 to L2 (1)
08	Polling compliance	2F	L2 entry
09	Configuration link width start (0)	30	L2 idle (0)
0A	Configuration link width start (1)	31	L2 idle (1)
0B	Configuration link width accept (0)	32	Recovery lock (0)
0C	Configuration link width accept (1)	33	Recovery lock (1)
0D	Configuration lane number wait (0)	34	Recovery lock (2)
0E	Configuration lane number wait (1)	35	Recovery cfg (0)
0F	Configuration lane number wait (2)	36	Recovery cfg (1)
10	Configuration lane number wait (3)	37	Recovery idle (0)
11	Configuration lane number accept	38	Recovery idle (1)
12	Configuration complete (0)	39	Recovery to configuration
13	Configuration complete (1)	3A	Recovery cfg to configuration
14	Configuration idle (0)	3F	L0 no training

Table continues on the next page...

Table 16-317. PEX_LTSSM_STAT Status Codes (continued)

Status Code (Hex)	LTSSM State Description	Status Code (Hex)	LTSSM State Description
15	Configuration idle (1)	7F	Detect quiet EI
16	L0	49	Configuration link width start-RC
17	TX L0; RX L0s entry	4A	Configuration link width accept-RC
18	TX L0; RX L0s idle	4B	Configuration lane number wait-RC
19	TX L0; RX L0s fast training sequence (FTS)	4C	Configuration lane number accept-RC
1A	TX L0s entry (0); RX L0	60	Loopback slave active (0)
1B	TX L0s entry (0); RX L0s idle	61	Loopback slave active (1)
1C	TX L0s entry (0); RX L0s FTS	62	Loopback slave exit
1D	TX L0s entry (1); RX L0	68	Hot reset (0)
1E	TX L0s entry (1); RX L0s idle	69	Hot reset (1)
1F	TX L0s entry (1); RX L0s FTS	6A	Hot reset (0)-RC
20	TX L0s idle; RX L0	6B	Hot reset (1)-RC
21	TX L0s idle; RX L0s entry	75	Disabled (0)
22	TX L0s idle; RX L0s idle	71	Disabled (1)
23	TX L0s idle; RX L0s FTS	72	Disabled (2)
24	TX L0s FTS; RX L0	73	Disabled (3)
25	TX L0s FTS; RX L0s entry	74	Disabled (4)
26	TX L0s FTS; RX L0s idle	78	L0 to L1/L2-RC

Address: 404h



LTSSM_State_Status_Register field descriptions

Field	Description
31-7 -	This field is reserved. Reserved
6-0 Status_code	Status code. See Table 16-317 for encodings.

16.11.17 PCI Express Controller Core Clock Ratio Register (Controller_Core_Clock_Ratio_Register)

The PCI Express controller core clock ratio register is used to program the ratio of the actual PCI Express controller clock frequency to the default controller core frequency (333 MHz). This is required only when a PCI Express controller clock frequency other than the default 333 MHz has to be used.

As an example of programming PEX_GCLK_RATIO, consider the case where the actual PCI Express controller clock is 250 MHz, the ratio of the actual clock to the default clock (333 MHz) is 3:4. that is, the default core clock has to be multiplied by the ratio (3/4, which is equivalent to 12/16). So the register has to be programmed with the decimal numerator value 12 or 0x0000_000C.

Address: 440h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved																Clock_Ratio_Numerator															
W	Reserved																Clock_Ratio_Numerator															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0

Controller_Core_Clock_Ratio_Register field descriptions

Field	Description
31–6 -	This field is reserved. Reserved
5–0 Clock_Ratio_Numerator	The numerator of the ratio of the actual PCI Express controller clock frequency used to the default core clock frequency of 333 MHz . The denominator of the ratio is fixed at 16. The default value of this register is 0x10 (16 decimal), which corresponds to a ratio of 1:1 (or 16/16)

16.11.18 PCI Express Power Management Timer Register (Power_Management_Timer_Register)

The PCI Express power management timer register is used to program the time-in values for entering L0s and L1 power management states.

Address: 450h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved								L1_WAIT_PERIOD								L0s_TIME_IN															
W	Reserved								L1_WAIT_PERIOD								L0s_TIME_IN															
Reset	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	1	1	0	1	0	1	1	1	1	1	0	0	1	1	1	0

Power_Management_Timer_Register field descriptions

Field	Description
31–24 -	This field is reserved. Reserved

Table continues on the next page...

Power_Management_Timer_Register field descriptions (continued)

Field	Description
23–12 L1_WAIT_PERIOD	<p>Wait period (in PCI Express controller core clock cycles) before entering L1 power state after all functions are in a non-D0 power state. The value is calculated as:</p> <p>Time (in μsec) x PCI Express controller core clock frequency (in MHz)</p> <p>The time value must be less than 2 μsec; the default value (0x14D) is 1 μsec for the default clock frequency of 333 MHz .</p>
11–0 L0s_TIME_IN	<p>Time in value (in PCI Express controller core clock cycles) for entering L0s power state. The value is calculated as:</p> <p>Time (in μsec) x PCI Express controller core clock frequency (in MHz)</p> <p>The maximum time value is 7 μsec; the default value (0x7CE) is 6 μsec for the default clock frequency of 333 MHz .</p>

16.11.19 PCI Express PME Time-Out Register (PME_Time_Out_Register)

This register is supported only for EP mode.

The PCI Express PME time-out register is used to program the time-out value that the controller uses before re-sending a PME message to the host. If PME is requested by a function and the host does not clear the associated PME_STAT bit even after this time-out has expired, the PME message is sent again to the host by the PCI Express controller.

Address: 454h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R																																
W																																
Reset	0	0	0	0	0	0	0	1	1	1	1	1	1	1	0	0	0	0	0	1	1	1	1	0	0	0	1	0	0	0	0	0

PME_Time_Out_Register field descriptions

Field	Description
31–26 -	This field is reserved. Reserved
25–0 PME_TIMEOUT	<p>The PME time-out value specifies the interval before PME messages are resent by the controller, provided the PME_STAT bit in the PCI Express power management status and control register (offset 0x48) is not cleared by the host. The value for PME_TIMEOUT is specified in terms of PCI Express controller core clock cycles. The value is calculated as:</p> <p>Time (in μsec) x PCI Express controller core clock frequency (in MHz)</p> <p>The minimum time value is 100 msec; the default value (0x1FC1E20) is 100 msec for the default clock frequency of 333 MHz .</p>

16.11.20 PCI Express Subsystem Vendor ID Update Register (Subsystem_Vendor_ID_Update_Register)

This register is supported only for EP mode.

The PCI Express subsystem vendor ID update register is used to set the values for the Subsystem ID and Subsystem Vendor ID registers in the Type 0 configuration header.

When used as an endpoint, the controller's initialization software programs the desired subsystem ID and subsystem vendor ID values in PEX_SSVID_UPDATE before setting the CFG_READY bit in the PEX_CFG_READY register (see [Configuration Ready Register \(Configuration_Ready_Register\)](#)). That way, when the host begins system enumeration, the correct values are present in the Type 0 configuration header.

Address: 478h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R																																
W	SS_ID																SSV_ID															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Subsystem_Vendor_ID_Update_Register field descriptions

Field	Description
31–16 SS_ID	Subsystem ID [15-0] value
15–0 SSV_ID	Subsystem vendor ID [15-0] value

16.11.21 Configuration Ready Register (Configuration_Ready_Register)

The PCI Express configuration ready register is used to indicate configuration complete status to the transaction layer. The transaction layer handles configuration requests from external hosts only after the CFG_READY bit is set. All the configuration requests received from external hosts before the CFG_READY bit is set are completed with configuration request retry status (CRS). The CFG_READY bit in this register should be set after all relevant configuration registers have been programmed. This makes sure the external host reads the correct capability advertisements during enumeration.

PCI Express extended configuration space

Address: 4B0h

Bit	31	30	29	28	27	26	25	24		23	22	21	20	19	18	17	16
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8		7	6	5	4	3	2	1	0
R	Reserved															CFG_READY	
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	*

* Notes:

- CFG_READY field: Defined during POR. See description for details.

Configuration_Ready_Register field descriptions

Field	Description
31–1 -	This field is reserved. Reserved
0 CFG_READY	Configuration ready Note that the reset state of this bit is determined during POR. 1 The transaction layer accepts inbound configuration requests. 0 The transaction layer responds to all inbound configuration requests with retry (CRS)

16.11.22 Flow Control Update Timeout Register (Flow_Control_Update_Timeout_Register)

The PCI Express flow control update timeout register is used to program the timeout value for update-FC DLLP reception in terms of PCI Express controller core clock cycles.

Address: 4B8h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	
R	Reserved															FC_UPDATE_TIMEOUT	
W																	
Reset	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
R	FC_UPDATE_TIMEOUT																
W																	
Reset	0	0	0	0	0	1	0	0	0	0	1	0	1	0	0	0	

Flow_Control_Update_Timeout_Register field descriptions

Field	Description
31 DIS_LR	<p>Disable link retraining on FC update timeout error.</p> <p>The PCI Express protocol allows for optional link retraining by a device if FC updates are not received from the remote device within the timeout period programmed in bits [17-0] of this register. If the FC updates were lost due to link error issues, this retraining will help to restore normal operation. By default, this retraining is disabled. Since this is an optional feature, the user is allowed to enable or disable such automatic retraining, if desired.</p> <p>1 Disable link retraining on FC update timeout error. 0 Enable link retraining on FC update timeout error.</p>
30–18 -	This field is reserved. Reserved
17–0 FC_UPDATE_TIMEOUT	<p>FC update timeout value specifies the interval (in PCI Express controller core clock cycles) to wait for the reception of consecutive FC-update DLLPs before indicating a flow-control protocol error. The value is calculated as:</p> <p>Time (in μsec) x PCI Express controller core clock frequency (in MHz)</p> <p>The maximum time value is 200 μsec; the default value (0x10428) is 200 μsec for the default clock frequency of 333 MHz .</p>

16.11.23 Secondary Status Interrupt Mask Register (Secondary_Status_Interrupt_Mask_Register)

This register is supported only for RC mode.

The PCI Express secondary status interrupt mask register can be used to disable sideband interrupt generation when error bits in the PCI Express secondary status register are set. See [PCI Express Secondary Status Register \(Secondary_Status_Register\)](#), for more information. By default, interrupt generation due to secondary status errors is disabled.

Address: 5A0h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved										M_DPE	M_SSE	M_RMA	M_RTA	M_STA	M_MDPE
W	Reserved										1	1	1	1	1	1
Reset	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1

Secondary_Status_Interrupt_Mask_Register field descriptions

Field	Description
31–6 -	This field is reserved. Reserved
5 M_DPE	Mask detected parity error. Default value is 1 1 Mask enabled (interrupt cannot be generated) 0 Mask disabled
4 M_SSE	Mask signaled system error. Default value is 1 1 Mask enabled (interrupt cannot be generated) 0 Mask disabled
3 M_RMA	Mask received master abort. Default value is 1 1 Mask enabled (interrupt cannot be generated) 0 Mask disabled
2 M_RTA	Mask received target abort. Default value is 1 1 Mask enabled (interrupt cannot be generated) 0 Mask disabled

Table continues on the next page...

Secondary_Status_Interrupt_Mask_Register field descriptions (continued)

Field	Description
1 M_STA	Mask signaled target abort. Default value is 1 1 Mask enabled (interrupt cannot be generated) 0 Mask disabled
0 M_MDPE	Mask master data parity error. Default value is 1 1 Mask enabled (interrupt cannot be generated) 0 Mask disabled

16.12 Functional description

The PCI Express protocol relies on a requestor/completer relationship where one device requests that some desired action be performed by some target device and the target device completes the task and responds.

Usually the requests and responses occur through a network of links, but to the requestor and to the completer, the intermediate components are transparent.



Figure 16-327. Requestor/completer relationship

Each PCI Express device is divided into two halves-transmit (TX) and receive (RX), and each of these halves is further divided into three layers-transaction, data link, and physical-as shown in [Figure 16-328](#).

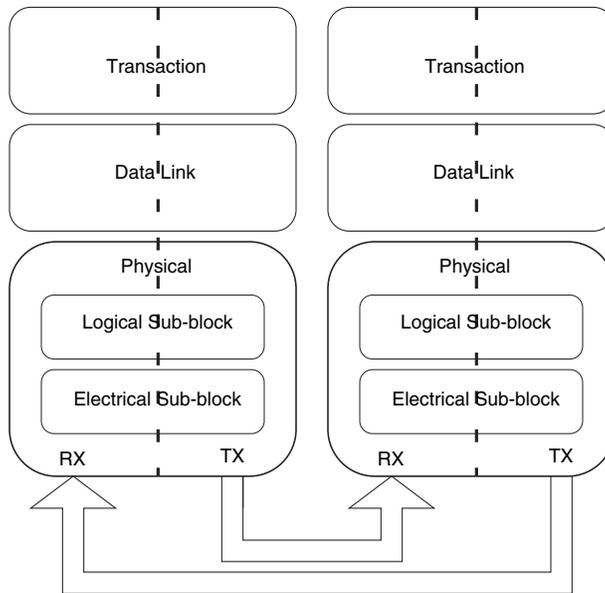


Figure 16-328. PCI Express high-level layering

Packets are formed in the transaction layer (TLPs) and data link layer (DLLPs), and each subsequent layer adds the necessary encodings and framing—as shown in [Figure 16-329](#). As packets are received, they are decoded and processed by the same layers but in reverse order, so they may be processed by the layer or by the device application software.

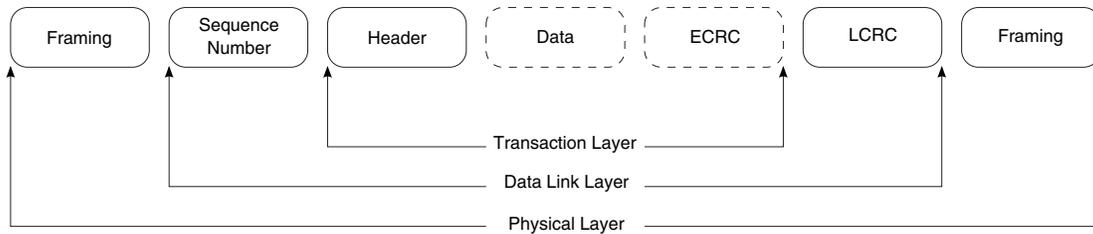


Figure 16-329. PCI Express packet flow

16.12.1 Architecture

This section describes implementation details of the PCI Express controller.

16.12.1.1 PCI Express transactions

Table 16-325 contains the list of transactions that the PCI Express controller supports as an initiator and a target.

Table 16-325. PCI Express transactions

PCI Express transaction	Supported as an Initiator	Supported as a target	Definition
Mrd	Yes	Yes	Memory read request
MRdLk	No	No	Memory read lock. As a target, CplLk with UR status is returned.
MWr	Yes	Yes	Memory write request to memory-mapped PCI-Express space
IORd	Yes (RC only)	No	I/O Read request. As a target, Cpl with UR status is returned.
IOWr	Yes (RC only)	No	I/O write request. As a target, Cpl with UR status is returned.
CfgRd0	Yes (RC only)	Yes	Configuration read type 0
CfgWr0	Yes (RC only)	Yes	Configuration write type 0
CfgRd1	Yes (RC only)	No	Configuration read type 1. As a target, Cpl with UR status is returned.
CfgWr1	Yes (RC only)	No	Configuration write type 1. As a target, Cpl with UR status is returned.
Msg	Yes	Yes	Message request
MsgD	Yes (RC only)	Yes (EP only)	Message request with data payload. Note that Set_Slot_Power_Limit is the only message with data that is supported and then only when the controller is an initiator and in RC mode or a target and in EP mode.
Cpl	Yes	Yes	Completion without data
CplD	Yes	Yes	Completion with data
CplLk	No	Yes	Completion for locked memory read without data. The only time that CplLk is returned with UR status is when the controller receives a MRdLk command.
CplDLk	No	No	Completion for locked memory read with data

16.12.1.2 Byte ordering

Whenever data cross a bridge between two busses, the byte ordering of data on the source and destination buses must be considered.

The internal platform bus of this device is inherently big endian and the PCI Express bus interface is inherently little endian.

There are two methods to handle ordering of data as it crosses a bridge—address invariance and data invariance. Address invariance preserves the addressing of bytes within a scalar data element, but not the relative significance of the bytes within that scalar. Conversely, data invariance preserves the relative significance of bytes within a scalar, but not the addressing of the individual bytes that make up a scalar.

This device uses address invariance as its byte ordering policy.

16.12.1.2.1 Address invariance

As stated above, address invariance preserves the byte address of each byte on an I/O interface as it is placed in memory or moved into a register.

This policy can have the effect of reversing the significance order of bytes (most significant to least significant and vice versa), but it has the benefit of preserving the format of general data structures. Provided that software is aware of the endianness and format of the data structure, it can correctly interpret the data on either side of the bridge.

The following figure shows the transfer of a 4-byte scalar, 0x4142_4344, from a big endian source across an address invariant bridge to a little endian destination.

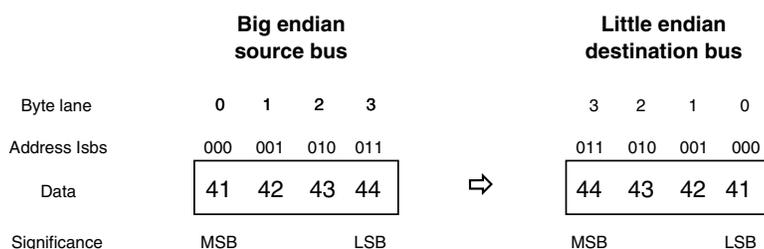


Figure 16-330. Address invariant byte ordering—4 bytes outbound

Note that although the significance of the bytes within the scalar have changed, the address of the individual bytes that make up the scalar have not changed. As long as software is aware that the source of the data used a big endian format, the data can be interpreted correctly.

The following figure shows data flowing the other way, from a little endian source to a big endian destination.

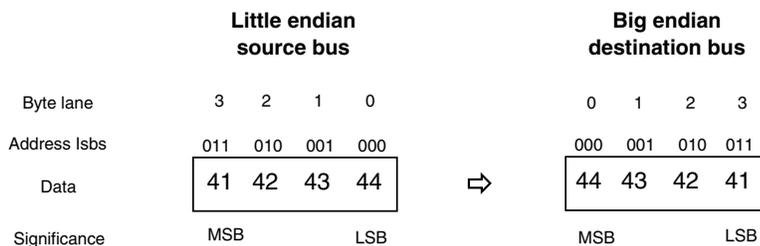


Figure 16-331. Address invariant byte ordering—4 bytes inbound

The following figure shows an outbound transfer of an 8-byte scalar, 0x5455_1617_CDCE_2728, using address invariance.

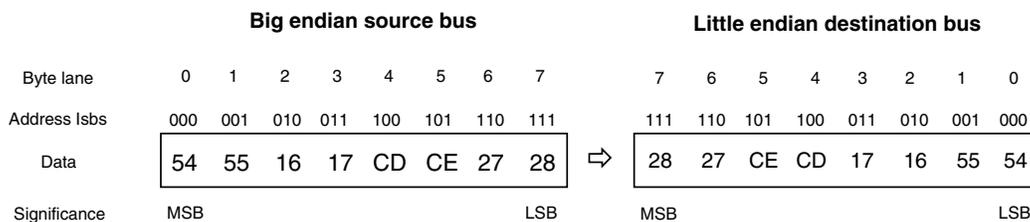


Figure 16-332. Address invariant byte ordering—8 bytes outbound

The following figure shows an inbound transfer of a 2-byte scalar, 0x5837, using address invariance.

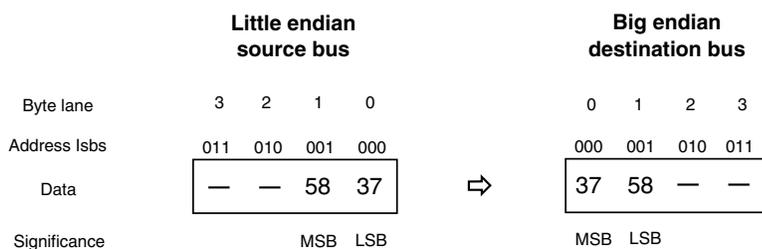


Figure 16-333. Address invariant byte ordering—2 bytes inbound

Note that in all of these examples, the original addresses of the individual bytes within the scalars (as created by the source) have been preserved.

16.12.1.2.2 Byte order for configuration transactions

All internal memory-mapped registers in the CCSR space use big endian byte ordering. However, the PCI Express specification defines PCI Express configuration registers as little endian.

All accesses to the PCI Express configuration port, PEX_CONFIG_DATA, including the those targeting the internal PCI Express configuration registers, use the address invariance policy as shown in [Figure 16-334](#). Therefore, software must access PEX_CONFIG_DATA with little-endian formatted data-either using the **lwbrx/stwbrx** instructions or by manipulating the data before writing to and after reading from PEX_CONFIG_DATA.

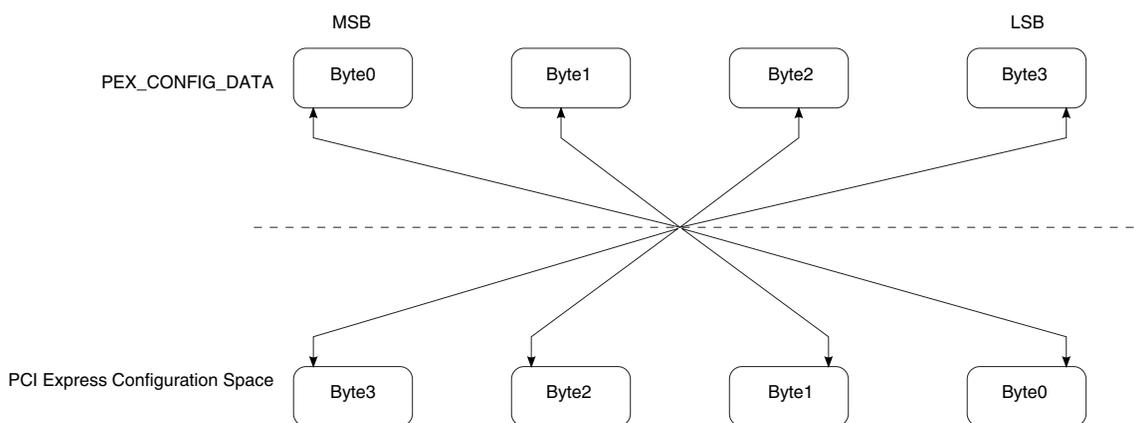


Figure 16-334. PEX_CONFIG_DATA byte ordering

16.12.1.3 Transaction ordering rules

In general, transactions are serviced in the order that they are received.

However, transactions can be reordered as they are sent due to a stalled condition such as a full internal buffer. [Table 16-326](#) describes the transaction ordering rules for this device.

Table 16-326. PCI Express controller TLP transaction ordering rules

Row pass column?		Posted request	Non-posted request		Completion	
		Memory write or message request (col 2)	Read request (Col 3)	I/O or Configuration write request (col 4)	Read completion (Col 5)	I/O or configuration write completion (col 6)
Posted request	Memory write or message request (row A)	a) No ¹ b) No ¹	Yes	Yes	a) Yes ² b) N/A ³	a) Yes ² b) N/A ³
Non-posted request	Read request (row B)	No	No	No	a) No ⁴ b) Yes ⁵	a) No ⁴ b) Yes ⁵
	I/O or configuration write request (Row C)	No	No	No	a) No ⁴ b) Yes ⁵	a) No ⁴ b) Yes ⁵
Completion	Read completion (row D)	a) No ⁶ b) Yes ⁷	Yes	Yes	a) No ⁸ b) No ⁸	No
	I/O or configuration write completion (row E)	a) No ⁶ b) Yes ⁷	Yes	Yes	No	No

1. Regardless of the setting of the relaxed ordering (RO) bit, a posted request cannot bypass another posted request.
2. Regardless of the setting of the relaxed ordering bit, a posted request can always bypass a completion.
3. N/A indicates that the original rules at these entries defined by the *PCI Express Base Specification* do not apply to RC or EP.
4. A non-posted request cannot bypass a completion if the relaxed ordering bit is cleared (that is, RO = 0).
5. A non-posted request can bypass a completion if the relaxed ordering bit is set (that is, RO = 1).
6. A read completion, I/O write completion, or configuration write completion cannot bypass a posted request if the relaxed ordering bit is cleared (that is, RO = 0).
7. A read completion, I/O write completion, or configuration write completion can bypass a posted request if the relaxed ordering bit is set (that is, RO = 1).
8. Regardless of the setting of the relaxed ordering bit, a read completion cannot bypass another read completion.

In general, the following points summarize the ordering rules for sending the next outstanding request:

- A posted request can bypass all other transactions except another posted request.
- A non-posted request cannot bypass posted or other non-posted requests, but it can bypass a completion if the relaxed ordering (RO) bit is set. (See [Table 16-326](#)).
- A completion can bypass posted requests if the relaxed ordering (RO) bit is set and can bypass non-posted transactions. However, a completion cannot bypass other completions.

16.12.1.4 PCI Express outbound ATMUs

Outbound transactions should hit into one of the outbound ranges defined by the outbound ATMUs. The outbound address translation windows must be aligned based on the granularity selected by the size fields. Outbound window misses use the default outbound register set (outbound ATMU window 0).

Overlapping outbound windows (1-4) are not supported and will cause undefined behavior. Note that for RC mode, all outbound transactions post ATMU must hit either into the memory base/limit range or the prefetchable memory base/limit range defined in the PCI Express type 1 header. For EP mode, there is no such requirement.

Note that in RC mode, there is no checking on whether the translated address actually hits into the memory base/limit range. It will just pass it through as is.

Figure 16-335 shows the outbound transaction flow.

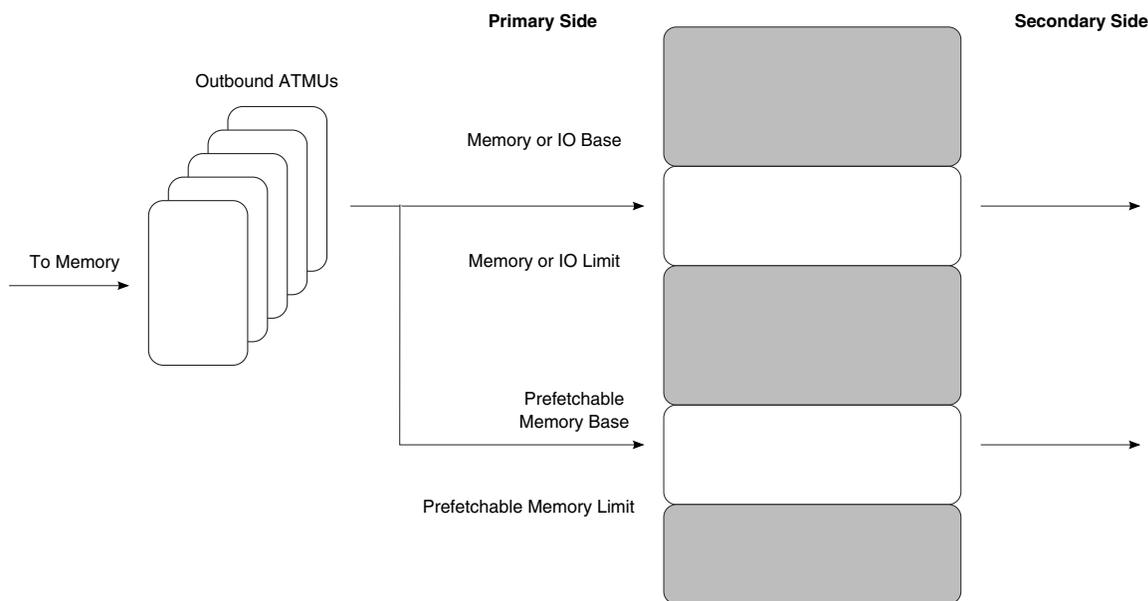


Figure 16-335. RC outbound transaction flow

16.12.1.5 PCI Express inbound ATMUs

There are differences between RC and EP implementations of the inbound ATMU registers as described in the following sections.

16.12.1.5.1 EP inbound ATMU implementation

All base address registers (BARs) reside in the PCI Express type 0 configuration header space which is accessible through the PEX_CONFIG_ADDR/PEX_CONFIG_DATA mechanism.

Note that host software must program these BARs using configuration type 0 cycles. There are 4 inbound BARs.

- Inbound window BAR0, also known as PEXCSRBAR, at configuration address 0x10 (32-bit). This is a fixed 1-Mbyte window used for inbound memory transactions that access memory-mapped registers.
- Inbound window BAR1 at configuration address 0x14 (32-bit)
- Inbound window BAR2 at configuration address 0x18-0x1c (64-bit)
- Inbound window BAR3 at configuration address 0x20-0x24 (64-bit)

The PCI Express controller does not implement a shadow of the inbound BARs in the memory-mapped register set. However, when there is a hit to the BAR(s), the PCI Express controller uses the corresponding translation and attribute registers in the memory-mapped register set (PEXITAR_n and PEXIWAR_n) for the translation. If the transaction hits multiple BARs, then the lowest-numbered BAR is used.

16.12.1.5.2 RC inbound ATMU implementation

In RC mode, inbound window BAR0/PEXCSRBAR is the only inbound BAR that resides in the PCI Express type 1 configuration header (at offset 0x10); the remaining inbound window BARs (PEXIWBAR[1-3]) reside outside of the type 1 header in the memory-mapped register space.

If the transaction hits any window, the translation is performed and then the transaction is sent to memory. If there is no hit to any one of the BARs, then an UR completion is returned for non-posted transactions. All posted transactions with no BAR hit are ignored.

Figure 16-336 shows the inbound transaction flow in RC mode.

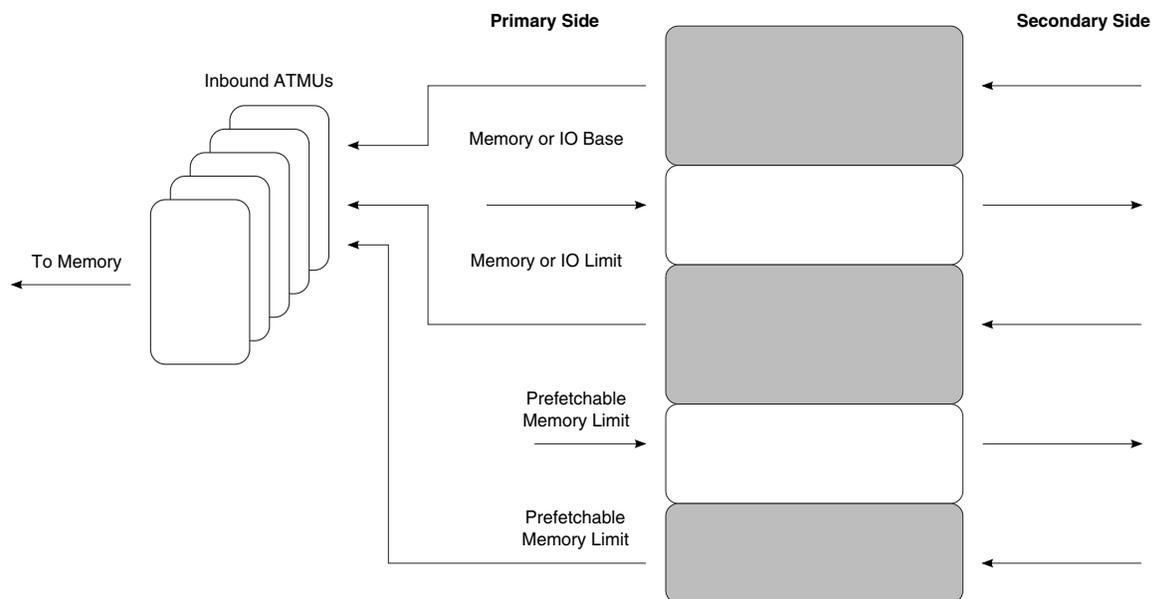


Figure 16-336. RC inbound transaction flow

16.12.1.6 Memory space addressing

A PCI Express memory transaction can address a 32- or 64-bit memory space.

The FMT[0] field in the PCI Express TLP header for a 32-bit address packet is 0; a 64-bit address packet has a FMT[0] = 1. The PCI Express TLP header for a memory read transaction has TYPE[4:0] = 00000 and FMT[1] = 0. A memory write transaction has TYPE[4:0] = 00000 and FMT[1] = 1. As an initiator, the controller is capable of sending 32- or 64-bit memory packets. Any transaction from the internal platform that (after passing through the translation mechanism) has a translated address greater than 4G is sent as a 64-bit memory packet. Otherwise, a 32-bit memory packet is sent. As a target device, the controller is capable of decoding 32- or 64-bit memory packets. This is done through two 32-bit inbound windows and two 64-bit inbound windows. All inbound addresses are translated to 36-bit internal platform addresses.

16.12.1.7 I/O space addressing

The controller does not support I/O transactions as a target. As an initiator, the controller can send I/O transactions in RC mode only.

This can be done by programming one of the outbound translation window's attribute to send I/O transactions. All I/O transactions only access 32-bit address I/O space. The PCI Express TLP header for an I/O read transaction has TYPE[4:0] = 00010 and FMT[1] = 0. The PCI Express TLP header for an I/O write transaction has TYPE[4:0] = 00010 and FMT[1] = 1.

16.12.1.8 Configuration space addressing

As an initiator, the controller supports both type 0 and type 1 configuration cycles when configured in RC mode. There are two methods of generating a configuration transaction; refer to [PCI Express configuration space access](#), for more information.

A configuration transaction can hit into the controller's internal configuration space, it can be sent out on the PCI Express link, or it can be internally terminated.

All configuration transactions sent on PCI Express require a response regardless whether they are read or a write configuration transactions. [Table 16-327](#) shows PCI Express TLP header configuration.

Table 16-327. PCI Express TLP header configuration

TLP header type	Configuration read transaction	Configuration write transaction
Type 0	TYPE[4:0]= 00100 and FMT[1] = 0	TYPE[4:0] = 00100 and FMT[1] = 1
Type 1	TYPE[4:0] = 00101 and FMT[1] = 0	TYPE[4:0] = 00101 and FMT[1] = 1

The controller does not generate configuration transactions in EP mode. Only inbound configuration transactions are supported in EP mode.

16.12.1.9 PCI Express configuration space access

PCI Express configuration space access depends on whether the device is in RC or EP mode, as described in the following sections.

16.12.1.9.1 RC configuration register access

In RC mode, there are two methods for accessing the external PCI Express configuration space:

- PCI Express configuration access registers (PEX_CONFIG_ADDR/ PEX_CONFIG_DATA)
- PCI Express outbound ATMU window

To access internal configuration space, software must rely on the PCI Express configuration access register (PEX_CONFIG_ADDR/ PEX_CONFIG_DATA) mechanism. To access external configuration space, software can either use configuration access registers or the outbound ATMU mechanism. For the configuration access register method, a value must be written to the PEX_CONFIG_ADDR register that specifies the targeted PCI Express bus, the targeted device on that bus, the targeted function within the device, and the configuration register in that device that should be accessed. The PCI Express controller's bus number is obtained from the PCI Express configuration header (type 1). Then either a write or a read to the PEX_CONFIG_DATA register triggers the actual write or read cycle to the configuration space.

NOTE

Accesses to the little-endian PCI Express configuration space must be properly formatted. See [Byte order for configuration transactions](#), for more information.

External configuration transactions should not be attempted until the link has successfully trained. Software can poll the LTSSM state status register (PEX_LTSSM_STAT) to check the status of link training before issuing external configuration requests.

Refer to [Configuration accesses and inbound writes to CCSR space](#), for special considerations regarding outbound configuration accesses and inbound writes to CCSR space.

16.12.1.9.1.1 PCI Express configuration access register mechanism

There are two types of configuration transactions (Type 0 and Type 1) needed to support hierarchical bridges.

- If the targeted bus number, and targeted device number equal to the PCI Express controller's bus number and device number, and the targeted function number is zero, then an internal PCI Express configuration cycle access is performed.
- If the targeted bus number does not equal the PCI Express controller's bus number, but does equal the secondary bus number (from the type 1 header) and the targeted device number is 0, then a Type 0 configuration transaction is sent to the PCI Express link.
- If the targeted bus number does not equal the PCI Express controller's bus number, and does not equal the secondary bus number (from the type 1 header), and the targeted bus number is less than or equal to the subordinate bus number (from the

type 1 header), then a Type 1 configuration transaction is sent to the PCI Express link.

- If none of the above conditions occur, then the PCI Express controller returns all 1s for reads and ignores writes. In this case, PEX_ERR_DR[ICCA] is set.

16.12.1.9.1.2 Outbound ATMU configuration mechanism (RC-only)

Software can also program one of the outbound ATMU windows to perform a configuration access. This is accomplished by programming the ReadTType or WriteTType field of the desired PEXOWAR to 0x2. Software must only issue 4-byte or less access to the ATMU configuration window and the access cannot cross a 4-byte boundary. The targeted bus number, targeted device number, targeted function number, register, and targeted extended register number sent are decoded from the outbound translated PCI Express address.

- targeted bus number[7:0] = PCI Express address[27:20]
- targeted device number[4:0] = PCI Express address[19:15]
- targeted function number[2:0] = PCI Express address[14:12]
- targeted extended register number[3:0] = PCI Express address[11:8]
- targeted register number[5:0] = PCI Express address[7:2]

A Type 0 configuration cycle is sent to the link if the targeted bus number equals the secondary bus number (from the type 1 header) and targeted device number is 0. A Type 1 configuration cycle is sent to the link if targeted bus number does not equal primary bus and secondary bus numbers and it is less than or equal to the subordinate bus number (from the type 1 header). For all other cases, the PCI Express controller squashes the write and read results in a response with error returned. In this case, PEX_ERR_DR[IACA] is set.

Note that the PCI Express controller does not support access to its internal configuration registers using the outbound ATMU mechanism. That is, the outbound ATMU mechanism must not be used to program the internal registers.

16.12.1.9.2 EP configuration register access

When the PCI Express controller is configured as an EP device it responds to remote host generated configuration cycles.

This is indicated by decoding the configuration command along with type 0 access in the packet. A remote host can access all of the PCI Express configuration area except the PCI Express controller internal CSR registers in the extended PCI Express configuration

space at offsets 0x400-0x6FF. The PCI Express controller internal CSR registers are not accessible by inbound PCI Express configuration transactions. Attempts to access these registers return all zeros.

While in EP mode, the PCI Express controller does not support generating configuration accesses as a master. All accesses to PEX_CONFIG_ADDR/PEX_CONFIG_DATA cause the device to access the internal configuration registers regardless of the targeted bus number or targeted device number programmed in the PEX_CONFIG_ADDR register. There is no configuration mechanism supported in EP mode using the ATMU window. If the outbound ATMU window is configured to issue a configuration transaction, all posted transactions hitting this window are ignored and all non-posted transactions get a response with an error.

16.12.1.10 Serialization of configuration and I/O writes

Configuration and I/O writes are serialized by the controller.

The logic after issuing a configuration write or IO write does not issue any new transactions until the outstanding configuration or I/O write is finished. This means that an acknowledgement packet from the link partner in the form of a CpL TLP packet must be seen or the transaction has timed out. If the CpL packet contains a CRS status, then the logic re-issues the configuration write transaction. It keeps retrying the request until either a status other than CRS is returned or the transaction times out.

Note that it is possible for outbound configuration read request to be requeued and be placed at the end of the request queue due to CRS condition.

16.12.1.11 Messages

Software message generation is supported in both RC and EP modes.

16.12.1.11.1 Outbound ATMU message generation

Software can choose to send a message by programming $PEXOWAR_n[WTT] = 0x5$.

A message is sent by writing a 4-byte transaction in big-endian format that hits in an outbound window configured to send messages.

Part of the 4-byte data is used to store information such as message code and routing information. [Table 16-328](#) describes the message data format.

Table 16-328. Internal platform Message data format

Bits	Name	Reset value	Description
0-15	-	-	Reserved
16-18	Routing	x	Routing mechanism. Contains the message's routing information
19-23	-	-	Reserved
24-31	Code	x	Message code. Contains the actual message type to be sent.

In addition to the outbound ATMU, the PEX PM Command register also provides the capability to send PME_Turn_Off message or PM_PME message by setting bits 31 or 29. See [PCI Express power management command register \(PEX_PEX_PMCR\)](#) for more information.

[Table 16-329](#) provides a complete list of supported outbound messages depending on whether RC or EP is configured.

Table 16-329. PCI Express ATMU outbound messages

Name	Code[7:0]	Routing[2:0]	RC	EP	Description
Assert_INTA	0010 0000	100	N/A	Yes	Sent upstream by endpoint
Assert_INTB	0010 0001	100	N/A	Yes	Sent upstream by endpoint
Assert_INTC	0010 0010	100	N/A	Yes	Sent upstream by endpoint
Assert_INTD	0010 0011	100	N/A	Yes	Sent upstream by endpoint
Deassert_INTA	0010 0100	100	N/A	Yes	Sent upstream by endpoint
Deassert_INTB	0010 0101	100	N/A	Yes	Sent upstream by endpoint
Deassert_INTC	0010 0110	100	N/A	Yes	Sent upstream by endpoint
Deassert_INTD	0010 0111	100	N/A	Yes	Sent upstream by endpoint
PM_Active_State_Nak	0001 0100	100	Yes	N/A	Terminate at receiver
PM_PME	0001 1000	000	N/A	Yes	Sent upstream by PME-requesting component
PME_Turn_Off	0001 1001	011	Yes	N/A	Broadcast downstream
PM_TO_Ack	0001 1011	101	N/A	Yes	Sent upstream by endpoint
ERR_COR	0011 0000	000	N/A	Yes	Sent by component when it detects a correctable error
ERR_NONFATAL	0011 0001	000	N/A	Yes	Sent by component when it detects a Non-fatal, uncorrectable error
ERR_FATAL	0011 0011	000	N/A	Yes	Sent by component when it detects a Fatal, uncorrectable error
Unlock	0000 0000	000	No	N/A	Not supported
Set_Slot_Power_Limit	0101 0000	100	Yes	N/A	Set slot power limit in upstream port
Vendor_Defined Type 0	0111 1110		No	No	Not supported
Vendor_Defined Type 1	0111 1111		No	No	Not supported
Attention_Indicator_On	0100 0001	100	Yes	N/A	Hot-plug message
Attention_Indicator_Blink	0100 0011	100	Yes	N/A	Hot-plug message

Table continues on the next page...

Table 16-329. PCI Express ATMU outbound messages (continued)

Name	Code[7:0]	Routing[2:0]	RC	EP	Description
Attention_Indicator_Off	0100 0000	100	Yes	N/A	Hot-plug message
Power_Indicator_On	0100 0101	100	Yes	N/A	Hot-plug message
Power_Indicator_Blink	0100 0111	100	Yes	N/A	Hot-plug message
Power_Indicator_Off	0100 0100	100	Yes	N/A	Hot-plug message
Attention_Button_Pressed	0100 1000	100	N/A	Yes	Hot-plug message

16.12.1.11.2 Inbound messages

Table 16-330 provides a complete list of supported inbound messages in RC mode.

Table 16-330. PCI Express RC inbound message handling

Name	Code[7:0]	Routing[2:0]	Action
Assert_INTA	0010 0000	100	Send to PIC
Assert_INTB	0010 0001	100	Send to PIC
Assert_INTC	0010 0010	100	Send to PIC
Assert_INTD	0010 0011	100	Send to PIC
De-assert_INTA	0010 0100	100	Send to PIC
De-assert_INTB	0010 0101	100	Send to PIC
De-assert_INTC	0010 0110	100	Send to PIC
De-assert_INTD	0010 0111	100	Send to PIC
PM_Active_State_Nak	0001 0100	100	No action taken
PM_PME	0001 1000	000	Generate interrupt to PIC if enabled
PME_Turn_Off	0001 1001	011	No action taken
PME_TO_Ack	0001 1011	101	Set PEX_PME_MES_DR[ENL23] bit and generate interrupt to PIC if enabled
ERR_COR	0011 0000	000	Generate interrupt to PIC if enabled
ERR_NONFATAL	0011 0001	000	Generate interrupt to PIC if enabled
ERR_FATAL	0011 0011	000	Generate interrupt to PIC if enabled
Unlock	0000 0000	000	No action taken
Set_Slot_Power_Limit	0101 0000	100	No action taken
Vendor_Defined Type 0	0111 1110		No action taken
Vendor_Defined Type 1	0111 1111		No action taken
Attention_Indicator_On	0100 0001	100	No action taken
Attention_Indicator_Blink	0100 0011	100	No action taken
Attention_Indicator_Off	0100 0000	100	No action taken
Power_Indicator_On	0100 0101	100	No action taken
Power_Indicator_Blink	0100 0111	100	No action taken
Power_Indicator_Off	0100 0100	100	No action taken
Attention_Button_Pressed	0100 1000	100	Set PEX_PME_MES_DR[ABP] bit and send interrupt if enabled.

Table 16-331 provides a complete list of supported inbound messages in EP mode.

Table 16-331. PCI Express EP inbound message handling

Name	Code[7:0]	Routing[2:0]	Action
Assert_INTA	0010 0000	100	No action taken
Assert_INTB	0010 0001	100	No action taken
Assert_INTC	0010 0010	100	No action taken
Assert_INTD	0010 0011	100	No action taken
Deassert_INTA	0010 0100	100	No action taken
Deassert_INTB	0010 0101	100	No action taken
Deassert_INTC	0010 0110	100	No action taken
Deassert_INTD	0010 0111	100	No action taken
PM_Active_State_Nak	0001 0100	100	No action taken
PM_PME	0001 1000	000	No action taken
PME_Turn_Off	0001 1001	011	Set PEX_PME_MES_DR[PTO] bit. Send interrupt if enabled.
PM_TO_Ack	0001 1011	101	No action taken
ERR_COR	0011 0000	000	No action taken
ERR_NONFATAL	0011 0001	000	No action taken
ERR_FATAL	0011 0011	000	No action taken
Unlock	0000 0000	000	No action taken
Set_Slot_Power_Limit	0101 0000	100	Update power value in PCI Express device capability register in configuration space.
Vendor_Defined Type 0	0111 1110		No action taken
Vendor_Defined Type 1	0111 1111		No action taken
Attention_Indicator_On	0100 0001	100	Set PEX_PME_MES_DR[AION] bit. Send interrupt if enabled.
Attention_Indicator_Blink	0100 0011	100	Set PEX_PME_MES_DR[AIB] bit. Send interrupt if enabled.
Attention_Indicator_Off	0100 0000	100	Set PEX_PME_MES_DR[AIOF] bit. Send interrupt if enabled.
Power_Indicator_On	0100 0101	100	Set PEX_PME_MES_DR[PION] bit. Send interrupt if enabled.
Power_Indicator_Blink	0100 0111	100	Set PEX_PME_MES_DR[PIB] bit. Send interrupt if enabled.
Power_Indicator_Off	0100 0100	100	Set PEX_PME_MES_DR[PIOF] bit. Send interrupt if enabled.
Attention_Button_Pressed	0100 1000	100	No action taken

16.12.1.12 Error handling

The PCI Express specification classifies errors as correctable and uncorrectable.

Correctable errors result in degraded performance, but uncorrectable errors generally result in functional failures. As shown in [Figure 16-337](#) uncorrectable errors can further be classified as fatal or non-fatal.

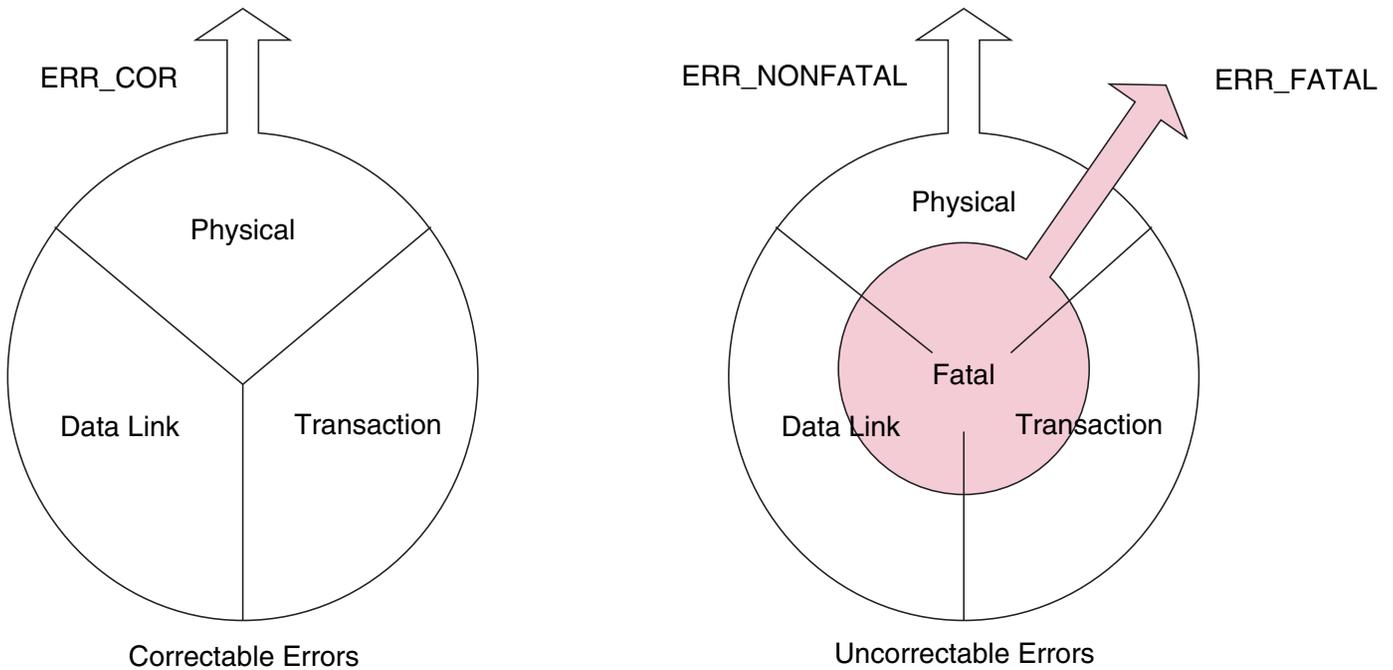
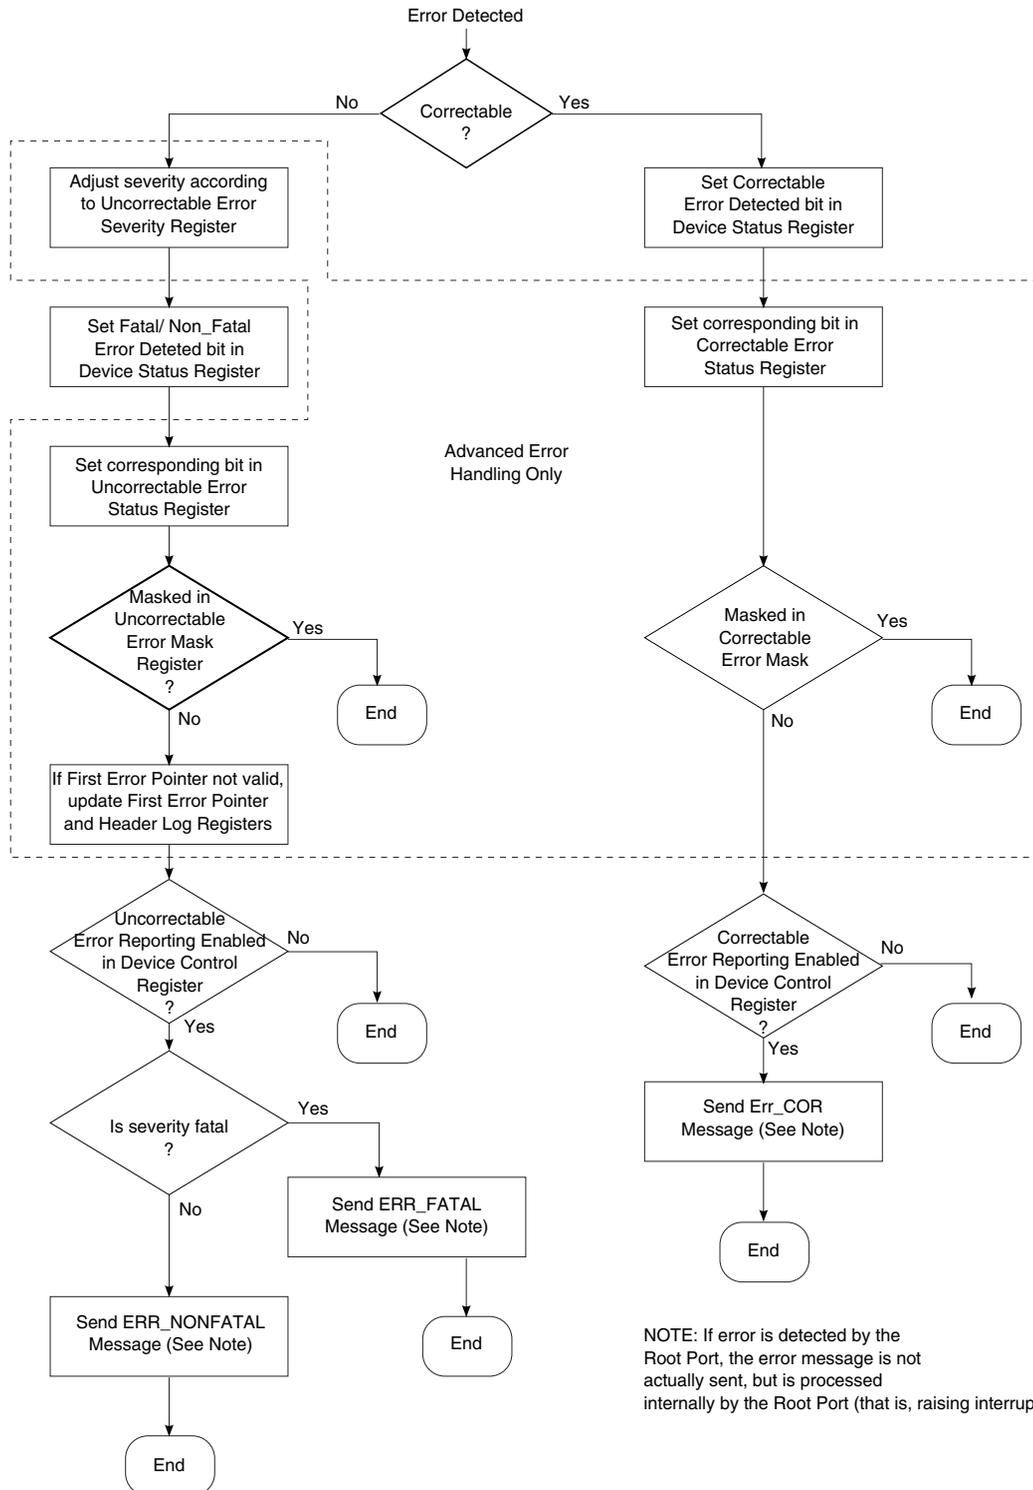


Figure 16-337. PCI Express error classification

16.12.1.12.1 PCI Express error logging and signaling

Figure 16-338 shows the PCI Express-defined sequence of operations related to signaling and logging of errors detected by a device.

Note that the PCI Express controller on this device supports the advanced error handling capabilities shown within the dotted lines.



P2020 QorIQ Integrated Processor Reference Manual, Rev. 2, 12/2012
 Figure 16-338. PCI Express device error signaling flowchart

16.12.1.12.2 PCI Express controller internal interrupt sources

Table 16-332 describes the sources of the PCI Express controller internal interrupt to the PIC and the preconditions for signaling the interrupt.

Table 16-332. PCI Express internal controller interrupt sources

Status register bit	Preconditions
Any bit in PEX_PME_MES_DR set	The corresponding interrupt enable bits must be set in PEX_PME_MES_IER
Any bit in PEX_ERR_DR set	The corresponding interrupt enable bits must be set in PEX_ERR_EN.
PCI Express Root Status Register[16] (PME status) is set	PCI Express Root Control Register [3] (PME interrupt enable) is set
PCI Express Root Error Status Register[6] (fatal error messages received) is set	PCI Express Root Error Command Register [2] (fatal error reporting enable) is set or PCI Express Root Control Register [2] (system error on fatal error enable) is set
PCI Express Root Error Status Register [5] (non-fatal error messages received) is set	PCI Express Root Error Command Register [1] (non-fatal error reporting enable) is set or PCI Express Root Control Register [1] (system error on non-fatal error enable) is set
PCI Express Root Error Status Register[0] (correctable error messages received) is set	PCI Express Root Error Command Register[0] (correctable error reporting enable) is set or PCI Express Root Control Register[0] (system error on correctable error enable) is set.
Any correctable error status bit in PCI Express Correctable Error Status Register is set	The corresponding error mask bit in PCI Express Correctable Error Mask Register is clear and PCI Express Root Error Command Register[0] (correctable error reporting enable) is set
Any fatal uncorrectable error status bit in PCI Express Uncorrectable Error Status Register is set. (The corresponding error is classified as fatal based on the severity setting in PCI Express Uncorrectable Error Severity Register.)	The corresponding error mask bit in PCI Express Uncorrectable Error Mask Register is clear and either PCI Express Device Control Register[2] (fatal error reporting) is set or PCI Express Command Register[8] (SERR) is set.
Any non-fatal uncorrectable error status bit in PCI Express Uncorrectable Error Status Register is set. (The corresponding error is classified as non-fatal based on the severity setting in PCI Express Uncorrectable Error Severity Register.)	The corresponding error mask bit in PCI Express Uncorrectable Error Mask Register is clear and either PCI Express Device Control Register[1] (non-fatal error reporting) is set or PCI Express Command Register[8] (SERR) is set.
PCI Express Secondary Status Register[8] (master data parity error) is set.	PCI Express Secondary Status Interrupt Mask Register[0] (mask master data parity error) is cleared and PCI Express Command Register[6] (parity error response) is set.
PCI Express Secondary Status Register[11] (signaled target abort) is set	PCI Express Secondary Status Interrupt Mask Register[1] (mask signaled target abort) is cleared.

Table continues on the next page...

Table 16-332. PCI Express internal controller interrupt sources (continued)

Status register bit	Preconditions
PCI Express Secondary Status Register[12] (received target abort) is set	PCI Express Secondary Status Interrupt Mask Register[2] (mask received target abort) is cleared.
PCI Express Secondary Status Register[13] (received master abort) is set	PCI Express Secondary Status Interrupt Mask Register[3] (mask received master abort) is cleared.
PCI Express Secondary Status Register[14] (signaled system error) is set.	PCI Express Secondary Status Interrupt Mask Register[4] (mask signaled system error) is cleared.
PCI Express Secondary Status Register[15] (detected parity error) is set	PCI Express Secondary Status Interrupt Mask Register[5] (mask detected parity error) is cleared.
PCI Express Slot Status Register[0] (attention button pressed) is set	PCI Express Slot Control Register[0] (attention button pressed enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1-0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[1] (power fault detected) is set	PCI Express Slot Control Register[1] (power fault detected enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1-0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[2] (MRL sensor changed) is set	PCI Express Slot Control Register[2] (MRL sensor changed enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1-0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[3] (presence detect changed) is set	PCI Express Slot Control Register[3] (presence detect changed enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set and either PCI Express PM Control Register[1-0] = 00 (the function power state is D0) or PCI Express PM Control Register[8] (PME enable) is set.
PCI Express Slot Status Register[4] (command completed) is set	PCI Express Slot Control Register[4] (command completed interrupt enable) is set and PCI Express Slot Control Register[5] (hot plug interrupt enable) is set.

16.12.1.12.3 Error conditions

Table 16-333 describes specific error types and the action taken for various transaction types.

Table 16-333. Error conditions

Transaction Type	Error Type	Action
Inbound response	PEX response time out. This case happens when the internal platform sends a non-posted request that did not get a response back after a specific amount of time specified in the outbound completion timeout register (PEX_OTB_CPL_TOR)	Log error (PEX_ERR_DR[PCT]) and send interrupt to PIC, if enabled.
Inbound response	Unexpected PEX response. This can happen if, after the response times out and the internal queue entry is deallocated, the response comes back.	Log unexpected completion error (PCI Express Uncorrectable Status Register[16]) and send interrupt to PIC, if enabled.
Inbound response	Unsupported request (UR) response status	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PEX_ERR_DR[CDNSC] and PCI Express Uncorrectable Status Register[20]) and send interrupt to PIC, if enabled.
Inbound response	Completer abort (CA) response status	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PEX_ERR_DR[PCAC, CDNSC] and PCI Express Uncorrectable Status Register[15]) and send interrupt to PIC, if enabled.
Inbound response	Poisoned TLP (EP=1)	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PCI Express Uncorrectable Status Register[12]) and send interrupt to PIC, if enabled.
Inbound response	ECRC error	Depending upon whether the initial internal request was broken up, the error is not sent until all responses come back for all portions of the internal request. Log the error (PCI Express Uncorrectable Status Register[19]) and send interrupt to PIC, if enabled.

Table continues on the next page...

Table 16-333. Error conditions (continued)

Transaction Type	Error Type	Action
Inbound response	Configuration Request Retry Status (CRS) timeout for a configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	<ol style="list-style-type: none"> 1. The controller always retries the transaction as soon as possible until a status other than CRS is returned. However, if a CRS status is returned after the configuration retry timeout (PEXCONF_RTY_TOR) timer expires, then the controller aborts the transaction and sends all 1s (0xFFFF_FFFF) data back to requester. 2. Log the error (PEX_ERR_DR[PCT]) and send interrupt to the PIC, if enabled.
Inbound response	UR response for configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	<ol style="list-style-type: none"> 1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PEX_ERR_DR[CDNSC] and PCI Express Uncorrectable Status Register[20]) and send interrupt to PIC, if enabled.
Inbound response	CA response for Configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	<ol style="list-style-type: none"> 1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PEX_ERR_DR[PCAC, CDNSC] and PCI Express Uncorrectable Status Register[15]) and send interrupt to PIC, if enabled.
Inbound response	Poisoned TLP (EP=1) response for Configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	<ol style="list-style-type: none"> 1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PCI Express Uncorrectable Status Register[12]) and send interrupt to PIC, if enabled.
Inbound response	ECRC error response for Configuration transaction that originates from PEX_CONFIG_ADDR/ PEX_CONFIG_DATA	<ol style="list-style-type: none"> 1. Send back all 1s (0xFFFF_FFFF) data. 2. Log the error (PCI Express Uncorrectable Status Register[19]) and send interrupt to PIC, if enabled.
Inbound response	Configuration Request Retry Status (CRS) response for Configuration transaction that originates from ATMU	<ol style="list-style-type: none"> 1. The controller always retries the transaction as soon as possible until a status other than CRS is returned. However, if a CRS status is returned after the configuration retry timeout (PEXCONF_RTY_TOR) timer expires, then the controller aborts the transaction. 2. Log the error (PEX_ERR_DR[CRST]) and send interrupt to the PIC, if enabled.
Inbound response	UR response for Configuration transaction that originates from ATMU	Log the error (PEX_ERR_DR[CDNSC] and PCI Express Uncorrectable Status Register[20]) and send interrupt to PIC, if enabled.
Inbound response	CA response for Configuration transaction that originates from ATMU	Log the error (PEX_ERR_DR[PCAC, CDNSC] and PCI Express Uncorrectable Status Register[15]) and send interrupt to PIC, if enabled.
Inbound response	Malformed TLP response	PCI Express controller does not pass the response back to the core. Therefore, a completion timeout error eventually occurs.

Table continues on the next page...

Table 16-333. Error conditions (continued)

Transaction Type	Error Type	Action
Inbound request	Poisoned TLP (EP=1)	<ol style="list-style-type: none"> 1. If it is a posted transaction, the controller drops it. 2. If it is a non-posted transaction, the controller returns a completion with UR status. 3. Release the proper credits
Inbound request	ECRC error	<ol style="list-style-type: none"> 1. If it is a posted transaction, the controller drops it. 2. If it is a non-posted transaction, the controller returns a completion with UR status. 3. Release the proper credits
Inbound request	PCI Express nullified request	The packet is dropped.
Outbound request	Outbound ATMU crossing	Log the error (PEX_ERR_DR[OAC]). The transaction is not sent out on the link.
Outbound request	Illegal message size	Log the error (PEX_ERR_DR[MIS]). The transaction is not sent out on the link.
Outbound request	Illegal I/O size	Log the error (PEX_ERR_DR[IOIS]). The transaction is not sent out on the link.
Outbound request	Illegal I/O address	Log the error (PEX_ERR_DR[IOIA]). The transaction is not sent out on the link.
Outbound request	Illegal configuration size	Log the error (PEX_ERR_DR[CIS]). The transaction is not sent out on the link.
Outbound response	Internal platform response with error (for example, an ECC error on a DDR read or the transaction maps to unknown address space).	Send poisoned TLP (EP=1) completion(s) for data that are known bad. If the poison data happens in the middle of the packet, the rest of the response packet(s) is also poisoned.

16.12.1.12.4 Error capture registers

The PCI Express error capture registers, PEX_ERR_CAP_R0 through PEX_ERR_CAP_R3, allow vital error information to be captured when an error occurs.

Different error information is reported depending on whether the error source is from an outbound transaction from an internal source or from an inbound transaction from an external source; the source of the captured error is reflected in PEX_ERR_CAP_STAT[GSID]. Note that after the initial error is captured, no further capturing is performed until the PEX_ERR_CAP_STAT[ECV] bit is clear.

16.12.1.12.4.1 Error capture registers (outbound error)

PEX_ERR_CAP_Rn for the case when the error is caused by an outbound transaction from an internal source and the error is due to timeout condition or PEX_CONFIG_ADDR/PEX_CONFIG_DATA access are shown in the following figure and tables.

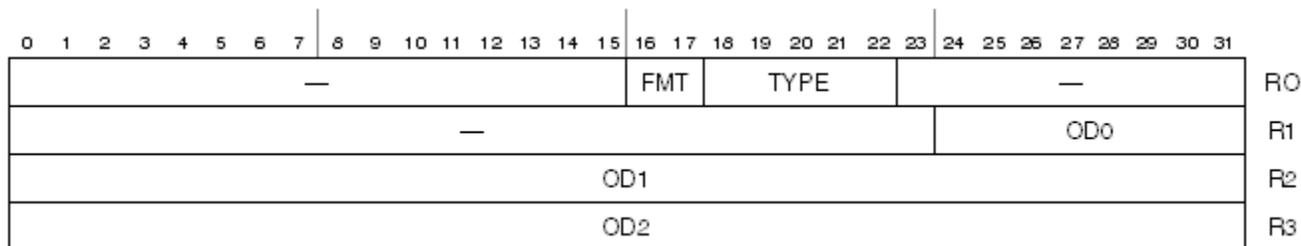


Figure 16-339. PCI Express error capture registers, outbound case

Table 16-334. PCI Express error capture register n field descriptions internal source, outbound transaction

ERR_CA P_n	Bits	Name	Description
R0	0-15	Reserved	-
	16-17	FMT	Header format
	18-22	TYPE	Transaction type
	23-31	Reserved	-
R1	0-23	Reserved	-
	24-31	OD0	Internal platform transaction information. Reserved for factory debug.
R2	0-31	OD1	Internal platform transaction information. Reserved for factory debug.
R3	0-31	OD2	Internal platform transaction information. Reserved for factory debug.

16.12.1.12.4.2 Error capture registers (inbound error)

PEX_ERR_CAP_Rn for the case when the error is caused by an inbound transaction from an external source are shown in the following figure and tables.

Inbound transactions that result in PEX_ERR_DR[PNM] or PEX_ERR_DR[PCAC] or PEX_ERR_DR[CDNSC] or PEX_ERR_DR[CRSNC] being set will result in PEX_ERR_CAP_STAT[GSID] indicating the controller itself was the source of the error.

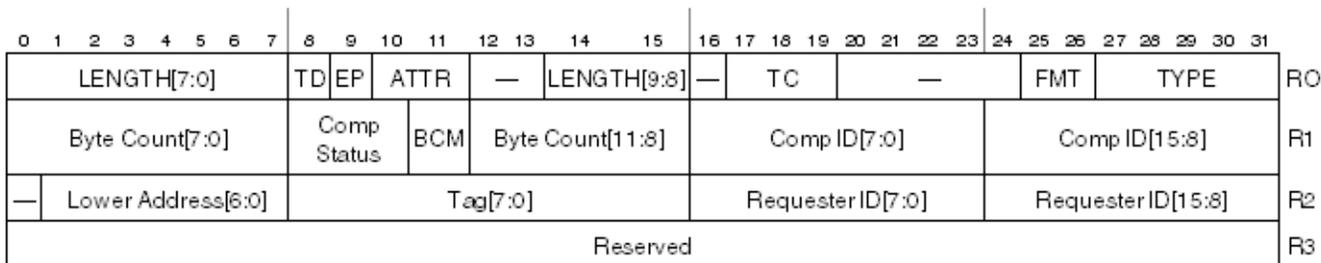


Figure 16-340. PCI Express error capture registers, inbound case, completion transaction

Table 16-335. PCI Express error capture register 1 field descriptions external source, inbound completion transaction

ERR_CAP_n	Bits	Name	Description
R0	24	Reserved	-
	25-26	FMT	Header format
	27-31	TYPE	Transaction type
	16	Reserved	-
	17-19	TC	Traffic class
	20-23	Reserved	-
	8	TD	TLP digest field present
	9	EP	Poisoned packet
	10-11	ATTR	Attributes
	12-13	Reserved	-
	14-15, 0-7	LENGTH[9:0]	Packet length
R1	24-31, 16-23	Comp ID[15:0]	Completion ID, bits 15:0
	8-10	Comp Status	Completion status
	11	BCM	Byte count modified
	12-15, 0-7	Byte Count[11:0]	Byte count, bits 11:0
R2	24-31, 16-23	Requester ID[15:0]	Requester ID, bits 15:0
	8-15	Tag[7:0]	Tag, bits 7:0
	0	Reserved	-
	1-7	Lower Address[6:0]	Lower address, bits 6:0
R3	0-31	Reserved	Not used

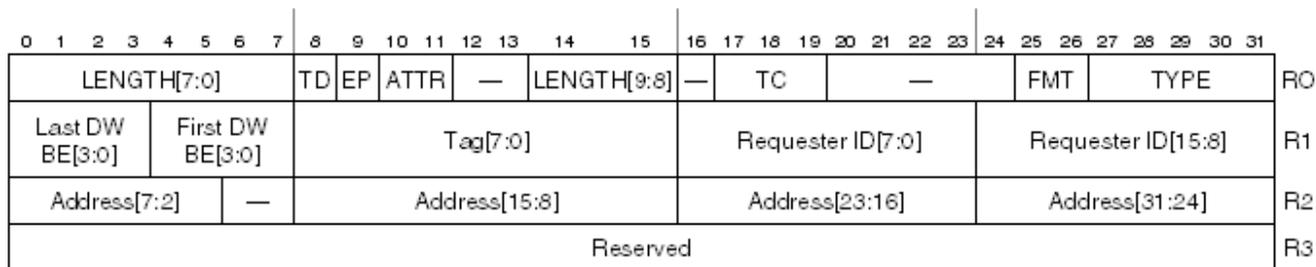


Figure 16-341. PCI Express error capture registers, inbound case, memory request transaction (3 DW)

Table 16-336. PCI Express error capture register 1 field descriptions external source, inbound memory request transaction (3 DW)

ERR_CAP_n	Bits	Name	Description
R0	24	Reserved	-
	25-26	FMT	Header format
	27-31	TYPE	Transaction type
	16	Reserved	-
	17-19	TC	Traffic class
	20-23	Reserved	-
	8	TD	TLP digest field present
	9	EP	Poisoned packet
	10-11	ATTR	Attributes
	12-13	Reserved	-
14-15, 0-7	LENGTH[9:0]	Packet length	
R1	24-31, 16-23	Requester ID[15:0]	Requester ID, bits 15:0
	8-15	Tag[7:0]	Tag, bits 7:0
	0-3	Last DW BE[3:0]	Last doubleword byte enables, bits 3:0
	4-7	First DW BE[3:0]	First doubleword byte enables, bits 3:0
R2	24-31, 16-23	Address[31:2]	Address, bits 31:2
	8-15, 0-5		
	6-7	Reserved	-
R3	0-31	Reserved	Not used

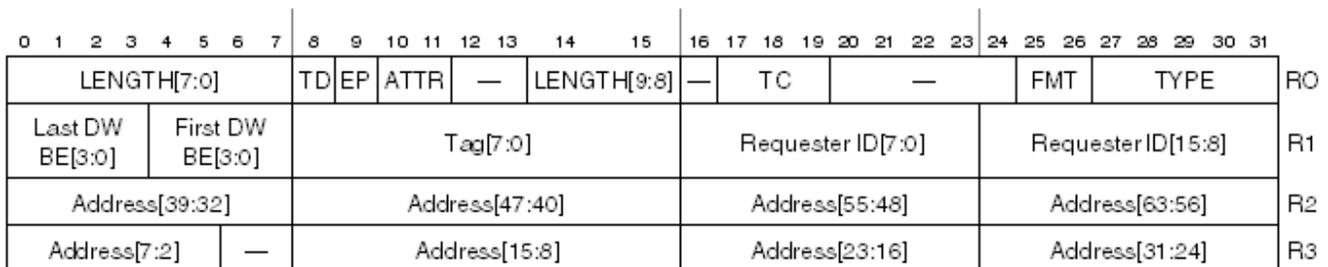


Figure 16-342. PCI Express error capture registers, inbound case, memory request transaction (4 DW)

Table 16-337. PCI Express error capture register 1 field descriptions external source, inbound memory request transaction (4 DW)

ERR_CAP_n	Bits	Name	Description
R0	24	Reserved	-
	25-26	FMT	Header format
	27-31	TYPE	Transaction type
	16	Reserved	-
	17-19	TC	Traffic class
	20-23	Reserved	-
	8	TD	TLP digest field present
	9	EP	Poisoned packet
	10-11	ATTR	Attributes
	12-13	Reserved	-
	14-15, 0-7	LENGTH[9:0]	Packet length
R1	24-31, 16-23	Requester ID[15:0]	Requester ID, bits 15:0
	8-15	Tag[7:0]	Tag, bits 7:0
	0-3	Last DW BE[3:0]	Last doubleword byte enables, bits 3:0
	4-7	First DW BE[3:0]	First doubleword byte enables, bits 3:0
R2	24-31, 16-23	Address[63:32]	Address, bits 63:32
	8-15, 0-7		
R3	24-31, 16-23	Address[31:2]	Address, bits 31:2
	8-15, 0-5		
	6-7	Reserved	-

16.12.2 Interrupts

Both INTx and message signaled interrupts (MSI) are supported; however there are differences depending on whether the PCI Express controller is configured as an RC or EP device.

16.12.2.1 EP interrupt generation

16.12.2.1.1 Hardware INTx message generation

Hardware INTx message generation is not supported in EP mode.

16.12.2.1.2 Hardware MSI generation

In EP mode, the PCI Express controller can be configured to automatically generate MSI transactions to the root complex when an interrupt event occurs.

The PCI Express controller uses *irq_out* (an internal version of the IRQ_OUT signal) to trigger the generation of the MSI. To trigger the MSI, interrupt events must be routed to *irq_out* by setting the EP (external pin) bit in the associated Interrupt Destination register in the PIC. Note that the IRQ_OUT signal should not be used for any other function if it is being used to trigger MSI transactions.

The remote root complex is expected to set up the MSI capability structure of all endpoints at system initialization by filling the Message Address and Message Data registers with appropriate values and setting the MSIE bit in the MSI Message Control register.

With the PCI Express controller properly configured, when it detects the leading edge of *irq_out* going active, it generates a PCI Express memory write transaction to the address specified in the MSI Message Address register (and MSI Message Upper Address register) with a data payload as specified in the MSI Message Data register (with leading zeros appended).

16.12.2.1.3 Software INTx message generation

Software can generate outbound assert or deassert INTx message transactions by using the outbound ATMU mechanism described in [Outbound ATMU message generation](#).

16.12.2.1.4 Software MSI generation

Host software has to set up the MSI capability registers to enable MSI mode, and have the correct values for the MSI address and data register.

Then local software has to read the MSI address in the MSI capability register and configure the outbound ATMU window to map the translated address to the MSI address. Software has to determine the number of allocated messages in the MSI capability register and allocates the appropriate data values to use. A write to the ATMU window containing the MSI address with the appropriate data value generates the desired MSI transaction to the remote RC.

16.12.2.2 RC handling of INTx message and MSI interrupts

16.12.2.2.1 INTx message handling

MSIs are the preferred interrupt signaling mechanism for PCI Express.

However, in RC mode, the PCI Express controller supports the INTx virtual-wire interrupt signaling mechanism (as described in the PCI Express specification). Whenever the controller receives an inbound INTx (INTA, INTB, INTC, or INTD) asserted or negated message, it asserts or negates an equivalent internal INTx signal (*inta*, *intb*, *intc*, or *intd*) to the interrupt controller.

The internal INTx signals from the PCI Express controller may be logically combined with the interrupt request (IRQ_n) input signals so that they share the same interrupt controlled by the associated $EIVPR_n$ and $EIDR_n$ registers in the interrupt controller. Refer to [Programmable interrupt controller \(PIC\)](#) for more information about sharing of PCI Express INTx interrupts and the external interrupt request (IRQ_n) signals.

If a PCI Express INTx interrupt is being used, then the interrupt controller must be configured so that external interrupts are level-sensitive ($EIVPR_n[S] = 1$).

16.12.2.2.2 MSI handling

An inbound MSI cycle must hit into the PEXCSRBAR window with the address offset that points to the MSIIR register in the PIC.

Note that it is the responsibility of the host software to configure each EP's MSI capability registers such that an MSI cycle generated from the EP device is routed to the MSIIR register in the PIC and for the appropriate interrupt to be generated to the core.

16.12.3 Initial credit advertisement

To prevent overflowing of the receiver's buffers and for ordering compliance purposes, the transmitter cannot send transactions unless it has enough flow control (FC) credits to send. Each device maintains an FC credit pool.

The FC information is conveyed between the two link partners by DLLPs during link training (initial credit advertisement). The transaction layer performs the FC accounting functions. One FC unit is four DWs (16-bytes) of data.

Table 16-338. Initial credit advertisement

Credit type	Initial credit advertisement
PH (Memory Write, Message Write)	4
PD (Memory Write, Message Write)	(256/16)x4=64
NPH (Memory Read, IO Read, Cfg Read, Cfg Write)	8
NPD (IO Write, Cfg Write)	2
CPLH (Memory Read Completion, IO R/W Completion, Cfg R/W Completion)	Infinite
CPLD (Memory Read Completion, IO Read Completion, Cfg Read Completion)	Infinite

16.12.4 Power management

All device power states are supported with the exception of D3cold with Vaux. In addition, all link power states are supported with the exception of L2 states.

Only L0s ASPM mode is supported if enabled by configuring the link control register's bits 1-0 in configuration space. Note that there is no power saving in the controller when the device is put into a non-D0 state. The only power saving is the I/O drivers when the controller is put into a non-L0 link state.

Table 16-339. Power management state supported

Component D-state	Permissible interconnect state	Action
D0	L0, L0s	In full operation.
D1	L0, L0s, L1	All outbound traffics are stalled. All inbound traffic is thrown away. The only exceptions are PME messages and configuration transactions. If the device is in RC mode, it is permissible to send a PM_Turn_Off message through the PEX power management command register.
D2	L0, L0s, L1	All outbound traffics are stalled. All inbound traffic is thrown away. The only exceptions are PME messages and configuration transactions. If the device is in RC mode, it is permissible to send a PM_Turn_Off message through the PEX power management command register.
D3hot	L0, L0s, L1, L2/ L3 ready	All outbound traffics are stalled. All inbound traffic is thrown away. The only exceptions are PME messages and configuration transactions. If the device is in RC mode, it is permissible to send a PM_Turn_Off message through the PEX power management command register. Note that if a transition of D3hot->D0 occurs, a reset is performed to the controller's configuration space. In addition, link training restarts.
D3cold	L3	Completely off.

16.12.4.1 L2/L3 ready link state

The L2/L3 ready link state is entered after the EP device is put into a D3hot state followed by a PME_Turn_Off/PME_TO_Ack message handshake protocol.

Exiting this state requires a POR reset or a WAKE_B signal from the EP device. The PCI Express controller (in EP mode) does not support the generation of beacon; therefore, as an alternative, the device can use one of the GPIO signals as an enable to an external tristate buffer to generate a WAKE_B signal, as shown in [Figure 16-343](#).

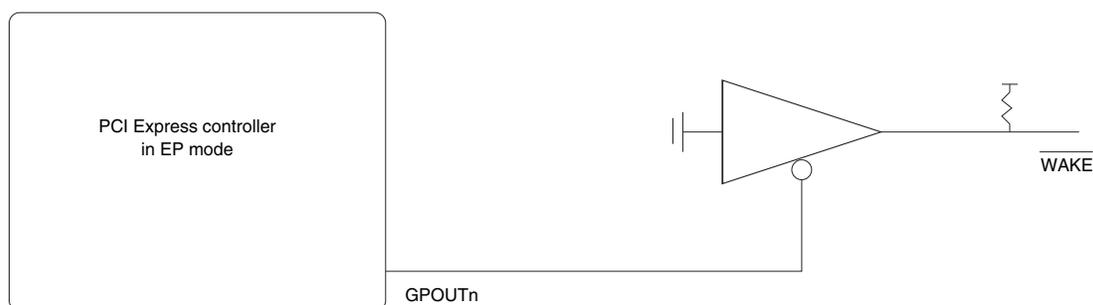


Figure 16-343. WAKE_B Generation example

In RC mode, the WAKE_B signal from the EP device can be connected to one of the external interrupt inputs to service the WAKE_B request.

16.12.5 Hot reset

When a hot reset condition occurs, the controller (in both RC and EP mode) initiates a clean-up of all outstanding transactions and returns to an idle state.

All configuration register bits that are non-sticky are reset. Link training takes place subsequently. The device is permitted to generate a hot reset condition on the bus when it is configured as an RC device by setting the "Secondary Bus Reset" bit in the Bridge Control Register in the configuration space. As an EP device, it is not permitted to generate a hot reset condition; it can only detect a hot reset condition and initiate the clean-up procedure appropriately.

16.12.6 Link down

Typically, a link down condition occurs after a hot reset event; however, it is possible for the link to go down unexpectedly without a hot reset event.

When this occurs, a link down condition is detected (`PEX_PME_MSG_DR[LDD]=1`). Link down is treated similarly to a hot reset condition.

Subsequently, while the link is down, all new posted outbound transactions are discarded. All new non-posted ATMU transactions are errored out. Non-posted configuration transactions issued using `PEX_CONFIG_ADDR/PEX_CONFIG_DATA` toward the link returns `0xFFFF_FFFF` (all 1s). As soon as the link is up again, the sending of transaction resumes.

Note that in EP mode, a link down condition causes the controller to reset all non-sticky bits in its PCI Express configuration registers as if it had been hot reset.

16.13 Initialization/application information

The following topics describe the initialization/application information for PCI Express.

16.13.1 EP Boot mode and inbound configuration transactions

In normal boot mode (`cfg_cpu_boot = 1`), the core is allowed to boot and configure the device.

During this time, the PCI Express interface retries all inbound PCI Express configuration transactions. When the core has configured the device to a state where it can accept inbound PCI Express configuration transactions, the boot code should set the `CFG_READY` bit in the `PEX_CFG_READY` register after which inbound PCI Express configuration transactions are accepted. Refer to [Configuration Ready Register \(Configuration_Ready_Register\)](#), for more information about the `CFG_READY` bit.

In boot hold-off mode (`cfg_cpu_boot = 0`), the core is prevented from fetching its first instruction by withholding its internal bus grant. During this time, the PCI Express interface accepts all inbound PCI Express configuration transactions which allows an external host/RC to configure the device. When the external host/RC has configured the device to a state where it can allow the core to fetch code from the boot vector, it sets the `EEBPCR[CPU_EN]` bit after which the core is granted the internal bus.

16.13.2 Automatic link retraining during initialization

During initialization, if the link partner is down or the device has not received DLLP updates within timeout interval specified in `PEX_FC_UPDATE_TOR[FC_UPDATE_TIMEOUT]`, it will fail link training.

It is recommended that software clear PEX_FC_UPDATE_TOR[DIS_LR] to allow for automatic retraining of the link. See [Flow Control Update Timeout Register \(Flow_Control_Update_Timeout_Register\)](#), for more information.

16.13.3 Configuration accesses and inbound writes to CCSR space

In RC mode, when the core issues an outbound configuration access to an external device using the PCI Express configuration access registers (PEX_CONFIG_ADDR/PEX_CONFIG_DATA), there is a potential deadlock if several inbound memory write transactions targeting the CCSR space are received before the completion for the configuration access is returned.

Inbound writes to CCSR space include MSIs and message interrupts to the MPIC. The deadlock also can occur when the endpoint generates a single write to CCSR followed by multiple writes to another target before the original configuration access by the RC completes. When the deadlock occurs the configuration access will time out and an error will be reported (if enabled). This deadlock can be avoided by any of the following methods:

1. Configure external agents to allow completions to bypass memory write requests. Many devices allow this when relaxed ordering (RO) is in effect; however, this ability is optional under the *PCI Express Base Specification* ordering rules, so it is possible that some devices may not support it.
2. For the RC, restrict external configuration accesses through PEX_CONFIG_ADDR/PEX_CONFIG_DATA to the enumeration of the bus and other initialization functions; do not use PEX_CONFIG_ADDR/PEX_CONFIG_DATA to access external devices during periods when PCI Express endpoints are generating inbound writes to CCSR space.
3. Govern the generation of write transactions by the endpoint device such that after a write to the CCSR space no other writes are generated until pending configuration accesses are completed or the write is guaranteed to have completed by the completion of a subsequent read request by the endpoint devices.
4. Use the PCI Express outbound ATMU window to generate outbound configuration accesses during periods when an endpoint device may be writing to CCSR space. See [RC configuration register access](#) for more information. The deadlock does not occur when the ATMU window is used for configuration accesses.

Chapter 17

Enhanced secure digital host controller (eSDHC)

The enhanced secure digital host controller (eSDHC) provides an interface between the host system and these types of memory cards.

17.1 eSDHC overview

The enhanced secure digital host controller (eSDHC) provides an interface between the host system and these types of memory cards:

- MultiMediaCard (MMC)

MMC is a universal low-cost data storage and communication medium designed to cover a wide area of applications including mobile video and gaming, which are available from either pre-loaded MMC cards or downloadable from cellular phones, WLAN, or other wireless networks. Old MMC cards are based on a seven-pin serial bus with a single data pin, while the new high-speed MMC communication is based on an advanced 11-pin serial bus designed to operate in a low voltage range.

- Secure digital (SD) card

The secure digital (SD) card is an evolution of old MMC technology. It is specifically designed to meet the security, capacity, performance, and environment requirements inherent in the emerging audio and video consumer electronic devices. The physical form factor, pin assignments, and data transfer protocol are forward-compatible with the old MMC.

The eSDHC acts as a bridge, passing host bus transactions to SD/MMC cards by sending commands and performing data accesses to or from the cards. It handles the SD/MMC protocol at the transmission level.

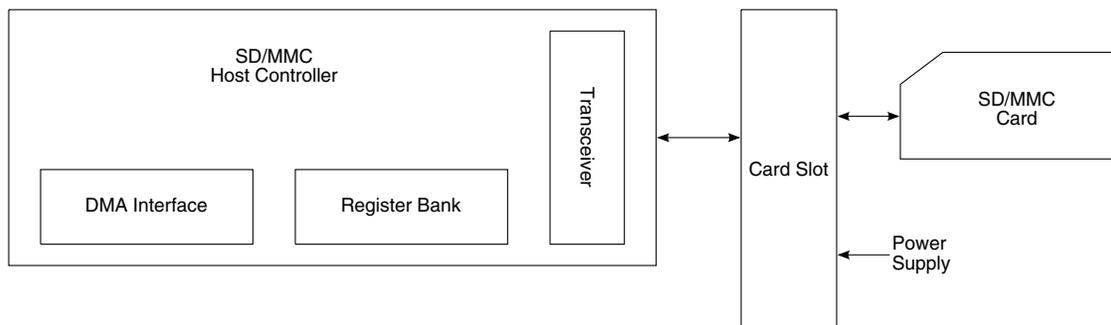


Figure 17-1. System connection of the eSDHC

This figure shows the eSDHC block diagram.

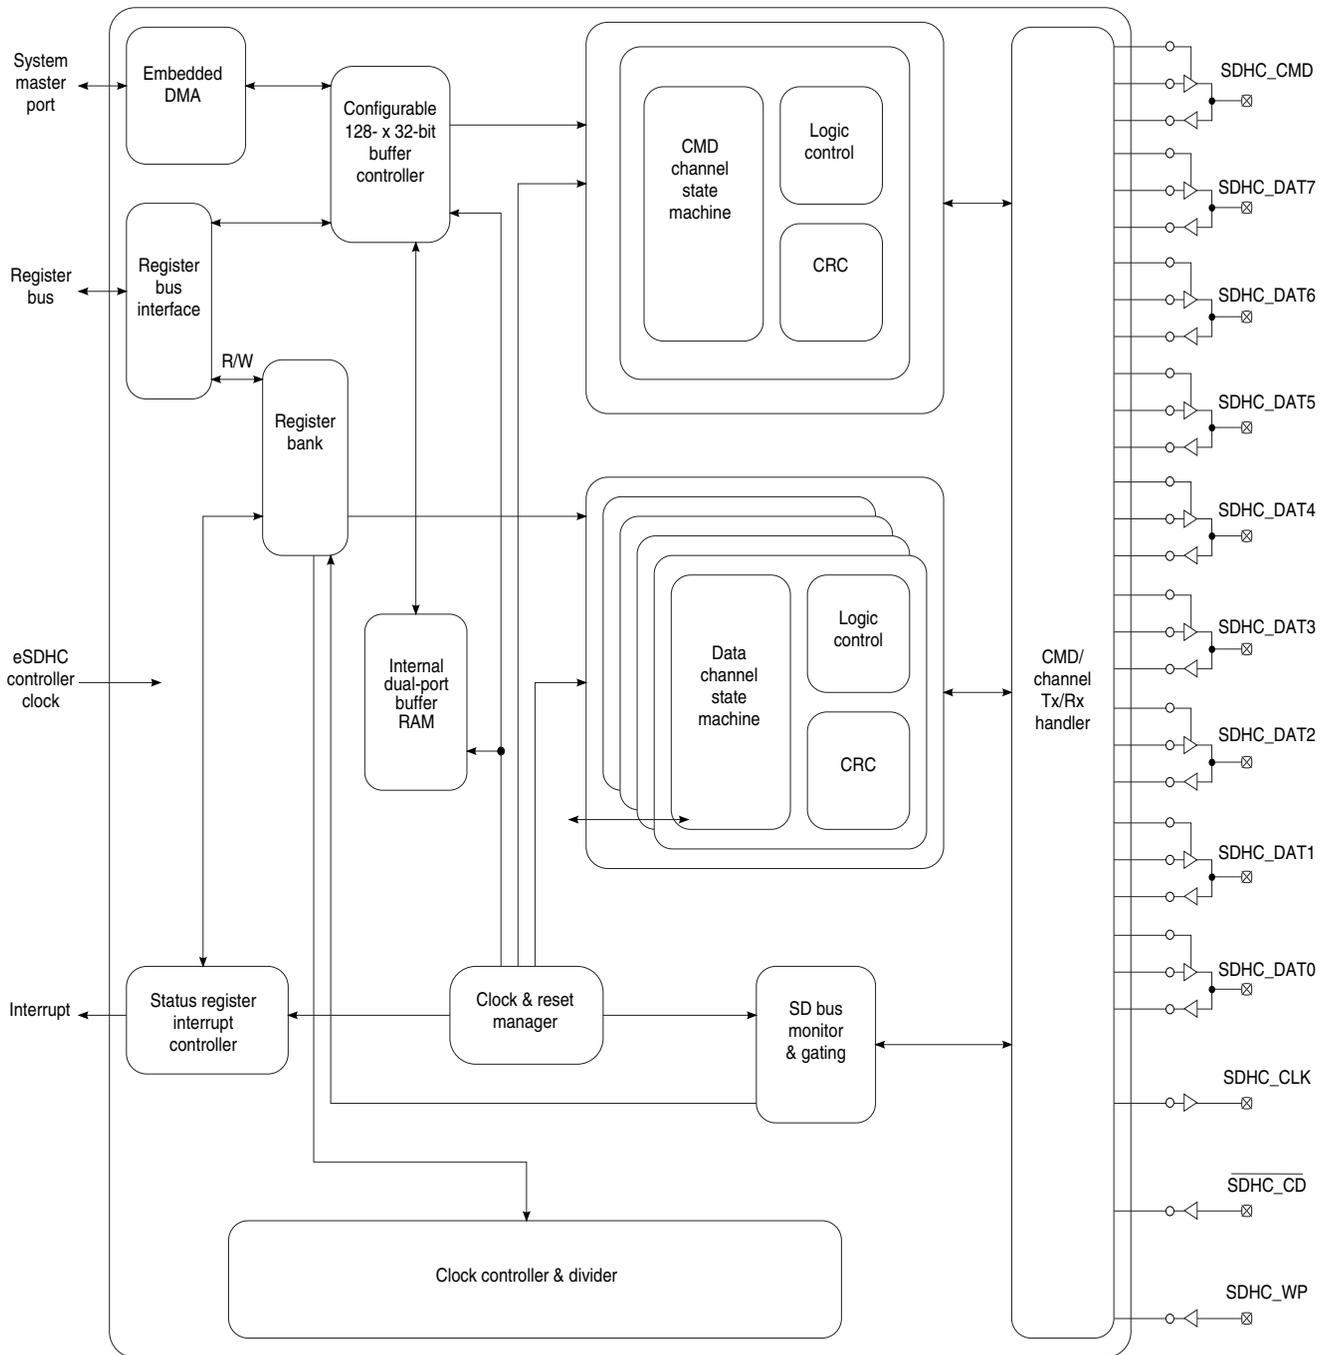


Figure 17-2. eSDHC block diagram

17.2 eSDHC features summary

The eSDHC includes the following features:

- Compatible with the following specifications:

- *SD Host Controller Standard Specification, Version 2.0* (<http://www.sdcard.org>) with test event register support
- *MultiMediaCard System Specification, Version 4.2* (<http://www.mmca.org>)
- *SD Memory Card Specification, Version 2.0* (<http://www.sdcard.org>)
- Designed to work with SD Memory, SDHC memory, miniSD Memory, MMC, MMC*plus*, and RS-MMC cards
- Card bus clock frequency up to 52 MHz
- Supports 1-/4-bit SD mode, 1-/4-/8-bit MMC modes
 - Up to 200 Mbps data transfer for SD/MMC cards using four parallel data lines
 - Up to 416 Mbps data transfer for MMC using 8 parallel data lines
- Single- and multi-block read and write
- Write-protection switch for write operations
- Synchronous abort
- Pause during the data transfer at a block gap
- Auto CMD12 for multi-block transfer
- Host can initiate non-data transfer commands while the data transfer is in progress
- Fully configurable 128 x 32-bit FIFO for read/write data
- Internal DMA capabilities
- Supports booting from eSDHC. See [eSDHC boot](#) for more detailed information

17.2.1 Data transfer modes

The eSDHC can select the following modes for data transfer:

- SD 1-bit
- SD 4-bit
- MMC 1-bit
- MMC 4-bit
- MMC 8-bit
- Identification mode (up to 400 kHz)
- MMC full-speed mode (up to 20 MHz)
- MMC high-speed mode (up to 52 MHz)
- SD full-speed mode (up to 25 MHz)
- SD high-speed mode (up to 50 MHz)

17.3 eSDHC external signal description

The eSDHC has 12 chip I/O signals.

- SDHC_CLK is the internally generated clock signal that drives the MMC or SD card.

- SDHC_CMD I/O sends commands and receives responses from the card.
- SDHC_DAT7 -SDHC_DAT0 perform data transfers between the eSDHC and the card. If the eSDHC is desired to support a 4-bit data transfer, SDHC_DAT7-SDHC_DAT4 can also be optional and tied high.
- SDHC_CD_B and SDHC_WP are card detection and write protection signals from the socket.

Table 17-1. Signal properties

Name	Port	Function	Reset State	Pull up/Pull down Required
SDHC_CLK	O	Clock for MMC/SD card	0	N/A
SDHC_CMD	I/O	Command line to card	High impedance	Pull up
SDHC_DAT7	I/O	8-bit mode: DAT7 line not used in other modes	High impedance	Pull up
SDHC_DAT6	I/O	8-bit mode: DAT6 line not used in other modes	High impedance	Pull up
SDHC_DAT5	I/O	8-bit mode: DAT5 line not used in other modes	High impedance	Pull up
SDHC_DAT4	I/O	8-bit mode: DAT4 line in not used in other modes	High impedance	Pull up
SDHC_DAT3	I/O	4-/8-bit mode: DAT3 line	High impedance	Pull up. Do not use DAT3 pin as a CD pin.
SDHC_DAT2	I/O	4-/8-bit mode: DAT2 line	High impedance	Pull up
SDHC_DAT1	I/O	8-bit mode: DAT1 line 4-bit mode: DAT1 line	High impedance	Pull up
SDHC_DAT0	I/O	8-bit mode: DAT0 line 4-bit mode: DAT0 line	High impedance	Pull up
SDHC_WP	I	Card write protect detect; if not used, tied low	N/A	N/A
SDHC_CD_B	I	Card detection pin; if not used, tie high. Low Card present High No card present	N/A	N/A

17.4 Enhanced Secure Digital Host Controller (eSDHC) Memory Map

The table below shows the memory mapped registers of the eSDHC module and their offsets. It lists the offset, name, and a cross-reference to the complete description of each register. Note that the full register address is comprised of CCSRBAR together with the eSDHC block base address and offset listed in the following table. Undefined 4-byte address spaces within offset 0x0000-0xFFFF are reserved.

NOTE

All eSDHC registers must be accessed as aligned 4-byte quantities. Accesses to the eSDHC registers that are less than 4-bytes are not supported.

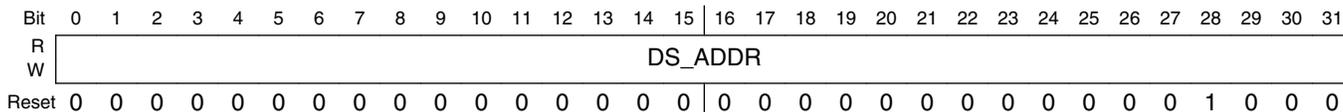
eSDHC memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/ page
2_E000	DMA system address (eSDHC_DSADDR)	32	R/W	0000_0008h	17.4.1/1487
2_E004	Block attributes (eSDHC_BLKATTR)	32	R/W	0000_0008h	17.4.2/1487
2_E008	Command argument (eSDHC_CMDARG)	32	R/W	0000_0000h	17.4.3/1488
2_E00C	Command transfer type (eSDHC_XFERTYP)	32	R/W	0000_0000h	17.4.4/1489
2_E010	Command response n (eSDHC_CMDRSP0)	32	R	0000_0000h	17.4.5/1492
2_E014	Command response n (eSDHC_CMDRSP1)	32	R	0000_0000h	17.4.5/1492
2_E018	Command response n (eSDHC_CMDRSP2)	32	R	0000_0000h	17.4.5/1492
2_E01C	Command response n (eSDHC_CMDRSP3)	32	R	0000_0000h	17.4.5/1492
2_E020	Data buffer access port (eSDHC_DATPORT)	32	R/W	0000_0000h	17.4.6/1493
2_E024	Present state (eSDHC_PRSTAT)	32	R	See section	17.4.7/1494
2_E028	Protocol control (eSDHC_PROCTL)	32	R/W	0000_0020h	17.4.8/1499
2_E02C	System control (eSDHC_SYSCTL)	32	R/W	0000_8000h	17.4.9/1502
2_E030	Interrupt status (eSDHC_IRQSTAT)	32	R/W	0000_0000h	17.4.10/1505
2_E034	Interrupt status enable (eSDHC_IRQSTATEN)	32	R/W	117F_013Fh	17.4.11/1510
2_E038	Interrupt signal enable (eSDHC_IRQSIGEN)	32	R/W	0000_0000h	17.4.12/1513
2_E03C	Auto CMD12 status (eSDHC_AUTO12ERR)	32	R	0000_0000h	17.4.13/1515
2_E040	Host controller capabilities (eSDHC_HOSTCAPBLT)	32	R	05E3_0000h	17.4.14/1518
2_E044	Watermark level (eSDHC_WML)	32	R/W	0010_0010h	17.4.15/1520
2_E050	Force event (eSDHC_FEVT)	32	W	0000_0000h	17.4.16/1521
2_E0FC	Host controller version (eSDHC_HOSTVER)	32	R	0000_1001h	17.4.17/1523
2_E40C	DMA control register (eSDHC_DCR)	32	R/W	0000_0000h	17.4.18/1524

17.4.1 DMA system address (eSDHC_DSADDR)

The DMA system address register contains the system memory address used for DMA transfers. Only access this register when no transactions are executing (after transactions have stopped). The host driver should wait until PRSSTAT[DLA] is cleared.

Address: 2_E000h base + 0h offset = 2_E000h



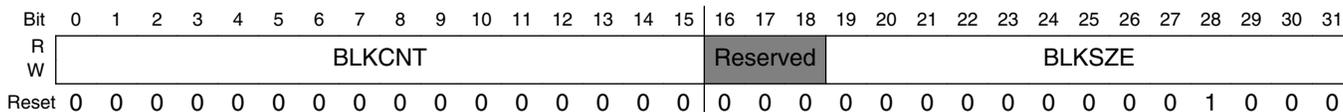
eSDHC_DSADDR field descriptions

Field	Description
0–31 DS_ADDR	DMA system address , lower 32 bits . When the eSDHC stops a DMA transfer, this register points to the system address of the next contiguous data position. The upper four bits of the DMA system address are stored in ECMCR[ESDHC_UPRADDR], see ECM control register (GUTS_ECMCR) . NOTE: The DS_ADDR must be aligned to a four-byte boundary; the two least-significant bits must be cleared.

17.4.2 Block attributes (eSDHC_BLKATTR)

The block attributes register configures the number of data blocks and the number of bytes in each block. Only access this register when no transactions are executing (after transactions have stopped). The host driver should wait until PRSSTAT[DLA] is cleared.

Address: 2_E000h base + 4h offset = 2_E004h



eSDHC_BLKATTR field descriptions

Field	Description
0–15 BLKCNT	Block count for current transfer. This field is enabled when XFERTYP[BCEN] is set and is valid only for multiple block transfers. The host driver should set this field to a value between 1 and the maximum block count. The eSDHC decrements the block count after each block transfer and stops when the count reaches zero. Clearing this field results in no data blocks being transferred. 0000 Stop count 0001 1 block 0002 2 blocks

Table continues on the next page...

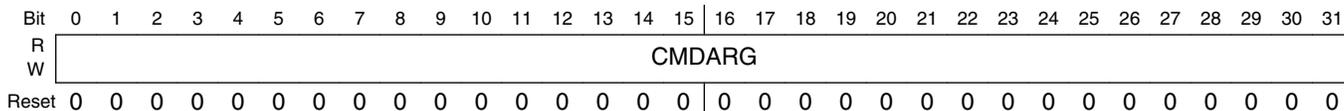
eSDHC_BLKATTR field descriptions (continued)

Field	Description
	... FFFF 65,535 blocks
16–18 -	This field is reserved.
19–31 BLKSZE	Transfer block size. Specifies the block size for block data transfers. Values can range from one byte up to the maximum buffer size. The DMA always writes at least four bytes to memory. Thus, software should allocate a buffer space rounded up to a 4-byte aligned size in order to avoid data corruption. 0000 No data transfer 0001 1 byte 0002 2 bytes 0003 3 bytes 0004 4 bytes ... 01FF 511 bytes 0200 512 bytes ... 0800 2048 bytes 1000 4096 bytes

17.4.3 Command argument (eSDHC_CMDARG)

The command argument register contains the SD/MMC command argument.

Address: 2_E000h base + 8h offset = 2_E008h



eSDHC_CMDARG field descriptions

Field	Description
0–31 CMDARG	Command argument. The SD/MMC command argument is specified as bits 39-8 of the command format in the SD or MMC Specification . If PRSSTAT[CIHB] is set, this register is write-protected.

17.4.4 Command transfer type (eSDHC_XFERTYP)

The transfer type register controls the operation of data transfers. The host driver should set this register before issuing a command followed by a data transfer. To prevent data loss, the eSDHC prevents a write to the bits that are involved in the data transfer of this register while the data transfer is active.

The host driver should check PRSSTAT[CDIHB] and PRSSTAT[CIHB] before writing to this register.

- If PRSSTAT[CDIHB] is set, any attempt to send a command with data by writing to this register is ignored.
- If PRSSTAT[CIHB] is set, any write to this register is ignored.

Table 17-6. Determination of transfer type

Multi/Single Block Select XFERTYP[MSBSEL]	Block Count Enable XFERTYP[BCEN]	Block Count BLKATTR[BLKCNT]	Function
0	X	X	Single Transfer
1	0	X	Infinite Transfer
1	1	Positive Number	Multiple Transfer
1	1	Zero	No Data Transfer

The table below shows how the response type can be determined by the command index check enable, command CRC check enable, and response type bits.

Table 17-7. Relation between parameters and name of response type

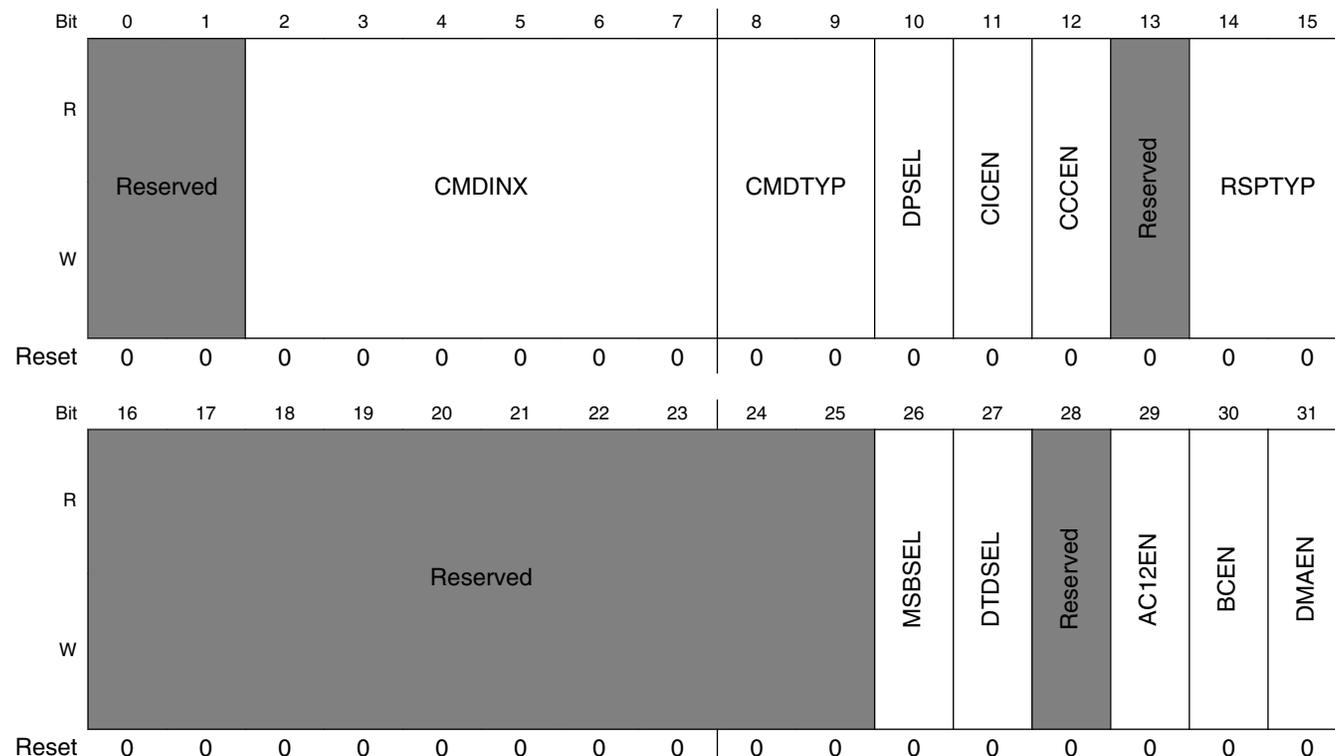
Response Type XFERTYP[RSPTYP]	Index Check Enable XFERTYP[CICEN]	CRC Check Enable XFERTYP[CCEN]	Response type
00	0	0	No Response
01	0	1	R2
10	0	0	R3, R4
10	1	1	R1, R5, R6
11	1	1	R1b, R5b

NOTE

The CRC field for R3 and R4 is expected to be all 1s. The CRC check should be disabled for these response types.

Enhanced Secure Digital Host Controller (eSDHC) Memory Map

Address: 2_E000h base + Ch offset = 2_E00Ch



eSDHC_XFERTYP field descriptions

Field	Description
0-1 -	This field is reserved.
2-7 CMDINX	Command index. These bits should be set to the command number (CMD0-63, ACMD0-63) that is specified in bits 45-40 of the command format in the <i>SD Memory Card Physical Layer Specification</i> .
8-9 CMDTYP	Command type. There is one type of special command-abort. Clear this bit field for all other commands. <ul style="list-style-type: none"> Abort command. If this command is set when executing a read transfer, the eSDHC stops reads to the buffer. If this command is set when executing a write transfer, the eSDHC stops driving the SDHC_DAT line. After issuing the abort command, the host driver should issue a software reset. (Abort transaction) 00 Normal-other commands 01 Reserved 10 Reserved 11 Abort-CMD12, CMD52 for writing I/O abort in CCCR
10 DPSEL	Data present select. Set to indicate that data is present and should be transferred using the SDHC_DAT line. It is cleared for the following: <ul style="list-style-type: none"> Commands using only the SDHC_CMD line (for example, CMD52) Commands with no data transfer but using busy signal on the SDHC_DAT[0] line (R1b or R5b, for example, CMD38) 0 No data present 1 Data present
11 CICEN	Command index check enable.

Table continues on the next page...

eSDHC_XFERTYP field descriptions (continued)

Field	Description
	<p>0 Disable. The index field is not checked.</p> <p>1 Enable. The eSDHC checks the index field in the response to see if it has the same value as the command index. If it is not, it is reported as a command index error.</p>
12 CCEN	<p>Command CRC check enable. The number of bits checked by the CRC field value changes according to the length of the response. (See RSPTYP[1:0] and Table 17-6.)</p> <p>0 Disable. The CRC field is not checked.</p> <p>1 Enable. The eSDHC checks the CRC field in the response if it contains the CRC field. If an error is detected, it is reported as a command CRC error.</p>
13 -	This field is reserved.
14–15 RSPTYP	<p>Response type select.</p> <p>00 No response</p> <p>01 Response length 136</p> <p>10 Response length 48</p> <p>11 Response length 48 check busy after response</p>
16–25 -	This field is reserved.
26 MSBSEL	<p>Multi/single block select. Enables multiple block SDHC_DAT line data transfers. For any other commands, this bit should be cleared. If this bit is cleared, it is not necessary to set the block count register. (See Table 17-7.)</p> <p>0 Single block</p> <p>1 Multiple blocks</p>
27 DTDSEL	<p>Data transfer direction select. Defines the direction of SDHC_DAT line data transfers. The bit is set by the host driver to transfer data from the SD card to the eSDHC and it is cleared for all other commands.</p> <p>0 Write (host to card)</p> <p>1 Read (card to host)</p>
28 -	This field is reserved.
29 AC12EN	<p>Auto CMD12 enable. Multiple block transfers for memory require CMD12 to stop the transaction. If this bit is set, the eSDHC issues CMD12 automatically when the last block transfer is completed. The host driver should not set this bit to issue commands that do not require CMD12 to stop a multiple block data transfer. In particular, secure commands defined in the Part 3 File Security specification do not require CMD12. In a single block transfer, the eSDHC ignores this bit.</p> <p>0 Disable</p> <p>1 Enable</p>
30 BCEN	<p>Block count enable. Enables the block attributes register, which is only relevant for multiple block transfers. When this bit is cleared, the block attributes register is disabled, which is useful in executing an infinite transfer.</p> <p>0 Disable</p> <p>1 Enable</p>
31 DMAEN	<p>DMA enable. Enables DMA functionality as described in DMA CCB interface. If this bit is set, a DMA operation should begin when the host driver writes to the CMDINX field of the transfer type register.</p>

Table continues on the next page...

eSDHC_XFERTYP field descriptions (continued)

Field	Description
0	Disable
1	Enable

17.4.5 Command response n (eSDHC_CMDRSPn)

The command response registers store the four parts of the response bits from the card.

The table below describes the mapping of command responses from the SD bus to the command response registers for each response type. In the table, R[] refers to a bit range within the response data as transmitted on the SD bus.

Table 17-9. Response bit definition for each response type

Response type	Meaning of response	Response field	Response register
R1,R1b (normal response)	Card status	R[39:8]	CMDRSP0
R1b (Auto CMD12 response)	Card status for Auto CMD12	R[39:8]	CMDRSP3
R2 (CID, CSD register)	CID/CSD register [127:8]	R[127:8]	{CMDRSP3[23:0], CMDRSP2, CMDRSP1, CMDRSP0}
R3 (OCR register)	OCR register for memory	R[39:8]	CMDRSP0
R4 (OCR register)	OCR register for I/O etc.	R[39:8]	CMDRSP0
R6 (publish RCA)	New published RCA[31:16] and card status[15:0]	R[39:8]	CMDRSP0

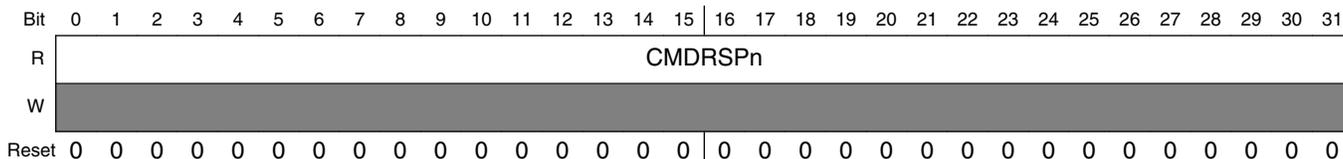
This table shows the following:

- Most responses with a length of 48 (R[47:0]) have 32 bits of the response data (R[39:8]) stored in the CMDRSP0 register.
- Responses of type R1b (Auto CMD12 responses) have response data bits R[39:8] stored in the CMDRSP3 register.
- Responses with length 136 (R[135:0]) have 120 bits of the response data (R[127:8]) stored in the CMDRSP0, 1, 2, and 3 registers.

To be able to read the response status efficiently, the eSDHC only stores part of the response data in the command response registers. This enables the host driver to efficiently read 32 bits of response data in one read cycle on a 32-bit bus system. Parts of the response, the index field, and the CRC are checked by the eSDHC (as specified by XFERTYP[CICEN, CCCEN]) and generate an error interrupt if any error is detected. The bit range for the CRC check depends on the response length. If the response length is 48, the eSDHC checks R[47:1], and if the response length is 136, the eSDHC checks R[119:1].

Since the eSDHC may have a multiple block data transfer executing concurrently with a CMD_wo_DAT (Command without data) command, the eSDHC stores the Auto CMD12 response in the CMDRSP3 register and the CMD_wo_DAT response is stored in CMDRSP0. This allows the eSDHC to avoid overwriting the Auto CMD12 response with the CMD_wo_DAT and vice versa. When the eSDHC modifies part of the command response registers it preserves the unmodified bits.

Address: 2_E000h base + 10h offset + (4d × i), where i=0d to 3d



eSDHC_CMDRSPn field descriptions

Field	Description
0–31 CMDRSPn	Command response. See Command Response 0-3 for the mapping of command responses from the SD bus to this register for each response type.

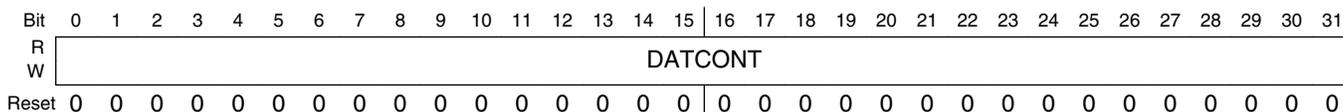
17.4.6 Data buffer access port (eSDHC_DATPORT)

The buffer data port register is a 32-bit data port register used to access the internal buffer.

NOTE

When the internal DMA is not enabled and a write transaction is in operation, DATPORT must not be read. DATPORT also must not be used to read (or write) data by the CPU or external DMA if the data will be written (or read) by the eSDHC internal DMA.

Address: 2_E000h base + 20h offset = 2_E020h



eSDHC_DATPORT field descriptions

Field	Description
0–31 DATCONT	Data content. The buffer data port register is for 32-bit data access by the CPU or an external DMA . When the internal DMA is enabled, any write to this register is ignored, and a read from this register always yields 0.

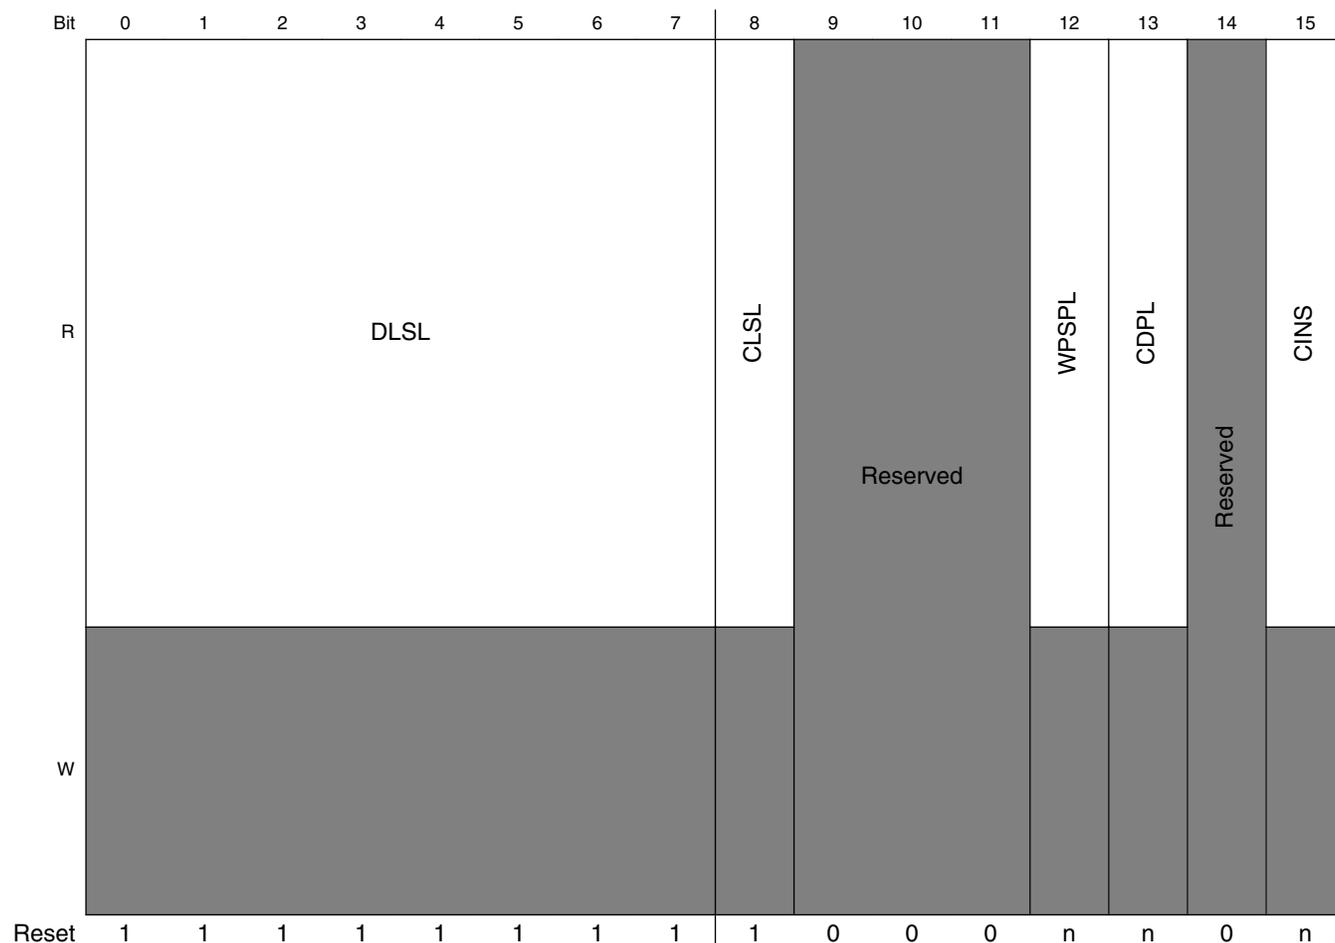
17.4.7 Present state (eSDHC_PRSTAT)

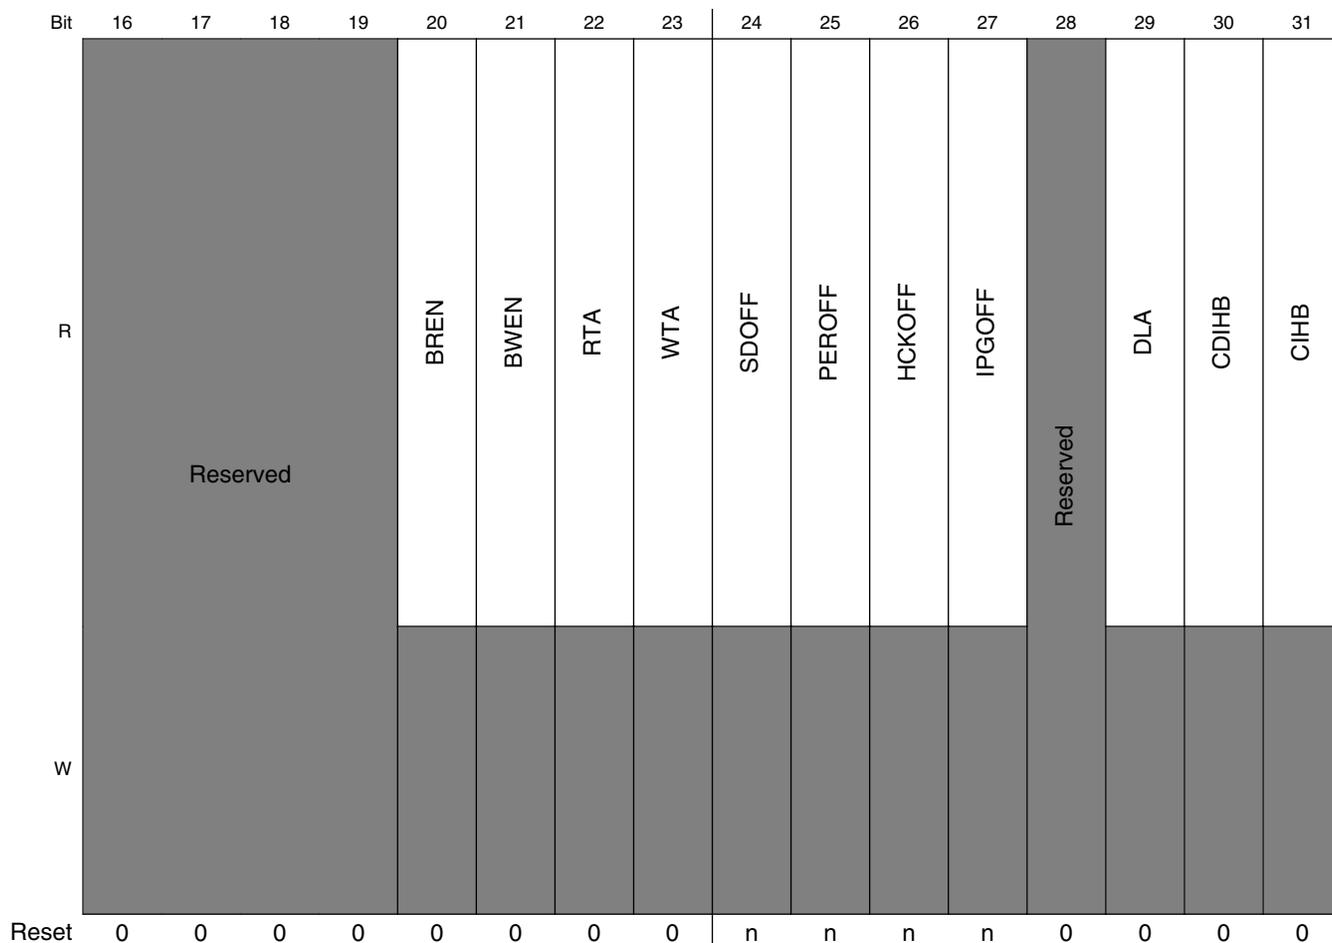
The present state register indicates the status of the eSDHC to the host driver.

NOTE

The host driver can issue CMD0, CMD12, CMD13 (for memory) when the SDHC_DAT lines are busy during a data transfer. These commands can be issued when PRSSTAT[CIHB] is cleared. Other commands should be issued when PRSSTAT[CDIHB] is cleared. Possible changes to the SD physical specification may add other commands to this list in the future.

Address: 2_E000h base + 24h offset = 2_E024h





eSDHC_PRSTAT field descriptions

Field	Description																
0–7 DLSL	<p>SDHC_DAT[7:0] line signal level. These bits are used to check the SDHC_DAT line level to recover from errors, and for debugging. This is especially useful in detecting the busy signal level from SDHC_DAT[0]. The reset value is affected by the external pull resistors. By default, read value of this bit field after reset is 0b11110111, when SDHC_DAT[3] is pull-down and other lines are pull-up.</p> <p>The meaning of PRSSTAT bits is as follows:</p> <p>PRSSTAT SDHC_DATn</p> <table border="0"> <tr><td>0</td><td>7</td></tr> <tr><td>1</td><td>6</td></tr> <tr><td>2</td><td>5</td></tr> <tr><td>3</td><td>4</td></tr> <tr><td>4</td><td>3</td></tr> <tr><td>5</td><td>2</td></tr> <tr><td>6</td><td>1</td></tr> <tr><td>7</td><td>0</td></tr> </table>	0	7	1	6	2	5	3	4	4	3	5	2	6	1	7	0
0	7																
1	6																
2	5																
3	4																
4	3																
5	2																
6	1																
7	0																
8 CLSL	<p>SDHC_CMD line signal level. This status is used to check the SDHC_CMD line level to recover from errors, and for debugging. The reset value is affected by the external pull resistor, by default, read value of this bit after reset is 1, when the command line is pull-up.</p>																

Table continues on the next page...

eSDHC_PRSTAT field descriptions (continued)

Field	Description
9–11 -	This field is reserved.
12 WPSPL	Write protect switch pin level. The write protect switch is supported for memory cards. This bit reflects the inverse of the SDHC_WP pin of the card socket. A software reset does not affect this bit. The reset value is affected by the external write protect switch. If the SDHC_WP pin is not used (PMUXCR[SDHC_WP] = 0, Alternate function signal multiplex control (GUTS_PMUXCR)), so that the reset value of this bit is 1 and write is enabled. 0 Write protected (SDHC_WP = 1) 1 Write enabled (SDHC_WP = 0)
13 CDPL	Card detect pin level. This bit reflects the inverse value of the SDHC_CD_B pin for the card socket. Debouncing is not performed on this bit. This bit may be valid, but it is not guaranteed because of a propagation delay. Use of this bit is limited to testing since it must be debounced by software. A software reset does not affect this bit. Write to the force event register does not affect this bit. The reset value is affected by the external card detection pin. If this bit is not used, it should be tied to 0. 0 No card present (SDHC_CD_B = 1) 1 Card present (SDHC_CD_B = 0)
14 -	This field is reserved.
15 CINS	Card inserted. Indicates if a card has been inserted. The eSDHC debounces this signal so that the host driver does not need to wait for it to stabilize. Changing from 0 to 1 generates a card-insertion interrupt in the interrupt status register and changing from 1 to 0 generates a card removal interrupt in the interrupt status register. A write to the force event register does not affect this bit. The software reset for all in the system control register does not affect this bit. A software reset does not affect this bit. 0 Power-on-reset or no card 1 Card inserted
16–19 -	This field is reserved.
20 BREN	Buffer read enable. This status is used for non-DMA read transfers. The eSDHC may implement multiple buffers to transfer data efficiently. This read-only flag indicates that a burst-length of valid data exists in the host-side buffer. When the buffer is read, this bit is cleared. When a burst length of data is ready in the buffer, this bit is set and a buffer read ready interrupt is generated (if the interrupt is enabled). 0 Buffer read disable 1 Buffer read enable
21 BWEN	Buffer write enable. This status is used for non-DMA write transfers. The eSDHC can implement multiple buffers to transfer data efficiently. This read-only flag indicates if space is available for a burst length of write data. When the buffer is written, this bit is cleared. When a burst length of data is written to the buffer, this bit is set and a buffer write ready interrupt is generated (if the interrupt is enabled). 0 Buffer write disable 1 Buffer write enable
22 RTA	Read transfer active. This status is used for detecting completion of a read transfer. This bit is set for either of the following conditions:

Table continues on the next page...

eSDHC_PRSTAT field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> After the end bit of the read command When writing a 1 to PROCTL[CREQ] to restart a read transfer <p>This bit is cleared for either of the following conditions:</p> <ul style="list-style-type: none"> When the last data block as specified by block length is transferred to the system When all valid data blocks have been transferred to the system and no current block transfers are being sent as a result of PROCTL[SABGREQ] being set. A transfer complete interrupt is generated when this bit changes to 0. <p>0 No valid data 1 Transferring data</p>
23 WTA	<p>Write transfer active. This status indicates a write transfer is active. If this bit is 0, it means no valid write data exists in eSDHC.</p> <p>This bit is set in either of the following cases:</p> <ul style="list-style-type: none"> After the end bit of the write command. When writing a 1 to PROCTL[CREQ] to restart a write transfer. <p>This bit is cleared in either of the following cases:</p> <ul style="list-style-type: none"> After getting the CRC status of the last data block, as specified by the transfer count (single and multiple) After getting the CRC status of any block where data transmission is about to be stopped by a stop-at-block-gap request. <p>During a write transaction, a IRQSTAT[BGE] interrupt is generated when this bit is changed to 0, as result of PROCTL[SABGREQ] being set. This status is useful for the host driver in determining when to issue commands during write busy.</p> <p>0 No valid data 1 Transferring data</p>
24 SDOFF	<p>SD clock gated off internally. Indicates the SD clock is internally gated off because of a buffer overrun, buffer underrun, or a read pause without read-wait assertion. This bit is for the host driver to debug data transaction on SD bus.</p> <p>This status bit resets to 0, but reflects the value of the automatic clock gating and may transition to 1 if the eSDHC is idle.</p>
25 PEROFF	<p>The internal bus clock gated off internally. This status bit indicates the internal bus clock is internally gated off. This bit is for the host driver to debug a transaction on SD bus.</p> <p>This status bit resets to 0, but reflects the value of the automatic clock gating and may transition to 1 if the eSDHC is idle.</p>
26 HCKOFF	<p>Master clock gated off internally. This status bit indicates master clock is internally gated off. This bit is for the host driver to debug a data transfer.</p> <p>This status bit resets to 0, but reflects the value of the automatic clock gating and may transition to 1 if the eSDHC is idle.</p>
27 IPGOFF	<p>Controller clock gated off internally. Indicates that the controller clock is internally gated off. This bit is for the host driver to debug.</p> <p>This status bit resets to 0, but reflects the value of the automatic clock gating and may transition to 1 if the eSDHC is idle.</p>
28 -	<p>This field is reserved.</p>

Table continues on the next page...

eSDHC_PRSTAT field descriptions (continued)

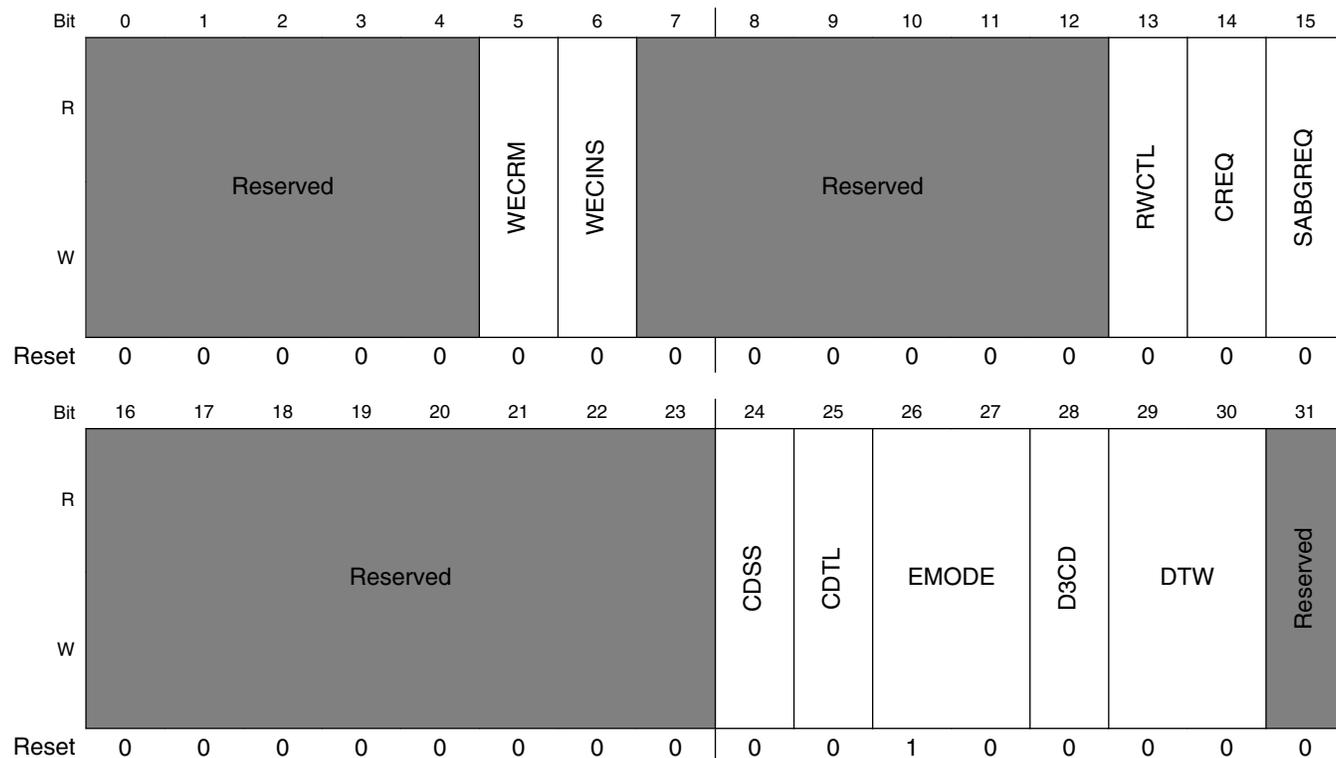
Field	Description
<p>29 DLA</p>	<p>Data line active. Indicates whether one of the SDHC_DAT line on SD bus is in use.</p> <p>For read transactions, this bit indicates if a read transfer is executing on the SD bus. Clearing this bit from 1 to 0 between data blocks generates a block gap event interrupt.</p> <p>This bit is set in either of the following cases:</p> <ul style="list-style-type: none"> • After the end bit of the read command • When writing a 1 to PROCTL[CREQ] to restart a read transfer <p>This bit is cleared in either of the following cases:</p> <ul style="list-style-type: none"> • When the end bit of the last data block is sent from the SD bus to the eSDHC • When beginning a read wait transfer initiated by a stop at block gap request <p>The eSDHC waits at the next block gap by driving read wait at the start of the interrupt cycle. If the read-wait signal is already driven (data buffer cannot receive data), the eSDHC can wait for current block gap by continuing to drive the read-wait signal. It is necessary to support read wait in order to use the suspend/resume function.</p> <p>For write transactions, this bit indicates that a write transfer is executing on the SD bus. Clearing this bit from 1 to 0 generates a transfer complete interrupt.</p> <p>This bit is set in any of the following cases:</p> <ul style="list-style-type: none"> • After the end bit of the write command • When writing a 1 to PROCTL[CREQ] to continue a write transfer <p>This bit is cleared in any of the following cases:</p> <ul style="list-style-type: none"> • When the SD card releases write-busy of the last data block, the eSDHC also detects if output is not busy. If the SD card does not drive the busy signal after CRC status is received, the eSDHC should consider the card drive not busy. • When the SD card releases write-busy prior to waiting for write transfer as a result of a stop at block gap request <p>0 SDHC_DAT line inactive 1 SDHC_DAT line active</p>
<p>30 CDIHB</p>	<p>Command inhibit (SDHC_DAT). This bit is set if the SDHC_DAT line is active, the read transfer active is set. If this bit is cleared, it indicates the eSDHC can issue the next SD/MMC command. Commands with busy signal belong to command inhibit (SDHC_DAT) (e.g. R1b and R5b type). Clearing from 1 to 0 generates a transfer complete interrupt.</p> <p>0 Can issue command which uses the SDHC_DAT line 1 Cannot issue command which uses the SDHC_DAT line</p>
<p>31 CIHB</p>	<p>Command inhibit (SDHC_CMD). This bit is cleared, if the SDHC_CMD line is not in use and the eSDHC can issue a SD/MMC command using the SDHC_CMD line.</p> <p>This bit is set immediately after the XFERTYP register is written. This bit is cleared when the command response is received. Even if the CDIHB bit is set, commands using only the SDHC_CMD line can be issued if this bit is cleared. Clearing from 1 to 0 generates a command complete interrupt.</p> <p>If the eSDHC cannot issue the command because of a command conflict error (see the command CRC error) or because of a command not issued by Auto CMD12 error, this bit remains set and IRQSTAT[CC] is not set. Status issuing Auto CMD12 is not read from this bit.</p> <p>0 Can issue command using only SDHC_CMD line 1 Cannot issue command</p>

17.4.8 Protocol control (eSDHC_PROCTL)

To restart the transfer after a stop at the block gap, the continue request should be used to restart the transfer.

Any time PROCTL[SABGREQ] stops the data transfer, the host driver should wait for IRQSTAT[TC] before attempting to restart the transfer. When restarting the data transfer by continue request, the host driver should clear PROCTL[SABGREQ] before or simultaneously.

Address: 2_E000h base + 28h offset = 2_E028h



eSDHC_PROCTL field descriptions

Field	Description
0-4 -	This field is reserved.
5 WECRM	Wake-up event enable on SD card removal. This bit enables wakeup event via card removal assertion in the IRQSTAT register. FN_WUS (wake-up support) in CIS does not affect this bit. 0 Disable 1 Enable
6 WECINS	Wake-up event enable on SD card insertion. This bit enables wakeup event via card insertion assertion in the IRQSTAT register. FN_WUS (wake-up support) in CIS does not affect this bit.

Table continues on the next page...

eSDHC_PROCTL field descriptions (continued)

Field	Description
	<p>NOTE: When this bit is set, the card insertion status and the eSDHC interrupt can be asserted without SD_CLK toggling. When the wakeup feature is not enabled, the SD_CLK must be active in order to assert the card insertion status and the eSDHC interrupt.</p> <p>0 Disable 1 Enable</p>
7–12 -	This field is reserved.
13 RWCTL	<p>Read wait control. Not supported on this device.</p> <p>If the card supports read wait, set this bit to enable the read wait protocol to stop read data using the SDHC_DAT[2] line. Otherwise, the eSDHC has to stop the SD clock to hold read data, which restricts command generation.</p> <p>If the card does not support read wait, this bit should never be set otherwise an SDHC_DAT line conflict may occur. If this bit is cleared, a stop-at-block-gap-during-read operation is also supported, but the eSDHC stops the SD clock to pause the reading operation.</p> <p>0 Disable read-wait control, and stop SD clock at block gap when the SABGREQ bit is set 1 Enable read-wait control, and assert read wait without stopping the SD clock at block gap when PROCTL[SABGREQ] is set</p>
14 CREQ	<p>Continue request. Restarts a transaction which was stopped using the stop-at-block-gap request. To cancel the request, clear SABGREQ and set this bit to restart the transfer.</p> <p>The eSDHC automatically clears this bit in either of the following cases:</p> <ul style="list-style-type: none"> • For a read transaction, the PRSSTAT[DLA] bit changes from 0 to 1 as a read transaction restarts. • For a write transaction, the PRSSTAT[WTA] bit changes from 0 to 1 as the write transaction restarts. <p>Therefore, it is not necessary for the host driver to clear. If SABGREQ and this bit are set, the continue request is ignored.</p> <p>0 No effect 1 Restart</p>
15 SABGREQ	<p>Stop at block gap request. Stops executing a transaction at the next block gap for both DMA and non-DMA transfers. Until the TC bit is set, indicating a transfer completion, the host driver should leave this bit set. Clearing SABGREQ and CREQ does not cause the transaction to restart.</p> <p>The eSDHC stops the SD bus clock to pause the read operation during the block gap.</p> <p>For write transfers in which the host driver writes data to the data port register, the host driver should set this bit after all block data is written. If this bit is set, the host driver should not write data to the DATPORT register after a block is sent. When this bit is set, the host driver should not clear this bit before IRQSTAT[TC] is set. Otherwise, the eSDHC behavior is undefined. Confirm that IRQSTAT[TC] is enabled.</p> <p>This bit affects PRSSTAT[RTA, WTA, DLA, CIHB].</p> <p>0 Transfer 1 Stop</p>
16–23 -	This field is reserved.
24 CDSS	<p>Card detect signal selection. Selects the source for card detection.</p> <p>0 Card detection level is selected (for normal purpose) 1 Card detection test level is selected (for test purpose)</p>

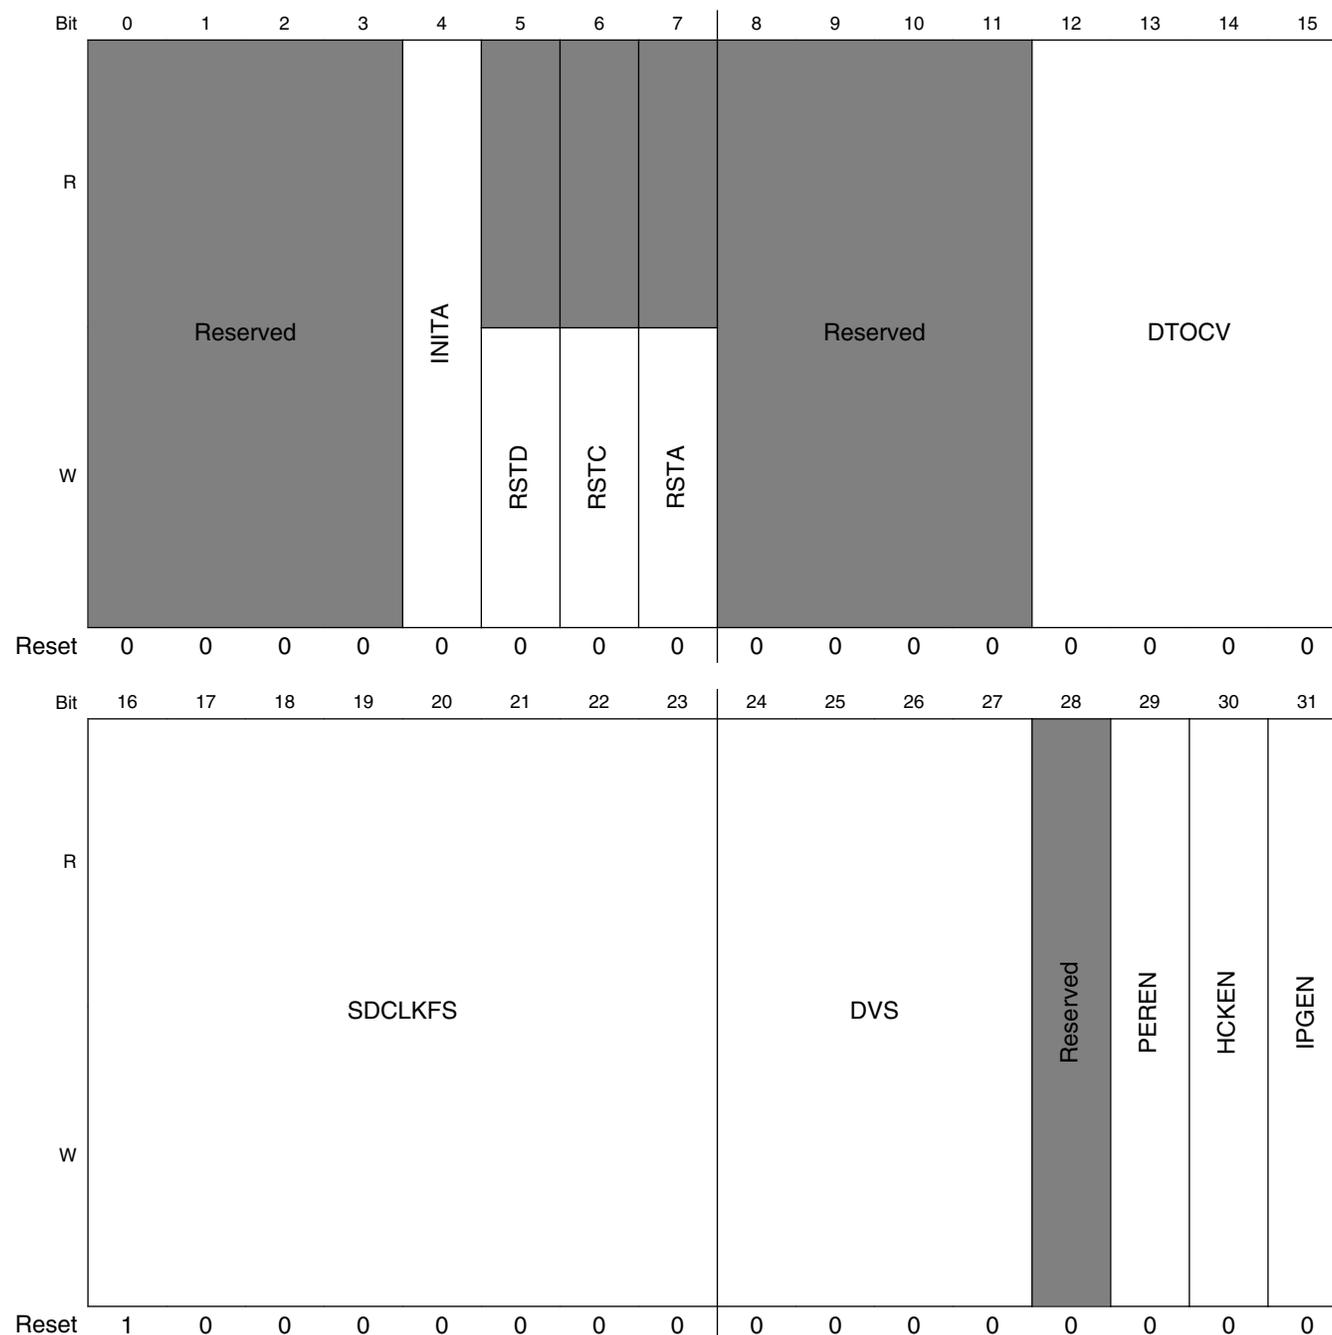
Table continues on the next page...

eSDHC_PROCTL field descriptions (continued)

Field	Description
25 CDTL	Card detect test level. Determines card insertion status when CDSS is set. 0 No card in the slot 1 Card is inserted
26–27 EMODE	Endian mode. eSDHC supports only address-invariant mode in data transfer. 00 Reserved 01 Reserved 10 Address-invariant mode. Each byte location in the main memory is mapped to the same byte location in the MMC/SD card. 11 Reserved
28 D3CD	SDHC_DAT3 as card detection pin. If this bit is set, SDHC_DAT3 should be pulled down to act as a card detection pin. Be cautious when using this feature, because SDHC_DAT3 is chip-select for SPI mode, and a pull-down on this pin and CMD0 may set the card into SPI mode, which the eSDHC does not support. 0 SDHC_DAT3 does not monitor card insertion 1 SDHC_DAT3 is card detection pin
29–30 DTW	Data transfer width. Selects the data width of the SD bus. The host driver should set it to match the data width of the card. 00 1-bit mode 01 4-bit mode 10 8-bit mode 11 Reserved
31 -	This field is reserved.

17.4.9 System control (eSDHC_SYSCTL)

Address: 2_E000h base + 2Ch offset = 2_E02Ch



eSDHC_SYSCTL field descriptions

Field	Description
0-3 -	This field is reserved.

Table continues on the next page...

eSDHC_SYSCTL field descriptions (continued)

Field	Description
4 INITA	<p>Initialization active. When this bit is written '1', 80 SD clocks are sent to the card. After the 80 clocks are sent, this bit is self-cleared. This bit is very useful during the card power-up period when 74 SD clocks are needed and clock auto-gating feature is enabled.</p> <p>Writing one to this bit when it is already set has no effect. Clearing this bit at any time does not affect it. When PRSSTAT[CIHB] or PRSSTAT[CDIHB] is set, writing a one to this bit is ignored. That is, when the command line or data line is active, writing to this bit is not allowed.</p>
5 RSTD	<p>Software reset for SDHC_DAT line. The DMA and part of the data circuit are reset. The following registers and bits are cleared by this bit:</p> <ul style="list-style-type: none"> • DATPORT register • Buffer is cleared and initialized; PRSSTAT register • PRSSTAT[BREN, BWEN, RTA, WTA, DLA, CDIHB] • PROCTL[CREQ, SABGREQ] • IRQSTAT[BRR, BWR, DINT, BGE, TC] <p>0 Work 1 Reset</p>
6 RSTC	<p>Software reset for SDHC_CMD line. Only part of the command circuit is reset. The following bits are cleared by this bit:</p> <ul style="list-style-type: none"> • PRSSTAT[CIHB] • IRQSTAT[CC] <p>0 Work 1 Reset</p>
7 RSTA	<p>Software reset for all. This reset affects the entire host controller except for the card-detection circuit. Register bits of type ROC, RW, RW1C, and RWAC are cleared.</p> <p>During its initialization, the host driver should set this bit to reset the eSDHC. The eSDHC should clear this bit when capabilities registers are valid and the host driver can read them. Additional use of the this bit does not affect the value of the capabilities registers. After this bit is set, it is recommended the host driver reset the external card and re-initialize it.</p> <p>0 Work 1 Reset</p>
8–11 -	This field is reserved.
12–15 DTCV	<p>Data timeout counter value. Determines the interval by which SDHC_DAT line timeouts are detected. See the data timeout error Interrupt status (eSDHC_IRQSTAT) , for information on factors that dictate timeout generation. Timeout clock frequency is generated by dividing the base clock SDHC_CLK value by this value. When setting this register, prevent inadvertent timeout events by clearing IRQSTATEN[DTOESEN].</p> <p>0000 SDHC_CLK x 2¹³ 0001 SDHC_CLK x 2¹⁴ ... 1110 SDHC_CLK x 2²⁷ 1111 Reserved</p>
16–23 SDCLKFS	<p>SDHC_CLK frequency select. This field, together with DVS, selects the frequency of SDHC_CLK pin. This bit holds the prescaler of the base clock frequency. Only the following settings are allowed:</p> <p>Multiple bits must not be set or the behavior of this prescaler is undefined.</p> <p>According to the SD Physical Specification version 2.0 , the maximum SD clock frequency is 52 MHz, and should never exceed this limit. The frequency of SDHC_CLK is set by the following formula:</p>

Table continues on the next page...

eSDHC_SYSCTL field descriptions (continued)

Field	Description
	<p>clock frequency = (base clock) ÷ [(SDCLKFS x 2) x (DVS + 1)]</p> <p>For example, if the base clock frequency is 96 MHz, and the target frequency is 25 MHz, then choosing the prescaler value of 0x1 and divisor value of 0x1 yields 24 MHz, which is the nearest frequency less than or equal to the target. Similarly, to approach a clock value of 400 kHz, the prescaler value of 0x8 and divisor value of 0xE yields the exact clock value of 400 kHz.</p> <p>The reset value of this bit field is 0x80. So, if the input base clock is about 96 MHz, the default SD clock after reset is 375 kHz.</p> <p>NOTE: The base clock frequency equals the platform clock/2.</p> <p>0x01 is multiplied by 2 to give a base clock divisor of 2 0x02 is multiplied by 2 to give a base clock divisor of 4 0x04 is multiplied by 2 to give a base clock divisor of 8 0x08 is multiplied by 2 to give a base clock divisor of 16 0x10 is multiplied by 2 to give a base clock divisor of 32 0x20 is multiplied by 2 to give a base clock divisor of 64 0x40 is multiplied by 2 to give a base clock divisor of 128 0x80 is multiplied by 2 to give a base clock divisor of 256</p>
24–27 DVS	<p>Divisor. Provides a more exact divisor to generate the desired SD clock frequency. The settings are as follows:</p> <p>0x0 after adding 1, gives a divisor of 1 0x1 after adding 1, gives a divisor of 2 ... 0xE after adding 1, gives a divisor of 15 0xF after adding 1, gives a divisor of 16</p>
28 -	This field is reserved.
29 PEREN	<p>Peripheral clock enable. If set, the peripheral clock is always active and no automatic gating is applied, thus SDHC_CLK is active only except auto gating-off during buffer danger. If cleared, the peripheral clock is automatically off when no transaction is on the SD bus. Clearing this bit does not stop SDHC_CLK immediately. The peripheral clock will be internally gated off, if none of the following factors are met:</p> <ul style="list-style-type: none"> • Command part is reset • Data part is reset • Soft reset • Command is about to send • Clock divisor is just updated • Continue request is just set • This bit is set • Card insertion is detected • Card removal is detected • Card external interrupt is detected • 80 clocks for initialization phase is ongoing <p>0 The peripheral clock is internally gated off 1 The peripheral clock is not automatically gated off</p>
30 HCKEN	<p>Master clock enable. If set, master clock is always active and no automatic gating is applied. If cleared, master clock is automatically off when no data transfer is on SD bus.</p> <p>NOTE: Master clock is the clock to the DMA engine and to the system bus interface logic.</p>

Table continues on the next page...

eSDHC_SYSCTL field descriptions (continued)

Field	Description
	0 Master clock is internally gated off 1 Master clock is not automatically gated off
31 IPGEN	Controller clock enable. If this bit is set, the controller clock is always active and no automatic gating is applied. The controller clock is internally gated off, if neither the following factors is met: <ul style="list-style-type: none"> • Command part is reset • Data part is reset • Soft reset • Command is about to send • Clock divisor is just updated • Continue request is just set • This bit is set • Card insertion is detected • Card removal is detected • Card external interrupt is detected • The internal bus clock is not gated off <p>NOTE: The controller clock is not auto-gated off if the peripheral clock is not gated off. So, clearing this bit only does not take effect if SYSCTL[PEREN] is not cleared.</p> 0 The controller clock is internally gated off 1 The controller clock is not automatically gated off

17.4.10 Interrupt status (eSDHC_IRQSTAT)

An interrupt is generated when one of the status bits and its corresponding interrupt enable bit are set. For all bits (except BGE), writing one to a bit clears it, while writing zero keeps the bit unchanged. More than one status can be cleared with a single register write.

The table below shows that command timeout error has higher priority than command complete. If both bits are set, it can be assumed that the response was not received correctly.

Table 17-23. Relation between command timeout error and command complete status

Command complete	Command timeout error	Meaning of the status
0	0	-
X	1	Response not received within 64 SDHC_CLK cycles
1	0	Response received

The table below shows that transfer complete has higher priority than the data timeout error. If both bits are set, the data transfer can be considered complete.

Table 17-24. Relation between data timeout error and transfer complete status

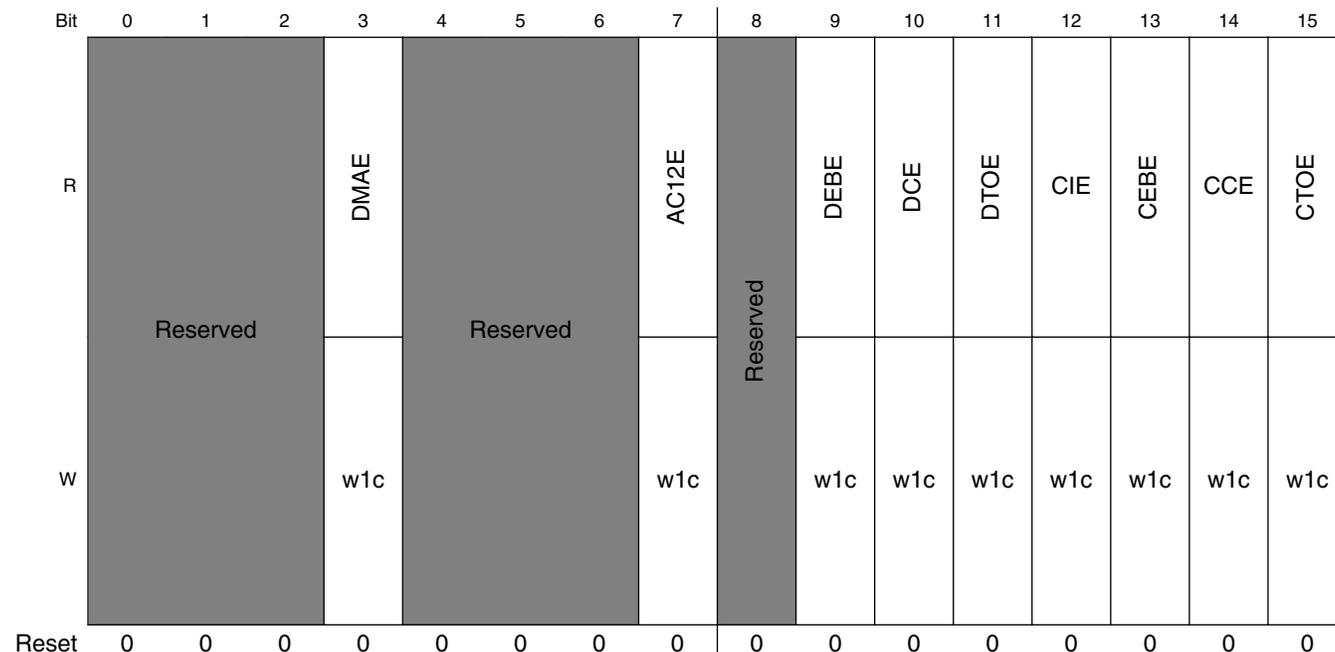
Transfer complete	Data timeout error	Meaning of the status
0	0	-
0	1	Timeout occur during transfer
1	X	Data transfer complete

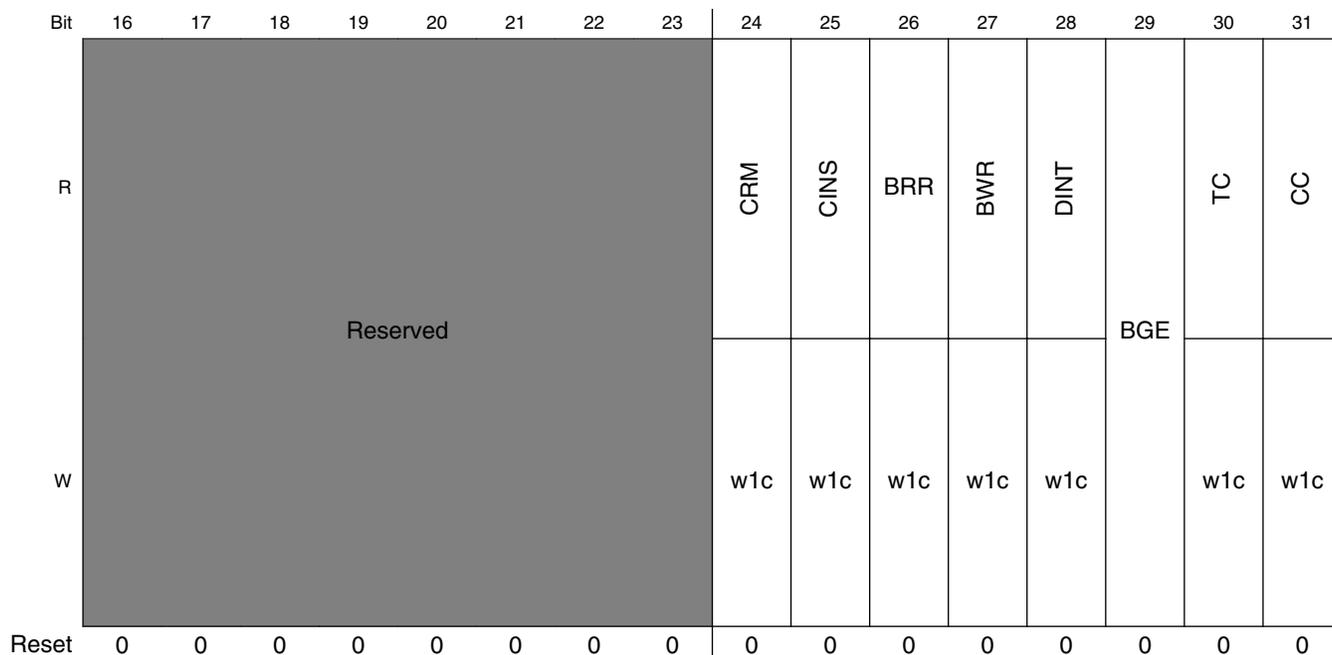
The relation between command CRC error and command timeout error is shown below.

Table 17-25. Relation between command CRC error and command timeout error

Command CRC error	Command timeout error	Meaning of the status
0	0	No error
0	1	Response timeout error
1	0	Response CRC error
1	1	SDHC_CMD line conflict

Address: 2_E000h base + 30h offset = 2_E030h





eSDHC_IRQSTAT field descriptions

Field	Description
0-2 -	This field is reserved.
3 DMAE	DMA error. Occurs when internal DMA transfer failed. This bit is set when some error occurs in the data transfer. The value in the DMA system address register is the next fetch address where the error occurs. Since any error corrupts the entire data block, the host driver should restart the transfer from the corrupted block boundary. The address of the block boundary can be calculated from the current DS_ADDR value or the remaining number of blocks and the block size. 0 No Error 1 Error
4-6 -	This field is reserved.
7 AC12E	Auto CMD12 error. Occurs when one of the bits in AUTOC12ERR is set. This bit is also set when Auto CMD12 is not executed due to a previous command error. 0 No Error 1 Error
8 -	This field is reserved.
9 DEBE	Data end bit error. Occurs when detecting 0 at the end bit position of read data on the SDHC_DAT line or at the end bit position of the CRC. 0 No Error 1 Error
10 DCE	Data CRC error. Occurs when detecting CRC error when transferring read data on the SDHC_DAT line or when detecting the write CRC status having a value other than 0b010.

Table continues on the next page...

eSDHC_IRQSTAT field descriptions (continued)

Field	Description
	0 No Error 1 Error
11 DIOE	Data timeout error. Occurs during one of following timeout conditions: <ul style="list-style-type: none"> • Busy timeout for R1b and R5b types • Busy timeout after write CRC status • Read data timeout 0 No error 1 Timeout
12 CIE	Command index error. Occurs if a command index error occurs in the command response. <p>NOTE: Under rare conditions, CIE may be set for a command without data while a command with data is in progress. On detecting a command CRC error (IRQSTAT[CCE]) or command index error (IRQSTAT[CIE]), perform error recovery and re-issue the command without data. If Auto CMD12 is enabled for data transfer then Auto CMD12 won't be issued by hardware, so software needs to issue it after data transfer completion.</p> 0 No error 1 Error
13 CEBE	Command end bit error. Occurs when the end bit of a command response is 0. 0 No error 1 End bit error generated
14 CCE	Command CRC error. A command CRC error is generated in two cases: <ul style="list-style-type: none"> • If a response is returned and IRQSTAT[CTOE] is cleared (indicating no timeout), this bit is set when detecting a CRC error in the command response. • The eSDHC detects a SDHC_CMD line conflict by monitoring the SDHC_CMD line when a command is issued. If the eSDHC drives the SDHC_CMD line to 1, but detects 0 on the SDHC_CMD line at the next SDHC_CLK edge, then the eSDHC aborts the command (stop driving SDHC_CMD line) and sets this bit. The CTOE bit is also set to distinguish the SDHC_CMD line conflict. <p>NOTE: Under rare conditions, CCE may be set for a command without data while a command with data is in progress. On detecting a command CRC error (IRQSTAT[CCE]) or command index error (IRQSTAT[CIE]), perform error recovery and re-issue the command without data. If Auto CMD12 is enabled for data transfer then Auto CMD12 won't be issued by hardware, so software needs to issue it after data transfer completion.</p> 0 No error 1 CRC error generated
15 CTOE	Command timeout error. Occurs if no response is returned within 64 SDHC_CLK cycles from the end bit of the command. Also, if eSDHC detects a SDHC_CMD line conflict, this bit is set along with IRQSTAT[CCE] as shown in Table 17-35 . 0 No error 1 Time out
16–23 -	This field is reserved.
24 CRM	Card removal. This bit is set if PRSSTAT[CINS] changes from 1 to 0. When the host driver writes 1 to this bit to clear it, the status of PRSSTAT[CINS] should be confirmed. Because the card-detect state may be changed when the host driver clears this bit, an interrupt event may not be generated.

Table continues on the next page...

eSDHC_IRQSTAT field descriptions (continued)

Field	Description
	<p>When this bit is cleared, it is set again if no card is inserted. To leave it cleared, clear IRQSTATEN[CRMSEN].</p> <p>0 Card state unstable or inserted 1 Card removed</p>
25 CINS	<p>Card insertion. This bit is set if PRSSTAT[CINS] changes from 0 to 1. When the host driver writes 1 to this bit to clear it, the status of PRSSTAT[CINS] should be confirmed. Because the card-detect state may be changed when the host driver clears this bit, an interrupt event may not be generated.</p> <p>When this bit is cleared, it is set again if a card has been inserted. To leave it cleared, clear IRQSTATEN[CINSEN].</p> <p>0 Card state unstable or removed 1 Card inserted</p>
26 BRR	<p>Buffer read ready. This bit is set if PRSSTAT[BREN] changes from 0 to 1.</p> <p>0 Not ready to read buffer 1 Ready to read buffer</p>
27 BWR	<p>Buffer write ready. This bit is set if PRSSTAT[BWEN] changes from 0 to 1.</p> <p>0 Not ready to write buffer 1 Ready to write buffer</p>
28 DINT	<p>DMA interrupt. Occurs when the internal DMA finishes the data transfer successfully. If errors occur during data transfer, this bit is not set. Instead, the DMAE bit is set.</p> <p>0 No DMA interrupt 1 DMA interrupt is generated</p>
29 BGE	<p>Block gap event. If PROCTL[SABGREQ] is set, this bit is set when a read or write transaction is stopped at a block gap. If PROCTL[SABGREQ] is cleared, this bit is not set.</p> <p>During a read transaction, this bit is set at the falling edge of the SDHC_DAT line active status (when the transaction is stopped at SD bus timing). Read wait must be supported to use this function.</p> <p>During a write transaction, this bit is set at the falling edge of PRSSTAT[WTA] (after reading the CRC status at SD bus timing).</p> <p>0 No block gap event 1 Transaction stopped at block gap</p>
30 TC	<p>Transfer complete. This bit is set when a read or write transfer is completed.</p> <p>For a read transaction, this bit is set at the falling edge of PRSSTAT[RTA]. There are two cases in which this interrupt is generated:</p> <ul style="list-style-type: none"> • When a data transfer completes, as specified by the data length (after the last data has been read to the host system). • When data has stopped at the block gap and completed the data transfer by setting PROCTL[SABGREQ] (after valid data has been read to the host system). <p>For a write transaction, this bit is set at the falling edge of PRSSTAT[DLA]. There are two cases in which this interrupt is generated:</p> <ul style="list-style-type: none"> • When the last data is written to the SD card, as specified by data length and the busy signal is released. • When data transfers are stopped at the block gap by setting PROCTL[SABGREQ] and data transfers have completed (after valid data is written to the SD card and the busy signal is released).

Table continues on the next page...

eSDHC_IRQSTAT field descriptions (continued)

Field	Description
31 CC	Command complete. This bit is set when the end bit of the command response is received (except Auto CMD12). See PRSSTAT[CIHB]. 0 No command complete 1 Command complete

17.4.11 Interrupt status enable (eSDHC_IRQSTATEN)

Setting the bits of IRQSTATEN enables the corresponding interrupt status bit to be set by the specified event. If any bit is cleared, the corresponding IRQSTAT bit is also cleared and is never set.

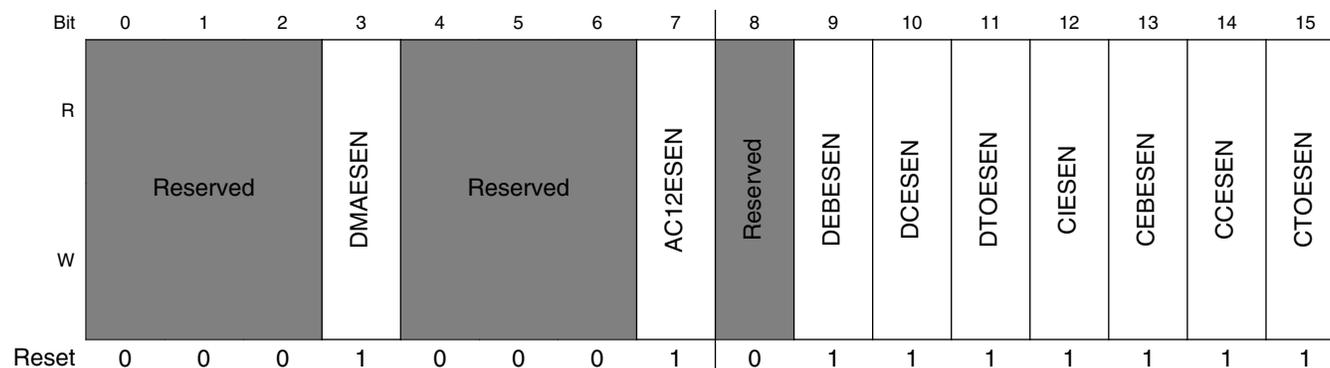
NOTE

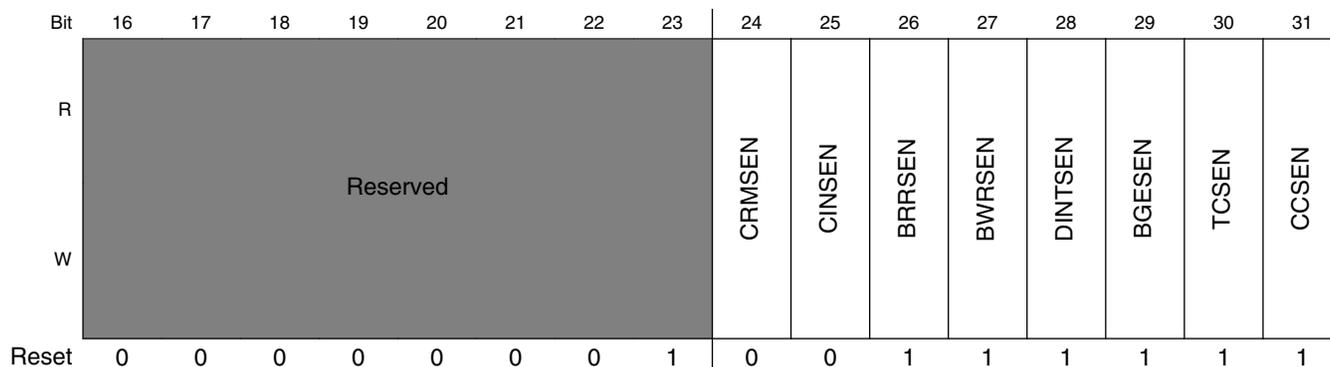
The eSDHC may sample the card interrupt signal during the interrupt period and hold its value in the flip-flop. As a result of synchronization, there is a delay in the card interrupt (which is asserted from the card) to the time the host system is informed.

NOTE

To detect a SDHC_CMD line conflict, the host driver must set both CTOESEN and CCESEN bits.

Address: 2_E000h base + 34h offset = 2_E034h





eSDHC_IRQSTATEN field descriptions

Field	Description
0-2 -	This field is reserved.
3 DMAESEN	DMA error status enable 0 Masked 1 Enabled
4-6 -	This field is reserved.
7 AC12ESEN	Auto CMD12 error status enable 0 Masked 1 Enabled
8 -	This field is reserved.
9 DEBESEN	Data end bit error status enable 0 Masked 1 Enabled
10 DCESEN	Data CRC error status enable 0 Masked 1 Enabled
11 DTESEN	Data timeout error status enable 0 Masked 1 Enabled
12 CIESEN	Command index error status enable 0 Masked 1 Enabled
13 CEBESEN	Command end bit error status enable 0 Masked 1 Enabled
14 CCESEN	Command CRC error status enable

Table continues on the next page...

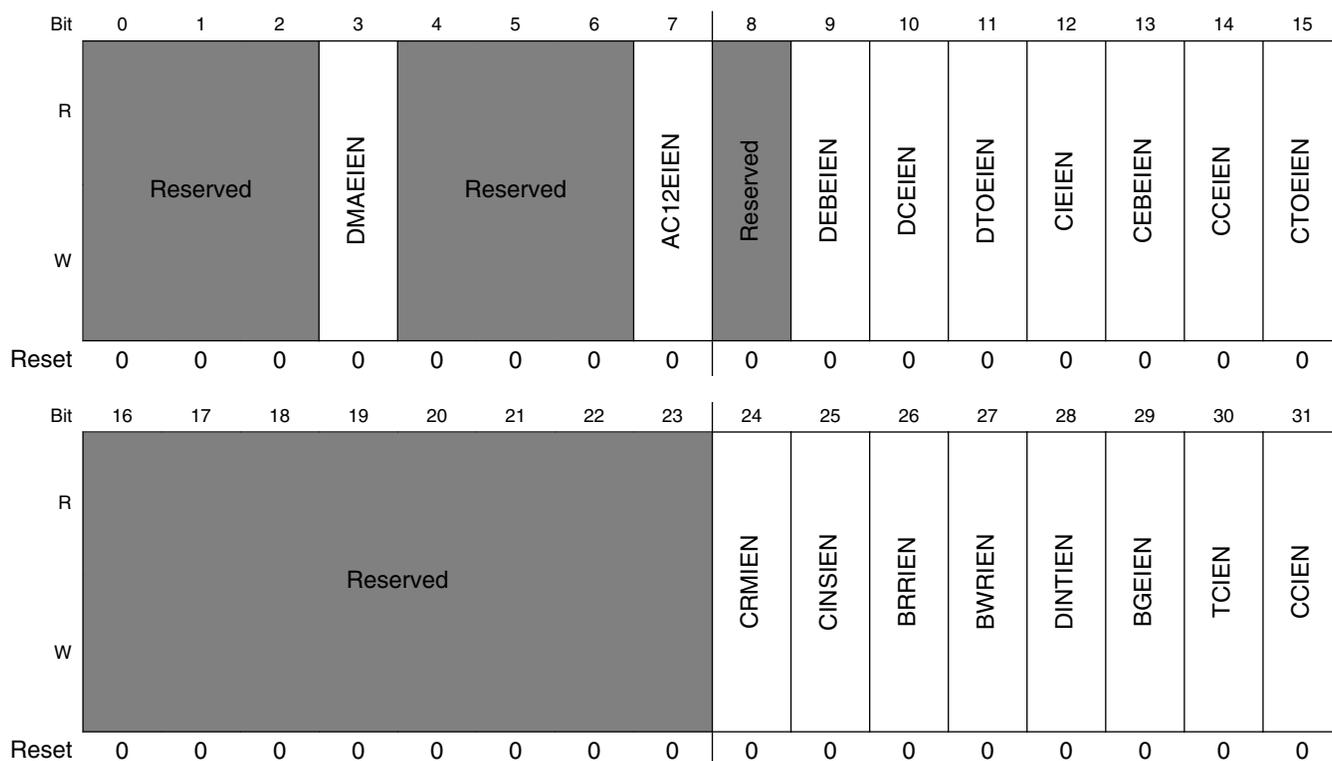
eSDHC_IRQSTATEN field descriptions (continued)

Field	Description
	0 Masked 1 Enabled
15 CTOESEN	Command timeout error status enable 0 Masked 1 Enabled
16–23 -	This field is reserved.
24 CRMSEN	Card removal status enable 0 Masked 1 Enabled
25 CINSEN	Card insertion status enable 0 Masked 1 Enabled
26 BRRSEN	Buffer read ready status enable 0 Masked 1 Enabled
27 BWRSEN	Buffer write ready status enable 0 Masked 1 Enabled
28 DINTSEN	DMA interrupt status enable 0 Masked 1 Enabled
29 BGESEN	Block gap event status enable 0 Masked 1 Enabled
30 TCSEN	Transfer complete status enable 0 Masked 1 Enabled
31 CCSEN	Command complete status enable 0 Masked 1 Enabled

17.4.12 Interrupt signal enable (eSDHC_IRQSIGEN)

IRQSIGEN selects which interrupt status is indicated to the host system as the interrupt. These status bits all share the same interrupt line. Setting any of these bits enables an interrupt generation. The corresponding status register bit generates an interrupt when the corresponding interrupt signal enable bit is set.

Address: 2_E000h base + 38h offset = 2_E038h



eSDHC_IRQSIGEN field descriptions

Field	Description
0–2 -	This field is reserved.
3 DMAEIEN	DMA error interrupt enable 0 Masked 1 Enabled
4–6 -	This field is reserved.
7 AC12EIEN	Auto CMD12 error interrupt enable

Table continues on the next page...

eSDHC_IRQSIGEN field descriptions (continued)

Field	Description
	0 Masked 1 Enabled
8 -	This field is reserved.
9 DEBEIEN	Data end bit error interrupt enable 0 Masked 1 Enabled
10 DCEIEN	Data CRC error interrupt enable 0 Masked 1 Enabled
11 DTOEIEN	Data timeout error interrupt enable 0 Masked 1 Enabled
12 CIEIEN	Command index error interrupt enable 0 Masked 1 Enabled
13 CEBEIEN	Command end bit error interrupt enable 0 Masked 1 Enabled
14 CCEIEN	Command CRC error interrupt enable 0 Masked 1 Enabled
15 CTOEIEN	Command timeout error interrupt enable 0 Masked 1 Enabled
16–23 -	This field is reserved.
24 CRMIEN	Card removal interrupt enable 0 Masked 1 Enabled
25 CINSIEN	Card insertion interrupt enable 0 Masked 1 Enabled
26 BRRIEN	Buffer read ready interrupt enable 0 Masked 1 Enabled
27 BWRIEN	Buffer write ready interrupt enable

Table continues on the next page...

eSDHC_IRQSIGEN field descriptions (continued)

Field	Description
	0 Masked 1 Enabled
28 DINTIEN	DMA interrupt enable 0 Masked 1 Enabled
29 BGEIEN	Block gap event interrupt enable 0 Masked 1 Enabled
30 TCIEN	Transfer complete interrupt enable 0 Masked 1 Enabled
31 CCIEN	Command complete interrupt enable 0 Masked 1 Enabled

17.4.13 Auto CMD12 status (eSDHC_AUTOC12ERR)

When IRQSTAT[AC12E] is set, the host driver checks this register to identify what kind of error Auto CMD12 indicated. This register is valid only when IRQSTAT[AC12E] is set.

Table 17-29. Relationship between command CRC error and command timeout error for Auto CMD12

Auto CMD12 CRC error	Auto CMD12 timeout error	Types of error
0	0	No error
0	1	Response timeout error
1	0	Response CRC error
1	1	SDHC_CMD line conflict

There are three scenarios when AUTOC12ERR can be changed:

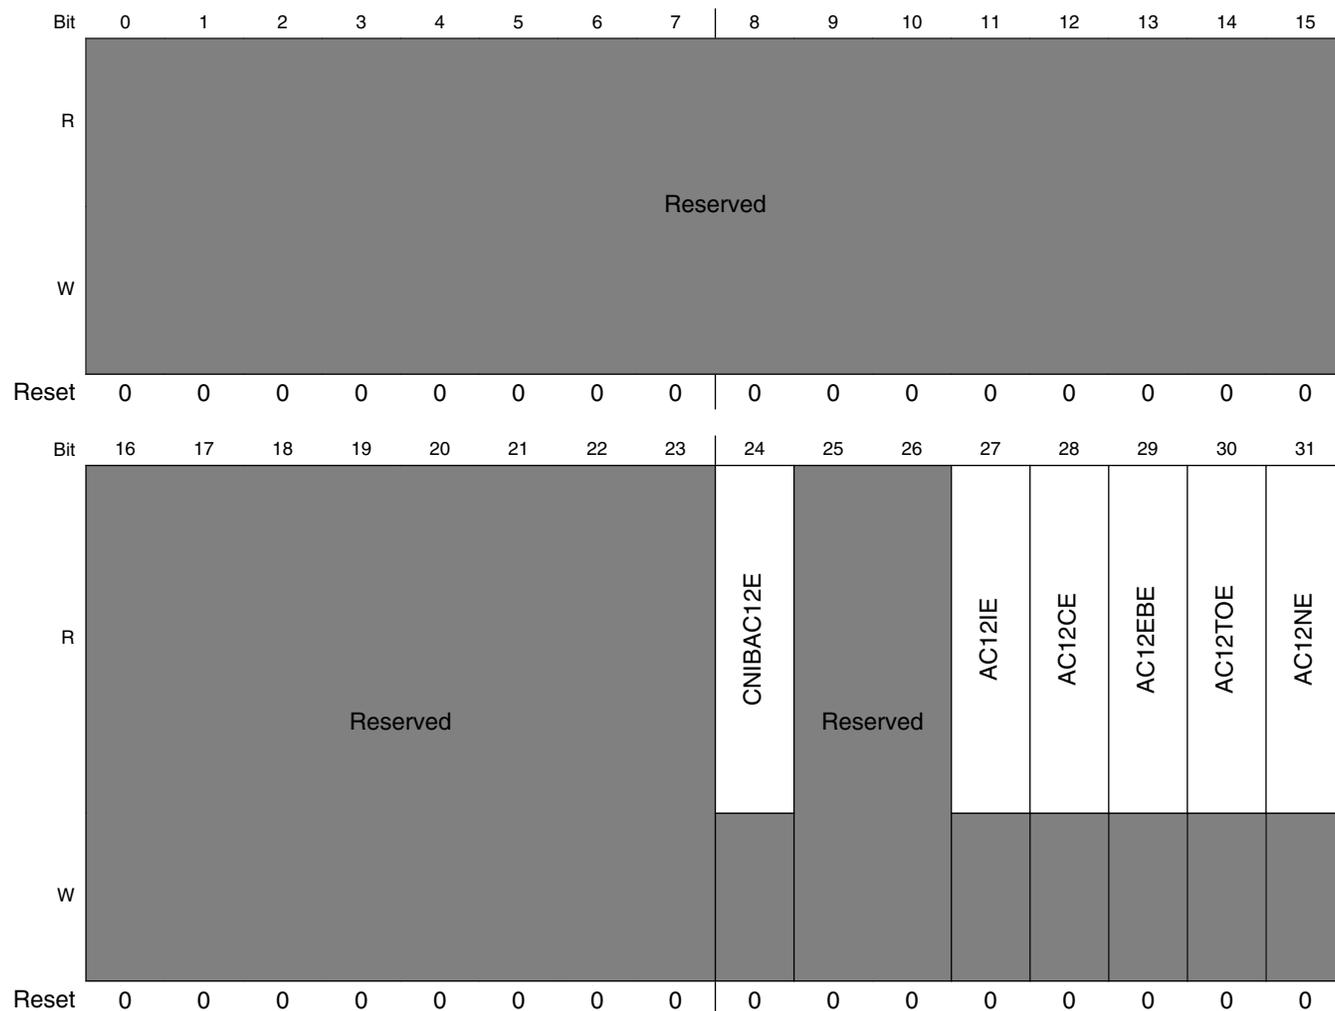
1. When eSDHC is going to issue Auto CMD12
 - Set AC12NE if Auto CMD12 cannot be issued due to an error in the previous command.
 - Clear AC12NE if Auto CMD12 is issued.
2. At the end bit of an Auto CMD12 response
 - Check received responses by checking the error bits 1-4.

Enhanced Secure Digital Host Controller (eSDHC) Memory Map

- Set if error is detected.
 - Clear if error is not detected.
3. Before reading AUTO12ERR[CNIBAC12E]
- Set CNIBAC12E if there is a command that cannot be issued
 - Clear CNIBAC12E if there is no command to issue

The timing of generating the Auto CMD12 error and writing to the command register is asynchronous. The command may be blocked by any Auto CMD12 error causing CNIBAC12E to be set. Therefore, it is suggested to read this register only when IRQSTAT[AC12E] is set. An Auto CMD12 error interrupt is generated when one of the error bits 0-4 is set. The CNIBAC12E error bit does not generate an interrupt.

Address: 2_E000h base + 3Ch offset = 2_E03Ch



eSDHC_AUTO12ERR field descriptions

Field	Description
0–23 -	This field is reserved.

Table continues on the next page...

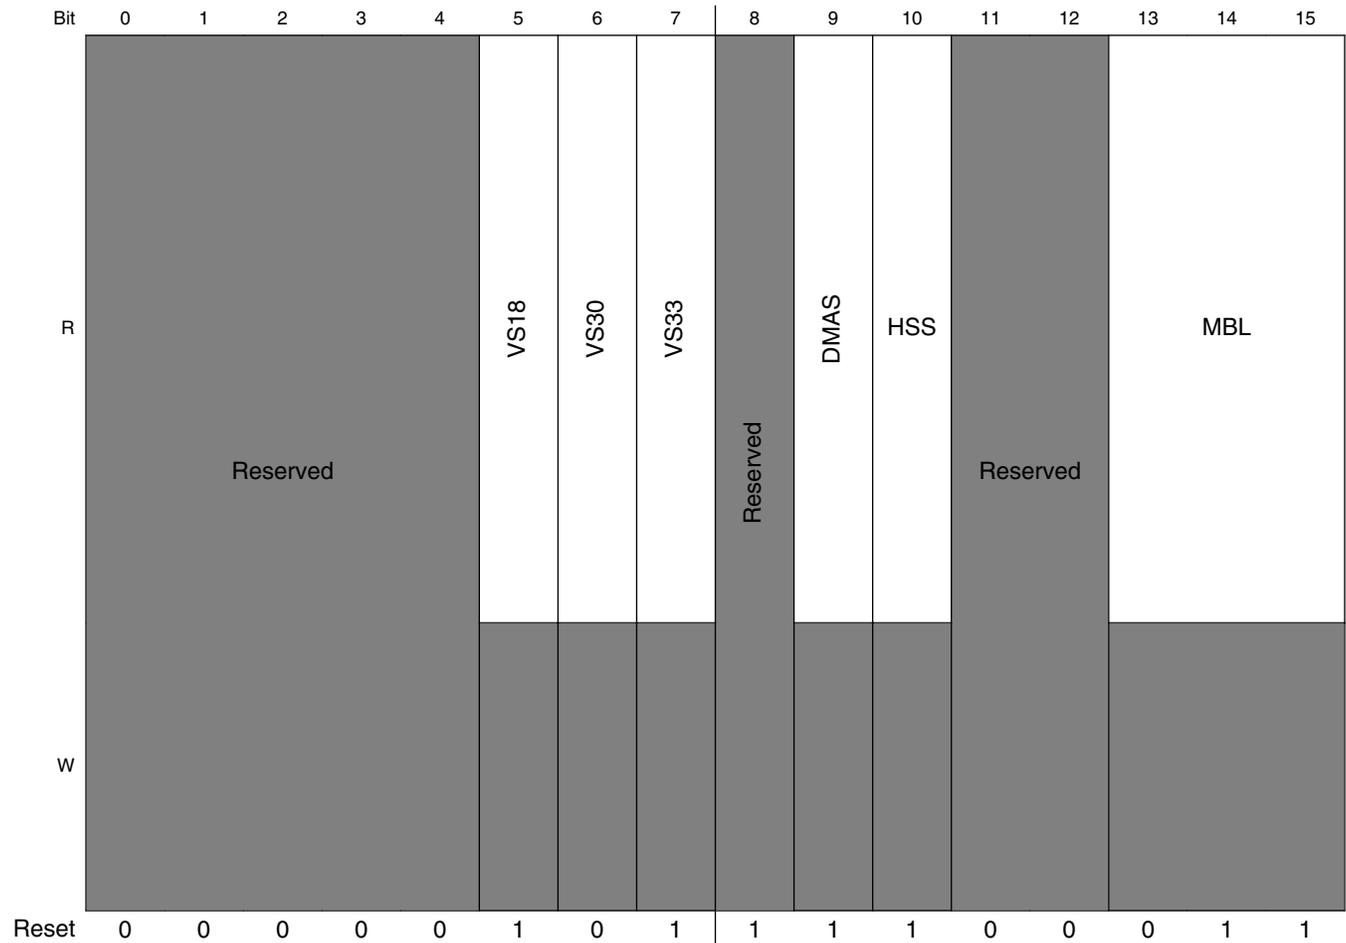
eSDHC_AUTOC12ERR field descriptions (continued)

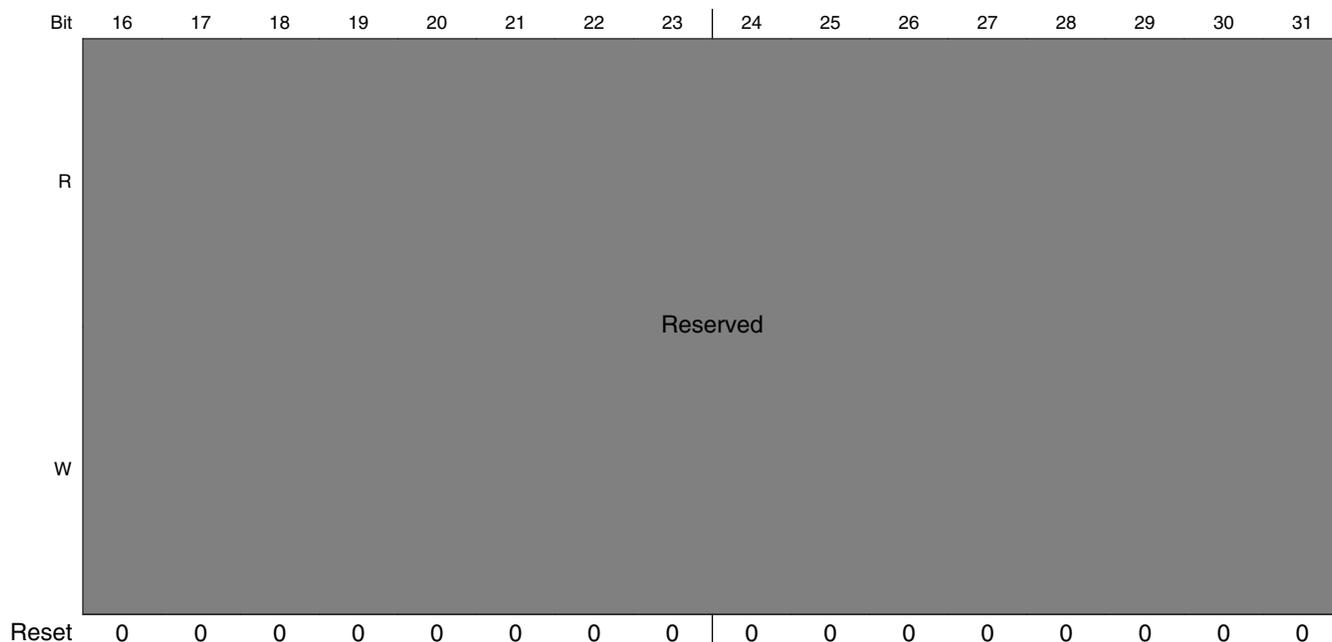
Field	Description
24 CNIBAC12E	Command not issued by Auto CMD12 error. This bit is set when CMD_wo_DAT is not executed due to an Auto CMD12 error (D04-D01). 0 No error 1 Not Issued
25–26 -	This field is reserved.
27 AC12IE	Auto CMD12 index error. Occurs if the command index error occurs in response to a command. 0 No error 1 Error, the CMD index in response is not CMD12
28 AC12CE	Auto CMD12 CRC error. Occurs when detecting a CRC error in the command response. 0 No CRC error 1 CRC error met in Auto CMD12 response
29 AC12EBE	Auto CMD12 end bit error. Occurs when detecting that the end bit of command response is 0 when it should be 1. 0 No error 1 End bit error generated
30 AC12TOE	Auto CMD12 timeout error. Occurs if no response is returned within 64 SDHC_CLK cycles from the end bit of the command. If this bit is set, the other error status bits (2-4) are meaningless. 0 No error 1 Time out
31 AC12NE	Auto CMD12 not executed. If a memory multiple block data transfer is not started due to command error, this bit is not set because it is not necessary to issue Auto CMD12. Setting this bit means eSDHC cannot issue Auto CMD12 to stop the memory multiple block data transfer due to some error. If this bit is set, the other error status bits (1-4) are meaningless. 0 Executed 1 Not executed

17.4.14 Host controller capabilities (eSDHC_HOSTCAPBLT)

The host controller capabilities provides the host driver with information specific to the eSDHC implementation. The value in this register does not change in a software reset, and any write to this register is ignored.

Address: 2_E000h base + 40h offset = 2_E040h





eSDHC_HOSTCAPBLT field descriptions

Field	Description
0–4 -	This field is reserved.
5 VS18	Voltage support 1.8 V. This bit depends on the host system ability. 0 1.8 V not supported 1 1.8 V supported
6 VS30	Voltage support 3.0 V. This bit depends on the host system ability. 0 3.0 V not supported 1 3.0 V supported
7 VS33	Voltage support 3.3 V. This bit depends on the host system ability. 0 3.3 V not supported 1 3.3 V supported
8 -	This field is reserved.
9 DMAS	DMA support. Indicates if eSDHC is capable of using internal DMA to transfer data between system memory and the data buffer directly. 0 DMA not supported 1 DMA supported
10 HSS	High speed support. Indicates if the eSDHC supports high speed mode and the host system can supply the SD clock frequency from 25 to 50 MHz. 0 High speed supported 1 High speed supported
11–12 -	This field is reserved.

Table continues on the next page...

eSDHC_HOSTCAPBLT field descriptions (continued)

Field	Description
13–15 MBL	Max block length. Indicates the maximum block size that the host driver can read and write to the buffer in the eSDHC. The buffer should transfer block size without wait cycles. 000 512 bytes 001 1024 bytes 010 2048 bytes 011 4096 bytes
16–31 -	This field is reserved.

17.4.15 Watermark level (eSDHC_WML)

Both write and read watermark levels are configurable. The value can be any number from 1-128 words.

Address: 2_E000h base + 44h offset = 2_E044h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0

eSDHC_WML field descriptions

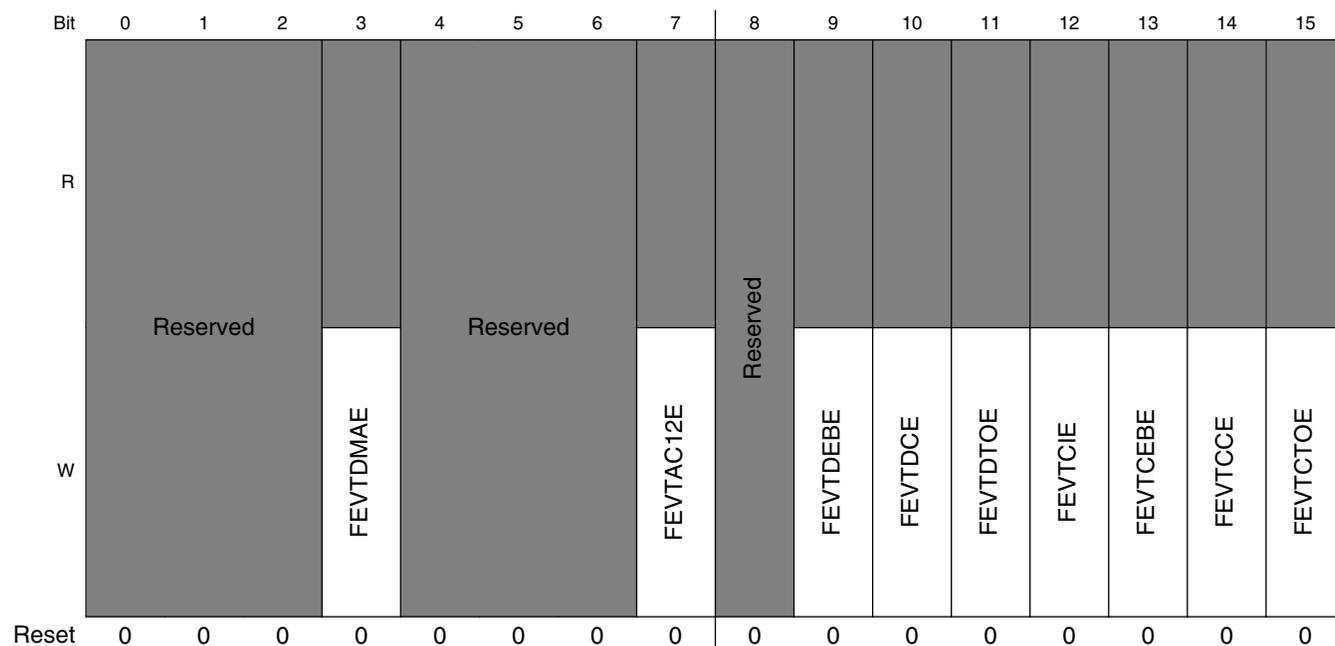
Field	Description
0–7 -	This field is reserved.
8–15 WR_WML	Write watermark level. Number of 32-bit words of watermark level in DMA write operation. Also, the number of words of write burst length. NOTE: The minimum value is 0x02, which represents 2 words (8 bytes).
16–23 -	This field is reserved.
24–31 RD_WML	Read watermark level. Number of 32-bit words of watermark level in DMA read operation. Also, the number of words of read burst length. NOTE: The minimum value for RD_WML is 0x02, which means 2 words (8 bytes), and the maximum value for RD_WML is 0x10, which means 16 words (64 bytes). Setting RD_WML to values outside this range results in non-predicted behavior.

17.4.16 Force event (eSDHC_FEVT)

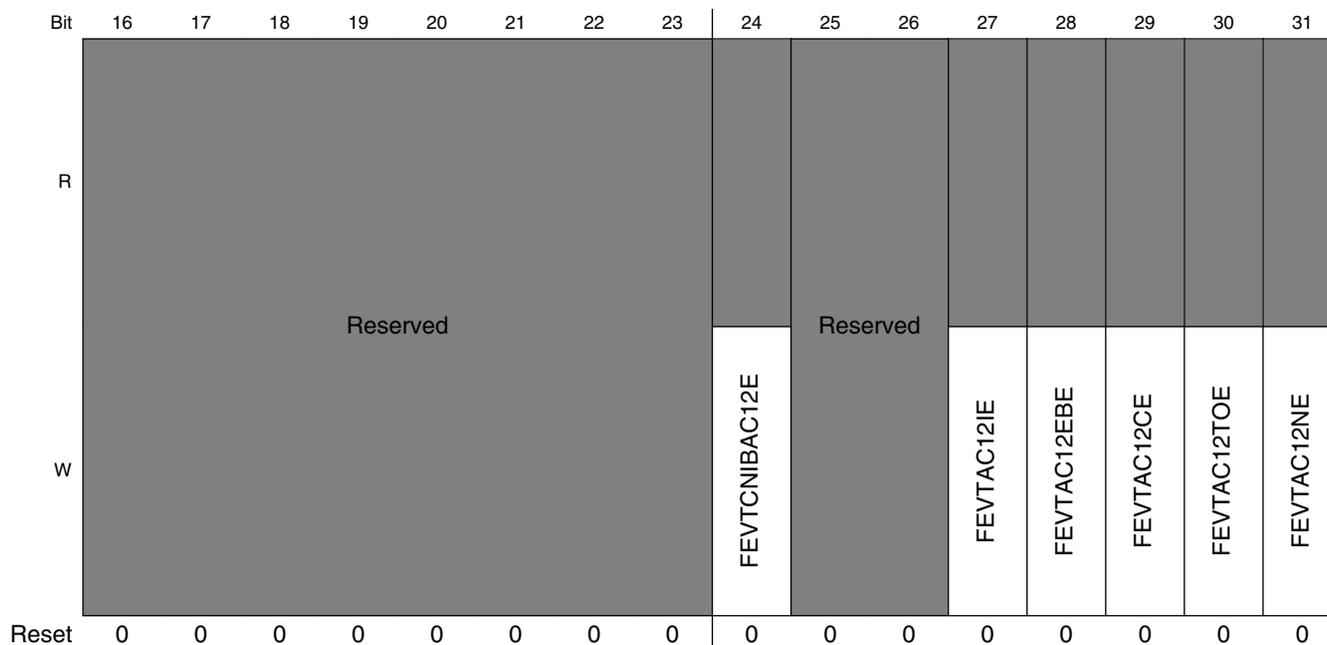
The force event register is not a physically implemented register. Rather, it is an address to which the IRQSTAT register can be written if the corresponding bit of IRQSTATEN is set. Therefore, this register is a write-only register and writing zero has no effect. Writing 1 to this register sets the corresponding bit of IRQSTAT. Reading from this register always returns zeroes.

Forcing a card interrupt generates a short pulse on the SDHC_DAT[1] line, and the driver may treat this interrupt as normal. The interrupt service routine may skip polling the card-interrupt source as the interrupt is self-cleared.

Address: 2_E000h base + 50h offset = 2_E050h



Enhanced Secure Digital Host Controller (eSDHC) Memory Map



eSDHC_FEVT field descriptions

Field	Description
0–2 -	This field is reserved.
3 FEVTDMAE	Force event DMA error. Forces IRQSTAT[DMAE] to set.
4–6 -	This field is reserved.
7 FEVTAC12E	Force event Auto CMD12 error. Forces IRQSTAT[AC12E] to set.
8 -	This field is reserved.
9 FEVTDEBE	Force event data end bit error. Forces IRQSTAT[DEBE] to set.
10 FEVTDCE	Force event data CRC error. Forces IRQSTAT[DCE] to set.
11 FEVTDTOE	Force event data time out error. Forces IRQSTAT[DTOE] to set.
12 FEVTCIE	Force event command index error. Forces IRQSTAT[CCE] to set.
13 FEVTCBE	Force event command end bit error. Forces IRQSTAT[CEBE] to set.
14 FEVTCCE	Force event command CRC error. Forces IRQSTAT[CCE] to set.
15 FEVTCIOE	Force event command time out error. Forces IRQSTAT[CTOE] to set.
16–23 -	This field is reserved.

Table continues on the next page...

eSDHC_FEVT field descriptions (continued)

Field	Description
24 FEVTCNIBAC12E	Force event command not executed by Auto CMD12 error. Forces AUTOC12ERR[CNIBAC12E] to set.
25–26 -	This field is reserved.
27 FEVTAC12IE	Force event Auto CMD12 index error. Forces AUTOC12ERR[AC12IE] to set.
28 FEVTAC12EBE	Force event Auto CMD12 end bit error. Forces AUTOC12ERR[AC12EBE] to set.
29 FEVTAC12CE	Force event Auto CMD12 CRC error. Forces AUTOC12ERR[AC12CE] to set.
30 FEVTAC12TOE	Force event Auto CMD12 time out error. Forces AUTOC12ERR[AC12TOE] to set.
31 FEVTAC12NE	Force event Auto CMD12 not executed. Forces AUTOC12ERR[AC12NE] to set.

17.4.17 Host controller version (eSDHC_HOSTVER)

The host controller version register contains the version for the vendor and the host controller. All the bits are read-only.

Address: 2_E000h base + FCh offset = 2_E0FCh

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															VVN				SVN												
W	Reserved															Reserved				Reserved												
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1

eSDHC_HOSTVER field descriptions

Field	Description
0–15 -	This field is reserved.
16–23 VVN	Vendor version number. The host driver should not use this status. The upper and the lower 4-bits indicate the version. Settings not shown are reserved. 0x00 Freescale eSDHC version 1.0 0x01 Freescale eSDHC version 1.1(used by P2020) 0x10 Freescale eSDHC version 2.0 0x11 Freescale eSDHC version 2.1 0x12 Freescale eSDHC version 2.2

Table continues on the next page...

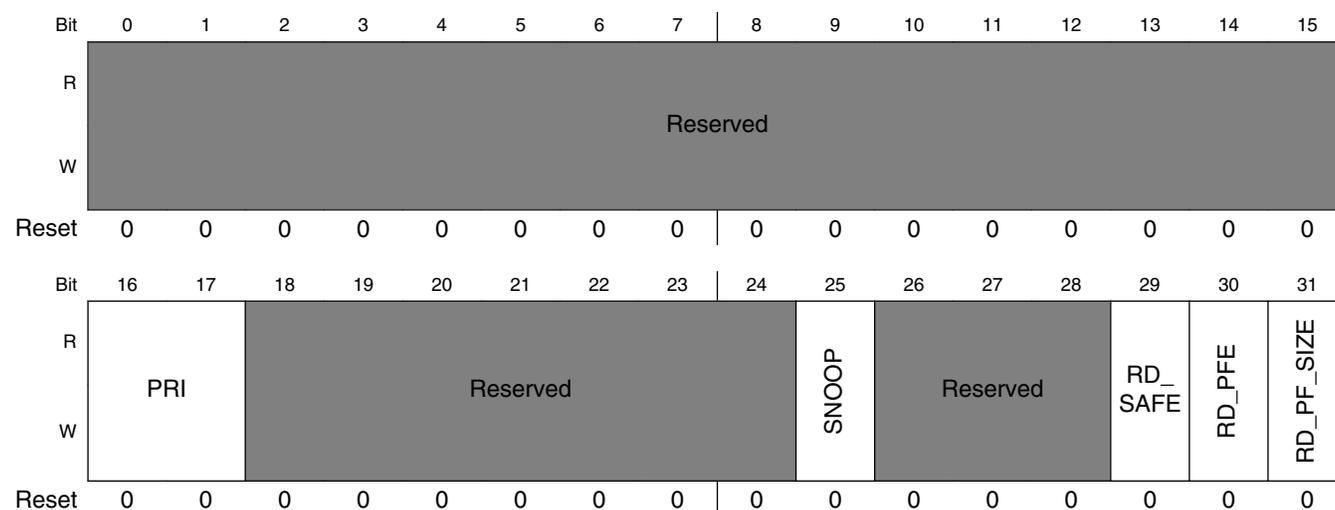
eSDHC_HOSTVER field descriptions (continued)

Field	Description
24–31 SVN	<p>Specification version number. Indicates for the host controller specification version. The upper and the lower 4-bits indicate the version.</p> <p>Settings not shown are reserved.</p> <p>0x00 SD Host Specification Version 1.0 0x01 SD Host Specification Version 2.0, supports the test event register .</p>

17.4.18 DMA control register (eSDHC_DCR)

The DMA control register controls various settings that affect the system response to DMA operations.

Address: 2_E000h base + 40Ch offset = 2_E40Ch



eSDHC_DCR field descriptions

Field	Description
0–15 -	This field is reserved.
16–17 PRI	<p>Priority. This field is used to present priority level for AHB arbitration for eSDHC DMA requests.</p> <p>00 Level 0 (lowest priority) 01 Level 1 10 Level 2 11 Level 3 (highest priority)</p>
18–24 -	This field is reserved.

Table continues on the next page...

eSDHC_DCR field descriptions (continued)

Field	Description
25 SNOOP	Snoop attribute. 0 DMA transactions are not snooped by the CPU data cache 1 DMA transactions are snooped by the CPU data cache
26–28 -	This field is reserved.
29 RD_SAFE	Read safe. This bit should be set only if the target of a read-DMA operation is a well-behaved memory that is not affected by the read operation and returns the same data if read again from the same location. This means that unaligned reading operation can be rounded up to enable more efficient read operations. 0 It is not safe to read more bytes that were intended 1 It is safe to read more bytes that were intended
30 RD_PFE	Read prefetch enable. This bit should be set if the target of read-DMA operation is a well-behaved memory that is not affected by the read operation and returns the same data if read again from the same location. This means that prefetch of data can be done by the internal bus units and it results in faster read completion. 0 It is not allowed to prefetch data on DMA read operation 1 It is allowed to prefetch data on DMA read operation
31 RD_PF_SIZE	Read prefetch size. Determines the prefetch byte count to be used if RD_PFE is set. 0 64 bytes prefetch 1 32 bytes prefetch

17.5 eSDHC functional description

The following sections provide a brief functional description of the major system blocks, including the data buffer, DMA CCB interface, register bank, register bus interface, dual-port memory wrapper, data/command controller, clock and reset manager, and clock generator.

17.5.1 Data buffer

The eSDHC uses one configurable data buffer so that data can be transferred between the internal system bus (register bus or CCB bus) and the SD card in an optimized manner to maximize throughput between the two clock domains (the IP peripheral clock and the master clock).

See the figure below for an illustration of the buffer scheme.

The buffer is used as temporary storage for data being transferred between the host system and the card. The burst lengths for read and write are both configurable and can be any value between 1 and 128 words.

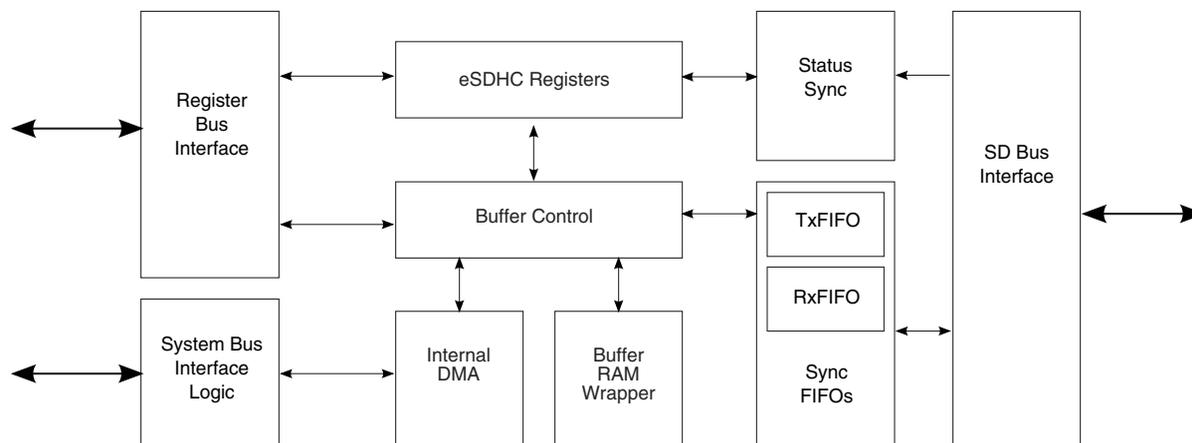


Figure 17-25. eSDHC buffer scheme

For a host read operation, when the amount of data exceeds the RD_WML value, the eSDHC sets PRSSTAT[BREN] and either:

- Issues a DMA request to inform the system to read the data
- Issues a DMA interrupt to inform the system to read the data
- When granted CCB access permission, the internal DMA burst-reads RD_WML number of words

Conversely, for a host write operation, when the amount of buffer spaces exceeds the WR_WML value, the eSDHC sets PRSSTAT[BWEN] and either:

- Issues a DMA request to inform the system to write data to the buffer
- Issues a DMA interrupt to inform the system to write data to the buffer
- When granted CCB crossbar access permission, the internal DMA burst-writes WR_WML number of words into the buffer

17.5.1.1 Write operation sequence

There are two ways to write data into the buffer when the user transfers data to the card:

- Processor core polling IRQSTAT[BWR] (interrupt or polling)
- Internal DMA

When the internal DMA is not used (XFERTYP[DMAEN] is not set when the command is sent), the eSDHC asserts an external DMA request when more than WML[WR_WML] number of empty buffer word slots are available and ready for receiving new data. At the same time, the eSDHC sets IRQSTAT[BWR]. The buffer write ready interrupt is generated if it is enabled by software.

When the internal DMA is used, the eSDHC does not inform the system before all the required number of bytes are transferred and no error is encountered. When an error occurs during the data transfer, the eSDHC aborts the data transfer and abandons the current block. The host driver should read the content of the DMA system address register to obtain the start address of abandoned data block. If the current data transfer is in multi-block mode, the eSDHC does not automatically send CMD12 even though XFERTYP[AC12EN] is set. Therefore, in this scenario, the host driver should send CMD12 and restart the write operation from that address. It is recommended that a software reset for data is applied before the transfer is restarted after error recovery.

The eSDHC does not start data transmission until the WML[WR_WML] number of words of data can be held in the buffer. If the buffer is empty and the host system does not write data in time, the eSDHC stops the SDHC_CLK to avoid a data buffer underrun situation.

17.5.1.2 Read operation sequence

There are two ways to read data from the buffer when transferring data to the card:

- Processor core polling IRQSTAT[BRR] (interrupt or polling)
- Internal DMA

When the internal DMA is not used (that is, XFERTYP[DMAEN] is not set when the command is sent), the eSDHC asserts a DMA request when more than WML[RD_WML] number of words are available and ready for the system to fetch the data. At the same time, the eSDHC sets the IRQSTAT[BRR] bit. The buffer read ready interrupt is generated if it is enabled by software.

When the internal DMA is used, the eSDHC does not inform the system before all the required number of bytes are transferred and no error is encountered. When an error occurs during the data transfer, the eSDHC aborts the data transfer and abandons the current block. The host driver should read the content of the DMA system address register to obtain the start address of the abandoned data block. If the current data transfer is in multiblock mode, the eSDHC does not automatically send CMD12, even though XFERTYP[AC12EN] is set. Therefore, in this scenario, the host driver should send CMD12 and restart the read operation from that address. It is recommended that a software reset for data is applied before the transfer is restarted after error recovery.

The eSDHC does not start data transmission until the WML[RD_WML] number of words of data are in the buffer. If the buffer is full and the host system does not read the data in time, the eSDHC stops the SDHC_CLK to avoid a data buffer overrun situation.

17.5.1.3 Data buffer size

To optimize use of the buffer, the buffer size must be known.

In the eSDHC, the data buffer can hold up to 128 32-bit words, and the read and write watermark levels are each configurable from one to 128 words. The host driver may configure the values according to the system situation and requirements.

During multiblock data transfer, the maximum block length is 4096 bytes, which can satisfy all the requirements from MMC and SD cards. Any block length less than this value is also allowed. The only restriction is from the external card since it may not support such a large block or a partial block access that is not an integer multiple of 512 bytes.

17.5.2 DMA CCB interface

The internal DMA implements a DMA engine and CCB master.

When the internal DMA is enabled (that is, XFERTYP[DMAEN] is set), the buffer interrupt status bits remain set if they are enabled. To avoid setting them, clear IRQSTATEN[BWRSEN, BRRSEN]. See [Figure 17-26](#) for an illustration of the DMA CCB interface block. Do not use the internal DMA to read (or write) data when the CPU or an external DMA writes (or reads) the data through the DATPORT register.

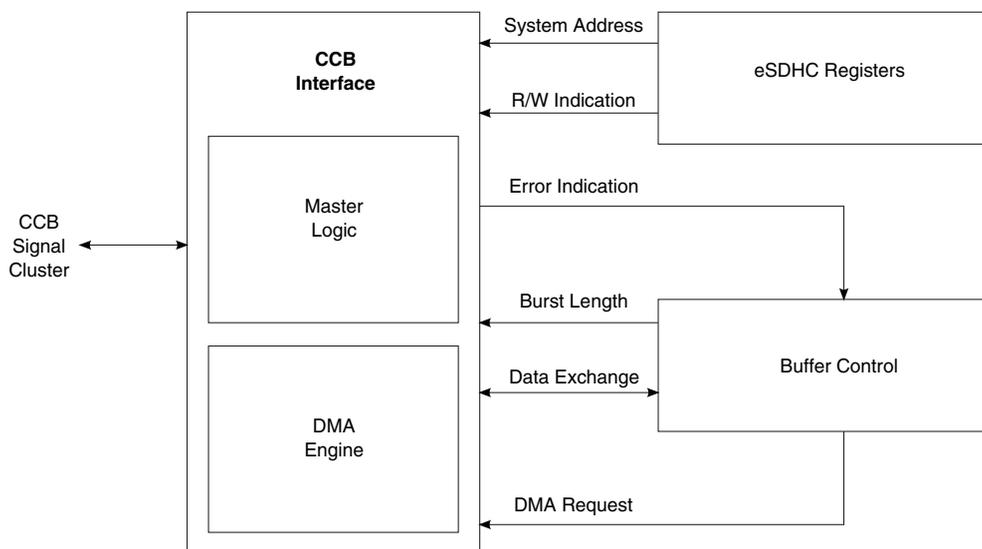


Figure 17-26. DMA system bus interface block

17.5.2.1 Internal DMA request

When the watermark level is met in the data transfer and the internal DMA is enabled, the data buffer block sends a DMA request to this block.

Meanwhile, the external DMA request is disabled. The delay of response from the internal DMA engine depends on the system bus loading and the priority assigned to the eSDHC. The DMA engine does not respond to the request during its burst transfer, and it is available as soon as the burst is over. When an access on the buffer is made, the data buffer deasserts the request. When the internal DMA accesses the buffer, the data buffer updates its internal buffer pointer. When the watermark level is satisfied, another DMA request is sent.

The data transfer is in the block unit and the last watermark level is always set to the remaining number of words. For example, for a multi-block data read with each block the size of 31 bytes, the burst length is set at six words. After the first burst transfer, if there are more than seven bytes in the buffer, which might include some data from the next block, another DMA read request is sent. This is because the remaining number of words to send for the current block is $(31 - 6 \times 4) \div 4 = 2$, and the eSDHC reads two words out of the buffer, with seven valid bytes and one stuff byte automatically added by the eSDHC.

17.5.2.2 DMA burst length

Maximum size is the only restriction on DMA burst length for the internal DMA engine.

The burst length for read or write can be one to 128 words. The actual DMA burst length depends on which is smaller: the configured watermark level or the remaining words of the current block.

Consider the example in [Internal DMA request](#). After six words are read, the burst length is two words to complete the 31-byte block. Then the burst length changes back to six words to prepare for the next 31-byte block. The host driver writer may take this variable burst length into account. For a constant burst length, the user may choose to configure the burst length as the divisor of block size.

17.5.2.3 CCB Master interface

It is possible that the internal DMA engine fails during data transfer. When an error occurs, the DMA engine stops the transfer and goes into idle state; at the same time, the internal data buffer stops working. `IRQSTAT[DMAE]` is set to inform the driver.

Once the `IRQSTAT[DMAE]` interrupt is received, software should send `CMD12` to abort the current transfer and read `DSADDR[DS_ADDR]` to obtain the start address of the corrupted block. After the DMA error is fixed, the software applies a data reset and restarts the transfer from this address to recover the corrupted block.

17.5.3 SD protocol unit

In addition to handling all SD protocol affairs, as well as other miscellaneous functions, the SD protocol unit performs the following:

- Acts as the bridge between the internal buffer and the SD bus
- Sends the command data and its argument serially
- Stores the serial response bit stream into corresponding registers
- Detects the bus state on the `SDHC_DAT[0]` line
- Gates off the SD clock when the buffer announces a danger status
- Detects the write-protect state

The SD protocol unit consists of the following four submodules:

- [SD transceiver](#)
- [SD clock and monitor](#)
- [Command agent](#)
- [Data agent](#)

17.5.3.1 SD transceiver

In the SD protocol unit, the transceiver is the main control module.

It consists of an FSM and the control module, from which the control signals for all other three modules are generated.

17.5.3.2 SD clock and monitor

This module monitors the signal level on all eight data lines and the command lines, directly route the level values into the register bank for the driver to debug with.

The transceiver reports the card insertion state according to the SDHC_CD_B state, or signal level on SDHC_DAT[3] line when PROCTL[D3CD] is set.

The module detects the SDHC_WP (write protect) line. With the information of SDHC_WP state, the register bank ignores the command accompanied by write operation, when the SD_WP switch is on.

If the internal data buffer is in danger and the SD clock must be gated off to avoid buffer over/underrun, this module asserts the gate of output SD clock to shut the clock off. When the buffer danger is eliminated when system access of the buffer catches up, the clock gate of this module is open and the SD clock is active again.

17.5.3.3 Command agent

The command agent deals with the transactions on the SDHC_CMD line.

See the figure below for an illustration of the structure for the command CRC shift register.

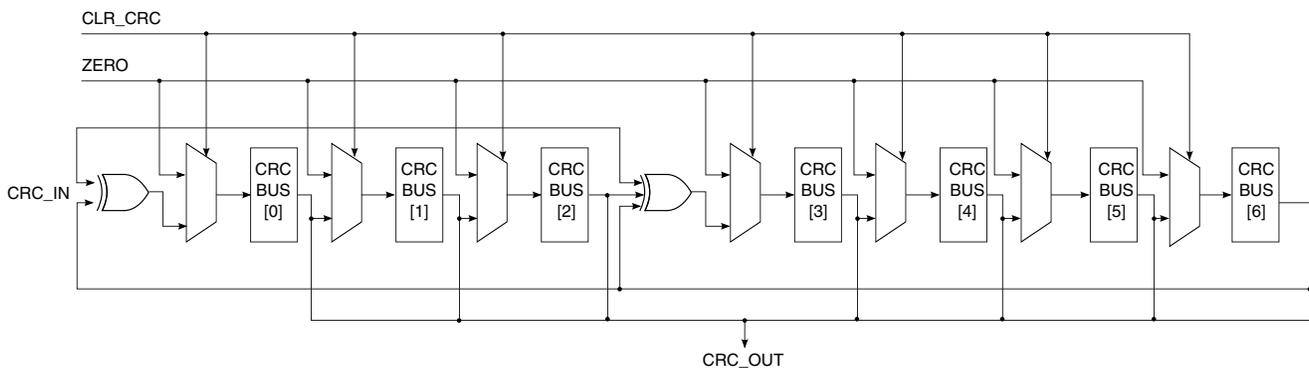


Figure 17-27. Command CRC shift register

The CRC polynomials for the SDHC_CMD are as follows:

1. Generator polynomial: $G(x) = x^7 + x^3 + 1$
 $M(x) = (\text{first bit}) \times x^n + (\text{second bit}) \times x^{n-1} + \dots + (\text{last bit}) \times x^0$
 $\text{CRC}[6:0] = \text{Remainder} [(M(x) \times x^7) \div G(x)]$

17.5.3.4 Data agent

The data agent handles the transactions on the eight data lines.

The CRC polynomials for the SDHC_DAT are as follows:

1. Generator polynomial: $G(x) = x^{16} + x^{12} + x^5 + 1$

$$M(x) = (\text{first bit}) \times x^n + (\text{second bit}) \times x^{n-1} + \dots + (\text{last bit}) \times x^0$$

$$\text{CRC}[15:0] = \text{Remainder} [(M(x) \times x^{16}) \div G(x)]$$

17.5.4 Clock and reset manager

This module controls the reset signals within the eSDHC.

Within eSDHC, there are four types of reset signals, all of which are fed into the clock and reset manager module: hardware reset, software reset for all, software reset for data, and software reset for command. Stable signals are generated inside the module in order to reset all other modules.

In addition, the clock and reset manager module:

- gates off all inside signals,
- monitors the activities of all other modules,
- supplies the clocks for other modules,
- and, when enabled, automatically gates off the corresponding clocks.

17.5.5 Clock generator

The clock generator generates the SDHC_CLK by dividing the internal bus clock into two stages.

Refer to [Figure 17-28](#) for the structure of the divider, in which the term base represents the frequency of the internal bus clock ($\text{ccb_clk}/2$). Refer to `SYSCCTL[SDCLKFS]` and `SYSCCTL[DVS]` (see [System control \(eSDHC_SYSCCTL\)](#)) to select the divisor values.

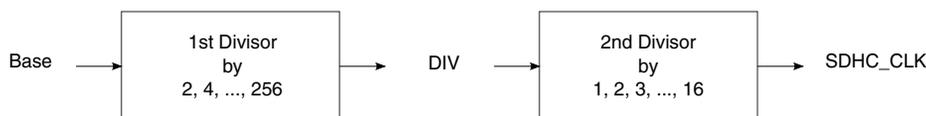


Figure 17-28. Two stages of clock divider

The first stage is a prescaler. The frequency of clock output from this stage, DIV, can be base, base/2, base/4, ..., or base/256.

The second stage outputs the actual clock, SDHC_CLK, as the driving clock for all sub-modules of SD protocol unit, and the sync FIFOs in [Figure 17-25](#) to synchronize with the data rate from the internal data buffer. It can be div, div/2, div/3,..., or div/16. Thus, the highest frequency of SDHC_CLK generated by the internal bus clock (ccb_clk/2) is base while the lowest frequency is base/4096.

17.5.6 Card insertion and removal detection

The eSDHC uses the SDHC_DAT[3] pin or the SDHC_CD_B pin to detect card insertion or removal.

- When the SDHC_DAT[3] pin is used for card detection, the user needs to pull down this pad as a default state. When there is no card on the MMC/SD bus, the SDHC_DAT[3] is pulled to a low voltage level by default. When any card is inserted to or removed from the socket, the eSDHC detects the logic value changes on the SDHC_DAT[3] pin and generates an interrupt.
- When the SDHC_DAT[3] pin is not used for card detection, SDHC_CD_B must be connected for card detection. It may be implemented by a GPIO. Whether or not SDHC_DAT[3] is configured for card detection, SDHC_CD_B is always a reference for card detection, either SDHC_DAT[3] or SDHC_CD_B reports the card inserted, the eSDHC informs the host system that a card is inserted, and the interrupt is sent if it is enabled.

Note that the SDHC_DAT[3] is not used for card detection on P2020.

17.5.7 Power management and wake-up events

When there is no operation between the eSDHC and the card through the SD bus, the user can completely disable the internal clocks in the chip-level clock-control module to save power. When using the eSDHC to communicate with the card, it can enable the clock and start the operation. This can be done by clearing the SCCR[SDHCCM] bits.

In some circumstances, when the clocks to eSDHC are disabled, or when the system is in low-power mode, there are events that require the user to enable the clock and handle the event. These events are called wake-up interrupts. The eSDHC can generate these interrupts even when there are no clocks enabled. The two interrupts that can be used as wake-up events are as follows:

- Card removal interrupt
- Card insertion interrupt

These two wake-up events (or wake-up interrupts) can also be used to wake the system from low-power modes.

NOTE

To make the interrupt as a wake-up event when all clocks to the eSDHC are disabled or when the entire system is in low-power mode, the corresponding wake-up enable bit must be set. See [Protocol control \(eSDHC_PROCTL\)](#) for more information on the eSDHC PROCTL register.

17.5.7.1 Setting wake-up events

For the eSDHC to respond to a wake-up event, the software must set the respective wake-up enable bit before the CPU enters sleep mode.

See [Protocol control \(eSDHC_PROCTL\)](#), for more information on the wake-up enable bits.

Before the software disables the host clock, it should ensure that all of the following conditions have been met:

- No read or write transfer is active
- Data and command lines are not active
- No interrupts are pending
- Internal data buffer is empty

17.6 Initialization/application information

All communication between the system and cards is controlled by the host.

The host sends the following types of commands: broadcast and addressed (point-to-point) commands.

Broadcast commands, such as GO_IDLE_STATE, SEND_OP_COND or ALL_SEND_CID, are intended for all cards. In broadcast mode, all cards are in the open-drain mode to avoid bus contention. See [Commands for MMC/SD](#), for the bc and bcr category commands.

After the broadcast command CMD3 is issued, the cards enter standby mode. From this point, addressed commands are used. In this mode, the SDHC_CMD/SDHC_DAT I/O pads turn to push-pull mode in order to provide the driving capability for maximum frequency operation. See [Commands for MMC/SD](#), for the ac and adtc category commands.

17.6.1 Command send and response receive basic operation

Assuming the data type WORD is an unsigned 32-bit integer, the following flow acts as a guideline for sending a command to the card(s):

```
send_command(cmd_index, cmd_arg, other requirements)
{
WORD wCmd; // 32-bit integer to make up the data to write into the XFERTYP register, it is
// recommended to implement in a bit-field manner
wCmd = (<cmd_index> & 0x3f) << 24; // set the first 8 bits as '00'+<cmd_index>
set CMDTYP, DPSEL, C ICEN, CCCEN, RSTTYP, and DTDSEL according to the command index;
// XFERTYP register bits
if (internal DMA is used) wCmd |= 0x1;
if (multi-block transfer) {
set XFERTYP[MSBSEL] bit;
if (finite block number) {
set XFERTYP[BCEN] bit;
if (autol2 command is to use) set XFERTYP[AC12EN] bit;
}
}
}
write_reg(CMDARG, <cmd_arg>); // configure the command argument
write_reg(XFERTYP, wCmd); // set XFERTYP register as wCmd value to issue the command
}
wait_for_response(cmd_index)
{
while (IRQSTAT[CC] is not set); // wait until command complete bit is set
read IRQSTAT register and check if any error bits about command are set;
if (any error bits are set) report error;
write 1 to clear IRQSTAT[CC] and all command error bits;
}
```

For the sake of simplicity, the function `wait_for_response` is implemented here by means of polling. For an effective and formal way, the response is usually checked after the command complete interrupt is received. This ensures the corresponding interrupt status bits are enabled.

For some scenarios, the response timeout is expected. For instance, after all cards respond to CMD3 and enter standby state, there is no response to the host when CMD2 is sent. The host driver should manage false errors similar to this with caution.

17.6.2 Card identification mode

When a card is inserted into the socket or the card is reset by the host, the host needs to validate the operation voltage range, identify the cards, and request the cards to publish the relative card address (RCA) or to set the RCA for the MMCs.

17.6.2.1 Card detect

This figure shows a flow diagram of the detection of MMC and SD cards using the eSDHC.

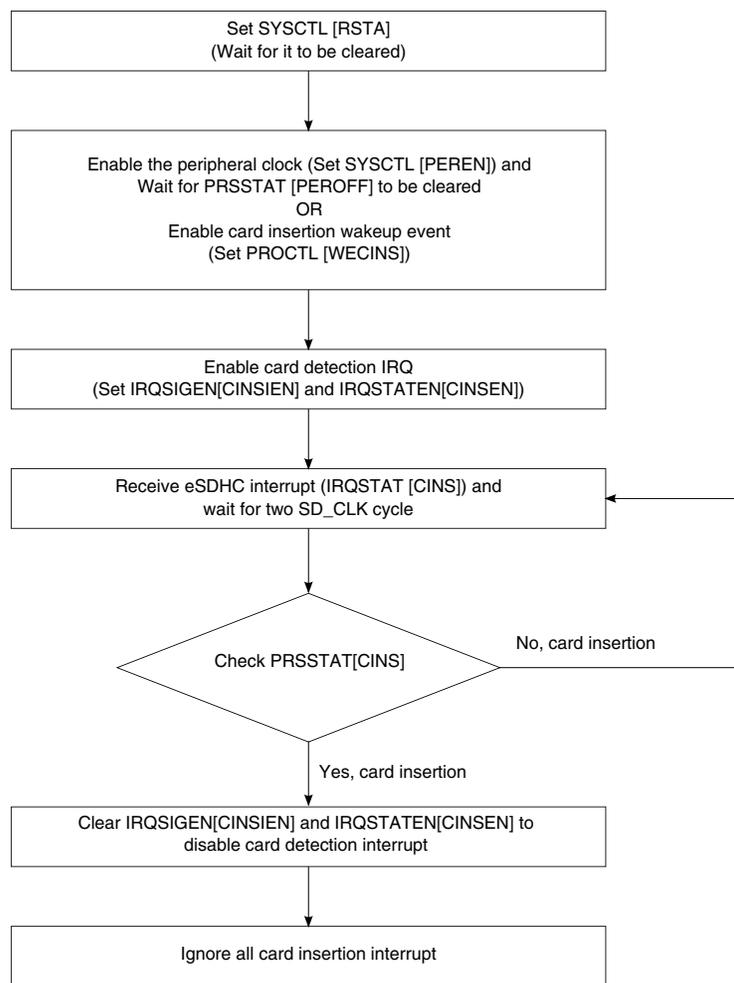


Figure 17-29. Flow diagram for card detection

1. Set SYSCTL[RSTA] and wait for it to be cleared.
2. Set SYSCTL[PEREN] to enable the clocks and wait for PRSSTAT[PEROFF] to be cleared or enable wakeup event (set PROCTL[WECINS]).
3. Set IRQSIGEN[CINSIEN] and IRQSTATEN[CINSEN] to enable card detection interrupt.
4. When an IRQSTAT[CINS] interrupt from eSDHC is received, wait for two SD_CLK cycle and then check PRSSTAT[CINS] to see if the interrupt is caused by card insertion.

5. Clear the IRQSIGEN[CINSIEN] and IRQSTATEN[CINSEN] to disable card detection interrupt and ignore all card insertion interrupt afterwards.

17.6.2.2 Reset

The host consists of the following types of reset:

- Hardware reset (card and host), which is driven by POR (power-on reset).
- Software reset (host only), which is proceeded by the write operation on the SYSCTL[RSTD], SYSCTL[RSTC], or SYSCTL[RSTA] bits to reset the data part, command part, or all parts of the host controller, respectively.
- Card reset (card only). The command CMD0, GO_IDLE_STATE, is the software reset command for all types of MMCs and SD memory cards. This command sets each card into idle state regardless of the current card state. The cards are initialized with a default relative card address (RCA = 0x0000) and with a default driver stage register setting (lowest speed, highest driving current capability).

After the card is reset, the host must validate the voltage range of the card. See [Figure 17-30](#) for the software flow chart.

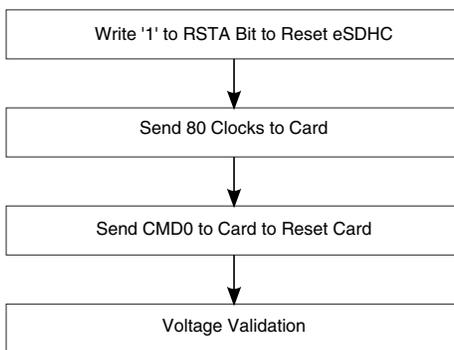


Figure 17-30. Flow chart for reset of eSDHC card

The following pseudocode shows how to initialize the eSDHC host controller and the memory card.

```

software_reset()
{
    Configure the I/O muxes to select SD signals;
    set_bit(SYSCTL, RSTA); // software reset the host
    set SYSCTL[DTOCV and SDCLKFS]; // get the SDHC_CLK of frequency
    around 400 kHz
    poll PRSSTAT[CIHB and CDIHB]; // wait until both bits are
    cleared
    set_bit(SYSCTRL, INTIA); // send 80 clock ticks for
    card to power-up
    If the card is SD/MMC
        send_command(CMD_GO_IDLE_STATE, <other parameters>); // reset the card with
  
```

```
CMD0
```

```
}
```

17.6.2.3 Voltage validation

All cards should be able to establish communication with the host by using any operation voltage in the maximum allowed voltage range, as specified in this standard.

However, the supported minimum and maximum values for V_{DD} are defined in the operation conditions register (OCR) and may not cover the whole range. Cards that store the CID (card identification) and CSD data in the preloaded memory are only able to communicate this information under data transfer V_{DD} conditions. This means that if the host and card have different V_{DD} ranges, the card is not able to complete the identification cycle, nor is it able to send CSD data.

Therefore, a special command is available, as follows:

- SEND_OP_CONT (CMD1 for MMC),
- SD_SEND_OP_CONT (ACMD41 for SD Memory)

The voltage validation procedure is designed to provide a mechanism to identify and reject cards that do not match the V_{DD} range(s) desired by the host. This is accomplished when the host sends the desired V_{DD} voltage window as the operand of this command. Cards that cannot perform data transfer in the specified range must discontinue any further bus operations and enter the inactive state.

By omitting the voltage range in the command, the host can query each card and determine the common voltage range before sending out-of-range cards into the inactive state. This query should be used if the host is able to select a common voltage range or if a notification should be sent to the system when a non-usable card in the stack is detected.

17.6.2.4 Card registry

Card registry on the MMC and SD cards differs.

SD card registry is as follows:

NOTE

The identification process starts at a clock rate lower than 400 kHz and a power voltage higher than 2.7 V, as defined by the card specification. At this time, the SDHC_CMD line output drives are push-pull drivers instead of open-drain.

1. After the bus is activated, the host requests the cards to send their valid operation conditions. The response to ACMD41 is the operation condition register of the card.
2. The host should send the same command to all of the new cards in the system. Incompatible cards are placed into the inactive state.
3. The host then issues the command, ALL_SEND_CID (CMD2), to each card to get its CID. Cards that are currently unidentified (that is, in ready state), send their CID number as the response.
4. After the CID is sent by the card, the card goes into the identification state.
5. The host then issues Send_Relative_Addr (CMD3), requesting the card publish a new relative card address (RCA) that is shorter than CID. This RCA is used to address the card for future data transfer operations.
6. Once the RCA is received, the card changes to standby state. At this point, if the host wants the card to have an alternative RCA number, it may ask the card to publish a new number by sending another Send_Relative_Addr command to the card. The last published RCA is the actual RCA of the card.
7. The host repeats the identification process with CMD2 and CMD3 for each card in the system until the last CMD2 gets no response from the cards in system.

MMC card registry is as follows:

NOTE

For MMC operation, the host starts the card identification process in open-drain mode with the identification clock rate lower than 400 kHz and the power voltage higher than 2.7 V. The open-drain driver stages on the SDHC_CMD line allow for parallel card operation during card identification.

1. After the bus is activated, the host requests the cards to send their valid operation conditions (CMD1). The response to CMD1 is the wired-OR operation on the condition restrictions of all cards in the system. Incompatible cards are sent into inactive state.
2. The host then issues the broadcast command All_Send_CID (CMD2), asking all cards for their unique CID number. All unidentified cards (the cards in ready state) simultaneously start sending their CID numbers serially, while bit-wise monitoring their outgoing bit stream. Those cards whose outgoing CID bits do not match the corresponding bits on the command line in any one of the bit periods stop sending their CID immediately and must wait for the next identification cycle. Since the CID is unique for each card, only one card can successfully send its full CID to the host. This card then goes into identification state.
3. Thereafter, the host issues Set_Relative_Addr (CMD3) to assign to this card a relative card address (RCA).

4. Once the RCA is received, the card state changes to standby state, the card does not react in further identification cycles, and its output driver switches from open-drain to push-pull.
5. The host repeats the process, namely CMD2 and CMD3, until the host receives a time-out condition to recognize completion of the identification process.

17.6.3 Card access

The following sections describe the supported access modes with external cards.

17.6.3.1 Block write

This section describes the process of writing data to external cards in block mode.

17.6.3.1.1 Normal write

During block write (CMD24-27), one or more blocks of data are transferred from the host to the card; the host appends a CRC to the end of each block.

If the CRC fails, the card should indicate the failure on the SDHC_DAT line (see below). The transferred data is discarded and not written, and all further transmitted blocks (in multi-block write mode) are ignored.

When the host uses partial blocks whose accumulated length is not block-aligned and block misalignment is not allowed (CSD parameter WRITE_BLK_MISALIGN is not set), the card detects the block misalignment error and aborts programming before the beginning of the first misaligned block. The card sets the ADDRESS_ERROR error bit in the status register, as defined in the MMC/SD specification, and then waits in the receive-data state for a stop command, ignoring all further data transfers. The write operation is also aborted when the host attempts to write over a write-protected area.

For MMC and SD cards, programming the CID and CSD registers does not require a previous block length setting. The transferred data is also CRC-protected. If a part of the CSD or CID register is stored in the ROM, this unchangeable section must match the corresponding section of the receive buffer. If this match fails, the card reports an error and does not change any register content.

Some cards may require a long and unpredictable period of time to write a block of data. After receiving a block of data and completing the CRC check, the card begins writing. When its write buffer is full and unable to accept new data from a new WRITE_BLOCK command, the card holds the SDHC_DAT line low. The host may poll the status of the card with a SEND_STATUS command (CMD13) cards, at any time and the card

responds with its status. The card status indicates whether the card can accept new data or if the write process is still in progress. The host may deselect the card by issuing CMD7 (to select a different card) to change the card into the standby state and release the SDHC_DAT line without interrupting the write operation. When re-selecting the card, the host reactivates the busy indication by pulling SDHC_DAT low when programming is still in progress and the write buffer is unavailable.

For the sake of simplicity, the software flow described below incorporates the internal DMA, and the write operation is a multi-block write with Auto CMD12 enabled. For the CPU polling status method and different transfer natures, remove the internal DMA part of the procedure and insert alternative steps.

1. Check the card status and wait until the card is ready for data.
2. Set the card block length.
 - For MMC/SD cards, use SET_BLOCKLEN (CMD16).
3. Set eSDHC BLKATTR[BLKSIZE] to the same block length set in the card in step 2.
4. Set eSDHC BLKATTR[BLKCNT] with the number of blocks to send.
5. Disable the buffer write ready interrupt, configure the DMA setting, and enable the eSDHC DMA when sending the command with data transfer. Set XFERTYP[AC12EN].
6. Wait for the transfer complete interrupt.
7. Check the status bit to see if a read CRC error or any other errors occurred between sending Auto CMD12 and receiving the response.

17.6.3.1.2 Write with pause

The write operation can be paused during the transfer.

Instead of stopping the SDHC_CLK at any time to pause all the operations which is also inaccessible to the host driver, the driver can set PROCTL[SABGREQ] to pause the transfer between the data blocks. Because there is no timeout condition in a write operation during the data blocks, a write operation to the cards can be paused in this way and if line SDHC_DAT0 is not required to de-assert to release busy state.

Similar to the flow described in [Normal write](#), the write with pause is shown with the same type of write operations:

1. Check the card status and wait until card is ready for data.
2. Set the card block length.
 - For MMC/SD cards, use SET_BLOCKLEN (CMD16)
3. Set the eSDHC BLKATTR[BLKSIZE] to the same as the block length set in the card in step 2.
4. Set eSDHC BLKATTR[BLKCNT] with the number of blocks to send.

5. Disable the buffer write ready interrupt, configure the DMA setting, and enable the eSDHC DMA when sending the command with data transfer. Set XFERTYP[AC12EN].
6. Set PROCTL[SABGREQ].
7. Wait for the transfer complete interrupt.
8. Clear PROCTL[SABGREQ].
9. Check the status bit to see if a read CRC error occurred.
10. Set PROCTL[CREQ] to continue the read operation.
11. Wait for the transfer complete interrupt.
12. Check the status bit to see if a read CRC error or any other errors occurred between sending Auto CMD12 and receiving the response.

The number of blocks left during the data transfer is accessible by reading the content of BLKATTR[BLKCNT]. Due to the data transfers and setting PROCTL[SABGREQ] are concurrent, along with the delay of register read and the register setting, the actual number of blocks left may not be the same as the value read earlier. The driver should read the value of BLKATTR[BLKCNT] after the transfer is paused and the transfer complete interrupt is received.

It is also possible that the transfer of the last block begins when the stop-at-block-gap request is sent to the buffer. In this case, the next block gap is the actual end of the transfer, and therefore, the request is ignored. The driver should treat this as a non-pause transfer and a common write operation.

When the write operation is paused, the data transfer inside the host system does not stop and the transfer remains active until the data buffer is full.

17.6.3.2 Block read

This section discusses normal read and read with pause.

17.6.3.2.1 Normal read

For block reads, the basic unit of a data transfer is a block whose maximum size is stored in areas defined in corresponding card specifications.

A CRC is appended to the end of each block, ensuring data transfer integrity. CMD17, CMD18, CMD53, and so on, can initiate a block read. After completing the transfer, the card returns to the transfer state.

For multiblock reads, data blocks are continuously transferred until a stop command is issued. If the host uses partial blocks whose accumulated length is not block aligned and block misalignment is not allowed, the card that does not support partial block length, should detect the block misalignment at the beginning of the first misaligned block and report the error, depending on its card type.

For simplicity, the software flow described below incorporates the internal DMA, and the read operation is a multiblock read with Auto CMD12 enabled. For the other method (CPU polling status) and different transfer nature, the internal DMA part should be removed and the alternative steps are straightforward.

1. Check the card status and wait until the card is ready for data.
2. Set the card block length.
 - For MMC/SD cards, use SET_BLOCKLEN (CMD16)
3. Set eSDHC BLKATTR[BLKSIZE] to the same as the block length set in the card in step 2.
4. Set eSDHC BLKATTR[BLKCNT] with the number of blocks to send.
5. Disable the buffer read ready interrupt, configure the DMA setting, and enable the eSDHC DMA when sending the command with data transfer. Set XFERTYP[AC12EN].
6. Wait for the transfer complete interrupt.
7. Check the status bit to see if a read CRC error or any other errors occurred between sending Auto CMD12 and receiving the response.

17.6.3.2.2 Read with pause

In general, the read operation is not able to pause.

Similarly to the flow described in [Normal read](#), the read with pause is shown with the following read operations:

1. Set PROCTL[RWCTL].
2. Check the card status and wait until the card is ready for data.
3. Set the card block length.
 - For MMC/SD cards, use SET_BLOCKLEN (CMD16)
4. Set eSDHC BLKATTR[BLKSIZE] to the same as the block length set in the card in Step 2.
5. Set eSDHC BLKATTR[BLKCNT] with the number of blocks to send.
6. Disable the buffer read ready interrupt, configure the DMA setting, and enable the eSDHC DMA when sending the command with data transfer. Set XFERTYP[AC12EN].
7. Set PROCTL[SABGREQ].
8. Wait for the transfer complete interrupt.

9. Clear PROCTL[SABGREQ].
10. Check the status bit to see if a read CRC error occurred.
11. Set PROCTL[CREQ] to continue the read operation.
12. Wait for the transfer complete interrupt.
13. Check the status bit to see if a read CRC error or any other errors occurred between sending Auto CMD12 and receiving the response.

Similar to the write operation, it is possible to meet the ending block of the transfer when paused. In this case, the eSDHC ignores the stop-at-block-gap request and treats it as a command read operation.

Unlike the write operation, there is no remaining data inside the buffer when the transfer is paused. All data received before the pause is transferred to the host system; the internal data buffer is not flushed.

17.6.3.3 Transfer error

This section discusses CRC errors, internal DMA errors, and Auto CMD12 errors.

17.6.3.3.1 CRC error

At the end of a block transfer, a write CRC status error or read CRC error may occur.

For this type of error, the last block received should be discarded because the integrity of the data block is not guaranteed. It is recommended that the following data blocks be discarded and the block be retransferred from the corrupted one. For a multiblock transfer, the host driver should issue CMD12 to abort the current process and start the transfer by a new data command. In this scenario, even when the XFERTYP[AC12EN, BCEN] are set, the eSDHC does not automatically send CMD12 because the last block is not transferred. On the other hand, if it is within the last block that CRC error occurs, Auto CMD12 is sent by the eSDHC. In this case, the driver should resend or re-obtain the last block with a single block transfer.

17.6.3.3.2 Internal DMA error

During the data transfer with the internal DMA, if the DMA engine encounters an error on the platform bus, the DMA operation is aborted and a DMA error interrupt is sent to the host system.

When acknowledged by such an interrupt, the driver should calculate the start address of the data block where the error occurred. The start address can be calculated by either of the following methods:

- Read the DSADDR[DSADDR] field. The error occurs during the previous burst. Therefore, by taking the block size, the previous burst length, and the start address of the next burst transfer into account, one can obtain the start address of the corrupted block.
- Read the BLKATTR[BLKCNT] field. The start address of the corrupted block can be calculated by the number of blocks left, the total number to transfer, the start address of transfer, and the size of each block. However, if BCEN is not set, the contents of the block attribute register does not change and this method does not work.

When a DMA error occurs, it is recommended to abort the current transfer by means of CMD12 (for multi-block transfer), apply a reset for data, and restart the transfer from the corrupted block to recover the error.

17.6.3.3.3 Auto CMD12 error

After the last block of a multi-block transfer is sent or received and XFERTYP[AC12EN] is set when the data transfer is initiated by the data command, the eSDHC automatically sends CMD12 to the card to stop the transfer.

When an error occurs at this point, the host driver should respond as follows:

Auto CMD12 response timeout	It is not certain whether the command has been accepted by the card or not. The driver should clear the Auto CMD12 error status bits and resend CMD12 until it is accepted by the card.
Auto CMD12 response CRC error	Because CMD12 has been received by the card, the card aborts the transfer. The driver may ignore the error and clear the error status bit.
Auto CMD12 conflict error or not sent	The command was not sent. Therefore, the driver should send CMD12 manually.

17.6.3.4 Card interrupt

The external cards can inform the host controller through the use of special signals.

17.6.4 Switch function

MMCs transferring data with a bus width other than one-bit wide is a new feature added to the MMC specification.

The high-speed timing mode for all card devices is also newly-defined in recent various card specifications. To enable these new features, a type of switch command should be issued by the host driver.

For SD cards, the high-speed mode is queried and enabled by CMD6 (with the mnemonic symbol as SWITCH_FUNC); for MMCs, the high-speed mode is queried by CMD8 and enabled by CMD6 (with the mnemonic symbol as SWITCH).

The 4-bit and 8-bit bus width of MMC is also enabled by the SWITCH command, but with a different argument.

These new functions can also be disabled by software reset, but such manner of restoring to normal mode is not recommended because a complete identification process is needed before the card is ready for data transfer.

17.6.4.1 Query, enable and disable SD high speed mode

```
enable_sd_high_speed_mode(void)
{
    set BLKATTR[BLKCNT] to 1 (block), set BLKATTR[BLKSIZE] to 64 (bytes);
    send CMD6, with argument 0xFFFFF1 and read 64 bytes of data accompanying the R1
        response;
    wait data transfer done bit is set;
    check if the bit 401 of received 512 bit is set;
    if (bit 401 is '0') report the SD card does not support high speed mode and return;
    send CMD6, with argument 0x80FFFF1 and read 64 bytes of data accompanying the R1
        response;
    check if the bit field 379~376 is 0xF;
    if (the bit field is 0xF) report the function switch failed and return;
    change clock divisor value or configure the system clock feeding into eSDHC to
generate
        the card_clk of around 50MHz;
    (data transactions like normal peers)
}
disable_sd_high_speed_mode(void)
{
    set BLKCNT field to 1 (block), set BLKSIZE field to 64 (bytes);
    send CMD6, with argument 0x80FFFF0 and read 64 bytes of data accompanying the R1
        response;
    check if the bit field 379~376 is 0xF;
    if (the bit field is 0xF) report the function switch failed and return;
    change clock divisor value or configure the system clock feeding into eSDHC to
generate
        the card_clk of the desired value below 25MHz;
    (data transactions like normal peers)
}
```

17.6.4.2 Query, enable and disable MMC high speed mode

```
enable_mmc_high_speed_mode(void)
{
    send CMD9 to get CSD value of MMC;
    check if the value of SPEC_VER field is 4 or above;
    if (SPEC_VER value is less than 4) report the MMC does not support high speed mode
and
```

```

        return;
    set BLKCNT field to 1 (block), set BLKSIZE field to 512 (bytes);
    send CMD8 to get EXT_CSD value of MMC;
    extract the value of CARD_TYPE field to check the 'high speed mode' in this MMC is
        26MHz or 52MHz;
    send CMD6 with argument 0x1B90100;
    send CMD13 to wait card ready (busy line released);
    send CMD8 to get EXT_CSD value of MMC;
    check if HS_TIMING byte (byte number 185) is 1;
    if (HS_TIMING is not 1) report MMC switching to high speed mode failed and return;
    change clock divisor value or configure the system clock feeding into eSDHC to
generate
        the card_clk of around 26MHz or 52MHz according to the CARD_TYPE;
    (data transactions like normal peers)
}
disable_mmc_high_speed_mode(void)
{
    send CMD6 with argument 0x2B90100;
    set BLKCNT field to 1 (block), set BLKSIZE field to 512 (bytes);
    send CMD8 to get EXT_CSD value of MMC;
    check if HS_TIMING byte (byte number 185) is 0;
    if (HS_TIMING is not 0) report the function switch failed and return;
    change clock divisor value or configure the system clock feeding into eSDHC to
generate
        the card_clk of the desired value below 20MHz;
    (data transactions like normal peers)
}

```

17.6.4.3 Set MMC bus width

```

change_mmc_bus_width(void)
{
    send CMD9 to get CSD value of MMC;
    check if the value of SPEC_VER field is 4 or above;
    if (SPEC_VER value is less than 4) report the MMC does not support multiple bit
width
        and return;
    send CMD6 with argument 0x3B70x00; (8-bit, x=2; 4-bit, x=1; 1-bit, x=0)
    send CMD13 to wait card ready (busy line released);
    (data transactions like normal peers)
}

```

17.6.5 Commands for MMC/SD

See the table below for a list of commands for the MMC/SD cards.

See the corresponding specifications for details about the command information.

Four kinds of commands control the MMC, as follows:

1. Broadcast commands (bc)-no response
2. Broadcast commands with response (bcr)-response from all cards simultaneously
3. Addressed (point-to-point) commands (ac)-no data transfer on SDHC_DAT
4. Addressed (point-to-point) data transfer commands (adtc)

Table 17-35. Commands for MMC/SD

CMD INDEX	Type	Argument	Resp	Abbreviation	Description ¹
CMD0	bc	[31:0] stuff bits	-	GO_IDLE_STATE	Resets all MMC and SD memory cards to idle state.
CMD1	bcr	[31:0] OCR without busy	R3	SEND_OP_COND	Asks all MMCs and SD memory cards in idle state to send their operation conditions register contents in the response on the SDHC_CMD line.
CMD2	bcr	[31:0] stuff bits	R2	ALL_SEND_CID	Asks all cards to send their CID numbers on the SDHC_CMD line.
CMD3 ⁽¹⁾	ac	[31:6] RCA [15:0] stuff bits	R1	SET/ SEND_RELATIVE_ADDR	Assigns relative address to the card.
CMD4	bc	[31:0] DSR [15:0] stuff bits	-	SET_DSR	Programs the DSR of all cards.
CMD6 ⁽²⁾	adtc	[31] Mode 0: Check function 1: Switch function [30:8] Reserved for function groups 6 ~ 3 (All 0 or 0xFFFF) [7:4] Function group1 for command system [3:0] Function group2 for access mode	R1	SWITCH_FUNC	Checks switch ability (mode 0) and switch card function (mode 1). Refer to SD Physical Specification version 2.0 for details.
CMD6 ⁽³⁾	ac	[31:26] Set to 0 [25:24] Access [23:16] Index [15:8] Value [7:3] Set to 0 [2:0] Cmd Set	R1b	SWITCH	Switches the mode of operation of the selected card or modifies the EXT_CSD registers. Refer to the MultiMediaCard System Specification version 4.0 final draft 2 for details.
CMD7	ac	[31:6] RCA [15:0] stuff bits	R1b	SELECT/DESELECT_CARD	Command toggles a card between the stand-by and transfer states or between the programming and disconnect states. In both cases, the card is selected by its own relative address and gets deselected by any other address; address 0 deselects all.
CMD8	bcr	[31:12] reserved bits [11:8] supply voltage(VHS) [7:0] check pattern	R7	SEND_IF_COND	Sends SD Memory Card interface condition, which includes host supply voltage information and asks the card whether card supports voltage. Reserved bits shall be set to '0'.
CMD8	adtc	[31:0] stuff bits	R1	SEND_EXT_CSD	The card sends its EXT_CSD register as a block of data, with block size of 512 bytes.

Table continues on the next page...

Table 17-35. Commands for MMC/SD (continued)

CMD INDEX	Type	Argument	Resp	Abbreviation	Description ¹
CMD9	ac	[31:6] RCA [15:0] stuff bits	R2	SEND_CSD	Addressed card sends its card-specific data (CSD) on the SDHC_CMD line.
CMD10	ac	[31:6] RCA [15:0] stuff bits	R2	SEND_CID	Addressed card sends its card-identification (CID) on the SDHC_CMD line.
CMD11	adtc	[31:0] data address	R1	READ_DAT_UNTIL_STOP	Reads data stream from the card starting at the given address until STOP_TRANSMISSION is received.
CMD12	ac	[31:0] stuff bits	R1b	STOP_TRANSMISSION	Forces the card to stop transmission.
CMD13	ac	[31:6] RCA [15:0] stuff bits	R1	SEND_STATUS	Addressed card sends its status register.
CMD14	Reserved				
CMD15	ac	[31:6] RCA [15:0] stuff bits	-	GO_INACTIVE_STATE	Sets the card to inactive state in order to protect the card stack against communication breakdowns.
CMD16	ac	[31:0] block length	R1	SET_BLOCKLEN	Sets the block length (in bytes) for all following block commands (read and write). Default block length is specified in the CSD.
CMD17	adtc	[31:0] data address	R1	READ_SINGLE_BLOCK	Reads a block of the size selected by the SET_BLOCKLEN command.
CMD18	adtc	[31:0] data address	R1	READ_MULTIPLE_BLOCK	Continuously transfers data blocks from card to host until interrupted by a stop command.
CMD19	Reserved				
CMD20	adtc	[31:0] data address	R1	WRITE_DAT_UNTIL_STOP	Writes data stream from the host starting at the given address until the STOP_TRANSMISSION command is received.
CMD21-23	Reserved				
CMD24	adtc	[31:0] data address	R1	WRITE_BLOCK	Writes a block of the size selected by the SET_BLOCKLEN command.
CMD25	adtc	[31:0] data address	R1	WRITE_MULTIPLE_BLOCK	Continuously writes blocks of data until the STOP_TRANSMISSION command is received.
CMD26	adtc	[31:0] stuff bits	R1	PROGRAM_CID	Programming of the card identification register. This command should be issued only once per card. The card contains hardware to prevent this operation after the first programming. Normally this command is reserved for the manufacturer.
CMD27	adtc	[31:0] stuff bits	R1	PROGRAM_CSD	Programming of the programmable bits of the CSD.

Table continues on the next page...

Table 17-35. Commands for MMC/SD (continued)

CMD INDEX	Type	Argument	Resp	Abbreviation	Description ¹
CMD28	ac	[31:0] data address	R1b	SET_WRITE_PROT	If the card has write-protection features, this command sets the write protection bit of the addressed group. The properties of write protection are coded in the card-specific data (WP_GRP_SIZE).
CMD29	ac	[31:0] data address	R1b	CLR_WRITE_PROT	If the card provides write-protection features, this command clears the write protection bit of the addressed group.
CMD30	adtc	[31:0] write protect data address	R1	SEND_WRITE_PROT	If the card provides write-protection features, this command asks the card to send the status of the write-protection bits.
CMD31	Reserved				
CMD32	ac	[31:0] data address	R1	TAG_SECTOR_START	Sets the address of the first sector of the erase group.
CMD33	ac	[31:0] data address	R1	TAG_SECTOR_END	Sets the address of the last write block of the continuous range to be erased.
CMD34	ac	[31:0] data address	R1	UNTAG_SECTOR	Removes one previously selected sector from the erase selection.
CMD35	ac	[31:0] data address	R1	TAG_ERASE_GROUP_START	Sets the address of the first erase group within a range to be selected for erase.
CMD36	ac	[31:0] data address	R1	TAG_ERASE_GROUP_END	Sets the address of the last erase group within a continuous range to be selected for erase.
CMD37	ac	[31:0] data address	R1	UNTAG_ERASE_GROUP	Removes one previously selected erase group from the erase selection.
CMD38	ac	[31:0] stuff bits	R1b	ERASE	Erase all previously selected sectors.
CMD39	ac	[31:0] RCA [15] register write flag [14:8] register address [7:0] register data	R4	FAST_IO	Used to write and read 8-bit (register) data fields. The command address a card and a register and provides the data for writing if the write flag is set. The R4 response contains data read from the address register. This command accesses application dependent registers which are not defined in the MMC standard.
CMD40	bcr	[31:0] stuff bits	R5	GO_IRQ_STATE	Sets the system into interrupt mode.
CMD41	Reserved				
CDM42	adtc	[31:0] stuff bits	R1b	LOCK_UNLOCK	Used to set/reset the password or lock/unlock the card. The size of the data block is set by the SET_BLOCK_LEN command.
CMD43-51	Reserved				

Table continues on the next page...

Table 17-35. Commands for MMC/SD (continued)

CMD INDEX	Type	Argument	Resp	Abbreviation	Description ¹
CMD52	ac	[31:0] stuff bits	R5	IO_RW_DIRECT	Access a single register within the total 128 Kbytes of register space in any I/O function.
CMD53	ac	[31:0] stuff bits	R5	IO_RW_EXTENDED	Access a multiple I/O register with a single command, it allows the reading or writing of a large number of I/O registers.
CMD54	Reserved				
CMD55	ac	[31:16] RCA [15:0] stuff bits	R1	APP_CMD	Indicates to the card that the next command is an application specific command rather than a standard command.
CMD56	adtc	[31:1] stuff bits [0]: RD/WR	R1b	GEN_CMD	Used either to transfer a data block to the card or to get a data block from the card for general-purpose or application-specific commands. The size of the data block is set by the SET_BLOCK_LEN command.
CMD57-63	Reserved				
ACMDs should be preceded with the APP_CMD command					
(Commands listed below are for SD cards only. Other SD commands not listed below are not supported by this module)					
ACMD6	ac	[31:2] stuff bits [1:0] bus width	R1	SET_BUS_WIDTH	Defines the data bus width (00 = 1 bit or 10 = 4 bit bus) to be used for data transfer. The allowed data bus widths are given in DCR register.
ACMD13	adtc	[31:0] stuff bits	R1	SD_STATUS	Send the SD memory card status.
ACMD18	adtc	[31:0] stuff bits	R1	SECURE_READ_MULTI_BLOCK	Protected Area Access Command: Reads continuously transfer data blocks from Protected Area of SD Memory Card. Refer to Security Specification Version 2.00 for more details.
ACMD22	adtc	[31:0] stuff bits	R1	SEND_NUM_WR_SECTORS	Send the number of the written (without errors) sectors. Responds with 32 bit + CRC data block.
ACMD23	ac	[31:23] stuff bits [22:0] number of blocks	R1	SET_WR_BLK_ERASE_COUNT	-
ACMD25	adtc	[31:0] stuff bits	R1	SECURE_WRITE_MULTI_BLOCK	Protected Area Access Command: Writes continuously transfer data blocks to Protected Area of SD memory card. Refer to Security Specification Version 2.00 for more details.

Table continues on the next page...

Table 17-35. Commands for MMC/SD (continued)

CMD INDEX	Type	Argument	Resp	Abbreviation	Description ¹
ACMD26	adtc	[31:0] stuff bits	R1	SECURE_WRITE_MKB	System Area Access Command: Overwrite the existing Media Key Block (MKB) on System Area of SD Memory Card with new MKB. This command is used in dynamic update MKB scheme. Refer to Security Specification Version 2.00 for more details.
ACMD38	ac	[31:0] stuff bits	R1b	SECURE_ERASE	Protected Area Access Command: Erase a specified region of the Protected Area of SD Memory Card. Refer to Security Specification Version 2.00 for more details.
ACMD41	bcr	[31:0] OCR	R3	SD_APP_OP_COND	Asks the accessed card to send its operating condition register (OCR) content in the response on the SDHC_CMD line.
ACMD42	ac	-	R1	SET_CLR_CARD_DETECT	-
ACMD43	adtc	[31:24]Unit_Count: [23:16] MKB_ID: [15:0]Unit_Offset:	R1	GET_MKB	Reads Media Key Block from System Area of SD Memory Card. <ul style="list-style-type: none"> Unit_Count specifies the Number of units to read. (Here, a unit = 512 bytes (fixed).) MKB_ID specifies the application unique number. Unit_Offset specifies the start address(offset) to read. Refer to Security Specification Version 2.00 for more details.
ACMD44	adtc	[31:0] stuff bits	R1	GET_MID	Reads Media ID from the System Area of SD Memory Card. Refer to Security Specification Version 2.00 for more details.
ACMD45	adtc	[31:0] stuff bits	R1	SET_CER_RN1	AKE Command: Writes random number RN1 as challenge1 in AKE process. Refer to Security Specification Version 2.00 for more details.
ACMD46	adtc	[31:0] stuff bits	R1	GET_CER_RN2	AKE Command: Reads random number RN2 as challenge2 in AKE process. Refer to Security Specification Version 2.00 for more details.

Table continues on the next page...

Table 17-35. Commands for MMC/SD (continued)

CMD INDEX	Type	Argument	Resp	Abbreviation	Description ¹
ACMD47	adtc	[31:0] stuff bits	R1	SET_CER_RES2	AKE Command: Writes RES2 as response2 to RN2 in AKE process. Refer to Security Specification Version 2.00 for more details.
ACMD48	adtc	[31:0] stuff bits	R1	GET_CER_RES1	AKE Command: Reads RES1 as response1 to RN1 in AKE process. Refer to Security Specification Version 2.00 for more details.
ACMD49	ac	[31:0] stuff bits	R1b	CHANGE_SECURE_AREA	Protected Area Access Command: Change size of Protected Area. Refer to Security Specification Version 2.00 for more details.
ACMD51	adtc	[31:0] stuff bits	R1	SEND_SCR	Reads the SD Configuration Register (SCR)

1. Registers mentioned in this table are SD card registers.

NOTE

- CMD3 differs for MMC and SD cards. For MMC cards, CMD3 is referred to as SET_RELATIVE_ADDR and has a response type R1. For SD cards, CMD3 is referred to as SEND_RELATIVE_ADDR and has a response type R6, with RCA inside.
- CMD6 differs completely between high-speed MMC cards and high-speed SD cards. Command SWITCH_FUNC is used for high speed SD cards.
- Command SWITCH is for high-speed MMC cards. The index field can contain any value from 0-255, but only values 0-191 are valid. If the index value is in the 192-255 range, the card does not perform any modification and the status bit EXT_CSD[SWITCH_ERROR] is set.

Table 17-36. EXT_CSD access modes

Bits	Access name	Operation
00	Command set	The command set is changed according to the command set field of the argument
01	Set bits	The bits in the pointed byte are set, according to the set bits in the value field.
10	Clear bits	The bits in the pointed byte are cleared, according to the set bits in the value field.
11	Write byte	The value field is written into the pointed byte.

[other]

- CMD8 differs for MMC and SD cards. For SD cards, CMD8 is referred to as SEND_IF_COND. For MMC cards, CMD8 is referred to as SEND_EXT_CSD.

17.6.6 Software restrictions

When polling read or write, once the software begins a buffer read or write, it must access exactly the number of times as set in the watermark level register, as if a DMA burst occurred.

When the internal DMA is not enabled and a write transaction is in operation, DATPORT (described in [Data buffer access port \(eSDHC_DATPORT\)](#)) must not be read.

Chapter 18

Universal Serial Bus Interface

This chapter describes the universal serial bus (USB) interface of the device.

18.1 Introduction

This chapter describes the universal serial bus (USB) interface of the device.

The USB interface implements many industry standards. However, it is beyond the scope of this document to document the intricacies of these standards. Instead, it is left to the reader to refer to the governing specifications.

The following documents are available from the USB Implementers Forum web page at <http://www.usb.org/developers/docs/>.

- *Universal Serial Bus Revision 2.0 Specification*

The following documents are available from the Intel USB Specifications web page at <http://www.intel.com/technology/usb/spec.htm>.

- *Enhanced Host Controller Interface (EHCI) Specification for Universal Serial Bus, Revision 1.0*

The following documents are available from the ULPI web page at <http://www.ulpi.org>.

- *UTMI+ Specification, Revision 1.0*
- *UTMI+ Low Pin Interface (ULPI) Specification, Revision 1.0*

The device implements a dual-role (DR) USB module. This module may be connected to an external port. Collectively the module and external port are called the USB interface. The USB interface is shown in the figure below.

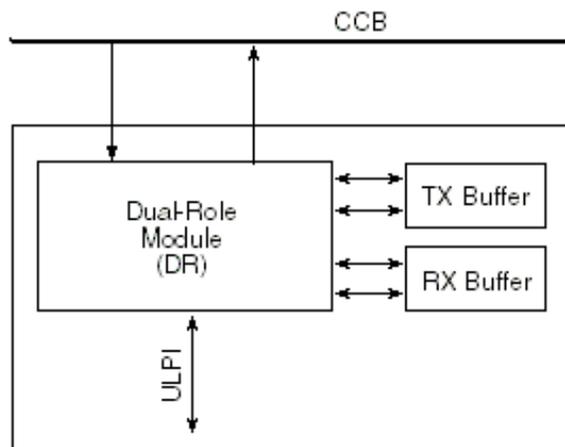


Figure 18-1. USB interface block diagram

18.1.1 Overview

The USB DR module is a USB 2.0-compliant serial interface engine for implementing a USB interface.

The registers and data structures for the module are based on the *Enhanced Host Controller Interface Specification for Universal Serial Bus* (EHCI) from Intel Corporation. The DR module can act as a device or host controller.

The DR module supports the required signaling for and UTMI+ low pin interface (ULPI) transceivers (PHYs). The , and the PHY interfacing to the ULPI is an external PHY.

The module contains a chaining DMA (direct memory access) engine that reduces the interrupt load on the application processor and reduces the total system bus bandwidth that must be dedicated to servicing the USB interface requirements.

18.1.2 Features

The USB DR module includes the following features:

- Complies with USB specification rev 2.0
- Supports operation as a standalone USB host controller
 - Supports enhanced host controller interface (EHCI)
- Supports high-speed (480 Mbps), full-speed (12 Mbps), and low-speed (1.5 Mbps) operation. Low speed is only supported in host mode.
- Supports external PHY with ULPI (UTMI+ low pin interface)
- Supports operation as a standalone USB device

- Supports one upstream facing port
- Supports six bidirectional USB endpoints
- Host and device support

18.1.3 Modes of operation

The USB DR module has two basic operating modes: host and device .

NOTE

Only high-speed and full-speed operations are supported in device mode.

18.2 USB external signals

This section contains detailed descriptions of all the USB controller signals.

18.2.1 ULPI interface

The ULPI (UTMI+ low pin interface) is a reduced pin-count (12 signals) extension of the UTMI+ specification.

Pin count is reduced by converting relatively static signals to register bits, and providing a bidirectional, generic data bus that carries USB and register data. This interface minimizes pin count requirements for external PHYs. The table below describes the signals for the ULPI interface.

Table 18-1. ULPI signal descriptions

Signal	I/O	Description	
USB_DIR	I	Direction. USB_DIR controls the direction of the data bus. When the PHY has data to transfer to USB port, it drives USB_DIR high to take ownership of the bus. When the PHY has no data to transfer it drives USB_DIR low and monitors the bus for link activity. The PHY pulls USB_DIR high whenever the interface cannot accept data from the link.	
		State Meaning	Asserted-PHY has data to transfer to the link. Negated-PHY has no data to transfer.
		Timing	Synchronous to PHY_CLK.

Table continues on the next page...

Table 18-1. ULPI signal descriptions (continued)

Signal	I/O	Description
USB_NXT	I	Next data. The PHY asserts USB_NXT to throttle the data. When USB port is sending data to the PHY, USB_NXT indicates when the current byte has been accepted by the PHY. The USB port places the next byte on the data bus in the following clock cycle. When the PHY is sending data to USB port, USB_NXT indicates when a new byte is available for USB port to consume.
		State Meaning Asserted-PHY is ready to transfer byte. Negated-PHY is not ready.
		Timing Synchronous to PHY_CLK.
USB_STP	O	Stop. USB_STP indicates the end of a transfer on the bus.
		State Meaning Asserted-USB asserts this signal for 1 clock cycle to stop the data stream currently on the bus. If USB port is sending data to the PHY, USB_STP indicates the last byte of data was previously on the bus. If the PHY is sending data to USB port, USB_STP forces the PHY to end its transfer, negate USB_DIR and relinquish control of the data bus to the USB port. Negated-Indicates normal operation.
		Timing Synchronous to PHY_CLK.
USB_PWRFAULT	I	Power fault. USB_PWRFAULT indicates whether a power fault occurred on the USB port Vbus.
		State Meaning Asserted-Indicates that a Vbus fault occurred. Applications that support power switching must shut down Vbus power. Negated-Indicates normal operation.
		Timing Synchronous to PHY_CLK.
USB_PCTL0	O	Port control 0. USB_PCTL0 controls the port status indicator LED 0 when in host mode.
		State Meaning Asserted-LED on. Negated-LED off.
		Timing Synchronous to PHY_CLK.
USB_PCTL1	O	Port control 1. USB_PCTL1 controls the port status indicator LED 1 when in host mode.
		State Meaning Asserted-LED on. Negated-LED off.
		Timing Synchronous to PHY_CLK.
USB_D[7:0]	I/O	Data bit <i>n</i> . USB_D _{<i>n</i>} is bit <i>n</i> of the 8-bit (USB_D7-USB_D0), uni-directional data bus used to carry USB register, and interrupt data between the PHY and the USB controller.
		State Meaning Asserted-Data bit <i>n</i> is 1. Negated-Data bit <i>n</i> is 0.
		Timing Synchronous to PHY_CLK.

18.2.2 PHY clocks

The USB_CLK input provides the clocking signal for the ULPI PHY interface .

The clock is 60 MHz. Detailed clock specifications are given in the appropriate hardware specifications document.

NOTE

A write to registers in the USB controller memory map may cause the system to hang if `PORTSC[PHCD]=0` when no USB PHY clock is applied.

18.3 USB memory map/register definition

This section provides the memory map and detailed descriptions of all USB interface registers.

The following sections provide details about the registers in the USB memory map.

NOTE

Memory may be viewed from either a big-endian or little-endian byte ordering perspective depending on the processor configuration. In big-endian mode, the most-significant byte of word 0 is located at address 0 and the least-significant byte of word 0 is located at address 3. In little-endian mode, the least-significant byte of word 0 is located at address 0 and the most-significant byte of word 0 is located at address 3. Within registers, bits are numbered within a word starting with bit 31 as the most-significant bit. By convention USB registers use little-endian byte ordering. In the USB DR module, these are the registers from offsets 0x00 to 0x1FF. The registers associated with the internal system interface (0x400 and above) use big-endian byte ordering.

USB memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_2000	Identification register (USB_ID)	32	R	E241_AF05h	18.3.1/1561
2_2100	Capability register length (USB_CAPLENGTH)	8	R	40h	18.3.2/1562
2_2102	Host controller interface version number (USB_HCIVERSION)	16	R	0100h	18.3.3/1562
2_2104	Host controller structural parameters (USB_HCSPARAMS)	32	R	0111_0011h	18.3.4/1563
2_2108	Host controller capability parameters (USB_HCCPARAMS)	32	R	0000_0006h	18.3.5/1565
2_2120	Device controller interface version number (USB_DCIVERSION)	16	R	0001h	18.3.6/1567

Table continues on the next page...

USB memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_2124	Device controller capability parameters (USB_DCCPARAMS)	32	R	0000_0186h	18.3.7/1568
2_2140	USB command (USB_USBCMD)	32	R/W	0008_0B00h	18.3.8/1568
2_2144	USB status (USB_USBSTS)	32	R/W	0000_0000h	18.3.9/1572
2_2148	USB interrupt enable (USB_USBINTR)	32	R/W	0000_0000h	18.3.10/1576
2_214C	USB frame index (USB_FRINDEX)	32	R/W	0000_0000h	18.3.11/1578
2_2154	Periodic frame list base address [host mode] (USB_PERIODICLISTBASE)	32	R/W	0000_0000h	18.3.12/1579
2_2154	USB device address [device mode] (USB_DEVICEADDR)	32	R/W	0000_0000h	18.3.13/1580
2_2158	Next asynchronous list addr [host mode] (USB_ASYNCLISTADDR)	32	R/W	0000_0000h	18.3.14/1580
2_2158	Address at endpoint list [device mode] (USB_ENDPOINTLISTADDR)	32	R/W	0000_0000h	18.3.15/1581
2_2160	Master interface data burst size (USB_BURSTSIZE)	32	R/W	0000_1010h	18.3.16/1582
2_2164	Transmit FIFO tuning controls (USB_TXFILLTUNING)	32	R/W	0002_0000h	18.3.17/1582
2_2170	ULPI register access (USB_ULPI_VIEWPORT)	32	R/W	0000_0000h	18.3.18/1584
2_2180	Configured flag register (USB_CONFIGFLAG)	32	R	0000_0001h	18.3.19/1586
2_2184	Port status/control (USB_PORTSC)	32	R/W	9C00_0000h	18.3.20/1587
2_21A8	USB device mode (USB_USBMODE)	32	R/W	0000_0000h	18.3.21/1594
2_21AC	Endpoint setup status (USB_ENDPTSETUPSTAT)	32	w1c	0000_0000h	18.3.22/1595
2_21B0	Endpoint initialization (USB_ENDPOINTPRIME)	32	R/W	0000_0000h	18.3.23/1596
2_21B4	Endpoint flush (USB_ENDPTFLUSH)	32	R/W	0000_0000h	18.3.24/1597
2_21B8	Endpoint status (USB_ENDPTSTATUS)	32	R	0000_0000h	18.3.25/1597
2_21BC	Endpoint complete (USB_ENDPTCOMPLETE)	32	w1c	0000_0000h	18.3.26/1598
2_21C0	Endpoint control 0 (USB_ENDPTCTRL0)	32	R/W	0080_0080h	18.3.27/1600
2_21C4	Endpoint control n (USB_ENDPTCTRL1)	32	R/W	0000_0000h	18.3.28/1602
2_21C8	Endpoint control n (USB_ENDPTCTRL2)	32	R/W	0000_0000h	18.3.28/1602

Table continues on the next page...

USB memory map (continued)

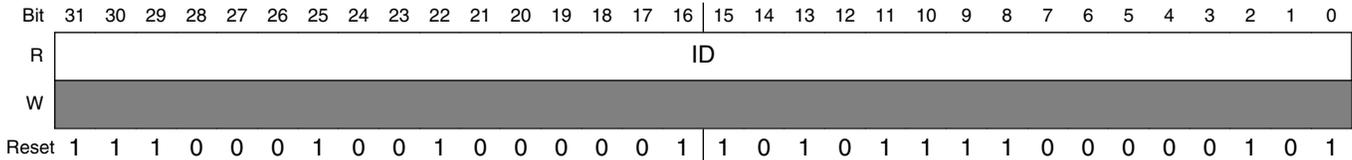
Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
2_21CC	Endpoint control n (USB_ENDPTCTRL3)	32	R/W	0000_0000h	18.3.28/1602
2_21D0	Endpoint control n (USB_ENDPTCTRL4)	32	R/W	0000_0000h	18.3.28/1602
2_21D4	Endpoint control n (USB_ENDPTCTRL5)	32	R/W	0000_0000h	18.3.28/1602
2_2400	Snoop n (USB_SNOOP1)	32	R/W	0000_0000h	18.3.29/1604
2_2404	Snoop n (USB_SNOOP2)	32	R/W	0000_0000h	18.3.29/1604
2_2408	Age count threshold (USB_AGE_CNT_THRESH)	32	R/W	0000_0000h	18.3.30/1605
2_240C	Priority control (USB_PRI_CTRL)	32	R/W	0000_0000h	18.3.31/1607
2_2410	System interface control (USB_SI_CTRL)	32	R/W	0000_0000h	18.3.32/1608
2_2500	Control (USB_CONTROL)	32	R/W	0000_0000h	18.3.33/1609

18.3.1 Identification register (USB_ID)

The identification registers are used to declare the slave interface presence and include a table of the hardware configuration parameters. These registers are not defined by the EHCI specification.

The ID register provides a simple way to determine if the USB DR module is provided in the system. The ID register identifies the USB DR and its revision.

Address: 2_2000h base + 0h offset = 2_2000h



USB_ID field descriptions

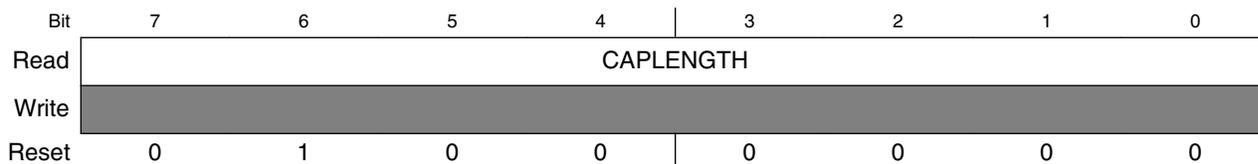
Field	Description
31-0 ID	Configuration number.

18.3.2 Capability register length (USB_CAPLENGTH)

The capability registers specify the software limits, restrictions, and capabilities of the host/device controller implementation. Most of these registers are defined by the EHCI specification. Registers that are not defined by the EHCI specification are noted in their descriptions.

CAPLENGTH is used as an offset to add to the register base address to find the beginning of the operational register space, that is, the location of the USBCMD register.

Address: 2_2000h base + 100h offset = 2_2100h



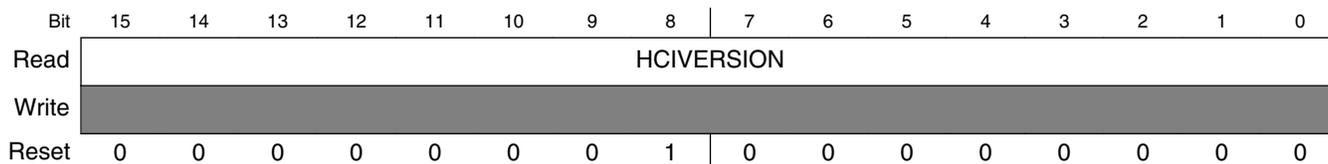
USB_CAPLENGTH field descriptions

Field	Description
7-0 CAPLENGTH	Indicate which offset to add to the register base address at the beginning of the Operational Register. Value is 0x40.

18.3.3 Host controller interface version number (USB_HCIVERSION)

HCIVERSION contains a BCD encoding of the EHCI revision number supported by this host controller. The most-significant byte of the register represents a major revision and the least-significant byte is the minor revision.

Address: 2_2000h base + 102h offset = 2_2102h



USB_HCIVERSION field descriptions

Field	Description
15-0 HCIVERSION	EHCI revision number. Value is 0x0100 indicating version 1.0.

18.3.4 Host controller structural parameters (USB_HCSPARAMS)

HCSPARAMS contains structural parameters such as the number of downstream ports

Address: 2_2000h base + 104h offset = 2_2104h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Reserved				N_TT				N_PTT				Reserved			PI
W	Reserved				Reserved				Reserved				Reserved			Reserved
Reset	0	0	0	0	0	0	0	1	0	0	0	1	0	0	0	1
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	N_CC				N_PCC				Reserved			PPC	N_PORTS			
W	Reserved				Reserved				Reserved			Reserved	Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	1

USB_HCSPARAMS field descriptions

Field	Description
31–28 -	This field is reserved. Reserved, should be cleared.
27–24 N_TT	Number of transaction translators. This is a non-EHCI field. This field indicates the number of embedded transaction translators associated the module. See Embedded transaction translator function .
23–20 N_PTT	Ports per transaction translator. This is a non-EHCI field. The number of ports assigned to each transaction translator. This is equal to N_PORTS.
19–17 -	This field is reserved. Reserved, should be cleared.
16 PI	Port indicators. Indicates whether the ports support port indicator control. 1 The port status and control registers include a R/W field for controlling the state of the port indicator.
15–12 N_CC	Number of companion controllers associated with the DR controller. Always 0.
11–8 N_PCC	Number ports per CC. This field indicates the number of ports supported per internal companion controller. Always 0.
7–5 -	This field is reserved. Reserved, should be cleared.
4 PPC	Power port control. Indicates whether the host controller supports port power control. ULPI Mode: 0 USBDR will write 0 for DRVVBUS bit of OTG Control register in PHY. 1 USBDR will write 1 for DRVVBUS bit of OTG Control register in PHY. The OTG Control register is defined in ULPI specification. UTMI Mode: In UTMI mode, USBDR will control the DRVVBUS signal defined in UTMI+ specification.

Table continues on the next page...

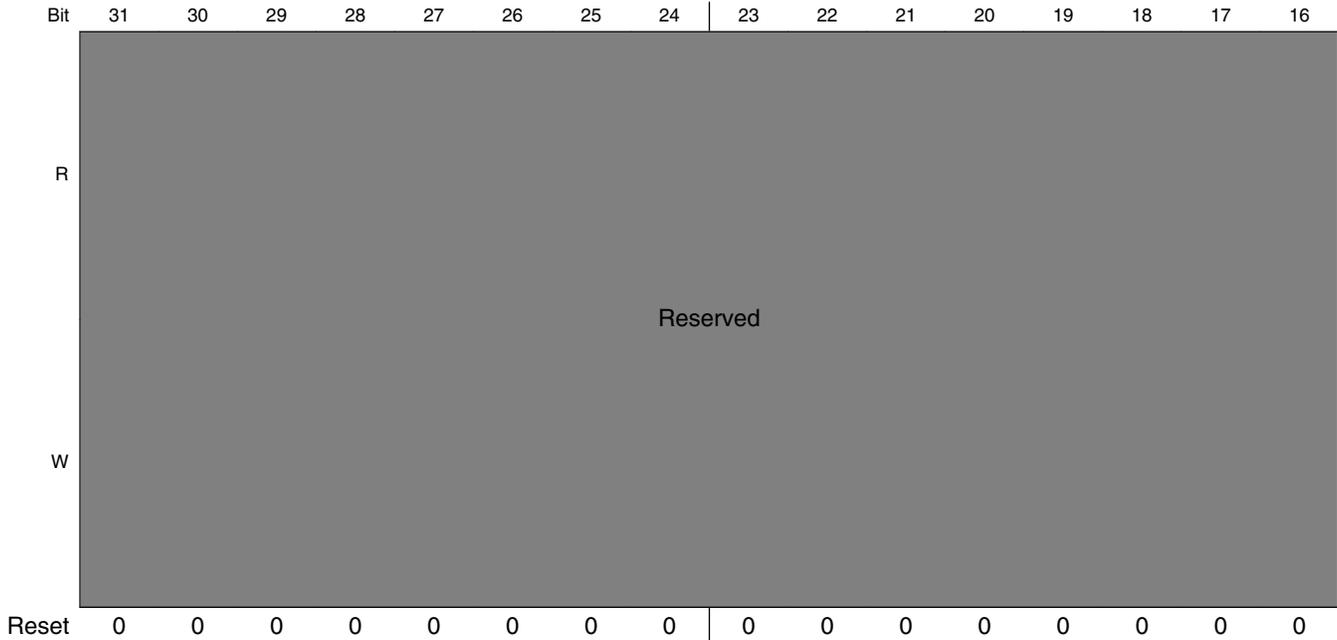
USB_HCSPARAMS field descriptions (continued)

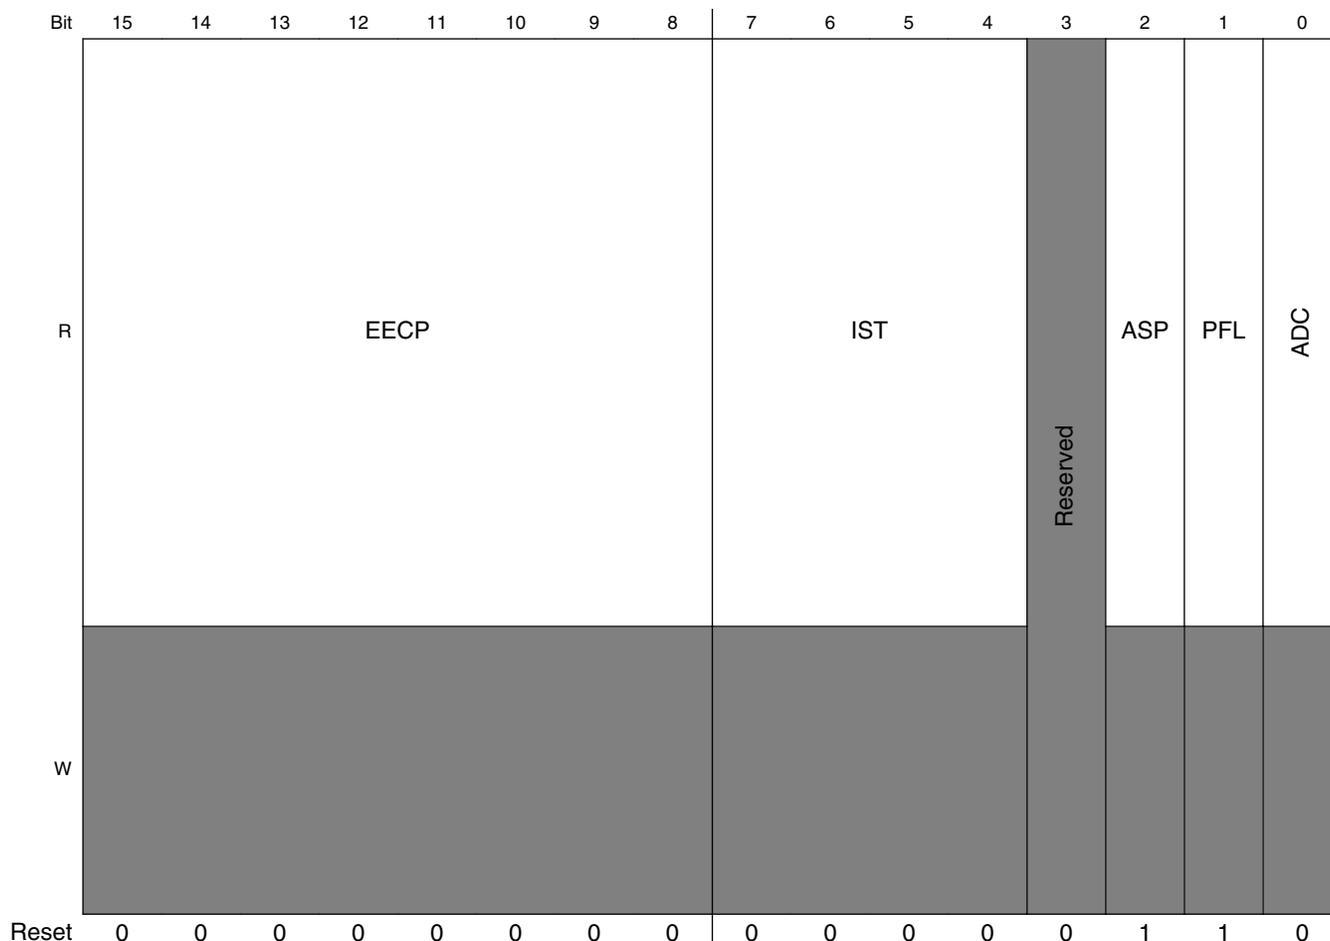
Field	Description
	0 USBDR will drive utmi DRVVBUS signal LOW. 1 USBDR will drive utmi DRVVBUS signal HIGH. The operation of Port Power bit is independent of USBDR mode (host or device / USBMODE[CM]).
3-0 N_PORTS	Number of ports. Number of physical downstream ports implemented for host applications. The value of this field determines how many port registers are addressable in the operational register.

18.3.5 Host controller capability parameters (USB_HCCPARAMS)

HCCPARAMS identifies multiple mode control (time-base bit functionality) addressing capability.

Address: 2_2000h base + 108h offset = 2_2108h





USB_HCCPARAMS field descriptions

Field	Description
31–16 -	This field is reserved. Reserved, should be cleared.
15–8 EECP	EHCI extended capabilities pointer. Indicates the existence of a capabilities list. A value of 0x00 indicates no extended capabilities are implemented. A non-zero value in this register indicates the offset in PCI configuration space of the first EHCI extended capability. The pointer value must be 0x40 or greater if implemented to maintain the consistency of the PCI header defined for this class of device. This field is always 0.
7–4 IST	Isochronous scheduling threshold. Indicates, relative to the current position of the executing host controller, where software can reliably update the isochronous schedule. When bit 7 is zero, the value of the least significant 3 bits indicates the number of microframes a host controller can hold a set of isochronous data structures (one or more) before flushing the state. When bit 7 is a one, then host software assumes the host controller may cache an isochronous data structure for an entire frame. This field is always 0.
3 -	This field is reserved. Reserved, should be cleared.
2 ASP	Asynchronous schedule park capability. Indicates whether the USB DR module supports the park feature for high-speed queue heads in the asynchronous schedule. The feature can be disabled or enabled and set to a specific level by using the <i>asynchronous schedule park mode enable</i> and <i>asynchronous schedule park mode count</i> fields in the USBCMD register.

Table continues on the next page...

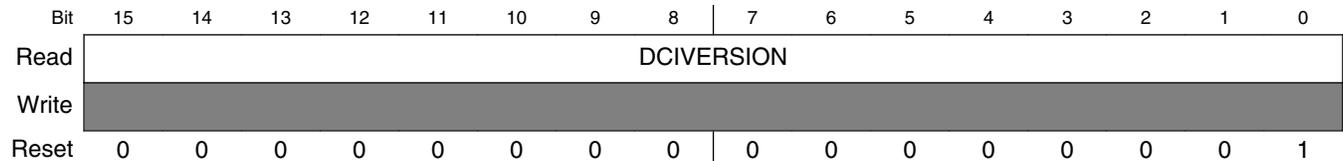
USB_HCCPARAMS field descriptions (continued)

Field	Description
	This field is always 1 (park feature supported).
1 PFL	Programmable frame list flag. Indicates whether system software can specify and use a frame list length less than 1024 elements. Frame list size is configured via the USBCMD register frame list size field. The frame list must always be aligned on a 4-K page boundary. This requirement ensures that the frame list is always physically contiguous. This field is always 1.
0 ADC	64-bit addressing capability. Always 0; 64-bit addressing is not supported. 0 Data structures use 32-bit address memory pointers

18.3.6 Device controller interface version number (USB_DCIVERSION)

This register is not defined in the EHCI specification. DCIVERSION is a two-byte register containing a BCD encoding of the device controller interface. The most-significant byte of the register represents a major revision and the least-significant byte is the minor revision.

Address: 2_2000h base + 120h offset = 2_2120h



USB_DCIVERSION field descriptions

Field	Description
15–0 DCIVERSION	Device interface revision number.

18.3.7 Device controller capability parameters (USB_DCCPARAMS)

This register is not defined in the EHCI specification. This register describes the overall host/device capability of the DR module.

Address: 2_2000h base + 124h offset = 2_2124h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
R	Reserved								HC	DC	Reserved			DEN			
W	Reserved										Reserved			Reserved			
Reset	0	0	0	0	0	0	0	1	1	0	0	0	0	0	1	1	0

USB_DCCPARAMS field descriptions

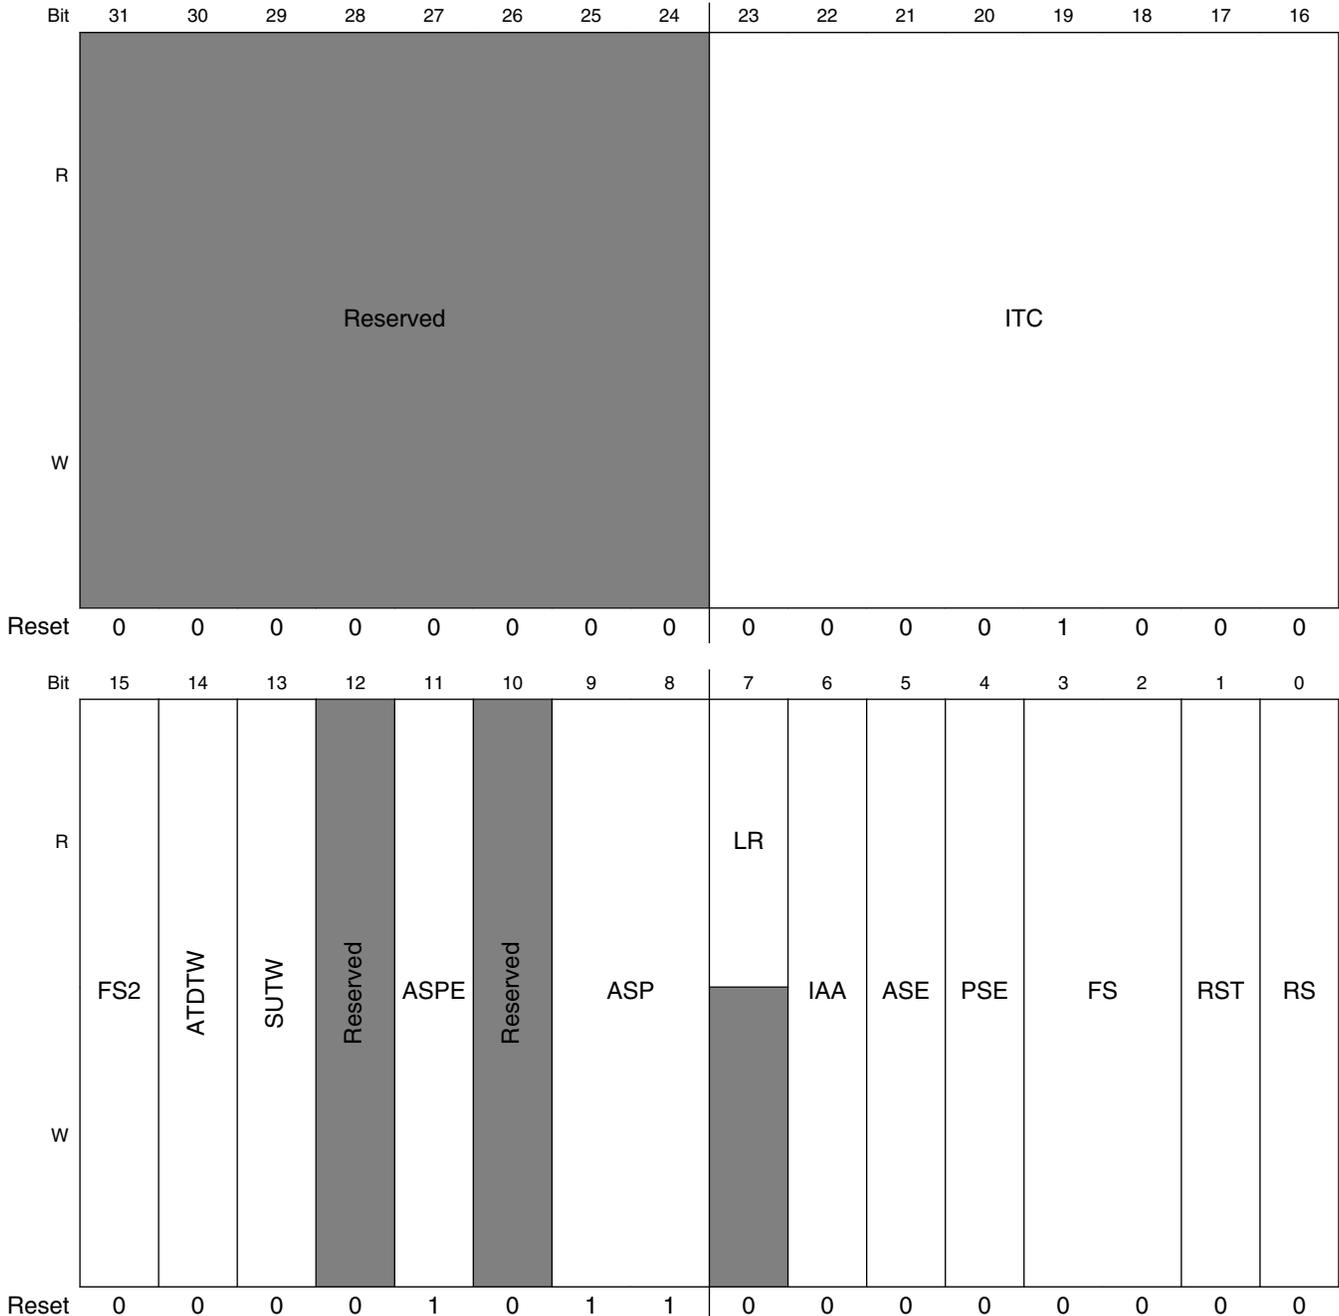
Field	Description
31–9 -	This field is reserved. Reserved, should be cleared.
8 HC	Host capable. Always 1, indicating the USB DR controller can operate as an EHCI compatible USB 2.0 host
7 DC	Device capable. Always 1, indicating the USB DR controller can operate as an USB 2.0 device. 1 Device capability. 0 No device capability (host only).
6–5 -	This field is reserved. Reserved, should be cleared.
4–0 DEN	Device endpoint number. Indicates the number of endpoints built into the device controller. Always 0x 6 .

18.3.8 USB command (USB_USBCMD)

The operational registers are comprised of dynamic control or status registers that may be read-only, read/write, or read/write-1-to-clear. The following sections define the operational registers.

The module executes the command indicated in this register.

Address: 2_2000h base + 140h offset = 2_2140h



USB_USBCMD field descriptions

Field	Description
31–24 -	This field is reserved. Reserved, should be cleared.
23–16 ITC	Interrupt threshold control. The system software uses this field to set the maximum rate at which the USB DR module issues interrupts. ITC contains the maximum interrupt interval measured in microframes. Valid values are shown below. 0x00 Immediate (no threshold)

Table continues on the next page...

USB_USBCMD field descriptions (continued)

Field	Description
	0x01 1 microframe 0x02 2 microframes 0x04 4 microframes 0x08 8microframes 0x10 16 microframes 0x20 32 microframes 0x40 40 microframes
15 FS2	See bits 3-2 below. This is a non-EHCI bit.
14 ATDTW	Add dTD TripWire. This is a non-EHCI bit. Used as a semaphore when a dTD is added to an active (primed) endpoint. This bit is set and cleared by software. This bit shall also be cleared by hardware when its state machine is in hazard region where adding a dTD to a primed endpoint may go unrecognized. More information on the use of this bit is described in Device operational model .
13 SUTW	Setup tripwire. This is a non-EHCI bit. Used as a semaphore when the 8 bytes of setup data read extracted from a QH by the DCD. If the setup lockout mode is off (See USB device mode (USB_USBMODE)) then there exists a hazard when new setup data arrives and the DCD is copying setup from the QH for a previous setup packet. This bit is set and cleared by software and will be cleared by hardware when a hazard exists. More information on the use of this bit is described in Device operational model .
12 -	This field is reserved. Reserved, should be cleared.
11 ASPE	Asynchronous schedule park mode enable. Software uses this bit to enable or disable park mode. 0 Disabled 1 Enabled
10 -	This field is reserved. Reserved, should be cleared.
9–8 ASP	Asynchronous schedule park mode count. It contains a count of the number of successive transactions the host controller is allowed to execute from a high-speed queue head on the Asynchronous schedule before continuing traversal of the Asynchronous schedule. Valid values are 0x1H to 0x3H. Software must not write a zero to this field when ASPE is set as this results in undefined behavior.
7 LR	Light host/device controller reset (OPTIONAL). Not implemented. Always 0.
6 IAA	<p>Interrupt on async advance doorbell. Used as a doorbell by software to tell the USB DR controller to issue an interrupt the next time it advances asynchronous schedule. Software must write a 1 to this bit to ring the doorbell.</p> <p>When the controller has evicted all appropriate cached schedule states, it sets USBSTS[AAI]. If USBINTR[AAE] is set, the host controller asserts an interrupt at the next interrupt threshold.</p> <p>The controller clears this bit after it has set USBSTS[AAI]. Software should not set this bit when the asynchronous schedule is inactive. Doing so yields undefined results.</p> <p>This bit is only used in host mode. Setting this bit when the USB DR module is in device mode is selected results in undefined results.</p>
5 ASE	<p>Asynchronous schedule enable. Controls whether the controller skips processing the asynchronous schedule. Only used in host mode.</p> <p>0 Do not process the asynchronous schedule 1 Use the ASYNCLISTADDR register to access the asynchronous schedule.</p>

Table continues on the next page...

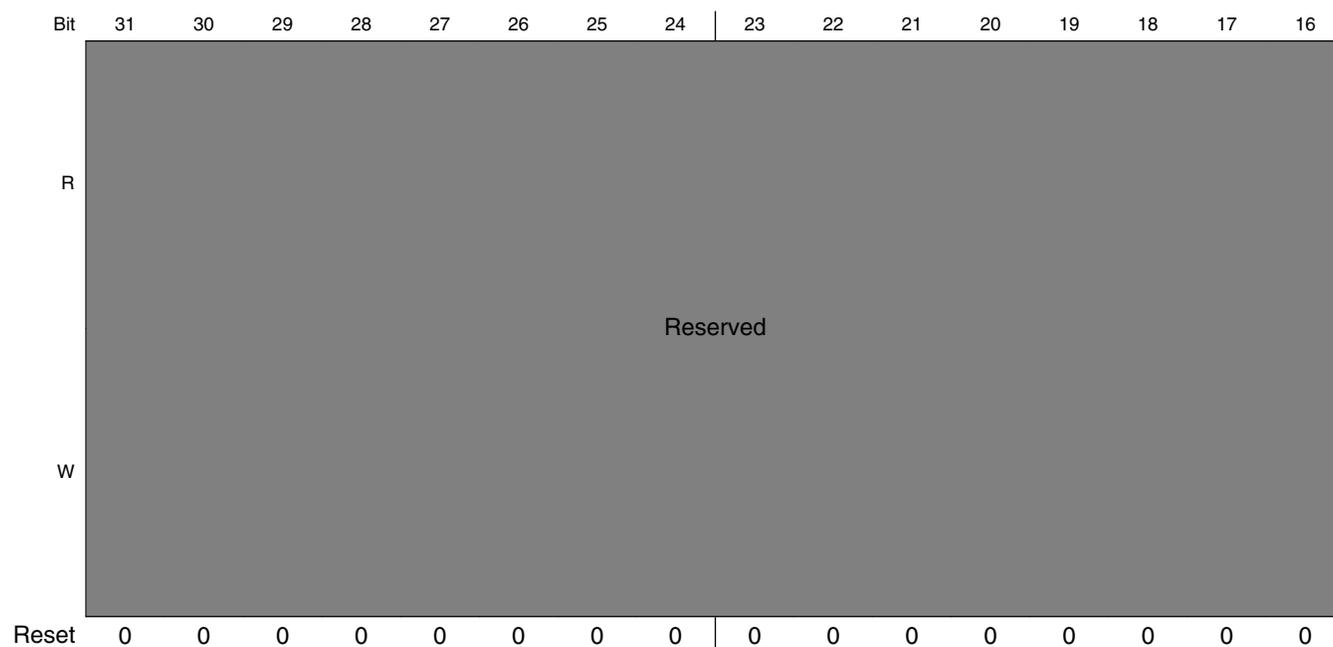
USB_USBCMD field descriptions (continued)

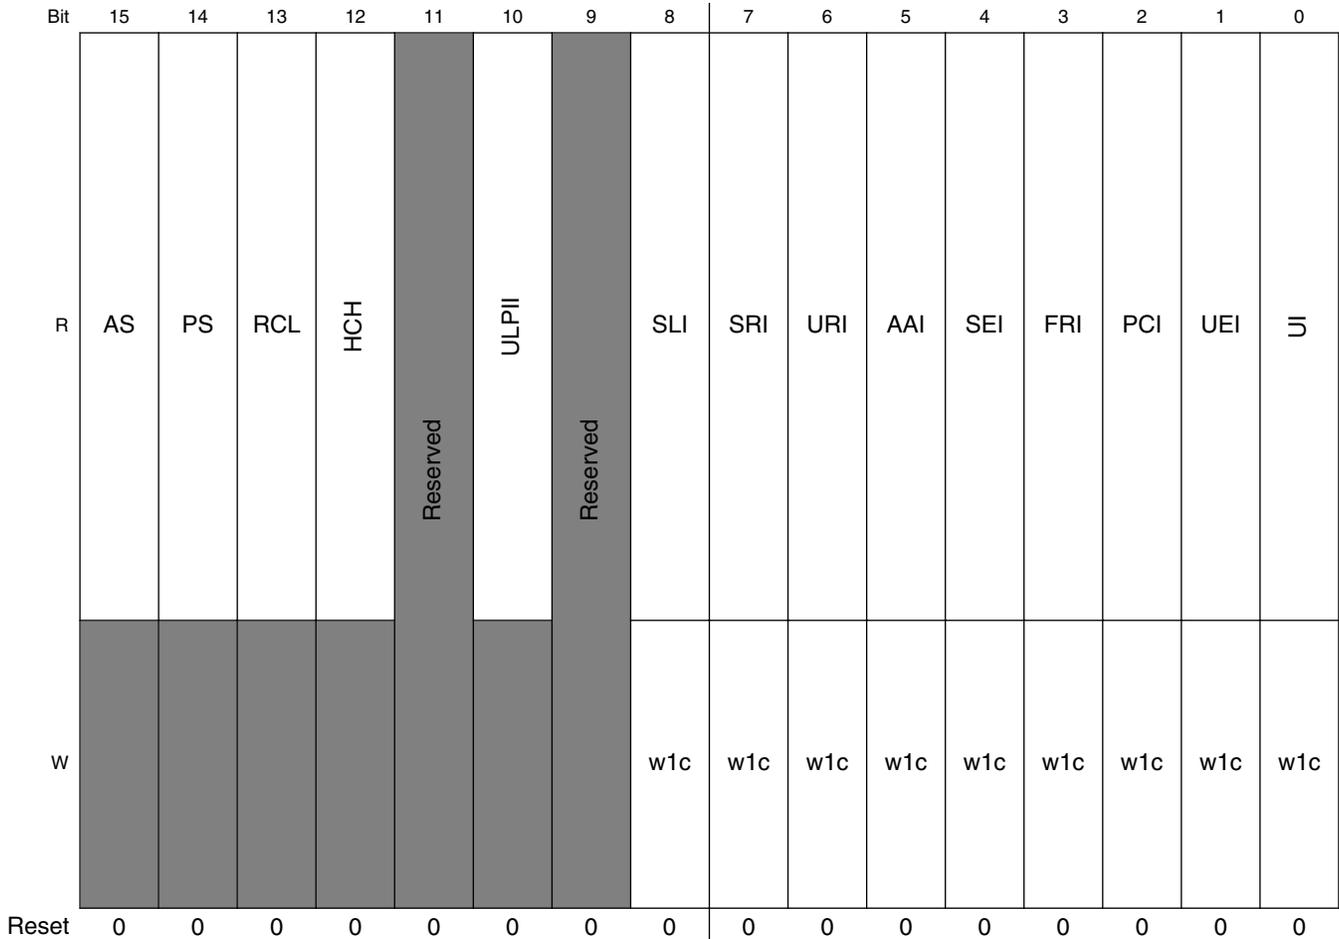
Field	Description
4 PSE	<p>Periodic schedule enable. Controls whether the controller skips processing the periodic schedule. Only used in host mode.</p> <p>0 Do not process the periodic schedule. 1 Use the PERIODICLISTBASE register to access the periodic schedule.</p>
3–2 FS	<p>Frame list size. Together with bit 15 these bits make the FS[2:0] field. This field is read/write only if programmable frame list flag in the HCCPARAMS registers is set to 1. This field specifies the size of the frame list that controls which bits in FRINDEX should be used for the frame list current index. Only used in host mode. Note that values below 256 elements are not defined in the EHCI specification.</p> <p>000 1024 elements (4096 bytes) 001 512 elements (2048 bytes) 010 256 elements (1024 bytes) 011 128 elements (512 bytes) 100 64 elements (256 bytes) 101 32 elements (128 bytes) 110 16 elements (64 bytes) 111 8 elements (32 bytes)</p>
1 RST	<p>Controller reset. Software uses this bit to reset the controller. This bit is cleared by the controller when the reset process is complete. Software cannot terminate the reset process early by writing a zero to this register.</p> <p>Host mode:</p> <ul style="list-style-type: none"> When software sets this bit, the host controller resets its internal pipelines, timers, counters, state machines and so on to their initial value. Any transaction currently in progress on USB is immediately terminated. A USB reset is not driven on downstream ports. Software should not set this bit when USBSTS[HCH] is a zero. Attempting to reset an actively running host controller results in undefined behavior. <p>Device mode:</p> <ul style="list-style-type: none"> When software sets this bit, the USB DR controller resets its internal pipelines, timers, counters, state machines and so on to their initial value. Any transaction currently in progress on USB is immediately terminated. Writing a one to this bit in device mode is not recommended.
0 RS	<p>Run/Stop.</p> <p>Host mode:</p> <ul style="list-style-type: none"> When this bit is set, the controller proceeds with the execution of the schedule. The controller continues execution as long as this bit is set. When this bit is set to 0, the host controller completes the current transaction on the USB and then halts. The USBSTS[HCH] bit indicates when the USB DR controller has finished the transaction and has entered the stopped state. Software should not write a one to this field unless the controller is in the halted state (that is, USBSTS[HCH] is a one). <p>Device mode:</p> <ul style="list-style-type: none"> Setting this bit causes the USB DR controller to enable a pull-up on D+ and initiate an attach event. This control bit is not directly connected to the pull-up enable, as the pull-up is disabled upon transitioning into high-speed mode. Software should use this bit to prevent an attach event before the controller has been properly initialized. Clearing this bit causes a detach event. <p>0 Stop 1 Run</p>

18.3.9 USB status (USB_USBSTS)

The USB status register indicates various states of the USB DR module and any pending interrupts. This register does not indicate status resulting from a transaction on the serial bus. Software clears certain bits in this register by writing a 1 to them (indicated by a w1c in the bit's W cell).

Address: 2_2000h base + 144h offset = 2_2144h





USB_USBSTS field descriptions

Field	Description
31–16 -	This field is reserved. Reserved, should be cleared.
15 AS	Asynchronous schedule status. Reports the current real status of the asynchronous schedule. The USB DR controller is not required to immediately disable or enable the asynchronous schedule when software transitions USB_CMD[ASE]. When this bit and USB_CMD[ASE] have the same value, the asynchronous schedule is either enabled (1) or disabled (0). Only used in host mode. 0 Disabled 1 Enabled
14 PS	Periodic schedule status. Reports the current real status of the periodic schedule. The USB DR controller is not required to immediately disable or enable the periodic schedule when software transitions USB_CMD[PSE]. When this bit and USB_CMD[PSE] have the same value, the periodic schedule is either enabled (1) or disabled (0). Only used in host mode. 0 Disabled 1 Enabled
13 RCL	Reclamation. Used to detect an empty asynchronous schedule. Only used by the host mode. 0 Non-empty asynchronous schedule 1 Empty asynchronous schedule

Table continues on the next page...

USB_USBSTS field descriptions (continued)

Field	Description
12 HCH	<p>HC halted. This bit is a zero whenever USBCMD[RS] is a one. The USB DR controller sets this bit to one after it has stopped executing because of USBCMD[RS] being cleared, either by software or by the host controller hardware (for example, internal error). Only used in host mode.</p> <p>0 Running 1 Halted</p>
11 -	<p>This field is reserved. Reserved, should be cleared.</p>
10 ULPII	<p>ULPI interrupt. An event completion to the viewport register sets this bit. If the ULPI enables the USBINTR[ULPIE] to be set, the USB interrupt (UI) occurs.</p>
9 -	<p>This field is reserved. Reserved, should be cleared.</p>
8 SLI	<p>DC Suspend. This is a non-EHCI bit. When a device controller enters a suspend state from an active state, this bit is set. The device controller clears the bit upon exiting from a suspend state. Only used by the device controller.</p> <p>0 Active 1 Suspended</p>
7 SRI	<p>Host mode:</p> <ul style="list-style-type: none"> This is a non-EHCI status bit. In host mode, this bit is set every 125 μs, provided the PHY clock is present and running (for example, the port is NOT suspended), and can be used by the host controller driver as a time base. <p>Device mode:</p> <ul style="list-style-type: none"> SOF received. When the USB DR controller detects a Start Of (Micro)Frame, this bit is set. When a SOF is extremely late, the DR controller automatically sets this bit to indicate that an SOF was expected. Therefore, this bit is set roughly every 1 msec in device FS mode and every 125 μs in HS mode and is synchronized to the actual SOF that is received. Because the controller is initialized to FS before connect, this bit is set at an interval of 1 msec during the prelude to the connect and chirp. <p>Software writes a 1 to this bit to clear it.</p>
6 URI	<p>USB reset received. This is a non-EHCI bit. When the USB DR controller detects a USB reset and enters the default state, this bit is set. Software can write a 1 to this bit to clear the USB reset received status bit. Only used by the device mode.</p> <p>0 No reset received 1 Reset received</p>
5 AAI	<p>Interrupt on async advance. System software can force the controller to issue an interrupt the next time the USB DR controller advances the asynchronous schedule by writing a one to USBCMD[IAA]. This status bit indicates the assertion of that interrupt source. Only used by the host mode.</p> <p>0 No async advance interrupt 1 Async advance interrupt</p>
4 SEI	<p>System error. This bit is set whenever an error is detected on the system bus. If USBINTR[SEE] is set, an interrupt is generated. The interrupt and status bits remain asserted until cleared by writing a 1 to this bit. Additionally, when in host mode, USBCMD[RS] is cleared, effectively disabling the USB DR controller. For the USB DR controller in device mode, an interrupt is generated, but no other action is taken.</p> <p>0 Normal operation 1 Error</p>

Table continues on the next page...

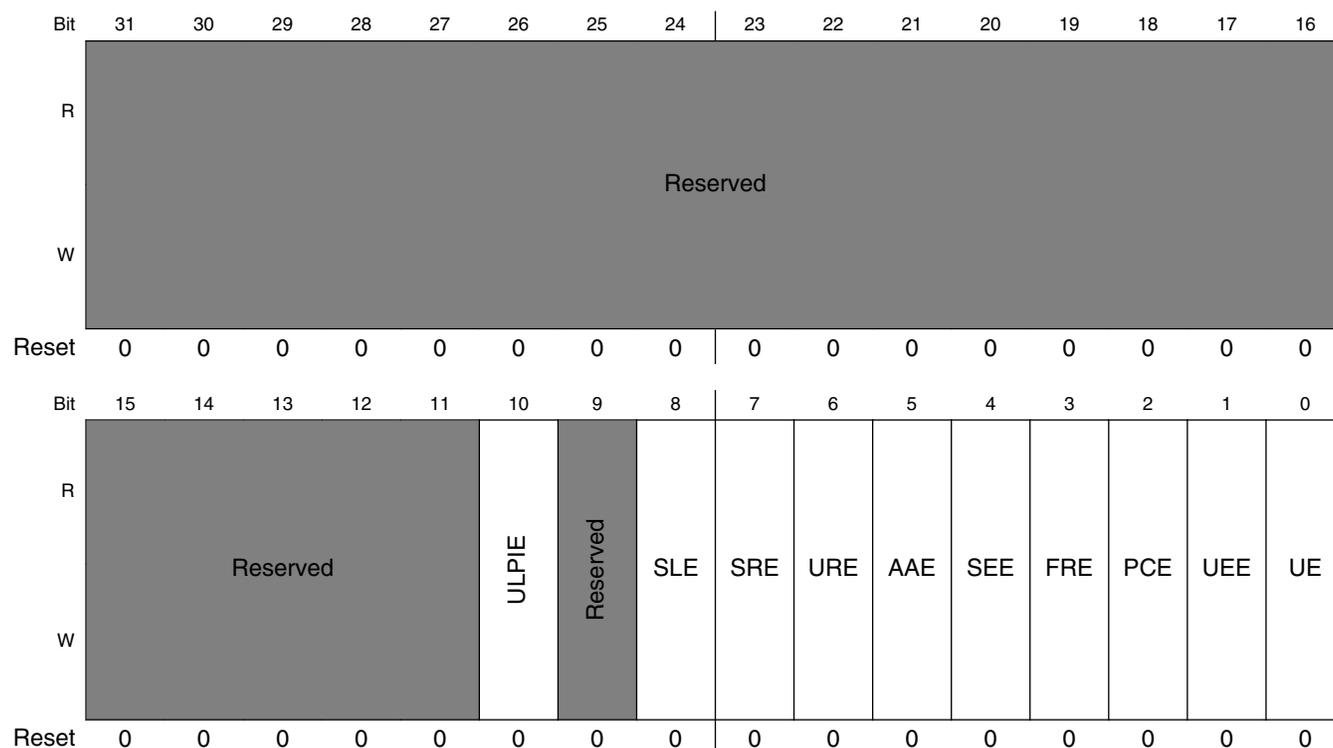
USB_USBSTS field descriptions (continued)

Field	Description
3 FRI	Frame list rollover. The controller sets this bit to a one when the frame list index rolls over from its maximum value to zero. The exact value at which the rollover occurs depends on the frame list size. For example, if the frame list size (as programmed in USBCMD[FS]) is 1024, FRINDEX rolls over every time FRINDEX [13] toggles. Similarly, if the size is 512, the USB DR controller sets this bit to a one every time FRINDEX [12] toggles. Only used by the host mode.
2 PCI	Host mode: <ul style="list-style-type: none"> • Port change detect. The controller sets this bit when a connect status occurs on any port, a port enable/disable change occurs, an over current change occurs, or PORTSC[FPR] is set as the result of a J-K transition on the suspended port. Device mode: <ul style="list-style-type: none"> • The USB DR controller sets this bit when it enters the full or high-speed operational state. When the port controller exits the full or high-speed operation states due to reset or suspend events, the notification mechanisms are USBSTS[URI] and USBSTS[SLI], respectively. This bit is not EHCI compatible.
1 UEI	USB error interrupt (USBERRINT). When completion of a USB transaction results in an error condition, this bit is set by the controller. This bit is set along with the UI, if the TD on which the error interrupt occurred also had its interrupt on complete (IOC) bit set. See Section 4.15.1 in EHCI for a complete list of host error interrupt conditions. Also see Table 18-104 in this chapter for more information on device error matrix. For the USB DR controller in device mode, only resume signaling is detected, all others are ignored. 0 No error 1 Error detected
0 UI	USB interrupt (USBINT). This bit is set by the controller when the cause of an interrupt is a completion of a USB transaction where the transfer descriptor (TD) has an interrupt on complete (IOC) bit set. This bit is also set by the controller when a short packet is detected. A short packet is when the actual number of bytes received was less than the expected number of bytes.

18.3.10 USB interrupt enable (USB_USBINTR)

The interrupts to software are enabled with the USB interrupt enable register. An interrupt is generated when a bit is set and the corresponding interrupt is active. The USB status register (USBSTS) still shows interrupt sources even if they are disabled by the USBINTR register, allowing polling of interrupt events by the software.

Address: 2_2000h base + 148h offset = 2_2148h



USB_USBINTR field descriptions

Field	Description
31–11 -	This field is reserved. Reserved, should be cleared.
10 ULPIE	ULPI interrupt enable. An event completion to the viewport register sets the USBSTS[ULPII]. If the ULPI enables ULPIE bit to be set, then the USBINT (USBSTS[U]) occurs. 0 Disable 1 Enable
9 -	This field is reserved. Reserved, should be cleared.
8 SLE	Sleep enable. This is a non-EHCI bit. When this bit is a one, and USBSTS[SLI] transitions, the USB DR controller issues an interrupt. The interrupt is acknowledged by software writing a one to USBSTS[SLI]. Only used in device mode.

Table continues on the next page...

USB_USBINTR field descriptions (continued)

Field	Description
	0 Disable 1 Enable
7 SRE	SOF received enable. This is a non-EHCI bit. When this bit is a one, and USBSTS[SRI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[SRI]. 0 Disable 1 Enable
6 URE	USB reset enable. This is a non-EHCI bit. When this bit is a one, and USBSTS[URI] is a one, the device controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[URI] bit. Only used in device mode. 0 Disable 1 Enable
5 AAE	Interrupt on async advance enable. When this bit is a one, and USBSTS[AAI] is a one, the controller issues an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[AAI]. Only used in host mode. 0 Disable 1 Enable
4 SEE	System error enable. When this bit is a one, and USBSTS[SEI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[SEI]. 0 Disable 1 Enable
3 FRE	Frame list rollover enable. When this bit is a one, and USBSTS[FRI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[FRI]. Only used by the host mode. 0 Disable 1 Enable
2 PCE	Port change detect enable. When this bit is a one, and USBSTS[PCI] is a one, the controller issues an interrupt. The interrupt is acknowledged by software clearing USBSTS[PCI]. 0 Disable 1 Enable
1 UEE	USB error interrupt enable. When this bit is a one, and USBSTS[UEI] is a one, the controller issues an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[UEI]. 0 Disable 1 Enable
0 UE	USB interrupt enable. When this bit is a one, and USBSTS[UI] is a one, the DR controller issues an interrupt at the next interrupt threshold. The interrupt is acknowledged by software clearing USBSTS[UI]. 0 Disable 1 Enable

18.3.11 USB frame index (USB_FRINDEX)

In host mode, the frame index register is used by the controller to index the periodic frame list. The register updates every 125 microseconds (once each microframe). Bits N-3 are used to select a particular entry in the periodic frame list during periodic schedule execution. The number of bits used for the index depends on the size of the frame list as set by system software in USBCMD[FS].

This register must be written as a DWord. Byte writes produce undefined results. This register cannot be written unless the USB DR controller is in the Halted state as indicated by the USBSTS[HCH]. A write to this register while USBCMD[RS] is set produces undefined results. Writes to this register also affect the SOF value.

In device mode, this register is read-only and the USB DR controller updates the FRINDEX[13-3] register from the frame number indicated by the SOF marker. Whenever a SOF is received by the USB bus, FRINDEX[13-3] is checked against the SOF marker. If FRINDEX[13-3] is different from the SOF marker, FRINDEX[13-3] is set to the SOF value and FRINDEX[2-0] is cleared (that is, SOF for 1 msec frame). If FRINDEX[13-3] is equal to the SOF value, FRINDEX[2-0] is incremented (that is, SOF for 125- μ sec microframe).

The table below illustrates values of N based on the value of the Frame List Size in the USBCMD register, when used in host mode.

Table 18-13. FRINDEX N Values

USBCMD[FS]	Frame List Size	FRINDEX N value
000	1024 elements (4096 bytes)	12
001	512 elements (2048 bytes)	11
010	256 elements (1024 bytes)	10
011	128 elements (512 bytes)	9
100	64 elements (256 bytes)	8
101	32 elements (128 bytes)	7
110	16 elements (64 bytes)	6
111	8 elements (32 bytes)	5

Address: 2_2000h base + 14Ch offset = 2_214Ch

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved																FRINDEX															
W	Reserved																FRINDEX															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

USB_FRINDEX field descriptions

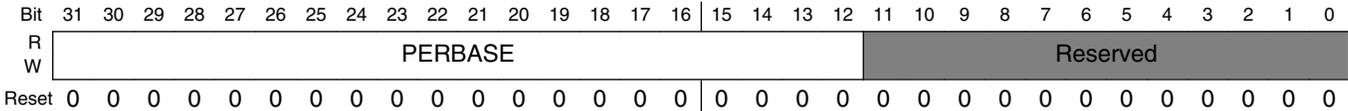
Field	Description
31–14 -	This field is reserved. Reserved, should be cleared.
13–0 FRINDEX	Frame index. The value in this register increments at the end of each time frame (for example, microframe). Bits N-3 are used for the Frame List current index. This means that each location of the frame list is accessed 8 times (frames or microframes) before moving to the next index. In device mode, the value is the current frame number of the last frame transmitted. It is not used as an index. In either mode, bits 2-0 indicate the current microframe.

18.3.12 Periodic frame list base address [host mode] (USB_PERIODICLISTBASE)

This register contains the beginning address of the Periodic frame list in the system memory. The host controller driver loads this register prior to starting the schedule execution by the controller. The memory structure referenced by this physical memory pointer is assumed to be 4-Kbyte aligned. The contents of this register are combined with the frame index register (FRINDEX) to enable the controller to step through the Periodic Frame List in sequence.

Note that this register is shared between the host and device mode functions. In host mode, it is the PERIODICLISTBASE register; in device mode, it is the DEVICEADDR register. See [USB device address \[device mode\] \(USB_DEVICEADDR\)](#) , for more information.

Address: 2_2000h base + 154h offset = 2_2154h



USB_PERIODICLISTBASE field descriptions

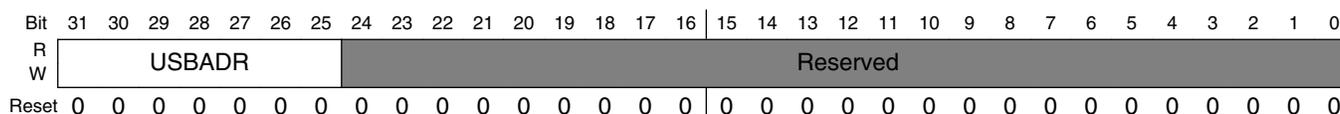
Field	Description
31–12 PERBASE	Base address. Correspond to memory address signal [31:12]. Only used in the host mode.
11–0 -	This field is reserved. Reserved, should be cleared.

18.3.13 USB device address [device mode] (USB_DEVICEADDR)

The device address register is not defined in the EHCI specification. In device mode, the upper seven bits of this register represent the device address. After any controller reset or a USB reset, the device address is set to the default address (0). The default address will match all incoming addresses. Software shall reprogram the address after receiving a SET_ADDRESS descriptor.

Note that this register is shared between the host and device mode functions. In device mode, it is the DEVICEADDR register; in host mode, it is the PERIODICLISTBASE register. See [Periodic frame list base address \[host mode\] \(USB_PERIODICLISTBASE\)](#) , for more information.

Address: 2_2000h base + 154h offset = 2_2154h



USB_DEVICEADDR field descriptions

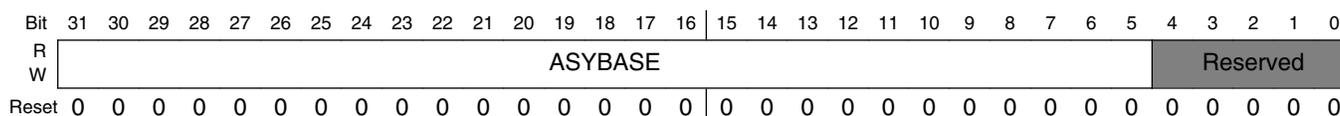
Field	Description
31–25 USBADR	Device address. This field corresponds to the USB device address.
24–0 -	This field is reserved. Reserved, should be cleared.

18.3.14 Next asynchronous list addr [host mode] (USB_ASYNC_LISTADDR)

This 32-bit register contains the address of the next asynchronous queue head to be executed by the host. Bits 4-0 of this register cannot be modified by the system software and always return zeros when read.

Note that this register is shared between the host and device mode functions. In host mode, it is the ASYNCLISTADDR register; in device mode, it is the ENDPOINTLISTADDR register. See [Address at endpoint list \[device mode\] \(USB_ENDPOINTLISTADDR\)](#) , for more information.

Address: 2_2000h base + 158h offset = 2_2158h



USB_ASYNC_LIST_ADDR field descriptions

Field	Description
31–5 ASYBASE	Link pointer low (LPL). These bits correspond to memory address signals [31:5]. This field may only reference a queue head (QH). Only used by the host controller.
4–0 -	This field is reserved. Reserved, should be cleared.

18.3.15 Address at endpoint list [device mode] (USB_ENDPOINT_LIST_ADDR)

The endpoint list address register is not defined in the EHCI specification. In device mode, this register contains the address of the top of the endpoint list in system memory. Bits 10–0 of this register cannot be modified by the system software and always return zeros when read. The memory structure referenced by this physical memory pointer is assumed to be 64-bytes. The queue head is actually a 48-byte structure, but must be aligned on 64-byte boundary. However, the ENDPOINT_LIST_ADDR[EPBASE] has a granularity of 2 Kbytes, so in practice the queue head should be 2-Kbyte aligned.

Note that this register is shared between the host and device mode functions. In device mode, it is the ENDPOINT_LIST_ADDR register; in host mode, it is the ASYNC_LIST_ADDR register. See [Next asynchronous list addr \[host mode\] \(USB_ASYNC_LIST_ADDR\)](#), for more information.

Address: 2_2000h base + 158h offset = 2_2158h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	EPBASE																Reserved															
W	EPBASE																Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

USB_ENDPOINT_LIST_ADDR field descriptions

Field	Description
31–11 EPBASE	Endpoint list address. Address of the top of the endpoint list.
10–0 -	This field is reserved. Reserved, should be cleared.

18.3.16 Master interface data burst size (USB_BURSTSIZE)

The master interface data burst size register is not defined in the EHCI specification. This register is used to control and dynamically change the burst size used during data movement on the initiator (master) interface.

Address: 2_2000h base + 160h offset = 2_2160h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved																TXPBURST						RXPBURST									
W	Reserved																TXPBURST						RXPBURST									
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0	0	0	0

USB_BURSTSIZE field descriptions

Field	Description
31–16 -	This field is reserved. Reserved, should be cleared.
15–8 TXPBURST	Programmable TX burst length. This register represents the maximum length of a burst in 32-bit words while moving data from system memory to the USB bus. Must not be set to greater than 16.
7–0 RXPBURST	Programmable RX burst length. This register represents the maximum length of a burst in 32-bit words while moving data from the USB bus to system memory. Must not be set to greater than 16.

18.3.17 Transmit FIFO tuning controls (USB_TXFILLTUNING)

The transmit FIFO tuning controls register is not defined in the EHCI specification. This register is used to control and dynamically change the burst size used during data movement on device DMA transfers. It is only used in host mode.

The fields in this register control performance tuning associated with how the USB DR module posts data to the TX latency FIFO before moving the data onto the USB bus. The specific areas of performance include how much data to post into the FIFO and an estimate for how long that operation should take in the target system.

Definitions:

T_0 = Standard packet overhead

T_1 = Time to send data payload

T_s = Total Packet Flight Time (send-only) packet ($T_s = T_0 + T_1$)

T_{ff} = Time to fetch packet into TX FIFO up to specified level.

T_p = Total Packet Time (fetch and send) packet ($T_p = T_{ff} + T_s$)

Upon discovery of a transmit (OUT/SETUP) packet in the data structures, host controller checks to ensure T_p remains before the end of the [micro]frame. If so it proceeds to pre-fill the TX FIFO. If at any time during the pre-fill operation the time remaining the [micro]frame is $< T_s$ then the packet attempt ceases and the packet is tried at a later time. Although this is not an error condition and the module eventually recovers, a mark is made in the scheduler health counter to note the occurrence of a back-off event. When a back-off event is detected, the partial packet fetched may need to be discarded from the latency buffer to make room for periodic traffic that will begin after the next SOF. Too many back-off events can waste bandwidth and power on the system bus and thus should be minimized (not necessarily eliminated). Back-offs can be minimized with use of the TSCHEALTH (T_{ff}) parameter described below.

Address: 2_2000h base + 164h offset = 2_2164h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
R	Reserved																TXFIFOTHRES				Reserved				TXSCHHEALTH				TXSCHOH				
W	Reserved																TXFIFOTHRES				Reserved				TXSCHHEALTH				TXSCHOH				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

USB_TXFILLTUNING field descriptions

Field	Description
31–22 -	This field is reserved. Reserved, should be cleared.
21–16 TXFIFOTHRES	FIFO burst threshold. Control the number of data bursts that are posted to the TX latency FIFO in host mode before the packet begins on to the bus. The minimum value is 2 and this value should be as low as possible to maximize USB performance. A higher value can be used in systems with unpredictable latency and/or insufficient bandwidth where the FIFO may underrun because the data transferred from the latency FIFO to USB occurs before it can be replenished from system memory. This value is ignored if USBMODE[SDIS] (stream disable bit) is set. When USBMODE[SDIS] is set, the host controller behaves as if TXFIFOTHRES is set to the maximum value.
15–13 -	This field is reserved. Reserved, should be cleared.
12–8 TXSCHHEALTH	Scheduler health counter. Increment when the host controller fails to fill the TX latency FIFO to the level programmed by TXFIFOTHRES before running out of time to send the packet before the next Start-Of-Frame. This health counter measures the number of times this occurs to provide feedback to selecting a proper TXSCHOH. Writing to this register clears the counter and this counter stops counting after reaching the maximum of 31.
7–0 TXSCHOH	Scheduler overhead. These bits add an additional fixed offset to the schedule time estimator described above as T_{ff} . As an approximation, the value chosen for this register should limit the number of back-off events captured in the TXSCHHEALTH to less than 10 per second in a highly utilized bus. Choosing a value that is too high for this register is not desired as it can needlessly reduce USB utilization. <ul style="list-style-type: none"> The time unit represented in this register is 1.267μs when a device is connected in high-speed mode. The time unit represented in this register is 6.333μs when a device is connected in low-/full-speed mode. <p>For most applications, TXSCHOH can be set to 4 or less. A good value to begin with is: $TXFIFOTHRES \times (BURSTSIZE \times 4 \text{ bytes-per-word}) \div (40 \times \text{TimeUnit})$, always rounded to the next higher integer. TimeUnit is either 1.267 or 6.333 as noted earlier in this description. For example, if TXFIFOTHRES is 5 and</p>

Table continues on the next page...

USB_TXFILLTUNING field descriptions (continued)

Field	Description
	<p>BURSTSIZE is 8, then set TXSCHOH to $5 \times (8 \times 4) \div (40 \times 1.267) = 4$ for a high-speed link. If this value of TXSCHOH results in a TXSCHHEALTH count of 0 per second, try lowering the value by 1 if optimizing performance is desired. If TXSCHHEALTH exceeds 10 per second, try raising the value by 1.</p> <p>If streaming mode is disabled via the USBMODE register, treat TXFIFOTHRES as the maximum value for purposes of the TXSCHOH calculation.</p>

18.3.18 ULPI register access (USB_ULPI_VIEWPORT)

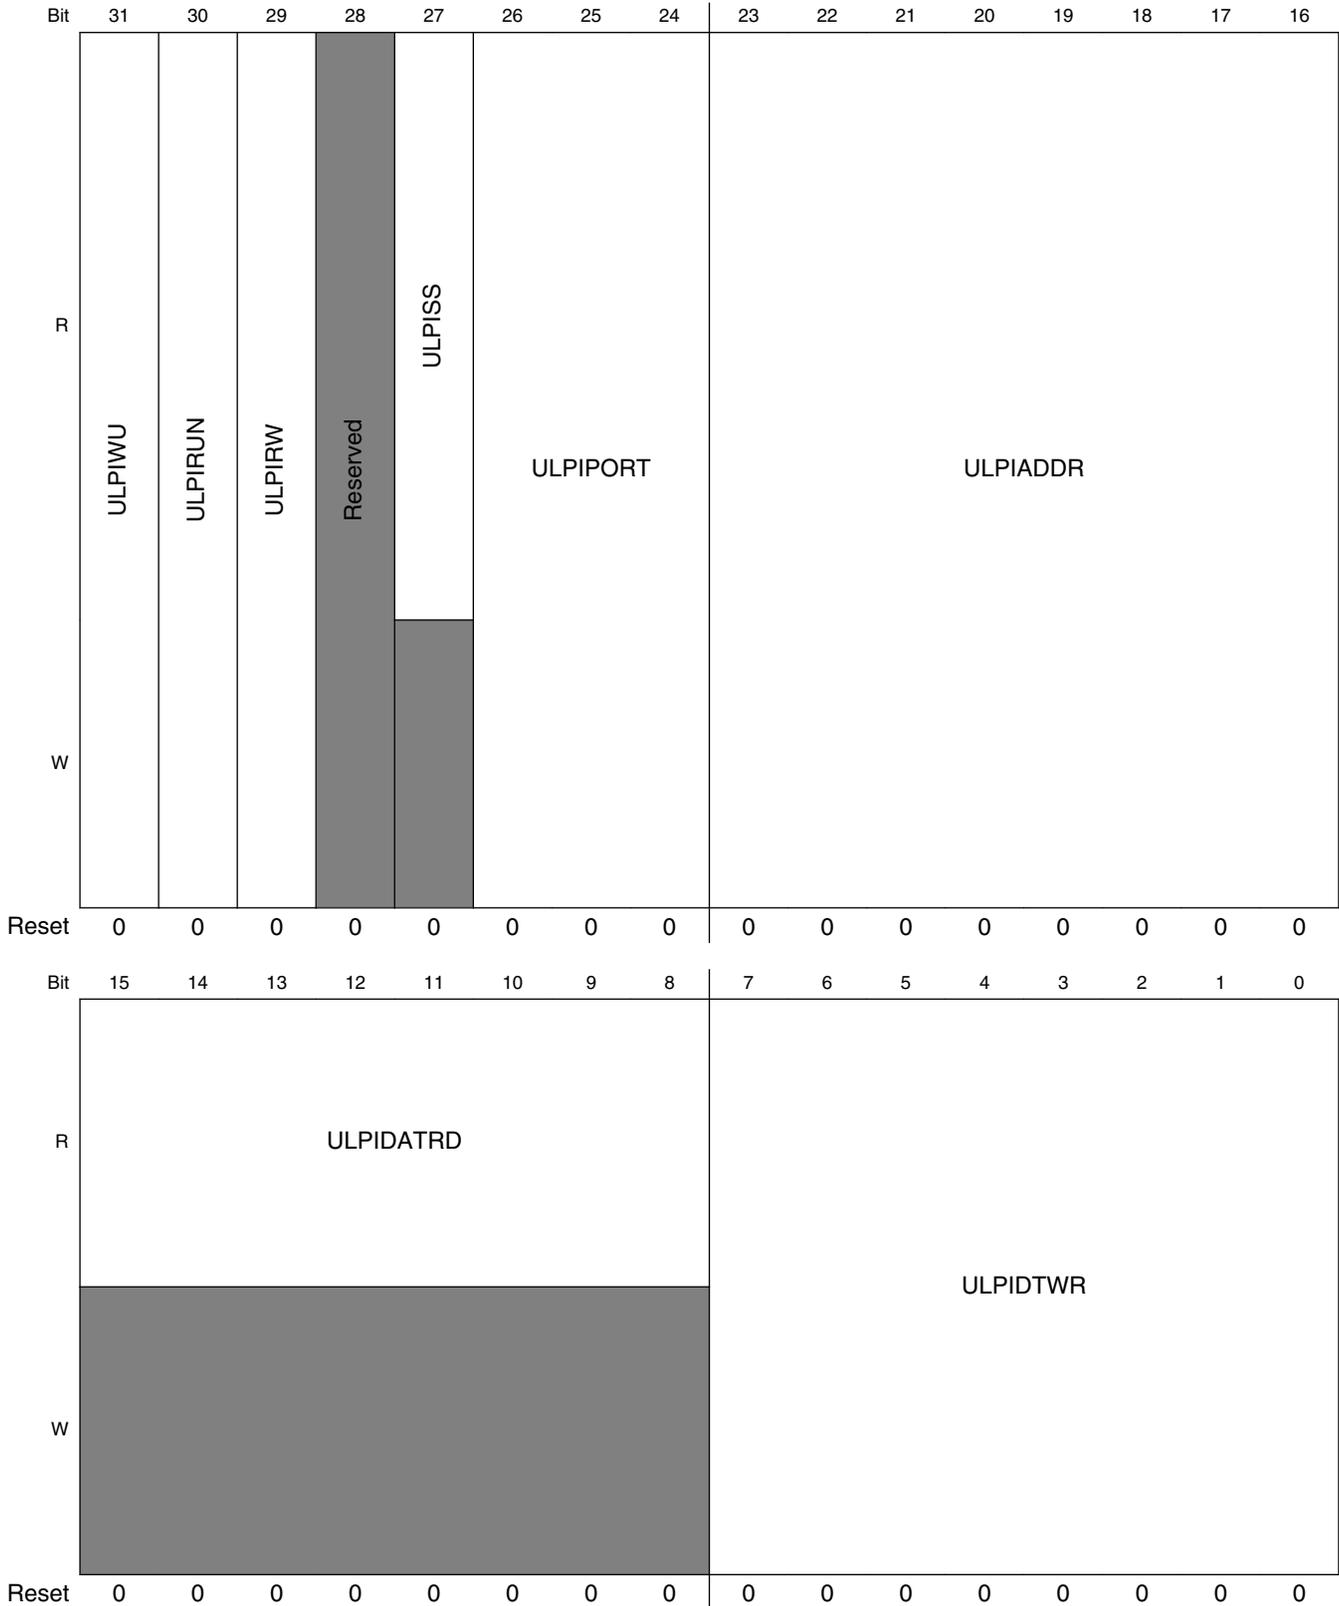
The ULPI register access provides indirect access to the ULPI PHY register set. Although the controller modules perform access to the ULPI PHY register set, there may be extraordinary circumstances where software may need direct access. Be advised that writes to the ULPI through the ULPI viewport can substantially harm standard USB operations. Currently no usage model has been defined where software should need to execute writes directly to the ULPI. Note that executing read operations through the ULPI viewport should have no harmful side effects to standard USB operations. Also note that if the ULPI interface is not enabled, this register will always read zeros.

There are two operations that can be performed with the ULPI viewport, wakeup and read/write operations. The wakeup operation is used to put the ULPI interface into normal operation mode and re-enable the clock if necessary. A wakeup operation is required before accessing the registers when the ULPI interface is operating in low power mode, serial mode, or carkit mode. The ULPI state can be determined by reading the sync state bit (ULPISS). If this bit is set, then the ULPI interface is running in normal operation mode and can accept read/write operations. If the ULPISS is cleared, then read/write operations will not be able execute. Undefined behavior results if a read or write operation is performed when ULPISS is cleared. To execute a wakeup operation, write all 32-bits of the ULPI Viewport where ULPIPORT is constructed appropriately and the ULPIWU bit is set and the ULPIRUN bit is cleared. Poll the ULPI Viewport until ULPIWU is cleared for the operation to complete.

To execute a read or write operation, write all 32-bits of the ULPI Viewport where ULPIDATWR, ULPIADDR, ULPIPORT, ULPIRW are constructed appropriately and the ULPIRUN bit is set. Poll the ULPI Viewport until ULPIRUN is cleared for the operation to complete. For read operations, ULPIDATRD is valid once ULPIRUN is cleared.

The polling method above can be replaced with interrupts using the ULPI interrupt defined in the USBSTS and USBINTR registers. When a wakeup or read/write operation completes, the ULPI interrupt is set.

Address: 2_2000h base + 170h offset = 2_2170h



USB_ULPI_VIEWPORT field descriptions

Field	Description
31 ULPIWU	ULPI Wake Up. Writing 1 to this bit begins the wakeup operation. This bit automatically transitions to 0 after the wakeup is complete. Once this bit is set, it can not be cleared by software. NOTE: The driver must never execute a wakeup and a read/write operation at the same time.
30 ULPIRUN	ULPI Run. Writing 1 to this bit begins a read/write operation. This bit automatically transitions to 0 after the read/write is complete. Once this bit is set, it can not be cleared by software. NOTE: The driver must never execute a wakeup and a read/write operation at the same time.
29 ULPIRW	This bit selects between running a read or write operation to the ULPI. 0 Read 1 Write
28 -	This field is reserved. Reserved, should be cleared.
27 ULPISS	This bit represents the state of the ULPI interface. Before reading this bit, the ULPIPORT field should be set accordingly if used with the multi-port host. Otherwise, this field should always remain 0. 0 Any other state (that is, carkit, serial, low power). 1 Normal Sync State.
26–24 ULPIPORT	For wakeup or read/write operations this value selects the port number to which the ULPI PHY is attached. Valid values are 0 and 1.
23–16 ULPIADDR	When a read or write operation is commanded, the address of the operation is written to this field.
15–8 ULPIDATRD	After a read operation completes, the result is placed in this field.
7–0 ULPIDTWR	When a write operation is commanded, the data to be sent is written to this field.

18.3.19 Configured flag register (USB_CONFIGFLAG)

This EHCI register is not used in this implementation. A read from this register returns a constant of a 0x0000_0001 to indicate that all port routings default to this host controller.

Address: 2_2000h base + 180h offset = 2_2180h

Bit	31	30	29	28	27	26	25	24		23	22	21	20	19	18	17	16
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8		7	6	5	4	3	2	1	0
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	1

USB_CONFIGFLAG field descriptions

Field	Description
31–1 -	This field is reserved. Reserved.
0 CF	Configure flag. Always 1. Indicating all port routings default to this host.

18.3.20 Port status/control (USB_PORTSC)

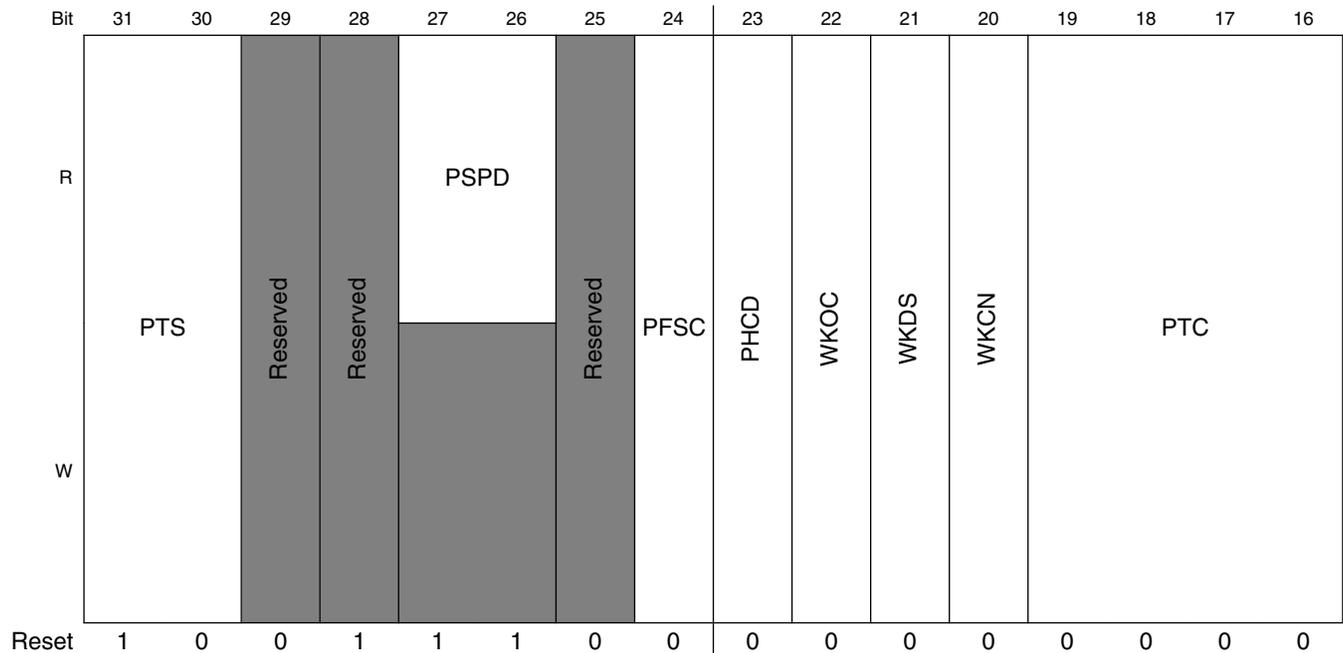
The port status and control (PORTSC) register is only reset when power is initially applied or in response to a controller reset. The initial conditions of a port are:

- No device connected
- Port disabled

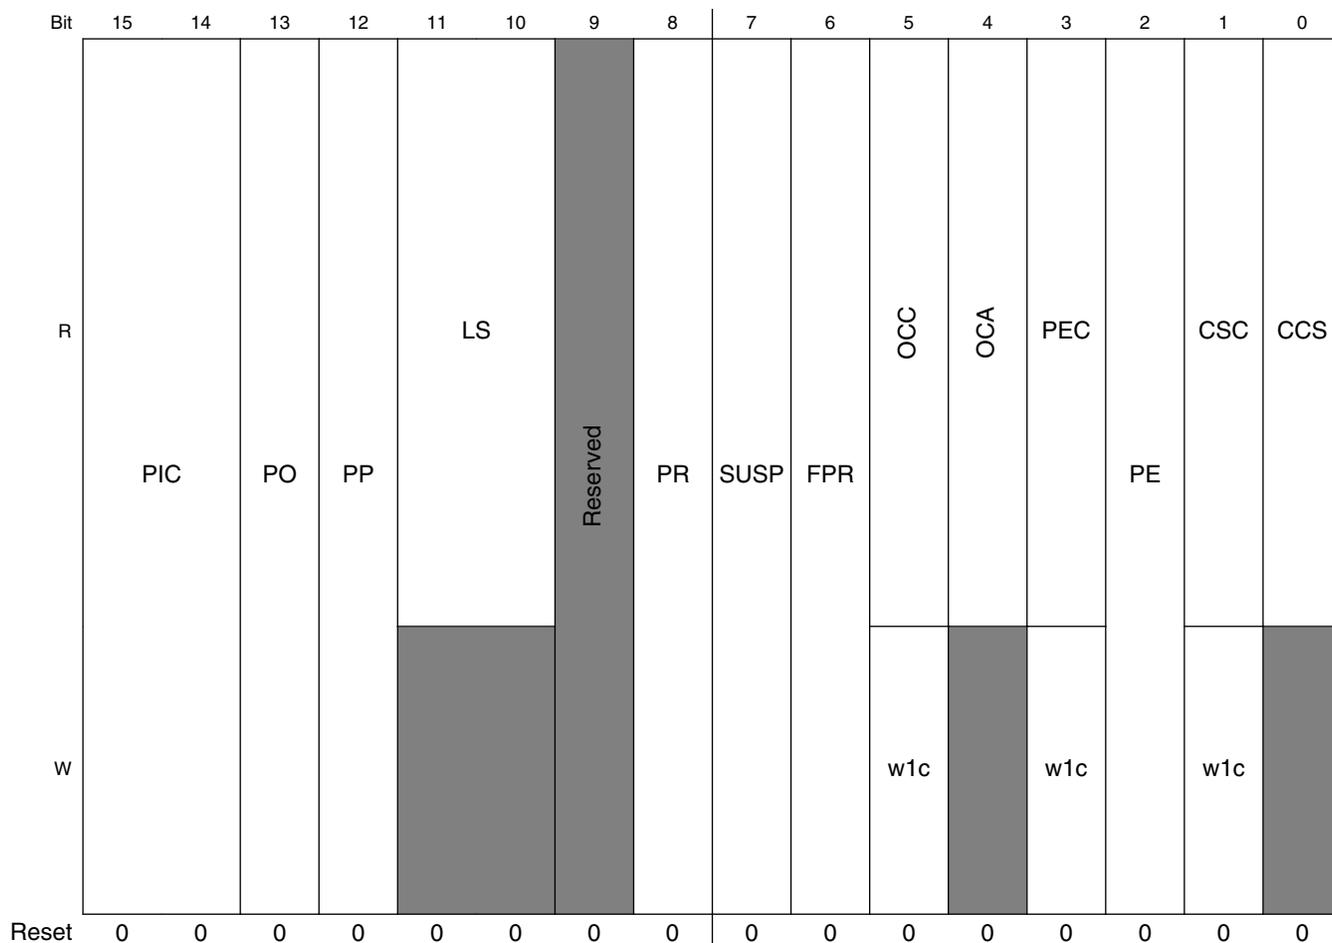
If the port has port power control, this state remains until software applies power to the port by setting port power to one.

In device mode, the USB DR controller does not support power control. Port control in device mode is only used for status port reset, suspend, and current connect status. It is also used to initiate test mode or force signaling and allows software to put the PHY into low power suspend mode and disable the PHY clock.

Address: 2_2000h base + 184h offset = 2_2184h



USB memory map/register definition



USB_PORTSC field descriptions

Field	Description
31–30 PTS	Port transceiver select. This register bit is used to control which parallel transceiver interface is selected. This bit is not defined in the EHCI specification. 00 Reserved 01 Reserved, should be cleared 10 ULPI parallel interface 11 Reserved
29 -	This field is reserved. Reserved, should be cleared
28 -	This field is reserved. Reserved, should be cleared
27–26 PSPD	Port speed. This read-only register field indicates the speed at which the port is operating. This bit is not defined in the EHCI specification. 00 Full-speed 01 Low-speed 10 High-speed 11 Undefined

Table continues on the next page...

USB_PORTSC field descriptions (continued)

Field	Description
25 -	This field is reserved. Reserved, should be cleared
24 PFSC	Port force full-speed connect. Used to disable the chirp sequence that allows the port to identify itself as a HS port. This is useful for testing FS configurations with a HS host, hub or device. This bit is not defined in the EHCI specification. This bit is for debugging purposes. 0 Allow the port to identify itself as high speed. 1 Force the port to only connect at full speed.
23 PHCD	PHY low power suspend. This bit is not defined in the EHCI specification. Host mode: <ul style="list-style-type: none"> The PHY can be put into low power suspend - when the downstream device has been put into suspend mode or when no downstream device is connected. Low power suspend is completely under the control of software. Device mode: <ul style="list-style-type: none"> The PHY can be put into low power suspend - when the device is not running (USBCMD[RS] = 0b) or suspend signaling is detected on the USB. Low power suspend is cleared automatically when the resume signaling has been detected or when forcing port resume. Reading this bit indicates the status of the PHY. NOTE: If there is no clock connected to the USB_CLK signals, PHCD must be set and the following registers should not be written: DEVICE_ADDR/PERIODICLISTBASE, PORTSC, ENDPTCTRL0, ENDPTCTRL1, ENDPTCTRL2, ENDPTCTRL3, ENDPTCTRL4, ENDPTCTRL5 . 0 Normal PHY operation. 1 Signal the PHY to enter low power suspend mode
22 WKOC	Wake on over-current enable. Writing this bit to a one enables the port to be sensitive to over-current conditions as wake-up events. This field is zero if Port Power (PP) is zero. This bit is used only in Host mode , and is used by an external power control circuit .
21 WKDS	Wake on disconnect enable. Writing this bit to a one enables the port to be sensitive to device disconnects as wake-up events. This field is zero if Port Power(PP) is zero or in device mode. This bit is used only in Host mode , and is used by an external power control circuit .
20 WKCEN	Wake on connect enable. Writing this bit to a one enables the port to be sensitive to device connects as wake-up events. This field is zero if Port Power(PP) is zero or in device mode. This bit is used only in Host mode , and is used by an external power control circuit .
19–16 PTC	Port test control. Any other value than zero indicates that the port is operating in test mode. Refer to Chapter 7 of the USB Specification Revision 2.0 [3] for details on each test mode. 0000 Not Enabled 0001 J_STATE 0010 K_STATE 0011 SEQ_NAK

Table continues on the next page...

USB_PORTSC field descriptions (continued)

Field	Description
	0100 Packet 0101 FORCE_ENABLE 0110 - 1111 Reserved, should be cleared
15–14 PIC	<p>Port indicator control. Control the link indicator signals. These signals are valid for host mode only. Refer to the USB Specification Revision 2.0 [3] for a description on how these bits are to be used. This field is output from the module on the USB port control signals for use by an external LED driving circuit.</p> <p>NOTE: The PORTSC[14] is mapped to PCTL0 and PORTSC[15] is mapped to PCTL1.</p> 00 Off 01 Amber 10 Green 11 Undefined
13 PO	<p>Port owner. Unconditionally goes to a 0 when the configured bit in the CONFIGFLAG register makes a 0 to 1 transition. This bit unconditionally goes to 1 whenever the Configured bit is zero. System software uses this field to release ownership of the port to a selected module (in the event that the attached device is not a high-speed device). Software writes a one to this bit when the attached device is not a high-speed device. A one in this bit means that an internal companion controller owns and controls the port.</p> <p>Port owner hand-off is not implemented in this design, therefore this bit is always 0.</p>
12 PP	<p>Port power. Represents the current setting of the switch (0=off, 1=on). When power is not available on a port (that is, PP equals a 0), the port is non-functional and will not report attaches, detaches, and so on.</p> <p>When an over-current condition is detected on a powered port, the PP bit in each affected port is transitioned by the host controller driver from a one to a zero (removing power from the port).</p> <p>This feature is implemented in the host controller (PPC = 1).</p> <p>In a device-only implementation port power control is not necessary, thus PPC and PP = 0.</p>
11–10 LS	<p>Line status. Reflect the current logical levels of the USB D+ (bit 11) and D- (bit 10) signal lines. The use of line status by the host controller driver is not necessary (unlike EHCI), because the connection of FS and LS is managed by hardware.</p> 00 SE0 10 J-state 01 K-state 11 Undefined
9 -	<p>This field is reserved. Reserved, should be cleared</p>
8 PR	<p>Port reset.</p> <p>Host mode:</p> <ul style="list-style-type: none"> When software writes a one to this bit the bus-reset sequence as defined in the USB Specification Revision 2.0 is started. This bit will automatically change to zero after the reset sequence is complete. This behavior is different from EHCI where the host controller driver is required to set this bit to a zero after the reset duration is timed in the driver. <p>Device mode:</p> <ul style="list-style-type: none"> This bit is a read only status bit. Device reset from the USB bus is also indicated in the USBSTS register.

Table continues on the next page...

USB_PORTSC field descriptions (continued)

Field	Description
	This field is zero if Port Power(PP) is zero. 1 Port is in reset. 0 Port is not in reset.
7 SUSP	Suspend. Host mode: <ul style="list-style-type: none"> • The port enabled bit (PE) and suspend (SUSP) bit define the port states as follows: 0x Disable 10 Enable 11 Suspend <ul style="list-style-type: none"> • When in suspend state, downstream propagation of data is blocked on this port, except for port reset. The blocking occurs at the end of the current transaction if a transaction was in progress when this bit was written to 1. In the suspend state, the port is sensitive to resume detection. Note that the bit status does not change until the port is suspended and that there may be a delay in suspending a port if there is a transaction currently in progress on the USB. • The module unconditionally sets this bit to zero when software clears the FPR bit. A write of zero to this bit is ignored by the host controller. If host software sets this bit to a one when the port is not enabled (that is, port enabled bit is a zero) the results are undefined. • This field is zero if Port Power (PP) is zero in host mode. Device mode: 1 Port in suspend state. 0 Port not in suspend state. Default. In device mode this bit is a read-only status bit.
6 FPR	Force port resume. This bit is not-EHCI compatible. 1 Resume detected/driven on port. 0 No resume (K-state) detected/driven on port. Host mode: <ul style="list-style-type: none"> • Software sets this bit to one to drive resume signaling. The controller sets this bit to one if a J-to-K transition is detected while the port is in the Suspend state. When this bit transitions to a one a J-to-K transition is detected, USBSTS[PCI] (port change detect) is also set. This bit will automatically change to zero after the resume sequence is complete. This behavior is different from EHCI where the host controller driver is required to set this bit to a zero after the resume duration is timed in the driver. • Note that when the controller owns the port, the resume sequence follows the defined sequence documented in the USB Specification Revision 2.0. The resume signaling (Full-speed 'K') is driven on the port as long as this bit remains a one. This bit will remain a one until the port has switched to the high-speed idle. Writing a zero has no affect because the port controller will time the resume operation and clear this bit when the port control state switches to HS or FS idle. • This field is zero if Port Power (PP) is zero in host mode. Device mode: <ul style="list-style-type: none"> • After the device has been in Suspend State for 5 msec or more, software must set this bit to one to drive resume signaling before clearing. The USB DR controller will set this bit to one if a J-to-K transition is detected while the port is in the Suspend state. The bit is cleared when the device returns to normal operation. Also, when this bit transitions to a one because a J-to-K transition detected, USBSTS[PCI] is also set.

Table continues on the next page...

USB_PORTSC field descriptions (continued)

Field	Description
5 OCC	<p>Over-current change. This bit gets set when there is a change to over-current active. Software clears this bit by writing a one to this bit position.</p> <p>Host mode:</p> <ul style="list-style-type: none"> The user can provide over-current detection to the USB <i>n</i>_PWRFAULT signal for this condition. <p>Device mode:</p> <ul style="list-style-type: none"> This bit must always be 0. <p>1 Over current detect. 0 No over current.</p>
4 OCA	<p>Over-current active. This bit will automatically transition from one to zero when the over current condition is removed.</p> <p>Host mode:</p> <ul style="list-style-type: none"> The user can provide over-current detection to the USB <i>n</i>_PWRFAULT signal for this condition. <p>Device mode:</p> <ul style="list-style-type: none"> This bit must always be 0. <p>1 Port currently in over-current condition. 0 Port not in over-current condition.</p>
3 PEC	<p>Port enable/disable change.</p> <p>For the root hub, this bit gets set only when a port is disabled due to disconnect on the port or due to the appropriate conditions existing at the EOF2 point (See Chapter 11 of the USB Specification). Software clears this by writing a one to it.</p> <p>In device mode:</p> <ul style="list-style-type: none"> The device port is always enabled. (This bit is zero.) <p>This field is zero if Port Power(PP) is zero.</p> <p>1 Port disabled. 0 No change.</p>
2 PE	<p>Port enabled/disabled.</p> <p>Host mode:</p> <ul style="list-style-type: none"> Ports can only be enabled by the controller as a part of the reset and enable. Software cannot enable a port by writing a one to this field. Ports can be disabled by either a fault condition (disconnect event or other fault condition) or by the host software. Note that the bit status does not change until the port state actually changes. There may be a delay in disabling or enabling a port due to other host and bus events. When the port is disabled, (0) downstream propagation of data is blocked except for reset. This field is zero if Port Power(PP) is zero in host mode. <p>Device mode:</p> <ul style="list-style-type: none"> The device port is always enabled. (This bit is one.)
1 CSC	<p>Connect change status.</p> <p>Host mode:</p> <ul style="list-style-type: none"> This bit indicates a change has occurred in the port's Current Connect Status. the controller sets this bit for all changes to the port device connect status, even if system software has not cleared an

Table continues on the next page...

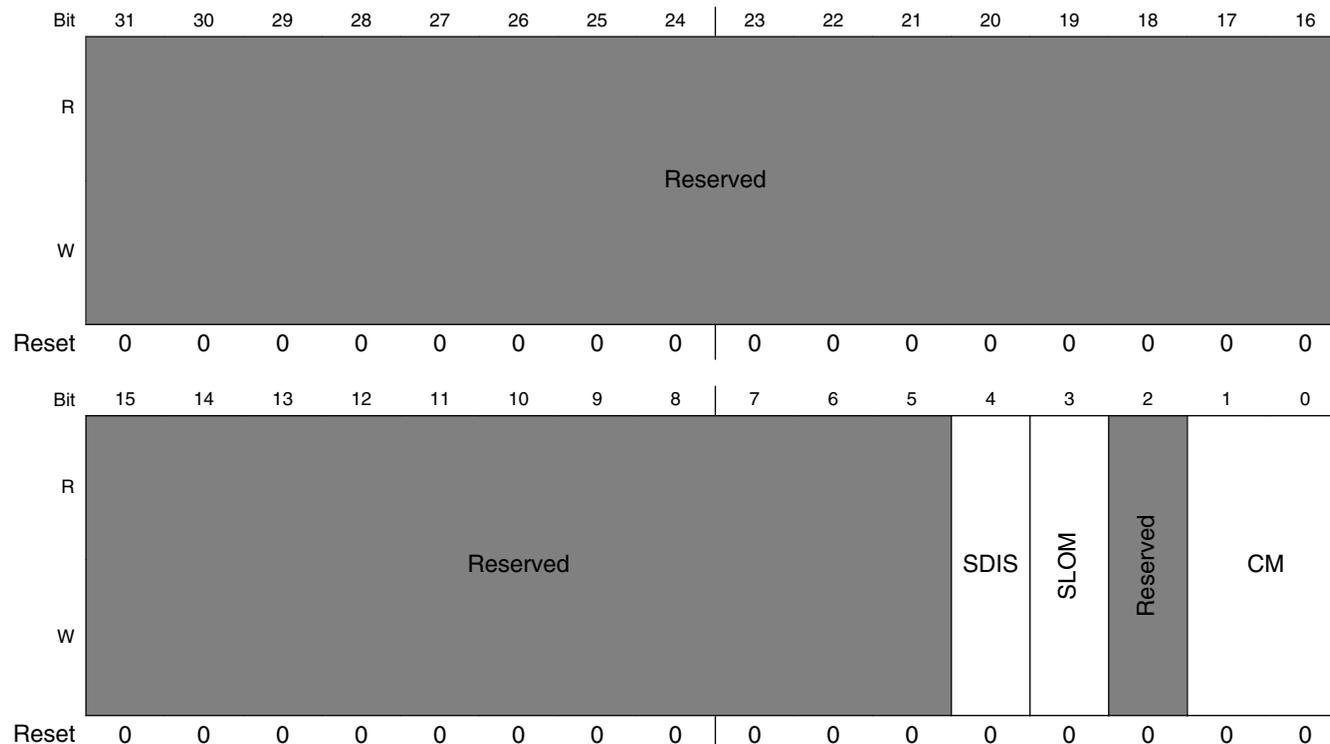
USB_PORTSC field descriptions (continued)

Field	Description
	<p>existing connect status change. For example, the insertion status changes twice before system software has cleared the changed condition, hub hardware is 'setting' an already-set bit (that is, the bit will remain set). Software clears this bit by writing a one to it.</p> <p>1 Connect Status has changed.</p> <p>0 No change.</p> <ul style="list-style-type: none"> • This field is zero if Port Power(PP) is zero. <p>Device mode:</p> <ul style="list-style-type: none"> • This bit is undefined.
<p>0 CCS</p>	<p>Current connect status.</p> <p>Host mode:</p> <p>1 Device is present</p> <p>0 No device present.</p> <p>This field is zero if Port Power(PP) is zero in host mode.</p> <p>In device mode:</p> <p>1 Attached</p> <p>0 Not attached.</p> <p>A one indicates that the device successfully attached and is operating in either high-speed or full-speed as indicated by the PSPD field in this register. A zero indicates that the device did not attach successfully or was forcibly disconnected by the software writing a zero to USBCMD[RS] (run bit). It does not state the device being disconnected or suspended.</p>

18.3.21 USB device mode (USB_USBMODE)

The USB mode register is not defined in the EHCI specification. This register controls the operating mode of the module.

Address: 2_2000h base + 1A8h offset = 2_21A8h



USB_USBMODE field descriptions

Field	Description
31–5 -	This field is reserved. Reserved, should be cleared.
4 SDIS	Stream disable In host mode, setting this bit ensures that overruns/underruns of the latency FIFO are eliminated for low bandwidth systems where the RX and TX buffers are sufficient to contain the entire packet. Enabling stream disable also has the effect of ensuring the TX latency is filled to capacity or a complete packet is stored in FIFO before the packet is launched onto the USB. Note that time duration to pre-fill the FIFO becomes significant when stream disable is active. See Transmit FIFO tuning controls (USB_TXFILLTUNING) to characterize the adjustments needed for the scheduler when using this feature. Also note that in systems with high system bus utilization, setting this bit will ensure no overruns or underruns during operation, at the expense of link utilization. For those who desire optimal link performance, SDIS can be left clear, and the rules used under the description of the TXFILLTUNING register to limit underruns/overruns.

Table continues on the next page...

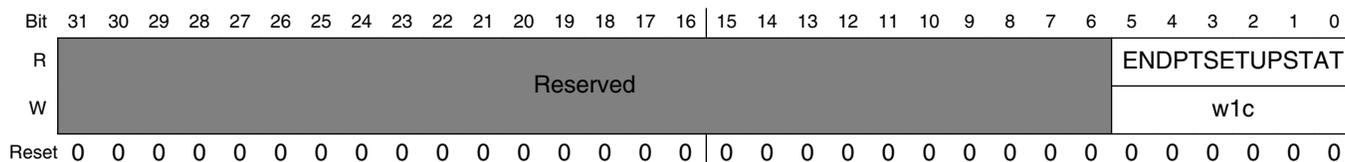
USB_USBMODE field descriptions (continued)

Field	Description
	<p>In device mode, setting this bit disables double priming on both RX and TX for low bandwidth systems. This mode ensures that when the RX and TX buffers are sufficient to contain an entire packet that the standard double buffering scheme is disabled to prevent overruns/underruns in bandwidth limited systems.</p> <p>Note that in high-speed mode, all packets received are responded to with a NYET handshake when stream disable is active.</p> <p>1 Active. 0 Inactive.</p>
3 SLOM	<p>Setup lockout mode. In device mode, this bit controls behavior of the setup lock mechanism. See Control endpoint operation model .</p> <p>1 Setup lockouts off. DCD requires use of setup data buffer tripwire in USBCMD (SUTW). 0 Setup lockouts on</p>
2 -	<p>This field is reserved. Reserved, should be cleared.</p>
1-0 CM	<p>Controller mode</p> <p>This register can only be written once after reset. If it is necessary to switch modes, software must reset the controller by writing to USBCMD[RST] before reprogramming this register.</p> <p>Defaults to the idle state and needs to be initialized to the desired operating mode after reset.</p> <p>00 Idle (default) . 01 Reserved, should be cleared. 10 Device controller . 11 Host controller .</p>

18.3.22 Endpoint setup status (USB_ENDPTSETUPSTAT)

The endpoint setup status register is not defined in the EHCI specification. This register contains the endpoint setup status. It is only used in device mode.

Address: 2_2000h base + 1ACh offset = 2_21ACh



USB_ENDPTSETUPSTAT field descriptions

Field	Description
31-6 -	<p>This field is reserved. Reserved, should be cleared.</p>

Table continues on the next page...

USB_ENDPTSETUPSTAT field descriptions (continued)

Field	Description
5–0 ENDPTSETUPSTAT	Setup endpoint status. For every setup transaction that is received, a corresponding bit in this register is set. Software must clear or acknowledge the setup transfer by writing a one to a respective bit after it has read the setup data from queue head. The response to a setup packet as in the order of operations and total response time is crucial to limit bus time outs while the setup lockout mechanism is engaged. This register is only used in device mode.

18.3.23 Endpoint initialization (USB_ENDPOINTPRIME)

The endpoint initialization register is not defined in the EHCI specification. This register is used to initialize endpoints. It is only used in device mode.

Address: 2_2000h base + 1B0h offset = 2_21B0h

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

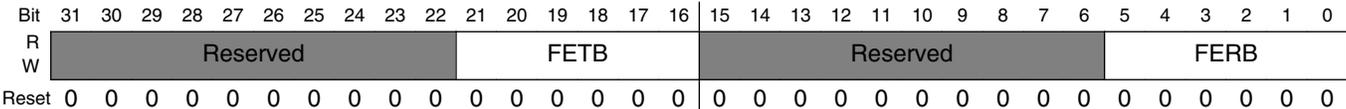
USB_ENDPOINTPRIME field descriptions

Field	Description
31–22 -	This field is reserved. Reserved, should be cleared.
21–16 PETB	Prime endpoint transmit buffer. For each endpoint a corresponding bit is used to request that a buffer prepared for a transmit operation in order to respond to a USB IN/INTERRUPT transaction. Software should write a one to the corresponding bit when posting a new transfer descriptor to an endpoint. Hardware will automatically use this bit to begin parsing for a new transfer descriptor from the queue head and prepare a transmit buffer. Hardware will clear this bit when the associated endpoint(s) is (are) successfully primed. PETB[5] (bit 21 of the register) corresponds to endpoint 5 . Note that these bits are momentarily set by hardware during hardware re-priming operations when a dTD is retired, and the dQH is updated.
15–6 -	This field is reserved. Reserved, should be cleared.
5–0 PERB	Prime endpoint receive buffer. For each endpoint, a corresponding bit is used to request a buffer prepare for a receive operation in order to respond to a USB OUT transaction. Software should write a one to the corresponding bit whenever posting a new transfer descriptor to an endpoint. Hardware will automatically use this bit to begin parsing for a new transfer descriptor from the queue head and prepare a receive buffer. Hardware will clear this bit when the associated endpoint(s) is (are) successfully primed. PERB[5] corresponds to endpoint 5 . Note that these bits are momentarily set by hardware during hardware re-priming operations when a dTD is retired, and the dQH is updated.

18.3.24 Endpoint flush (USB_ENDPTFLUSH)

The endpoint flush register is not defined in the EHCI specification. This register is only used in device mode.

Address: 2_2000h base + 1B4h offset = 2_21B4h



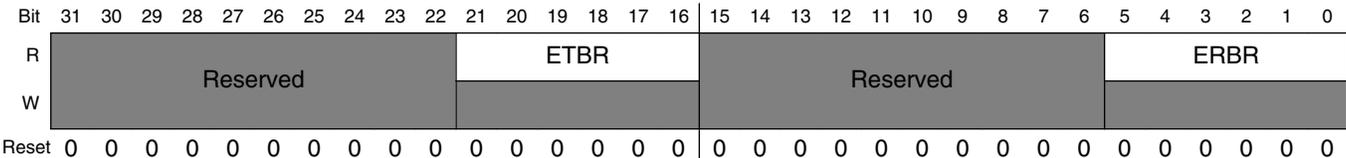
USB_ENDPTFLUSH field descriptions

Field	Description
31–22 -	This field is reserved. Reserved, should be cleared.
21–16 FETB	Flush endpoint transmit buffer. Writing a one to a bit(s) in this register will cause the associated endpoint(s) to clear any primed buffers. If a packet is in progress for one of the associated endpoints, then that transfer will continue until completion. Hardware will clear this register after the endpoint flush operation is successful. FETB[5] (bit 21 of the register) corresponds to endpoint 5 .
15–6 -	This field is reserved. Reserved, should be cleared.
5–0 FERB	Flush endpoint receive buffer. Writing a one to a bit(s) will cause the associated endpoint(s) to clear any primed buffers. If a packet is in progress for one of the associated endpoints, then that transfer will continue until completion. Hardware will clear this register after the endpoint flush operation is successful. FERB[5] corresponds to endpoint 5 .

18.3.25 Endpoint status (USB_ENDPTSTATUS)

The endpoint status register is not defined in the EHCI specification. This register is only used in device mode.

Address: 2_2000h base + 1B8h offset = 2_21B8h



USB_ENDPTSTATUS field descriptions

Field	Description
31–22 -	This field is reserved. Reserved, should be cleared
21–16 ETBR	Endpoint transmit buffer ready. One bit for each endpoint indicates status of the respective endpoint buffer. This bit is set by the hardware as a response to receiving a command from a corresponding bit in the ENDPTPRIME register. There will always be a delay between setting a bit in the ENDPTPRIME register and endpoint indicating ready. This delay time varies based upon the current USB traffic and the number of bits set in the ENDPTPRIME register. Buffer ready is cleared by USB reset, by the USB DMA system, or through the ENDPTFLUSH register. ETBR[5] (bit 21 of the register) corresponds to endpoint 5 . Note that these bits are momentarily cleared by hardware during hardware endpoint re-priming operations when a dTD is retired, and the dQH is updated.
15–6 -	This field is reserved. Reserved, should be cleared
5–0 ERBR	Endpoint receive buffer ready. One bit for each endpoint indicates status of the respective endpoint buffer. This bit is set by the hardware as a response to receiving a command from a corresponding bit in the ENDPTPRIME register. There will always be a delay between setting a bit in the ENDPTPRIME register and endpoint indicating ready. This delay time varies based upon the current USB traffic and the number of bits set in the ENDPTPRIME register. Buffer ready is cleared by USB reset, by the USB DMA system, or through the ENDPTFLUSH register. ERBR[5] corresponds to endpoint 5 . Note that these bits are momentarily cleared by hardware during hardware endpoint re-priming operations when a dTD is retired, and the dQH is updated.

18.3.26 Endpoint complete (USB_ENDPTCOMPLETE)

The endpoint complete register is not defined in the EHCI specification. This register is only used in device mode.

Address: 2_2000h base + 1BCh offset = 2_21BCh

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved											ETCE					Reserved											ERCE				
W	Reserved											w1c					Reserved											w1c				
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

USB_ENDPTCOMPLETE field descriptions

Field	Description
31–22 -	This field is reserved. Reserved, should be cleared
21–16 ETCE	Endpoint transmit complete event. Each bit indicates a transmit event (IN/INTERRUPT) occurred and software should read the corresponding endpoint queue to determine the endpoint status. If the corresponding IOC bit is set in the Transfer Descriptor, then this bit is set simultaneously with the USBINT. Writing a one will clear the corresponding bit in this register. ETCE[5] (bit 21 of the register) corresponds to endpoint 5 .

Table continues on the next page...

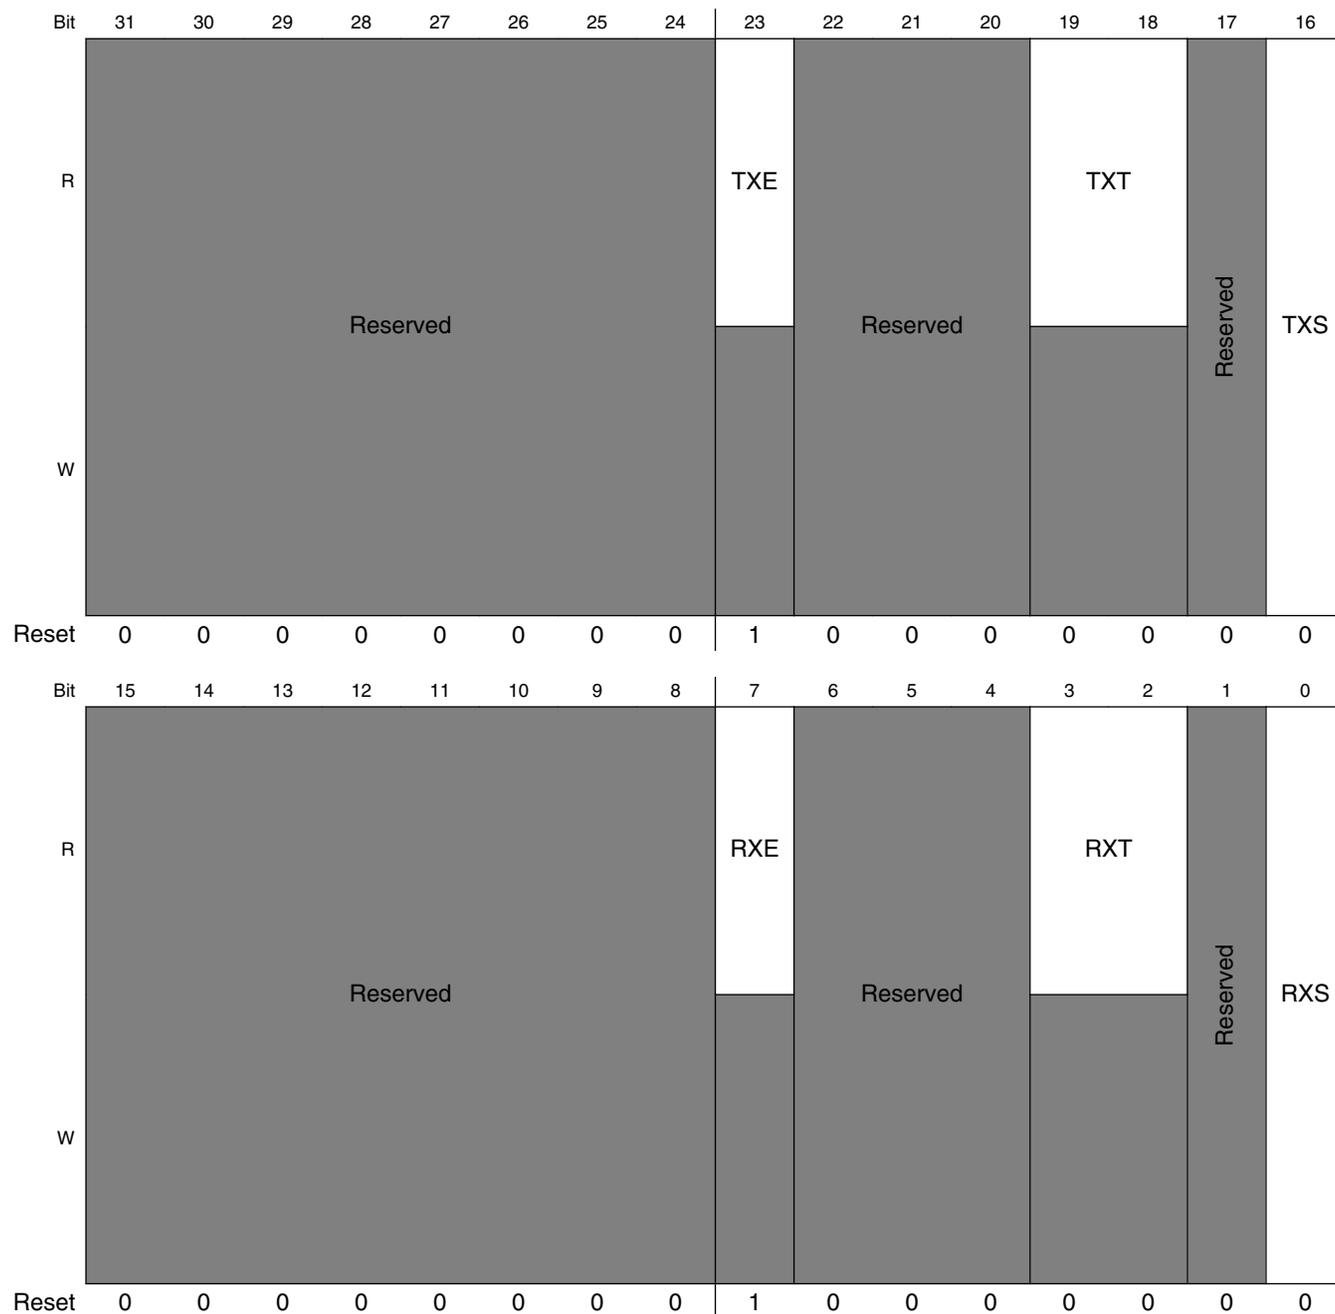
USB_ENDPTCOMPLETE field descriptions (continued)

Field	Description
15–6 -	This field is reserved. Reserved, should be cleared
5–0 ERCE	Endpoint receive complete event. Each bit indicates a received event (OUT/SETUP) occurred and software should read the corresponding endpoint queue to determine the transfer status. If the corresponding IOC bit is set in the Transfer Descriptor, then this bit is set simultaneously with the USBINT. Writing a one will clear the corresponding bit in this register. ERCE[5] corresponds to endpoint 5 .

18.3.27 Endpoint control 0 (USB_ENDPTCTRL0)

Endpoint control register 0 is not defined in the EHCI specification. Every device will implement endpoint 0 as a control endpoint.

Address: 2_2000h base + 1C0h offset = 2_21C0h



USB_ENDPTCTRL0 field descriptions

Field	Description
31–24 -	This field is reserved. Reserved, should be cleared.
23 TXE	TX endpoint enable. Endpoint zero is always enabled. 0 Disable 1 Enable
22–20 -	This field is reserved. Reserved, should be cleared.
19–18 TXT	TX endpoint type. Endpoint zero is always a control endpoint (00).
17 -	This field is reserved. Reserved, should be cleared.
16 TXS	TX endpoint stall. Software can write a one to this bit to force the endpoint to return a STALL handshake to the Host. It will continue returning STALL until the bit is cleared by software or it will automatically be cleared upon receipt of a new SETUP request. 1 Endpoint stalled 0 Endpoint OK
15–8 -	This field is reserved. Reserved, should be cleared.
7 RXE	RX endpoint enable. Endpoint zero is always enabled. 0 Disabled 1 Enabled
6–4 -	This field is reserved. Reserved, should be cleared.
3–2 RXT	RX endpoint type. Endpoint zero is always a control endpoint (00).
1 -	This field is reserved. Reserved, should be cleared.
0 RXS	RX endpoint stall Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue returning STALL until the bit is cleared by software or it will automatically be cleared upon receipt of a new SETUP request. 1 Endpoint stalled 0 Endpoint OK

18.3.28 Endpoint control n (USB_ENDPTCTRLn)

The endpoint control n registers are not defined in the EHCI specification. There is an ENDPTCTRL n register of each endpoint in a device.

Address: 2_2000h base + 1C4h offset + (4d × i), where i=0d to 4d

Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
R	Reserved								TXE	TXR	TXI	Reserved	TXT		TXD	TXS
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
R	Reserved								RXE	RXR	RXI	Reserved	RXT		RXD	RXS
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

USB_ENDPTCTRLn field descriptions

Field	Description
31–24 -	This field is reserved. Reserved, should be cleared
23 TXE	TX endpoint enable 0 Disabled 1 Enabled
22 TXR	TX data toggle reset. Whenever a configuration event is received for this endpoint, software must write a one to this bit in order to synchronize the data PID's between the Host and device.
21 TXI	TX data toggle inhibit. Used only for test and should always be written as zero. Writing a one to this bit will cause this endpoint to ignore the data toggle sequence and always transmit DATA0 for a data packet. 0 PID sequencing enabled 1 PID sequencing disabled
20 -	This field is reserved. Reserved, should be cleared

Table continues on the next page...

USB_ENDPTCTRLn field descriptions (continued)

Field	Description
19–18 TXT	TX endpoint type NOTE: When only one endpoint (RX or TX, but not both) of an endpoint pair is used, the unused endpoint should be configured as a bulk type endpoint. 00 Control 01 Isochronous 10 Bulk 11 Interrupt
17 TXD	TX endpoint data source. This bit should always be written as 0, which selects the dual port memory/DMA engine as the source.
16 TXS	TX endpoint stall. This bit is set automatically upon receipt of a SETUP request if this endpoint is not configured as a control endpoint. It is cleared automatically upon receipt of a SETUP request if this endpoint is configured as a control endpoint. Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue to returning STALL until this bit is either cleared by software or automatically cleared as above. 0 Endpoint OK 1 Endpoint stalled
15–8 -	This field is reserved. Reserved, should be cleared
7 RXE	RX endpoint enable 0 Disabled 1 Enabled
6 RXR	RX data toggle reset. Whenever a configuration event is received for this endpoint, software must write a one to this bit in order to synchronize the data PID's between the Host and device.
5 RXI	RX data toggle inhibit. This bit is only used for test and should always be written as zero. Writing a one to this bit will cause this endpoint to ignore the data toggle sequence and always accept data packets regardless of their data PID. 1 PID sequencing enabled 0 PID sequencing disabled
4 -	This field is reserved. Reserved, should be cleared
3–2 RXT	RX endpoint type. NOTE: When only one endpoint (RX or TX, but not both) of an endpoint pair is used, the unused endpoint should be configured as a bulk type endpoint. 00 Control 01 Isochronous 10 Bulk 11 Interrupt
1 RXD	RX endpoint data sink. This bit should always be written as 0, which selects the dual port memory/DMA engine as the sink.
0 RXS	RX endpoint stall. This bit is set automatically upon receipt of a SETUP request if this endpoint is not configured as a control endpoint. It is cleared automatically upon receipt a SETUP request if this endpoint is configured as a control endpoint,

Table continues on the next page...

USB_ENDPTCTRLn field descriptions (continued)

Field	Description
	Software can write a one to this bit to force the endpoint to return a STALL handshake to the host. It will continue to returning STALL until this bit is either cleared by software or automatically cleared as above,
1	Endpoint stalled
0	Endpoint OK

18.3.29 Snoop n (USB_SNOOPn)

The SNOOP1 and SNOOP2 registers are shown below. Note that these registers use big-endian byte ordering and are not defined in the EHCI specification. The SNOOP1 and SNOOP2 registers provide snooping control and address range selection function. Transactions that hit a snooping window will generate cache coherent transactions on the internal CCB bus. When the five lower bits (SNOOP *n* [27-31]) are equal to 00000, snooping is always disabled on the CCB for all DMA transfers. When SNOOP *n* [27-31] is 01011 through 11110, the twenty upper bits (SNOOP *n* [0-19]) provide the starting base address for which transactions are snooped. These twenty bits are compared to the twenty upper bits of the address provided by the DMA block of the USB controller. When a match occurs, the five lower bits are decoded as shown below. This provides a snooping region of 4 Kbytes to 2 Gbytes within each starting base address that is programmed by the core. The SNOOP *n* [20-26] are not used.

Address: 2_2000h base + 400h offset + (4d × i), where i=0d to 1d

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																Reserved							Snoop_Enables									
W	Snoop_address																															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

USB_SNOOPn field descriptions

Field	Description
0–19 Snoop_address	The starting base address for which transactions are snooped.
20–26 -	This field is reserved. Reserved, should be cleared
27–31 Snoop_Enables	0x00 Snooping disabled 0x0B 4-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-19] 0x0C 8-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-18] 0x0D 16-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-17] 0x0E 32-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-16] 0x0F 64-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-15] 0x10 128-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-14]

Table continues on the next page...

USB_SNOOP_n field descriptions (continued)

Field	Description
0x11	256-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-13]
0x12	512-Kbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-12]
0x13	1-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-11]
0x14	2-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-10]
0x15	4-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-9]
0x16	8-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-8]
0x17	16-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-7]
0x18	32-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-6]
0x19	64-M byte snoop range starting at the value defined by SNOOP <i>n</i> [0-5]
0x1A	128-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-4]
0x1B	256-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-3]
0x1C	512-Mbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-2]
0x1D	1-Gbyte snoop range starting at the value defined by SNOOP <i>n</i> [0-1]
0x1E	2-Gbyte snoop range starting at the value defined by SNOOP <i>n</i> [0]

18.3.30 Age count threshold (USB_AGE_CNT_THRESH)

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The age count threshold (AGE_CNT_THRESH) register provides the aging counter threshold value used to determine the priority state of the USB DR controller's internal system interface. This is used to increase the priority state of the module's system interface from zero to one. The actual priority level on the system bus for each state is defined by the PRI_CTRL register. The threshold value is in units of *ccb_clk* cycles. This register should be written during system initialization or during normal system operation when the system bus interface is idle. It can be read at any time.

If the aging counter is less than the AGE_CNT_THRESH value, priority state zero is chosen. If the aging counter is greater than or equal to the AGE_CNT_THRESH value, priority state one is chosen.

The aging counter begins to count from zero when a bus access is requested. It increments every bus cycle until the bus transaction completes. At the completion of a bus transaction, the counter is synchronously reset to zero. If there are any outstanding bus requests, the aging counter will then begin counting immediately.

The AGE_CNT_THRESH is compared against the value of the aging counter during each clock cycle of the current transaction. If AGE_CNT_THRESH is equal to zero, priority state one is always chosen. If the aging counter is less than the AGE_CNT_THRESH value, priority state zero is selected. If the aging counter is greater than or equal to the AGE_CNT_THRESH value, priority state one is selected.

USB memory map/register definition

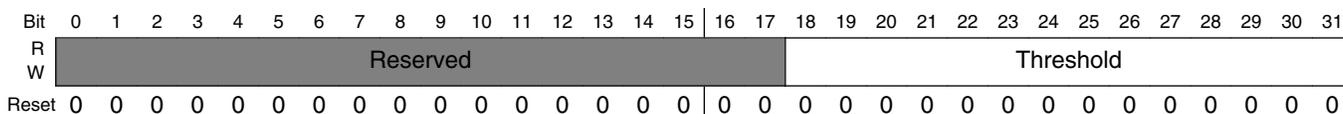
The two priority states of the aging counter function each have corresponding register bits which are programmed by the CPU. Thus, when the aging counter function is at priority state zero, PRI_CTRL[30-31] are selected and used to drive bus priority levels. When the aging counter function is at priority state one, PRI_CTRL[28-29] are selected and used to drive the priority.

The setting of AGE_CNT_THRESH is highly dependent on both the mix of other controllers operating on the system bus as well as the kind of traffic moving through the USB controller. A recommended approach is first to try leaving the aging mechanism disabled and see if the USB meets performance requirements. If USB performance does not meet application requirements, try the following settings:

- Set PRI_CTRL[pri_lvl0] to 0.
- Set PRI_CTRL[pri_lvl1] to 3.
- Set AGE_CNT_THRESH to 40.

This combination works for a wide variety of applications. If this combination still does not meet application requirements, try lowering AGE_CNT_THRESH by 5. On the contrary, the setting 40 may be too conservative for some applications. If USB performance is acceptable at 40, try raising the value in increments of 5. Raising AGE_CNT_THRESH benefits the other controllers on the system bus by reducing the frequency that this USB controller raises its priority to the arbiter.

Address: 2_2000h base + 408h offset = 2_2408h



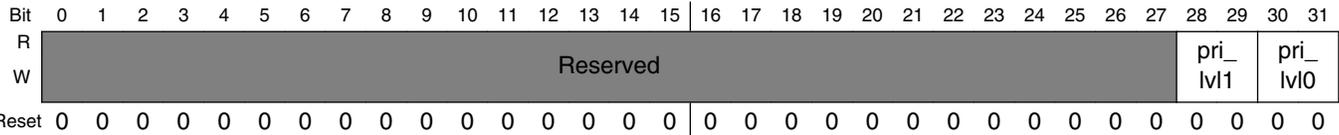
USB_AGE_CNT_THRESH field descriptions

Field	Description
0–17 -	This field is reserved. Reserved, should be cleared
18–31 Threshold	Aging counter threshold value.

18.3.31 Priority control (USB_PRI_CTRL)

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The priority control (PRI_CTRL) register sets the priority level for each of two priority states. The priority state is determined by the value programmed in the AGE_CNT_THRESH register and the number of *ccb_clk* cycles that a particular transaction takes to complete .

Address: 2_2000h base + 40Ch offset = 2_240Ch



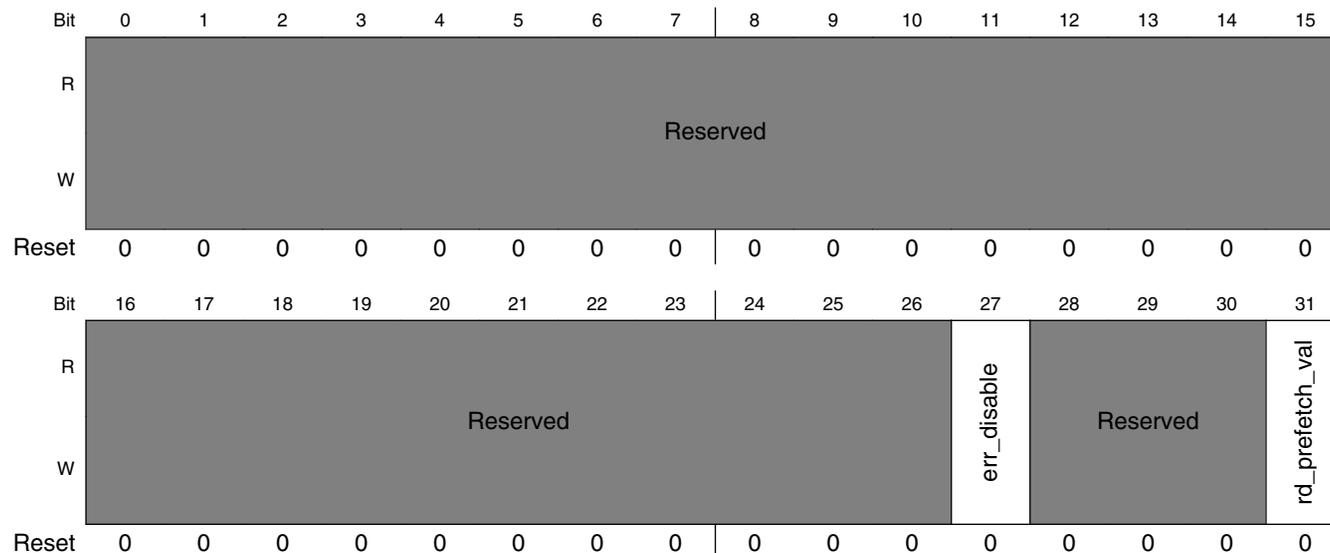
USB_PRI_CTRL field descriptions

Field	Description
0–27 -	This field is reserved. Reserved, should be cleared
28–29 pri_lvl1	Priority level for priority state 1. The highest priority is 2'h3 and the lowest priority is 2'b0.
30–31 pri_lvl0	Priority level for priority state 0. The highest priority is 2'h3 and the lowest priority is 2'b0.

18.3.32 System interface control (USB_SI_CTRL)

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The system interface control register (SI_CTRL) controls various functions pertaining to the internal system interface.

Address: 2_2000h base + 410h offset = 2_2410h



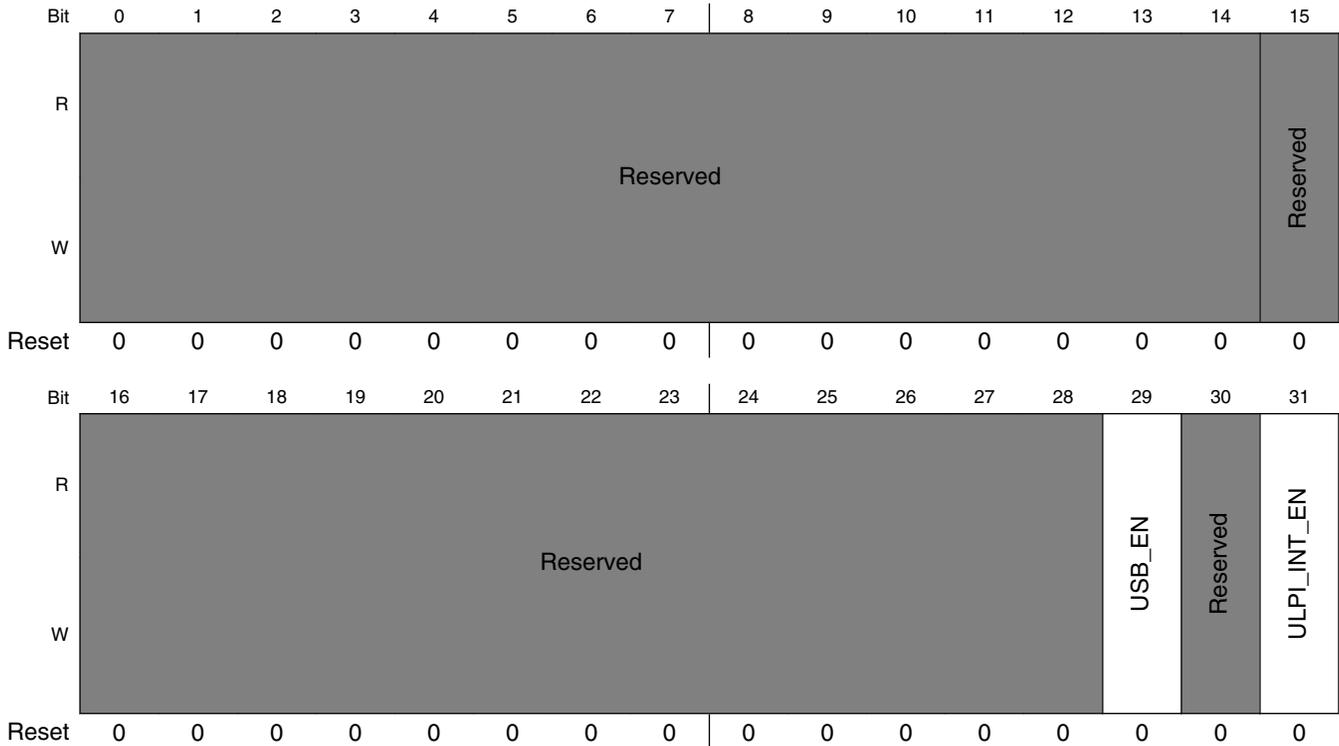
USB_SI_CTRL field descriptions

Field	Description
0–26 -	This field is reserved. Reserved, should be cleared
27 err_disable	When this bit is set, it causes the controller to ignore system bus errors. If cleared the controller responds according to the values set in USBSTS[SEI] and USBINT[SEE]. 0 enable 1 disable
28–30 -	This field is reserved. Reserved, should be cleared
31 rd_prefetch_val	Selects whether 32 bytes or 64 bytes are fetched during burst read transactions at the system interface. When this input is ZERO 64 bytes are fetched, and when it is ONE 32 bytes are fetched. The setting of rd_prefetch_val must match the setting of the larger of TXPBURST and RXPBURST fields in the BURSTSIZE register. If either of these fields is 64 bytes, then rd_prefetch_val must be left cleared. Otherwise, this value should be set. 0 64-byte fetch 1 32-byte fetch

18.3.33 Control (USB_CONTROL)

Note that this register uses big-endian byte ordering and is not defined in the EHCI specification. The USB general purpose (CONTROL) register contains the general-purpose IP control register outputs.

Address: 2_2000h base + 500h offset = 2_2500h



USB_CONTROL field descriptions

Field	Description
0–14 -	This field is reserved. Reserved
15 -	This field is reserved. Reserved
16–28 -	This field is reserved. Reserved
29 USB_EN	UTMI mode: This bit is used to enable the USB interface. It must be set before setting RS bit in USB CMD register. 1 Enable 0 Disable

Table continues on the next page...

USB_CONTROL field descriptions (continued)

Field	Description
	<p>ULPI mode: In safe mode, all USB interface signals are put into input mode or driven inactive, except for SUSPEND_STP, which is driven high. Also, the input signal DIR is forced to appear high to the controller. This prevents any start-up problems that otherwise could occur if the PHY and the controller take significantly different times to complete power-on reset.</p> <p>1 Normal operation 0 Safe mode</p>
30 -	<p>This field is reserved. Reserved</p>
31 ULPI_INT_EN	<p>Used to enable the ULPI low power wakeup interrupt from the PHY when the PHY is in low power mode only.</p> <p>NOTE: PORTSC[PHCD] bit must be set.</p> <p>0 ULPI low power wakeup interrupt disabled 1 ULPI low power wakeup interrupt enabled</p>

18.4 Functional description

The USB DR module can be broken down into functional sub-blocks, which are described below.

18.4.1 System interface

The system interface block contains all the control and status registers that allow a processor to interface to the USB DR module.

These registers allow the processor to control the configuration of the module, ascertain the capabilities of the module, and control the module's operation. It also has registers to control snoopability and priority of the DMA interface.

18.4.2 DMA engine

The module contains a local DMA engine.

The DMA engine interfaces internally to the CCB . It is responsible for moving all of the data to be transferred over the USB between the module and buffers in system memory.

Like the system interface block, the DMA engine block uses a simple synchronous bus signaling protocol that eases connections to a number of different standard buses.

The DMA controller must access both control information and packet data from system memory. The control information is contained in link list-based queue structures. The DMA controller has state machines that are able to parse data structures defined in the EHCI specification. In host mode, the data structures are EHCI compliant and represent queues of transfers to be performed by the host controller, including the split-transaction requests that allow an EHCI controller to direct packets to FS and LS devices. In device mode, the data structures are designed to be similar to those in the EHCI specification and are used to allow device responses to be queued for each of the active pipes in the device.

18.4.3 FIFO RAM controller

The FIFO RAM controller is used for context information and to control FIFOs between the protocol engine and the DMA controller.

These FIFOs decouple the system processor/memory bus requests from the extremely tight timing required by USB.

The use of the FIFO buffers differs between host and device mode operation. In host mode, a single data channel is maintained in each direction through the buffer memory. In device mode, multiple FIFO channels are maintained for each of the active endpoints in the system.

In host mode, the USB DR module uses a 512-byte Tx buffer and a 512-byte Rx buffer. Device operation uses a single 512-byte Rx buffer and a 512-byte Tx buffer for each endpoint. The 512-byte buffers allow the module to buffer a complete HS bulk packet.

18.4.4 PHY interface

The USB DR module interfaces to any or ULPI-compatible PHY.

The primary function of the port controller block is to isolate the rest of the module from the transceiver, and to move all of the transceiver signaling into the primary clock domain of the module.

This allows the module to run synchronously with the system processor and its associated resources.

Due to pincount limitations the module only supports certain combinations of PHY interfaces and USB functionality. Refer to the table below for more information.

Table 18-43. Supported PHY interfaces

PHY	Function
ULPI	Host/Device

18.5 Host data structures

This section defines the interface data structures used to communicate control, status, and data between HCD (software) and the Enhanced Host Controller (hardware).

The data structure definitions in this section support a 32-bit memory buffer address space. The interface consists of a periodic schedule, periodic frame list, asynchronous schedule, isochronous transaction descriptors, split-transaction isochronous transfer descriptors, queue heads, and queue element transfer descriptors.

The periodic frame list is the root of all periodic (isochronous and interrupt transfer type) support for the host controller interface. The asynchronous list is the root for all the bulk and control transfer type support. Isochronous data streams are managed using isochronous transaction descriptors. Isochronous split-transaction data streams are managed with split-transaction isochronous transfer descriptors. All interrupt, control, and bulk data streams are managed with queue heads and queue element transfer descriptors. These data structures are optimized to reduce the total memory footprint of the schedule and to reduce (on average) the number of memory accesses needed to execute a USB transaction.

Note that software must ensure that no interface data structure reachable by the EHCI host controller spans a 4K-page boundary.

The data structures defined in this section are (from the host controller's perspective) a mix of read-only and read/writable fields. The host controller must preserve the read-only fields on all data structure writes.

18.5.1 Periodic frame list

The figure below shows the organization of the periodic schedule. This schedule is for all periodic transfers (isochronous and interrupt).

The periodic schedule is referenced from the operational registers space using the PERIODICLISTBASE address register and the FRINDEX register. The periodic schedule is based on an array of pointers called the periodic frame list. The

Frame list link pointers always reference memory objects that are 32-byte aligned. The referenced object may be an isochronous transfer descriptor for high-speed devices, a split-transaction isochronous transfer descriptor (for full-speed isochronous endpoints), or a queue head (used to support high-, full- and low-speed interrupt). System software should not place non-periodic schedule items into the periodic schedule. The least-significant bits in a frame list pointer are used to key the host controller in as to the type of object the pointer is referencing.

The least-significant bit is the T bit (bit 0). When this bit is set, the host controller never uses the value of the frame list pointer as a physical memory pointer. The Typ field indicates the exact type of data structure being referenced by this pointer. The value encodings for the Typ field are given in the table below.

Table 18-45. Typ field encodings

Typ	Description
00	Isochronous transfer descriptor
01	Queue head
10	Split transaction isochronous transfer descriptor
11	Frame span traversal node

18.5.2 Asynchronous list queue head pointer

The asynchronous transfer list (based at the ASYNCLISTADDR register) is where all the control and bulk transfers are managed.

Host controllers use this list only when it reaches the end of the periodic list, the periodic list is disabled, or the periodic list is empty. The figure below shows the asynchronous schedule organization.

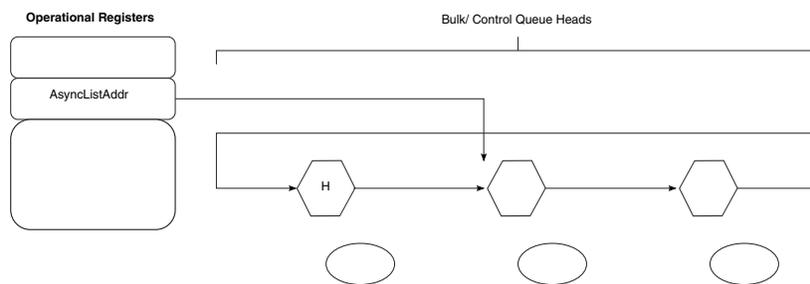


Figure 18-43. Asynchronous schedule organization

The asynchronous list is a simple circular list of queue heads. The ASYNCLISTADDR register is simply a pointer to the next queue head. This implements a pure round-robin service for all queue heads linked into the asynchronous list.

18.5.3 Isochronous (high-speed) transfer descriptor (iT D)

The table below shows the format of an isochronous transfer descriptor.

This structure is used only for high-speed isochronous endpoints.

All other transfer types should use queue structures. Isochronous TDs must be aligned on a 32-byte boundary.

Table 18-46. Isochronous transaction descriptor (iT D)

3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	15	1	1	1	11	1	9	8	7	6	5	4	3	2	1	0	offset
1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6		4	3	2		0										
Next link pointer																								00	Typ	T	0x00				
Status ¹		Transaction 0 Length ¹												ioc	PG ²	Transaction 0 Offset ²								0x04							
Status ¹		Transaction 1 length ¹												ioc	PG ²	Transaction 1 offset ²								0x08							
Status ¹		Transaction 2 length ¹												ioc	PG ²	Transaction 2 offset ²								0x0C							
Status ¹		Transaction 3 length ¹												ioc	PG ²	Transaction 3 offset ²								0x10							
Status ¹		Transaction 4 length ¹												ioc	PG ²	Transaction 4 offset ²								0x14							
Status ¹		Transaction 5 length ¹												ioc	PG ²	Transaction 5 offset ²								0x18							
Status ¹		Transaction 6 length ¹												ioc	PG ²	Transaction 6 offset ²								0x1C							
Status ¹		Transaction 7 length ¹												ioc	PG ²	Transaction 7 offset ²								0x20							
Buffer pointer (page 0)																		EndPt	R	Device address							0x24				
Buffer pointer (page 1)																		I/O	Maximum packet size							0x28					
Buffer pointer (Page 2)																		Reserved							Mult	0x2C					
Buffer pointer (page 3)																		Reserved							0x30						
Buffer pointer (page 4)																		Reserved							0x34						
Buffer pointer (page 5)																		Reserved							0x38						
Buffer pointer (page 6)																		Reserved							0x3C						

1. Host controller read/write; all others read-only.
2. These fields may be modified by the host controller if the I/O field indicates an OUT.

18.5.3.1 Next link pointer-iTD

The first DWord of an iTD is a pointer to the next schedule data structure, as shown in the table below.

Table 18-47. Next schedule element pointer

Bits	Name	Description
31-5	Link Pointer	Correspond to memory address signals [31:5], respectively. This field points to another isochronous transaction descriptor (iTD/siTD) or queue head (QH).
4-3	-	Reserved, should be cleared. These bits are reserved and their value has no effect on operation. Software should initialize this field to zero.
2-1	Typ	Indicates to the host controller whether the item referenced is an iTD, siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. Value encodings are: 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)
0	T	Terminate 1 Link Pointer field is not valid. 0 Link Pointer field is valid.

18.5.3.2 iTD transaction status and control list

DWords 1-8 constitute eight slots of transaction control and status.

Each transaction description includes the following:

- Status results field
- Transaction length (bytes to send for OUT transactions and bytes received for IN transactions).
- Buffer offset. The PG and Transaction *n* Offset fields are used with the buffer pointer list to construct the starting buffer address for the transaction.

The host controller uses the information in each transaction description, plus the endpoint information contained in the first three DWords of the buffer page pointer list, to execute a transaction on the USB.

The table below shows the iTD transaction status and control fields.

Table 18-48. iTD transaction status and control

Bits	Name	Description
31-28	Status	<p>Records the status of the transaction executed by the host controller for this slot. This field is a bit vector with the following encoding:</p> <p>31 Active. Set by software to enable the execution of an isochronous transaction by the host controller. When the transaction associated with this descriptor is completed, the host controller clears this bit indicating that a transaction for this element should not be executed when it is next encountered in the schedule.</p> <p>30 Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overflow) or is unable to supply data fast enough during transmission (underflow). If an overflow condition occurs, no action is necessary.</p> <p>29 Babble detected. Set by the host controller during status update when "babble" is detected during the transaction generated by this descriptor.</p> <p>28 Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (Time-out, CRC, Bad PID, and so on). This bit may only be set for isochronous IN transactions.</p>
27-16	Transaction <i>n</i> Length	<p>For an OUT, this field is the number of data bytes the host controller will send during the transaction. The host controller is not required to update this field to reflect the actual number of bytes transferred during the transfer. For an IN, the initial value of the endpoint to deliver. During the status update, the host controller writes back the field is the number of bytes the host expects the number of bytes successfully received. The value in this register is the actual byte count (for example, 0 zero length data, 1 one byte, 2 two bytes, and so on). The maximum value this field may contain is 0xC00 (3072).</p>
15	ioc	<p>Interrupt on complete. If this bit is set, it specifies that when this transaction completes, the host controller should issue an interrupt at the next interrupt threshold.</p>
14-12	PG	<p>These bits are set by software to indicate which of the buffer page pointers the offset field in this slot should be concatenated to produce the starting memory address for this transaction. The valid range of values for this field is 0 to 6.</p>
11-0	Transaction <i>n</i> Offset	<p>This field is a value that is an offset, expressed in bytes, from the beginning of a buffer. This field is concatenated onto the buffer page pointer indicated in the adjacent PG field to produce the starting buffer address for this transaction.</p>

18.5.3.3 iTD buffer page pointer list (plus)

DWords 9-15 of an isochronous transaction descriptor are nominally page pointers (4K aligned) to the data buffer for this transfer descriptor.

This data structure requires the associated data buffer to be contiguous (relative to virtual memory), but allows the physical memory pages to be non-contiguous. Seven page pointers are provided to support the expression of eight isochronous transfers. The seven pointers allow for 3 (transactions) x 1024 (maximum packet size) x 8 (transaction records) = 24 576 bytes to be moved with this data structure, regardless of the alignment offset of the first page.

Since each pointer is a 4 K-aligned page pointer, the least-significant 12 bits in several of the page pointers are used for other purposes.

The following tables describe buffer pointer page *n*.

Table 18-49. Buffer pointer page 0 (Plus)

Bits	Name	Description
31-12	Buffer pointer (Page 0)	A 4K-aligned pointer to physical memory. Corresponds to memory address bits 31-12.
11-8	EndPt	Selects the particular endpoint number on the device serving as the data source or sink.
7	-	Reserved, should be cleared. Reserved for future use and should be initialized by software to zero.
6-0	Device Address	This field selects the specific device serving as the data source or sink.

Table 18-50. iTD buffer pointer page 1 (plus)

Bits	Name	Description
31-12	Buffer Pointer (Page 1)	This is a 4K aligned pointer to physical memory. Corresponds to memory address bits 31-12.
11	I/O	Direction (I/O). This field encodes whether the high-speed transaction should use an IN or OUT PID. 0 OUT 1 IN
10-0	Maximum Packet Size	This directly corresponds to the maximum packet size of the associated endpoint (wMaxPacketSize). This field is used for high-bandwidth endpoints where more than one transaction is issued per transaction description (for example, per microframe). This field is used with the Multi field to support high-bandwidth pipes. This field is also used for all IN transfers to detect packet babble. Software should not set a value larger than 1024 (0x400). Any value larger yields undefined results.

Table 18-51. Buffer pointer page 2 (plus)

Bits	Name	Description
31-12	Buffer pointer (page 2)	This is a 4K-aligned pointer to physical memory. Corresponds to memory address bits 31-12.
11-2	-	Reserved, should be cleared. This bit reserved for future use and should be cleared.
1-0	Mult	Indicates to the host controller the number of transactions that should be executed per transaction description (for example, per microframe). 00 Reserved, should be cleared. A zero in this field yields undefined results. 01 One transaction to be issued for this endpoint per microframe 10 Two transactions to be issued for this endpoint per microframe 11 Three transactions to be issued for this endpoint per microframe

Table 18-52. Buffer pointer page 3-6

Bits	Name	Description
31-12	Buffer Pointer	This is a 4K aligned pointer to physical memory. Corresponds to memory address bits 31-12.
11-0	-	Reserved, should be cleared. These bits reserved for future use and should be cleared.

18.5.4 Split transaction isochronous transfer descriptor (siTD)

All full-speed isochronous transfers through the internal transaction translator are managed using the siTD data structure.

This data structure satisfies the operational requirements for managing the split transaction protocol.

The table below shows the split-transaction isochronous transfer descriptor (siTD).

Table 18-53. Split-transaction isochronous transaction descriptor (siTD)

31	3	2	2	2	2	2	2	2	2	2	2	1	1	1	1	15	1	1	1	1	1	9	8	7	6	5	4	3	2	1	0	offset
Next Link Pointer																										00	Typ	T	0x00			
I/O	Port Number			0	Hub Address				0000				EndPt		0	Device Address				0x04												
0000_0000_0000_00000										μFrame C-mask						μFrame S-mask						0x08										
ioc	P ¹	0000			Total Bytes to Transfer ¹				μFrame C-prog-mask ¹				Status ¹				0x0C															
Buffer Pointer (Page 0)										Current Offset ¹						0x10																
Buffer Pointer (Page 1)										000_0000				TP ¹		T-count ¹		0x14														
Back Pointer																						0000				T		0x18				

1. Host controller read/write; all others read-only.

18.5.4.1 Next link pointer-siTD

DWord0 of a siTD is a pointer to the next schedule data structure.

The table below describes the next link pointer fields.

Table 18-54. Next link pointer

Bits	Name	Description
31-5	Next Link Pointer	This field contains the address of the next data object to be processed in the periodic list and corresponds to memory address signals [31:5], respectively.
4-3	-	Reserved, should be cleared. These bits must be written as zeros.
2-1	Typ	Indicates to the host controller whether the item referenced is an iTD/siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. Value encodings are: 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)

Table continues on the next page...

Table 18-54. Next link pointer (continued)

Bits	Name	Description
0	T	Terminate. 0 Link pointer is valid. 1 Link pointer field is not valid.

18.5.4.2 siTD endpoint capabilities/characteristics

DWords 1 and 2 specify static information about the full-speed endpoint, the addressing of the parent Companion Controller, and microframe scheduling control.

The table below describes the endpoint and transaction translator characteristics.

Table 18-55. Endpoint and transaction translator characteristics

Bits	Name	Description
31	I/O	Direction (I/O). This field encodes whether the full-speed transaction should be an IN or OUT. 0 OUT 1 IN
30-24	Port Number	This field is the port number of the recipient transaction translator.
23	-	Reserved, should be cleared. Bit reserved and should be cleared.
22-16	Hub Address	This field holds the device address of the companion controllers' hub.
15-12	-	Reserved, should be cleared. Field reserved and should be cleared.
11-8	EndPt	Endpoint Number. Selects the particular endpoint number on the device serving as the data source or sink.
7	-	Reserved, should be cleared. Bit is reserved for future use. It should be cleared.
6-0	Device Address	Selects the specific device serving as the data source or sink.

The table below describes the microframe schedule control.

Table 18-56. Microframe schedule control

Bits	Name	Description
31-16	-	Reserved, should be cleared. This field reserved for future use. It should be cleared.
15-8	µFrame C-mask	Split completion mask. This field (along with the Active and SplitX- state fields in the status byte) is used to determine during which microframes the host controller should execute complete-split transactions. When the criteria for using this field is met, an all-zeros value has undefined behavior. The host controller uses the value of the three low-order bits of the FRINDEX register to index into this bit field. If the FRINDEX register value indexes to a position where the µFrame C-Mask field is a one, this siTD is a candidate for transaction execution. There may be more than one bit in this mask set.

Table continues on the next page...

Table 18-56. Microframe schedule control (continued)

Bits	Name	Description
7-0	μFrame S-mask	Split start mask. This field (along with the Active and SplitX-state fields in the Status byte) is used to determine during which microframes the host controller should execute start-split transactions. The host controller uses the value of the three low-order bits of the FRINDEX register to index into this bit field. If the FRINDEX register value indexes to a position where the μFrame S-mask field is a one, then this siTD is a candidate for transaction execution. An all zeros value in this field, in combination with existing periodic frame list has undefined results.

18.5.4.3 siTD transfer state

DWords 3-6 manage the state of the transfer, as described in the table below.

Table 18-57. siTD transfer status and control

Bits	Name	Description
31	ioc	Interrupt on complete 0 Do not interrupt when transaction is complete. 1 Do interrupt when transaction is complete. When the host controller determines that the split transaction has completed it will assert a hardware interrupt at the next interrupt threshold.
30	P	Page select. Indicates which data page pointer should be concatenated with the CurrentOffset field to construct a data buffer pointer 0 Selects Page 0 pointer 1 Selects Page 1 pointer The host controller is not required to write this field back when the siTD is retired (Active bit transitioned from a one to a zero).
29-26	-	Reserved, should be cleared. This field reserved for future use and should be cleared.
25-16	Total Bytes to Transfer	This field is initialized by software to the total number of bytes expected in this transfer. Maximum value is 1023 (3FFh)
15-8	μFrame C-prog-mask	Split complete progress mask. This field is used by the host controller to record which split-completes have been executed.

Table continues on the next page...

Table 18-57. siTD transfer status and control (continued)

Bits	Name	Description	
7-0	Status	This field records the status of the transaction executed by the host controller for this slot. This field is a bit vector with the following encoding:	
		Status Bits	Definition
		7	Active. Set by software to enable the execution of an isochronous split transaction by the host controller.
		6	ERR. Set by the host controller when an ERR response is received from the companion controller.
		5	Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). In the case of an under run, the host controller will transmit an incorrect CRC (thus invalidating the data at the endpoint). If an overrun condition occurs, no action is necessary.
		4	Babble detected. Set by the host controller during status update when "babble" is detected during the transaction generated by this descriptor.
		3	Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (Time-out, CRC, Bad PID, and so on). This bit will only be set for IN transactions.
		2	Missed microframe. The host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction.
		1	Split transaction state (SplitXstate). The bit encodings are: 0 Do start split. This value directs the host controller to issue a Start split transaction to the endpoint when a match is encountered in the S-mask. 1 Do complete split. This value directs the host controller to issue a Complete split transaction to the endpoint when a match is encountered in the C-mask.
0	Reserved, should be cleared. Bit reserved for future use and should be cleared.		

18.5.4.4 siTD buffer pointer list (plus)

DWords 4 and 5 are the data buffer page pointers for the transfer.

This structure supports one physical page cross. The most-significant 20 bits of each DWord in this section are the 4K (page) aligned buffer pointers. The least-significant 12 bits of each DWord are used as additional transfer state.

The table below describes the siTD buffer pointer page 0.

Table 18-58. siTD buffer pointer page 0 (plus)

Bits	Name	Description
31-12	Buffer Pointer (Page 0)	Bits 31-12 are 4K page-aligned, physical memory addresses. These bits correspond to physical address bits 31-12 respectively. The field P specifies the current active pointer
11-0	Current Offset	The 12 least-significant bits of the Page 0 pointer is the current byte offset for the current page pointer (as selected with the page indicator bit (P field)). The host controller is not required to write this field back when the siTD is retired (Active bit transitioned from a one to a zero).

The table below describes the siTD buffer pointer page 1.

Table 18-59. siTD buffer pointer page 1 (plus)

Bits	Name	Description
31-12	Buffer Pointer (Page 1)	Bits 31-12 are 4K page-aligned, physical memory addresses. These bits correspond to physical address bits 31-12 respectively. The field P specifies the current active pointer
11-5	-	Reserved, should be cleared.
4-3	TP	Transaction position. This field is used with T-count to determine whether to send all, first, middle, or last with each outbound transaction payload. System software must initialize this field with the appropriate starting value. The host controller must correctly manage this state during the lifetime of the transfer. The bit encodings are: 00 All. The entire full-speed transaction data payload is in this transaction (that is, less than or equal to 188 bytes). 01 Begin. This is the first data payload for a full-speed transaction that is greater than 188 bytes. 10 Mid. This is the middle payload for a full-speed OUT transaction that is larger than 188 bytes. 11 End. This is the last payload for a full-speed OUT transaction that was larger than 188 bytes.
2-0	T-Count	Transaction count. Software initializes this field with the number of OUT start-splits this transfer requires. Any value larger than 6 is undefined.

18.5.4.5 siTD back link pointer

DWord 6 of a siTD is simply another schedule link pointer.

This pointer is always zero, or references a siTD. This pointer cannot reference any other schedule data structure.

The table below describes the siTD back link pointer.

Table 18-60. siTD back link pointer

Bits	Name	Description
31-5	Back Pointer	A physical memory pointer to an siTD
4-1	-	Reserved, should be cleared. This field is reserved for future use. It should be cleared.
0	T	Terminate 0 siTD Back Pointer field is valid 1 siTD Back Pointer field is not valid

18.5.5 Queue element transfer descriptor (qTD)

This data structure is only used with a queue head.

This data structure is used for one or more USB transactions. This data structure is used to transfer up to 20,480 (5 x 4096) bytes. The structure contains two structure pointers used for queue advancement, a DWord of transfer state, and a five-element array of data buffer pointers. This structure is 32 bytes (or one 32-byte cache line). This data structure must be physically contiguous.

The buffer associated with this transfer must be virtually contiguous. The buffer may start on any byte boundary. A separate buffer pointer list element must be used for each physical page in the buffer, regardless of whether the buffer is physically contiguous.

Host controller updates (host controller writes) to stand-alone qTDs only occur during transfer retirement. References in the following bit field definitions of updates to the qTD are to the qTD portion of a queue head.

The table below shows the queue element transfer descriptors.

Table 18-61. Queue element transfer descriptor (qTD)

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Next qTD Pointer																												0000	T	0x00		
Alternate Next qTD Pointer																												0000	T	0x04		
dt ¹ Total Bytes to Transfer ¹										ioc		C_Page ¹		Cerr ¹	PID Code	Status ¹										0x08						
Buffer Pointer (Page 0)										Current Offset ¹										0x0C												
Buffer Pointer (Page 1)										0000_0000_0000										0x10												
Buffer Pointer (Page 2)										0000_0000_0000										0x14												
Buffer Pointer (Page 3)										0000_0000_0000										0x18												
Buffer Pointer (Page 4)										0000_0000_0000										0x1C												

1. Host controller read/write; all others read-only.

Queue element transfer descriptors must be aligned on 32-byte boundaries.

18.5.5.1 Next qTD pointer

The first DWord of an element transfer descriptor is a pointer to another transfer element descriptor.

The table below describes the qTD next element transfer pointer.

Table 18-62. qTD next element transfer pointer (DWord 0)

Bits	Name	Description
31-5	Next qTD Pointer	This field contains the physical memory address of the next qTD to be processed and corresponds to memory address signals [31:5], respectively.
4-1	-	Reserved, should be cleared. These bits are reserved and their value has no effect on operation.

Table continues on the next page...

Table 18-62. qTD next element transfer pointer (DWord 0) (continued)

Bits	Name	Description
0	T	Terminate. Indicates to the host controller that there are no more valid entries in the queue. 0 Pointer is valid (points to a valid transfer element descriptor) 1 Pointer is invalid

18.5.5.2 Alternate next qTD pointer

The second DWord of a queue element transfer descriptor is used to support hardware-only advance of the data stream to the next client buffer on short packet.

To be more explicit the host controller will always use this pointer when the current qTD is retired due to short packet. The table below describes the alternate qTD next element transfer pointer.

Table 18-63. qTD alternate next element transfer pointer (DWord 1)

Bits	Name	Description
31-5	Alternate next qTD pointer	This field contains the physical memory address of the next qTD to be processed in the event that the current qTD execution encounters a short packet (for an IN transaction). The field corresponds to memory address signals [31:5], respectively.
4-1	-	Reserved, should be cleared. These bits are reserved and their value has no effect on operation.
0	T	Terminate. Indicates to the host controller that there are no more valid entries in the queue. 0 Pointer is valid (points to a valid transfer element descriptor) 1 Pointer is invalid

18.5.5.3 qTD token

The third DWord of a queue element transfer descriptor contains most of the information the host controller requires to execute a USB transaction (the remaining endpoint-addressing information is specified in the queue head).

Note that some of the field descriptions in the table below reference fields are defined in the queue head. See [Queue head](#), for more information on these fields.

Table 18-64. qTD token (DWord 2)

Bits	Name	Description
31	dt	Data toggle. This is the data toggle sequence bit. The use of this bit depends on the setting of the Data Toggle Control bit in the queue head.

Table continues on the next page...

Table 18-64. qTD token (DWord 2) (continued)

Bits	Name	Description	
30-16	Total bytes to transfer	Total bytes to transfer. This field specifies the total number of bytes to be moved with this transfer descriptor. This field is decremented by the number of bytes actually moved during the transaction, only on the successful completion of the transaction. The maximum value software may store in this field is 5 x 4K (0x5000). This is the maximum number of bytes 5 page pointers can access. If the value of this field is zero when the host controller fetches this transfer descriptor (and the active bit is set), the host controller executes a zero-length transaction and retires the transfer descriptor. It is not a requirement for OUT transfers that total bytes to transfer be an even multiple of QH[Maximum Packet Length]. If software builds such a transfer descriptor for an OUT transfer, the last transaction will always be less than QH[Maximum Packet Length]. Although it is possible to create a transfer up to 20K this assumes the page is 0. When the offset cannot be predetermined, crossing past the 5th page can be guaranteed by limiting the total bytes to 16K. Therefore, the maximum recommended transfer is 16K (0x4000).	
15	ioc	Interrupt on complete. If this bit is set, the host controller should issue an interrupt at the next interrupt threshold when this qTD is completed.	
14-12	C_Page	Current range. This field is used as an index into the qTD buffer pointer list. Valid values are in the range 0x0 to 0x4. The host controller is not required to write this field back when the qTD is retired.	
11-10	Cerr	Error counter. 2-bit down counter that keeps track of the number of consecutive errors detected while executing this qTD. If this field is programmed with a non-zero value during setup, the host controller decrements the count and writes it back to the qTD if the transaction fails. If the counter counts from one to zero, the host controller marks the qTD inactive, sets the Halted bit to a one, and error status bit for the error that caused Cerr to decrement to zero. An interrupt is generated if USBINTR[UEE] is set. If the host controller driver (HCD) software programs this field to zero during setup, the host controller will not count errors for this qTD and there is no limit on the retries of this qTD. Note that write-backs of intermediate execution state are to the queue head overlay area, not the qTD.	
		Error	Decrement Counter
		Transaction Error	Yes
		Data Buffer Error	No. Data buffer errors are host problems. They don't count against the device's retries. Note that software must not program Cerr to a value of zero when the EPS field is programmed with a value indicating a full- or low-speed device. This combination could result in undefined behavior.
		Stalled	No. Detection of babble or stall automatically halts the queue head. Thus, count is not decremented
		Babble Detected	No. Detection of babble or stall automatically halts the queue head. Thus, count is not decremented
		No Error	No. If the EPS field indicates a HS device or the queue head is in the asynchronous schedule (and PIDCode indicates an IN or OUT) and a bus transaction completes and the host controller does not detect a transaction error, then the host controller should reset Cerr to extend the total number of errors for this transaction. For example, Cerr should be reset with maximum value (0b11) on each successful completion of a transaction. The host controller must never reset this field if the value at the start of the transaction is 0b00.

Table continues on the next page...

Table 18-64. qTD token (DWord 2) (continued)

Bits	Name	Description														
9-8	PID Code	This field is an encoding of the token, which should be used for transactions associated with this transfer descriptor. Encodings are: 00 OUT Token generates token (E1H) 01 IN Token generates token (69H) 10 SETUP Token generates token (2DH) (undefined if endpoint is an Interrupt transfer type, for example. μ Frame S-mask field in the queue head is non-zero.) 11 Reserved, should be cleared														
7-0	Status	This field is used by the host controller to communicate individual command execution states back to the host controller driver (HCD) software. This field contains the status of the last transaction performed on this qTD. The bit encodings are:														
		<table border="1"> <thead> <tr> <th>Bits</th> <th>Status Field Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Active. Set by software to enable the execution of transactions by the host controller.</td> </tr> <tr> <td>6</td> <td>Halted. Set by the host controller during status updates to indicate that a serious error has occurred at the device/endpoint addressed by this qTD. This can be caused by babble, the error counter counting down to zero, or reception of the STALL handshake from the device during a transaction. Any time that a transaction results in the Halted bit being set, the Active bit is also cleared.</td> </tr> <tr> <td>5</td> <td>Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). If an overrun condition occurs, the host controller will force a time-out condition on the USB, invalidating the transaction at the source. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.</td> </tr> <tr> <td>4</td> <td>Babble detected. Set by the host controller during status update when babble is detected during the transaction. In addition to setting this bit, the host controller also sets the Halted bit to a one. Since babble is considered a fatal error for the transfer, setting the Halted bit to a one insures that no more transactions occur because of this descriptor.</td> </tr> <tr> <td>3</td> <td>Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (time-out, CRC, bad PID). If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.</td> </tr> <tr> <td>2</td> <td>Missed microframe. This bit is ignored unless the QH[EPS] field indicates a full- or low-speed endpoint and the queue head is in the periodic list. This bit is set when the host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.</td> </tr> </tbody> </table>	Bits	Status Field Description	7	Active. Set by software to enable the execution of transactions by the host controller.	6	Halted. Set by the host controller during status updates to indicate that a serious error has occurred at the device/endpoint addressed by this qTD. This can be caused by babble, the error counter counting down to zero, or reception of the STALL handshake from the device during a transaction. Any time that a transaction results in the Halted bit being set, the Active bit is also cleared.	5	Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). If an overrun condition occurs, the host controller will force a time-out condition on the USB, invalidating the transaction at the source. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.	4	Babble detected. Set by the host controller during status update when babble is detected during the transaction. In addition to setting this bit, the host controller also sets the Halted bit to a one. Since babble is considered a fatal error for the transfer, setting the Halted bit to a one insures that no more transactions occur because of this descriptor.	3	Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (time-out, CRC, bad PID). If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.	2	Missed microframe. This bit is ignored unless the QH[EPS] field indicates a full- or low-speed endpoint and the queue head is in the periodic list. This bit is set when the host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.
Bits	Status Field Description															
7	Active. Set by software to enable the execution of transactions by the host controller.															
6	Halted. Set by the host controller during status updates to indicate that a serious error has occurred at the device/endpoint addressed by this qTD. This can be caused by babble, the error counter counting down to zero, or reception of the STALL handshake from the device during a transaction. Any time that a transaction results in the Halted bit being set, the Active bit is also cleared.															
5	Data buffer error. Set by the host controller during status update to indicate that the host controller is unable to keep up with the reception of incoming data (overrun) or is unable to supply data fast enough during transmission (under run). If an overrun condition occurs, the host controller will force a time-out condition on the USB, invalidating the transaction at the source. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.															
4	Babble detected. Set by the host controller during status update when babble is detected during the transaction. In addition to setting this bit, the host controller also sets the Halted bit to a one. Since babble is considered a fatal error for the transfer, setting the Halted bit to a one insures that no more transactions occur because of this descriptor.															
3	Transaction error (XactErr). Set by the host controller during status update in the case where the host did not receive a valid response from the device (time-out, CRC, bad PID). If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.															
2	Missed microframe. This bit is ignored unless the QH[EPS] field indicates a full- or low-speed endpoint and the queue head is in the periodic list. This bit is set when the host controller detected that a host-induced hold-off caused the host controller to miss a required complete-split transaction. If the host controller sets this bit to a one, then it remains a one for the duration of the transfer.															

Table continues on the next page...

Table 18-64. qTD token (DWord 2) (continued)

Bits	Name	Description
		<p>1</p> <p>Split transaction state (SplitXstate). This bit is ignored by the host controller unless the QH[EPS] field indicates a full- or low-speed endpoint. When a full- or low-speed device, the host controller uses this bit to track the state of the split- transaction. The functional requirements of the host controller for managing this state bit and the split transaction protocol depends on whether the endpoint is in the periodic or asynchronous schedule. The bit encodings are:</p> <p>0 Do start split. This value directs the host controller to issue a start split transaction to the endpoint.</p> <p>1 Do complete split. This value directs the host controller to issue a Complete split transaction to the endpoint.</p>
		<p>0</p> <p>Ping state (P)/ERR. If the QH[EPS] field indicates a high-speed device and the PID Code indicates an OUT endpoint, then this is the state bit for the Ping protocol. The bit encodings are:</p> <p>0 Do OUT. This value directs the host controller to issue an OUT PID to the endpoint.</p> <p>1 Do Ping. This value directs the host controller to issue a PING PID to the endpoint.</p> <p>If the QH[EPS] field does not indicate a high-speed device, then this field is used as an error indicator bit. It is set by the host controller whenever a periodic split-transaction receives an ERR handshake.</p>

18.5.5.4 qTD buffer page pointer list

The last five DWords of a queue element transfer descriptor make up an array of physical memory address pointers.

These pointers reference the individual pages of a data buffer.

System software initializes the Current Offset field to the starting offset into the current page, where current page is selected with the value in the C_Page field.

The table below describes the qTD buffer pointer.

Table 18-65. qTD buffer pointer

Bits	Name	Description
31-12	Buffer pointer (page <i>n</i>)	Each element in the list is a 4K page aligned physical memory address. The lower 12 bits in each pointer are reserved (except for the first one), as each memory pointer must reference the start of a 4K page. The field C_Page specifies the current active pointer. When the transfer element descriptor is fetched, the starting buffer address is selected using C_Page (similar to an array index to select an array element). If a transaction spans a 4K buffer boundary, the host controller must detect the page-span boundary in the data stream, increment C_Page and advance to the next buffer pointer in the list, and conclude the transaction via the new buffer pointer.

Table continues on the next page...

Table 18-65. qTD buffer pointer (continued)

Bits	Name	Description
11-0	Current offset (Page 0)/ - (pages 1-4)	Reserved in all pointers except the first one (that is, page 0). The host controller should ignore all reserved bits. For the page 0 current offset interpretation, this field is the byte offset into the current page (as selected by C_Page). The host controller is not required to write this field back when the qTD is retired. Software should ensure the reserved fields are initialized to zeros.

18.5.6 Queue head

The table below shows the queue head structure.

Table 18-66. Queue head layout

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset						
Queue Head Horizontal Link Pointer																											00	Typ	T	0x00								
RL		C	Maximum Packet Length				H	drc	EPS	EndPt		I	Device Address				0x04																					
Mult		Port Number			Hub Addr			μFrame C-mask				μFrame S-mask				0x08 ¹																						
Current qTD Pointer																											00000				0x0C							
Next qTD Pointer ²																											0000				T ²				0x10			
Alternate Next qTD Pointer ²																											NakCnt ²				T ²				0x14 ^{3,4}			
dt ⁵		Total Bytes to Transfer ²				ioc ²		C_Page ²		Cerr ²		PID Code ²		Status ²				0x18 ^{3,4}																				
Buffer Pointer (Page 0) ²										Current Offset ²						0x1C ^{3,4}																						
Buffer Pointer (Page 1) ²										0000				C-prog-mask ²						0x20 ^{3,4}																		
Buffer Pointer (Page 2) ²										S-bytes ²						FrameTag ²				0x24 ^{3,4}																		
Buffer Pointer (Page 3) ²										0000_0000_0000						0x28 ³																						
Buffer Pointer (Page 4) ²										0000_0000_0000						0x2C ³																						

1. Offsets 0x04 through 0x0B contain the static endpoint state.
2. Host controller read/write; all others read-only.
3. Offsets 0x10 through 0x2F contain the transfer overlay.
4. Offsets 0x14 through 0x27 contain the transfer results.
5. Host controller read/write; all others read-only.

18.5.6.1 Queue head horizontal link pointer

The first DWord of a queue head contains a link pointer to the next data object to be processed after any required processing in this queue has been completed, as well as the control bits defined below.

This pointer may reference a queue head or one of the isochronous transfer descriptors. It must not reference a queue element transfer descriptor.

The table below describes the queue head.

Table 18-67. Queue head DWord 0

Bits	Name	Description
31-5	QHLP	Queue head horizontal link pointer. This field contains the address of the next data object to be processed in the horizontal list and corresponds to memory address signals [31:5], respectively.
4-3	-	Reserved, should be cleared. These bits must be written as zeros.
2-1	Typ	Indicates to the hardware whether the item referenced by the link pointer is an iTD, siTD or a QH. This allows the host controller to perform the proper type of processing on the item after it is fetched. 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)
0	T	Terminate. 1 Last QH (pointer is invalid). 0 Pointer is valid. If the queue head is in the context of the periodic list, a one bit in this field indicates to the host controller that this is the end of the periodic list. This bit is ignored by the host controller when the queue head is in the asynchronous schedule. Software must ensure that queue heads reachable by the host controller always have valid horizontal link pointers.

18.5.6.2 Endpoint capabilities/characteristics

The second and third DWords of a queue head specify static information about the endpoint.

This information does not change over the lifetime of the endpoint. There are three types of information in this region:

- Endpoint characteristics. These are the USB endpoint characteristics, which include addressing, maximum packet size, and endpoint speed.
- Endpoint capabilities. These are adjustable parameters of the endpoint. They affect how the endpoint data stream is managed by the host controller.
- Split transaction characteristics. This data structure manages full- and low-speed data streams for bulk, control, and interrupt with split transactions to USB 2.0 Hub transaction translator. Additional fields exist for addressing the hub and scheduling the protocol transactions (for periodic).

The host controller must not modify the bits in this region.

The following tables describe the endpoint characteristics.

Table 18-68. Endpoint characteristics: Queue head DWord 1

Bits	Name	Description
31-28	RL	Nak count reload. This field contains a value, which is used by the host controller to reload the Nak Counter field.
27	C	Control endpoint flag. If the QH[EPS] field indicates the endpoint is not a high-speed device, and the endpoint is a control endpoint, then software must set this bit to a one. Otherwise, it should always set this bit to a zero.
26-16	Maximum packet length	This directly corresponds to the maximum packet size of the associated endpoint (wMaxPacketSize). The maximum value this field may contain is 0x400 (1024).
15	H	Head of reclamation list flag. This bit is set by system software to mark a queue head as being the head of the reclamation list.
14	dtc	Data toggle control (DTC). Specifies where the host controller should get the initial data toggle on an overlay transition. 0 Ignore DT bit from incoming qTD. Host controller preserves DT bit in the queue head. 1 Initial data toggle comes from incoming qTD DT bit. Host controller replaces DT bit in the queue head from the DT bit in the qTD.
13-12	EPS	Endpoint speed. This is the speed of the associated endpoint. 00 Full-speed (12 Mbps) 01 Low-speed (1.5 Mbps) 10 High-speed (480 Mbps) 11 Reserved, should be cleared This field must not be modified by the host controller.
11-8	EndPt	Endpoint number. Selects the particular endpoint number on the device serving as the data source or sink.
7	I	Inactivate on next transaction. This bit is used by system software to request that the host controller set the Active bit to zero. This field is only valid when the queue head is in the periodic schedule and the EPS field indicates a full- or low-speed endpoint. Setting this bit when the queue head is in the asynchronous schedule or the EPS field indicates a high-speed device yields undefined results.
6-0	Device address	Selects the specific device serving as the data source or sink.

Table 18-69. Endpoint capabilities: Queue head DWord 2

Bits	Name	Description
31-30	Mult	High-bandwidth pipe multiplier. This field is a multiplier used to key the host controller as the number of successive packets the host controller may submit to the endpoint in the current execution. The host controller makes the simplifying assumption that software properly initializes this field (regardless of location of queue head in the schedules or other run time parameters). 00 Reserved, should be cleared. A zero in this field yields undefined results. 01 One transaction to be issued for this endpoint per microframe 10 Two transactions to be issued for this endpoint per microframe 11 Three transactions to be issued for this endpoint per microframe
29-23	Port number	This field is ignored by the host controller unless the EPS field indicates a full- or low-speed device. The value is the port number identifier on the USB 2.0 hub (for hub at device address Hub Addr below), below which the full- or low-speed device associated with this endpoint is attached. This information is used in the split-transaction protocol.

Table continues on the next page...

Table 18-69. Endpoint capabilities: Queue head DWord 2 (continued)

Bits	Name	Description
22-16	Hub addr	This field is ignored by the host controller unless the EPS field indicates a full-or low-speed device. The value is the USB device address of the USB 2.0 hub below which the full- or low-speed device associated with this endpoint is attached. This field is used in the split-transaction protocol.
15-8	μFrame C-mask	This field is ignored by the host controller unless the EPS field indicates this device is a low- or full-speed device and this queue head is in the periodic list. This field (along with the Active and SplitX-state fields) is used to determine during which microframes the host controller should execute a complete-split transaction. When the criteria for using this field are met, a zero value in this field has undefined behavior. This field is used by the host controller to match against the three low-order bits of the FRINDEX register. If the FRINDEX register bits decode to a position where the μFrame C- mask field is a one, then this queue head is a candidate for transaction execution. There may be more than one bit in this mask set.
7-0	μFrame S-mask	Interrupt schedule mask. This field is used for all endpoint speeds. Software should set this field to a zero when the queue head is on the asynchronous schedule. A non-zero value in this field indicates an interrupt endpoint. The host controller uses the value of the three low-order bits of the FRINDEX register as an index into a bit position in this bit vector. If the μFrame S-mask field has a one at the indexed bit position then this queue head is a candidate for transaction execution. If the EPS field indicates the endpoint is a high-speed endpoint, then the transaction executed is determined by the PID_Code field contained in the execution area. This field is also used to support split transaction types: Interrupt (IN/OUT). This condition is true when this field is non-zero and the EPS field indicates this is either a full- or low-speed device. A zero value in this field, in combination with existing in the periodic frame list has undefined results.

18.5.6.3 Transfer overlay

The nine DWords in this area represent a transaction working space for the host controller.

The general operational model is that the host controller can detect whether the overlay area contains a description of an active transfer. If it does not contain an active transfer, then it follows the queue head horizontal link pointer to the next queue head. The host controller will never follow the next transfer queue element or alternate queue element pointers unless it is actively attempting to advance the queue. For the duration of the transfer, the host controller keeps the incremental status of the transfer in the overlay area. When the transfer is complete, the results are written back to the original queue element.

The DWord3 of a queue head contains a pointer to the source qTD currently associated with the overlay. The host controller uses this pointer to write back the overlay area into the source qTD after the transfer is complete.

The table below describes the current qTD link pointer.

Table 18-70. Current qTD link pointer

Bits	Name	Description
31-5	Current qTD pointer	Current element transaction descriptor link pointer. This field contains the address Of the current transaction being processed in this queue and corresponds to memory address signals [31:5], respectively.
4-0	-	Reserved, should be cleared. These bits are ignored by the host controller when using the value as an address to write data. The actual value may vary depending on the usage.

The DWords 4-11 of a queue head are the transaction overlay area. This area has the same base structure as a queue element transfer descriptor. The queue head utilizes the reserved fields of the page pointers to implement tracking the state of split transactions.

This area is characterized as an overlay because when the queue is advanced to the next queue element, the source queue element is merged onto this area. This area serves an execution cache for the transfer.

The table below describes the host-controller rules for bits in overlay.

Table 18-71. Host-controller rules for bits in overlay (DWords 5, 6, 8, and 9)

DWord	QH Offset	Bits	Name	Description
5	0x14	4-1	NakCnt	Nak counter-RW. This field is a counter the host controller decrements whenever a transaction for the endpoint associated with this queue head results in a Nak or Nyet response. This counter is reloaded from RL before a transaction is executed during the first pass of the reclamation list (relative to an Asynchronous List Restart condition). It is also loaded from RL during an overlay.
6	0x18	31	dt	Data toggle. The Data toggle control controls whether the host controller preserves this bit when an overlay operation is performed.
6	0x18	15	ioc	Interrupt on complete. The ioc control bit is always inherited from the source qTD when the overlay operation is performed.
6	0x18	11-10	Cerr	Error counter. Copied from the qTD during the overlay and written back during queue advancement.
6	0x18	0	Status[0]	Ping state (P)/ERR. If the EPS field indicates a high-speed endpoint, then this field should be preserved during the overlay operation.
8	0x20	7-0	C-prog-mask	Split-transaction complete-split progress. Initialized to zero during any overlay. This field is used to track the progress of an interrupt split-transaction.
9	0x24	11-5	S-bytes	Software must ensure that the S-bytes field in a qTD is zero before activating the qTD. Keeps track of the number of bytes sent or received during an IN or OUT split transaction.
9	0x24	4-0	FrameTag	Split-transaction frame tag. Initialized to zero during any overlay. This field is used to track the progress of an interrupt split-transaction.

18.5.7 Periodic frame span traversal node (FSTN)

The periodic frame span traversal node (FSTN) data structure, shown in the table below, is to be used only for managing full- and low-speed transactions that span a host-frame boundary.

Software must not use an FSTN in the asynchronous schedule. An FSTN in the asynchronous schedule results in undefined behavior.

Table 18-72. Frame span traversal node structure

31	3	2	2	2	2	2	2	2	2	2	1	1	1	1	15	1	1	1	1	1	9	8	7	6	5	4	3	2	1	0	offset
	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6		4	3	2	1	0										
Normal path link pointer																											-	Typ	T	0x00	
Back path link pointer																											-	Typ	T	0x04	

NOTE

The host controller performs only read operations to the FSTN data structure.

18.5.7.1 FSTN normal path pointer

The first DWord of an FSTN contains a link pointer to the next schedule object.

This object can be of any valid periodic schedule data type. The table below describes the FSTN normal path pointer.

Table 18-73. FSTN normal path pointer

Bits	Name	Description
31-5	NPLP	Normal path link pointer. Contains the address of the next data object to be processed in the periodic list and corresponds to memory address signals [31:5], respectively.
4-3	-	Reserved, should be cleared. These bits must be written as 0s.
2-1	Typ	Indicates to the host controller whether the item referenced is a iTD/siTd, QH, or FSTN. This allows the host controller to perform the proper type of processing on the item after it is fetched. 00 iTD (isochronous transfer descriptor) 01 QH (queue head) 10 siTD (split transaction isochronous transfer descriptor) 11 FSTN (frame span traversal node)
0	T	Terminate. 0 Link pointer is valid. 1 Link pointer field is not valid.

18.5.7.2 FSTN back path link pointer

The second DWord of an FSTN node contains a link pointer to a queue head.

If the T-bit in this pointer is a zero, then this FSTN is a Save-Place indicator. Its Typ field must be set by software to indicate the target data structure is a queue head. If the T-bit in this pointer is set, then this FSTN is the Restore indicator. When the T-bit is a one, the host controller ignores the Typ field.

The table below describes the FSTN back path link pointer.

Table 18-74. FSTN back path link pointer

Bits	Name	Description
31-5	BPLP	Back path link pointer. Contains the address of a queue head. This field corresponds to memory address signals [31:5], respectively.
4-3	-	Reserved, should be cleared. These bits must be written as 0s.
2-1	Typ	Software must ensure this field is set to indicate the target data structure is a Queue Head (01). Any other value in this field yields undefined results.
0	T	Terminate. 0 Link pointer is valid (that is, the host controller may use bits 31-5 as a valid memory address). This value also indicates that this FSTN is a Save-Place indicator. 1 Link pointer field is not valid (that is, the host controller must not use bits 31-5 as a valid memory address). This value also indicates that this FSTN is a Restore indicator.

18.6 Host operations

The general operational model for the USB DR module in host mode is defined by the Enhanced Host Controller Interface (EHCI) Specification. The EHCI specification describes the register-level interface for a host controller for the USB Revision 2.0. It includes a description of the hardware/software interface between system software and host controller hardware. Information concerning the initialization of the USB module is included in the following section; however, the full details of the EHCI specification are beyond the scope of this document.

18.6.1 Host controller initialization

After initial power-on or host controller reset (hardware or through USBCMD[RST]), all of the operational registers are at their default values.

To configure the external ULPI PHY the following initialization sequence is required:

1. The UTMI PHY should remain disabled if the ULPI is being used.
2. Set the CONTROL[PHY_CLK_SEL] bits to select the ULPI PHY as the source of USB controller PHY clock.
3. Wait for PHY clock to become valid. This can be determined by polling the CONTROL[PHY_CLK_VALID] status bit. Note that this bit is not valid once the CONTROL[USB_EN] bit is set

Once the PHY clock is valid the user can proceed to the host controller initialization phase.

In order to initialize the USB DR module, software should perform the following steps:

1. Set the controller mode to host mode. Optionally set USBMODE[SDIS] (streaming disable)

NOTE

Transitioning from device mode to host mode requires a host controller reset before modifying USBMODE.

2. Optionally modify the BURSTSIZE register.
3. Program the PTS field of the PORTSC register if using a non-ULPI PHY.
4. Set CONTROL[USB_EN].
5. Write the appropriate value to the USBINTR register to enable the appropriate interrupts.
6. Write the base address of the periodic frame list to the PERIODICLIST BASE register. If there are no work items in the periodic schedule, all elements of the periodic frame list should have their T-Bits set.
7. Write the USBCMD register to set the desired interrupt threshold, frame list size (if applicable) and turn on the controller by setting the RS bit.

At this point, the USB DR module is up and running and the port registers begin reporting device connects. System software can enumerate a port through the reset process (where the port is in the enabled state). At this point, the port is active with SOFs occurring down the enabled high-speed ports, but the schedules have not yet been enabled. The EHCI host controller will not transmit SOFs to enabled Full- or Low-speed ports.

In order to communicate with devices via the asynchronous schedule, system software must write the ASYNDLISTADDR register with the address of a control or bulk queue head. Software must then enable the asynchronous schedule by writing a one to USBCMD[ASE]. In order to communicate with devices via the periodic schedule, system software must enable the periodic schedule by writing a one to USBCMD[PSE]. Note that the schedules can be turned on before the first port is reset (and enabled).

Any time the USBCMD register is written, system software must ensure the appropriate bits are preserved, depending on the intended operation.

18.6.2 Power port

The HCSPARAMS[PPC] bit indicates whether the USB 2.0 host controller has port power control.

When the PPC bit is set, the host controller supports port power switches. Each available switch has an output enable. PPE is controlled based on the state of the combination bits-PPC bit, EHCI Configured (CF)-bit and individual Port Power (PP) bit. The Configured Flag and Port Power Control bits are always 1'b1 in Host Mode. The PPE always follows the state of Port Power (PP) bit that is, if PP is 0, PPE will be 0 and if PP is 1, PPE will be 1.

18.6.3 Reporting over-current

Host ports by definition are power providers on USB.

Whether the ports are considered high- or low-powered is a platform implementation issue. The EHCI PORTSC register has an over-current status and over-current change bit. The functionality of these bits is specified in the USB Specification Revision 2.0.

The over current detection and limiting logic resides outside the DR logic. The over-current condition effects the following bits in the PORTSC register on the EHCI port:

- Over-current active bit (OCA) is set. When the over-current condition goes away, the OCA will transition from a one to a zero.
- Over-current change bit (OCC) is set. On every transition of OCA, the controller will set OCC to a one. Software sets OCC to a zero by writing a one to this bit.
- Port enabled/disabled bit (PE) is cleared. When this change bit gets set, USBSTS[PCI] (the port change detect bit) is set.
- Port power (PP) bit may optionally be cleared. There is no requirement in USB that a power provider shut off power in an over current condition. It is sufficient to limit the current and leave power applied. When OCC transitions from a zero to a one, the controller also sets USBSTS[PCI] to a one. In addition, if the Port Change Interrupt Enable bit, USBINTR[PCE], is a one, the controller issues an interrupt to the system. Refer to [Table 18-75](#) for summary of behavior for over-current detection when the controller is halted (suspended from a device component point of view).

18.6.4 Suspend/resume

The host controller provides an equivalent suspend and resume model as that defined for individual ports in a USB 2.0 hub.

Control mechanisms are provided to allow system software to suspend and resume individual ports. The mechanisms allow the individual ports to be resumed completely through software initiation. Other control mechanisms are provided to parameterize the host controller's response (or sensitivity) to external resume events. In this discussion, host-initiated, or software-initiated resumes are called Resume Events/Actions; bus-initiated resume events are called wake-up events. The classes of wakeup events are:

- Remote-wakeup enabled device asserts resume signaling. In similar kind to USB 2.0 hubs, when in host mode the host controller responds to explicit device resume signaling and wake up the system (if necessary).
- Port connect and disconnect and over-current events. Sensitivity to these events can be turned on or off by using the port control bits in the PORTSC register.

Selective suspend is a feature supported by the PORTSC register. It is used to place specific ports into a suspend mode. This feature is used as a functional component for implementing the appropriate power management policy implemented in a particular operating system. When system software intends to suspend the bus, it should suspend the enabled port, then shut off the controller by setting the USBCMD[RS] to a zero.

When a wake event occurs the system will resume operation and system software must set the RS bit to a one and resume the suspended port.

18.6.4.1 Port suspend/resume

System software places the USB into suspend mode by writing a one into the appropriate PORTSC Suspend bit.

Software must only set the Suspend bit when the port is in the enabled state (Port Enabled bit is a one).

The host controller may evaluate the Suspend bit immediately or wait until a microframe or frame boundary occurs. If evaluated immediately, the port is not suspended until the current transaction (if one is executing) completes. Therefore, there may be several microframes of activity on the port until the host controller evaluates the Suspend bit. The host controller must evaluate the Suspend bit at least every frame boundary.

System software can initiate a resume on the suspended port by writing a one to PORTSC[FPR]. Software should not attempt to resume a port unless the port reports that it is in the suspended state. If system software sets PORTSC[FPR] when the port is not in

the suspended state, the resulting behavior is undefined. In order to assure proper USB device operation, software must wait for at least 10 milliseconds after a port indicates that it is suspended (Suspend bit is a one) before initiating a port resume through PORTSC[FPR]. When PORTSC[FPR] is set, the host controller sends resume signaling down the port. System software times the duration of the resume (nominally 20 milliseconds) then clears PORTSC[FPR]. When the host controller receives the write to transition PORTSC[FPR] to zero, it completes the resume sequence as defined in the USB specification, and clears both PORTSC[FPR] and PORTSC[SUSP]. Software-initiated port resumes do not affect the port change detect bit (USBSTS[PCI]) nor do they cause an interrupt if USBINTR[PCE] (port change interrupt enable) is a one. When a wake event occurs on a suspended port, the resume signaling is detected by the port and the resume is reflected downstream within 100 μ sec. The port's PORTSC[FPR] bit is set and USBSTS[PCI] is set. If USBINTR[PCE] is a one, the host controller issues a hardware interrupt.

System software observes the resume event on the port, delays a port resume time (nominally 20 milliseconds), then terminates the resume sequence by clearing PORTSC[FPR] in the port. The host controller receives the write of zero to PORTSC[FPR], terminates the resume sequence and clears PORTSC[FPR] and PORTSC[SUSP]. Software can determine that the port is enabled (not suspended) by sampling the PORTSC register and observing that the SUSP and FPR bits are zero. Software must ensure that the host controller is running (that is, USBSTS[HCH] is a zero), before terminating a resume by clearing the port's PORTSC[FPR] bit. If HCH is a one when PORTSC[FPR] is cleared, then SOFs will not occur down the enabled port and the device will return to suspend mode in a maximum of 10 milliseconds.

The table below summarizes the wake-up events. Whenever a resume event is detected, USBSTS[PCI] is set. If USBINTR[PCE] (port change interrupt enable) is a one, the host controller also generates an interrupt on the resume event. Software acknowledges the resume event interrupt by clearing the USBSTS[PCI].

Table 18-75. Behavior during wake-up events

Port status and signaling type	Signaled port response	Device state	
		D0	Not D0
Port disabled, resume K-State received	No effect	N/A	N/A
Port suspended, Resume K-State received	Resume reflected downstream on signaled port. PORTSC[FPR] is set. USBSTS[PCI] is set.	[1], [2]	[2]
Port is enabled, disabled or suspended, and the port's PORTSC[WKDS], is set. A disconnect is detected.	Depending on the initial port state, the PORTSC Connect (CCS) and Enable (PE) status bits are cleared, and the Connect Change status bit (CSC) is set. USBSTS[PCI] is set.	[1], [2]	[2]
Port is enabled, disabled or suspended, and the port's PORTSC[WKDS], is cleared. A disconnect is detected.	Depending on the initial port state, the PORTSC Connect (CCS) and Enable (PE) status bits are cleared, and the Connect Change status bit (CSC) is set. USBSTS[PCI] is set.	[1], [3]	[3]

Table continues on the next page...

Table 18-75. Behavior during wake-up events (continued)

Port status and signaling type	Signaled port response	Device state	
		D0	Not D0
Port is not connected and the port's PORTSC[WKCN] bit is a one. A connect is detected.	PORTSC Connect Status (CCS) and Connect Status Change (CSC) bits are set. USBSTS[PCI] is set.	[1], [2]	[2]
Port is not connected and the port's PORTSC[WKCN] bit is a zero. A connect is detected.	PORTSC Connect Status (CCS) and Connect Status Change (CSC) bits are set. USBSTS[PCI] is set.	[1], [3]	[3]
Port is connected and the port's PORTSC[WKOC] bit is a one. An over-current condition occurs.	PORTSC Over-current Active (OCA), Over-current Change (OCC) bits are set. If Port Enable/Disable bit (PE) is a one, it is cleared. USBSTS[PCI] is set	[1], [2]	[2]
Port is connected and the port's PORTSC[WKOC] bit is a zero. An over-current condition occurs.	PORTSC Over-current Active (OCA), Over-current Change (OCC) bits are set. If Port Enable/Disable bit (PE) is a one, it is cleared. USBSTS[PCI] is set.	[1], [3]	[3]
¹ Hardware interrupt issued if USBINTR[PCE] (port change interrupt enable) is set. ² PME# asserted if enabled (Note: PME Status must always be set). ³ PME# not asserted.			

18.6.5 Schedule traversal rules

The host controller executes transactions for devices using a simple, shared-memory schedule.

The schedule is comprised of a few data structures, organized into two distinct lists. The data structures are designed to provide the maximum flexibility required by USB, minimize memory traffic and hardware/software complexity.

System software maintains two schedules for the host controller: a periodic schedule and an asynchronous schedule. The root of the periodic schedule is the [Periodic frame list base address \[host mode\] \(USB_PERIODICLISTBASE\)](#) register. The PERIODICLISTBASE register is the physical memory base address of the periodic frame list. The periodic frame list is an array of physical memory pointers. The objects referenced from the frame list must be valid schedule data structures as defined in [Host data structures](#). In each microframe, if the periodic schedule is enabled (see [Periodic schedule](#)) then the host controller must execute from the periodic schedule before executing from the asynchronous schedule. It will only execute from the asynchronous schedule after it encounters the end of the periodic schedule. The host controller traverses the periodic schedule by constructing an array offset reference from the PERIODICLISTBASE and the FRINDEX registers (see the figure below). It fetches the element and begins traversing the graph of linked schedule data structures.

The end of the periodic schedule is identified by a next link pointer of a schedule data structure having its T-bit set. When the host controller encounters a T-Bit set during a horizontal traversal of the periodic list, it interprets this as an End-Of-Periodic-List mark. This causes the host controller to cease working on the periodic schedule and transitions immediately to traversing the asynchronous schedule. Once this transition is made, the host controller executes from the asynchronous schedule until the end of the microframe.

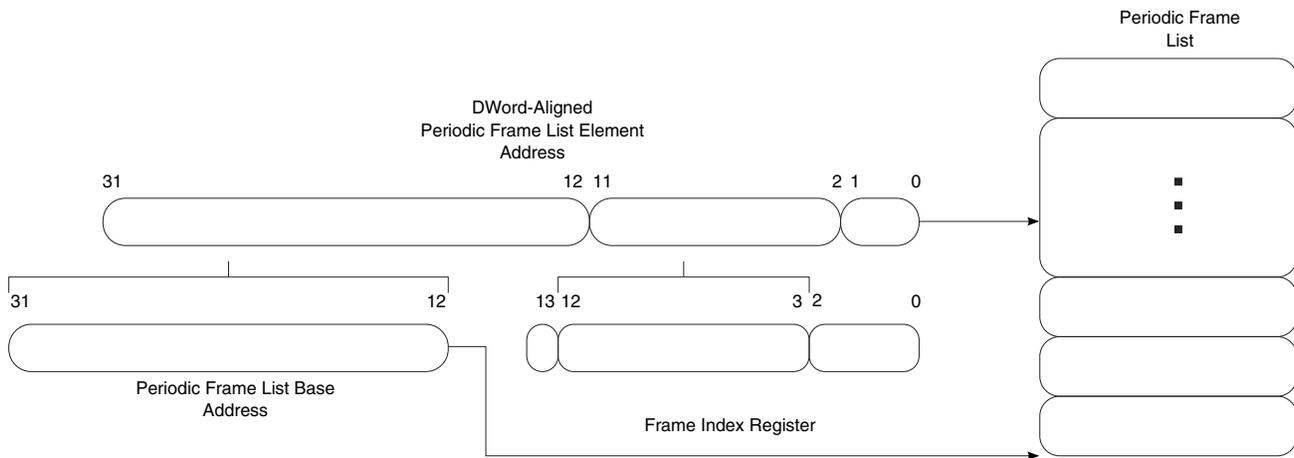


Figure 18-44. Derivation of pointer into frame list array

When the host controller determines that it is time to execute from the asynchronous list, it uses the operational register `ASYNCLISTADDR` to access the asynchronous schedule, as shown in the figure below.

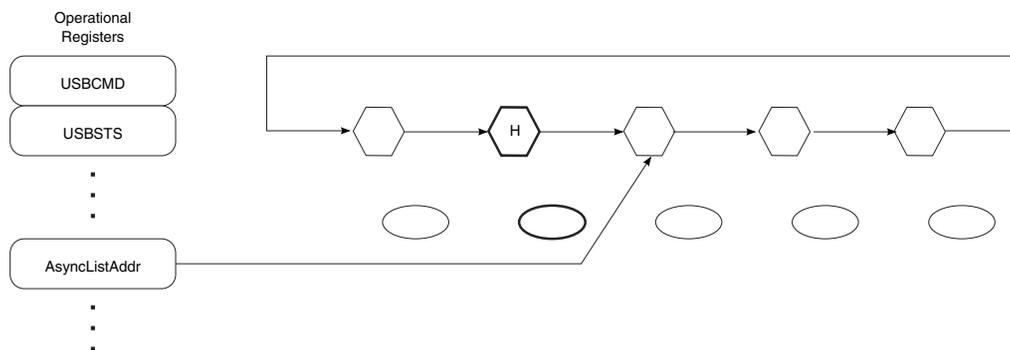


Figure 18-45. General format of asynchronous schedule list

The `ASYNCLISTADDR` register contains a physical memory pointer to the next queue head. When the host controller makes a transition to executing the asynchronous schedule, it begins by reading the queue head referenced by the `ASYNCLISTADDR` register. Software must set queue head horizontal pointer T-bits to a zero for queue heads in the asynchronous schedule.

18.6.6 Periodic schedule frame boundaries vs. bus frame boundaries

The USB Specification Revision 2.0 requires that the frame boundaries (SOF frame number changes) of the high-speed bus and the full- and low-speed bus(es) below USB 2.0 hubs be strictly aligned.

Super-imposed on this requirement is that USB 2.0 hubs manage full- and low-speed transactions via a microframe pipeline (see start- (SS) and complete- (CS) splits illustrated in the figure below). A simple, direct projection of the frame boundary model into the host controller interface schedule architecture creates tension (complexity for both hardware and software) between the frame boundaries and the scheduling mechanisms required to service the full- and low-speed transaction translator periodic pipelines.

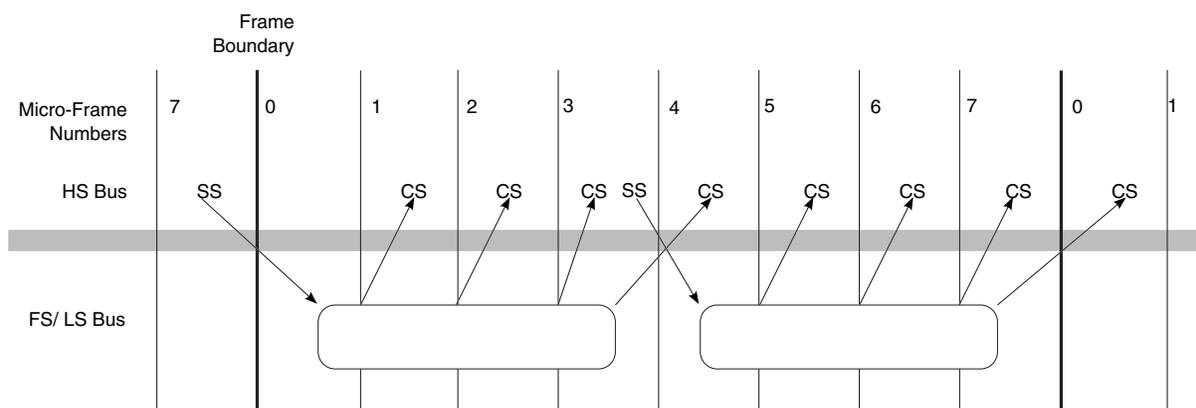


Figure 18-46. Frame boundary relationship between HS bus and FS/LS bus

The simple projection, as the figure above illustrates, introduces frame-boundary wrap conditions for scheduling on both the beginning and end of a frame. In order to reduce the complexity for hardware and software, the host controller is required to implement a one microframe phase shift for its view of frame boundaries. The phase shift eliminates the beginning of frame and frame-wrap scheduling boundary conditions.

The implementation of this phase shift requires that the host controller use one register value for accessing the periodic frame list and another value for the frame number value included in the SOF token. These two values are separate, but tightly coupled. The periodic frame list is accessed via the Frame List Index Register (FRINDEX). Bits FRINDEX[2-0], represent the microframe number. The SOF value is coupled to the value of FRINDEX[13-3]. Both FRINDEX[13-3] and the SOF value are incremented based on FRINDEX[2-0]. It is required that the SOF value be delayed from the FRINDEX value by one microframe. The one microframe delay yields a host controller periodic schedule

and bus frame boundary relationship as illustrated in the figure below. This adjustment allows software to trivially schedule the periodic start and complete-split transactions for full- and low-speed periodic endpoints, using the natural alignment of the periodic schedule interface.

The figure below illustrates how periodic schedule data structures relate to schedule frame boundaries and bus frame boundaries. To aid the presentation, two terms are defined. The host controller's view of the 1-millisecond boundaries is called H-Frames. The high-speed bus's view of the 1-millisecond boundaries is called B-Frames.

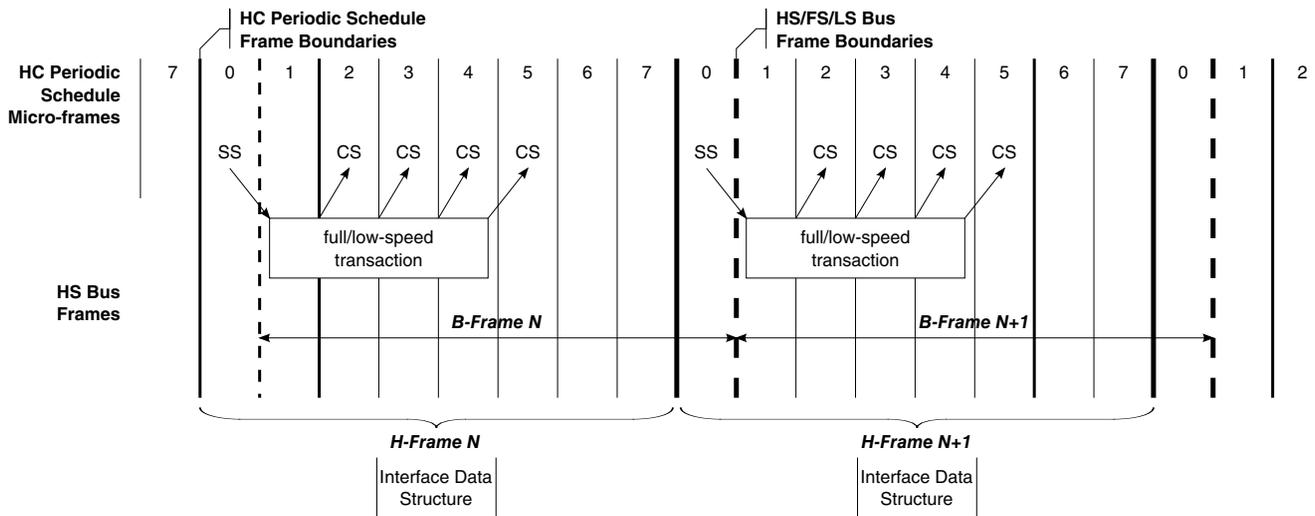


Figure 18-47. Relationship of periodic schedule frame boundaries to bus frame boundaries

H-Frame boundaries for the host controller correspond to increments of FRINDEX[13-3]. Microframe numbers for the H-Frame are tracked by FRINDEX[2-0]. B-Frame boundaries are visible on the high-speed bus via changes in the SOF token's frame number. Microframe numbers on the high-speed bus are only derived from the SOF token's frame number (that is, the high-speed bus will see eight SOFs with the same frame number value). H-Frames and B-Frames have the fixed relationship (that is, B-Frames lag H-Frames by one microframe time) illustrated in the figure above. The host controller's periodic schedule is naturally aligned to H-Frames. Software schedules transactions for full- and low-speed periodic endpoints relative the H-Frames. The result is these transactions execute on the high-speed bus at exactly the right time for the USB 2.0 hub periodic pipeline. As described in [USB frame index \(USB_FRINDEX\)](#), the SOF Value can be implemented as a shadow register (in this example, called SOFV), which lags the FRINDEX register bits [13-3] by one microframe count. The table below illustrates the required relationship between the value of FRINDEX and the value of

SOFV. This lag behavior can be accomplished by incrementing FRINDEX[13-3] based on carry-out on the 7 to 0 increment of FRINDEX[2-0] and incrementing SOFV based on the transition of 0 to 1 of FRINDEX[2-0].

Software is allowed to write to FRINDEX. [USB frame index \(USB_FRINDEX\)](#), provides the requirements that software should adhere when writing a new value in FRINDEX.

Table 18-76. Operation of FRINDEX and SOFV (SOF value register)

Current			Next		
FRINDEX[13-3]	SOFV	FRINDEX[2-0]	FRINDEX[13-3]	SOFV	FRINDEX[2-0]
N	N	111	N+1	N	000
N+1	N	000	N+1	N+1	001
N+1	N+1	001	N+1	N+1	010
N+1	N+1	010	N+1	N+1	011
N+1	N+1	011	N+1	N+1	100
N+1	N+1	100	N+1	N+1	101
N+1	N+1	101	N+1	N+1	110
N+1	N+1	110	N+1	N+1	111

18.6.7 Periodic schedule

The periodic schedule traversal is enabled or disabled through USBCMD[PSE] (periodic schedule enable).

If USBCMD[PSE] is cleared, then the host controller simply does not try to access the periodic frame list via the PERIODICLISTBASE register. Likewise, when USBCMD[PSE] is a one, then the host controller does use the PERIODICLISTBASE register to traverse the periodic schedule. The host controller will not react to modifications to USBCMD[PSE] immediately. In order to eliminate conflicts with split transactions, the host controller evaluates USBCMD[PSE] only when FRINDEX[2-0] is zero. System software must not disable the periodic schedule if the schedule contains an active split transaction work item that spans the 0b000 microframe. These work items must be removed from the schedule before USBCMD[PSE] is cleared. USBSTS[PS] (periodic schedule status) indicates status of the periodic schedule. System software enables (or disables) the periodic schedule by setting (or clearing) USBCMD[PSE]. Software then can poll USBSTS[PS] to determine when the periodic schedule has made the desired transition. Software must not modify USBCMD[PSE] unless the value of USBCMD[PSE] equals that of USBSTS[PS].

The periodic schedule is used to manage all isochronous and interrupt transfer streams. The base of the periodic schedule is the periodic frame list. Software links schedule data structures to the periodic frame list to produce a graph of scheduled data structures. The graph represents an appropriate sequence of transactions on the USB. The figure below illustrates isochronous transfers (using iTDs and siTDs) with a period of one are linked directly to the periodic frame list. Interrupt transfers (are managed with queue heads) and isochronous streams with periods other than one are linked following the period-one iTD/siTDs. Interrupt queue heads are linked into the frame list ordered by poll rate. Longer poll rates are linked first (for example, closest to the periodic frame list), followed by shorter poll rates, with queue heads with a poll rate of one, on the very end.

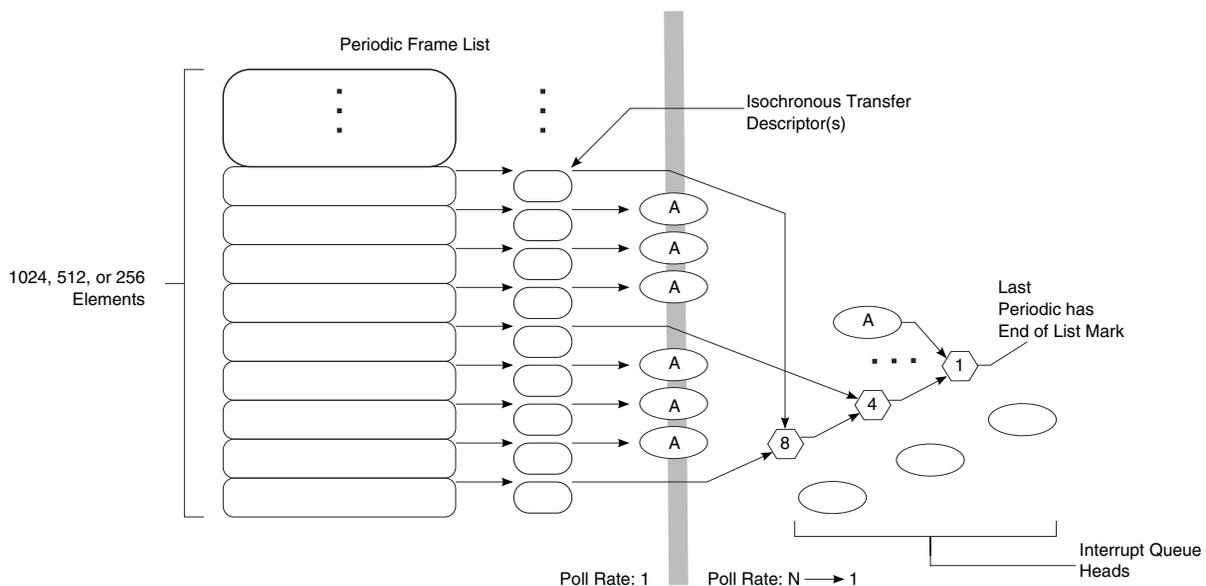


Figure 18-48. Example periodic schedule

18.6.8 Managing isochronous transfers using iTDs

The structure of an iTD is presented in isochronous (high-speed) transfer descriptor (iTD).

There are four distinct sections to an iTD:

- Next link pointer
 - This is the first field.
 - This field is for schedule linkage purposes only.
- Transaction description array

- This is an eight-element array.
- Each element represents control and status information for one microframe's worth of transactions for a single high-speed isochronous endpoint.
- Buffer page pointer array
 - This is a 7-element array of physical memory pointers to data buffers.
 - These are 4K aligned pointers to physical memory.
- Endpoint capabilities
 - This area utilizes the unused low-order 12 bits of the buffer page pointer array.
 - Its fields are used across all transactions executed for this iTD, including endpoint addressing, transfer direction, maximum packet size and high-bandwidth multiplier.

18.6.8.1 Host controller operational model for iTDs

The host controller uses FRINDEX register bits 12-3 to index into the periodic frame list.

This means that the host controller visits each frame list element eight consecutive times before incrementing to the next periodic frame list element. Each iTD contains eight transaction descriptions, which map directly to FRINDEX register bits 2-0. Each iTD can span 8 microframes worth of transactions. When the host controller fetches an iTD, it uses FRINDEX register bits 2-0 to index into the transaction description array. When the first iTD in the periodic list is traversed after periodic schedule is enabled, the value of FRINDEX[2:0] may be other than 0, so the first transaction issued by the controller may be any of the eight available active transactions. If the active bit in the Status field of the indexed transaction description is cleared, the host controller ignores the iTD and follows the Next pointer to the next schedule data structure.

When the indexed active bit is a one the host controller continues to parse the iTD. It stores the indexed transaction description and the general endpoint information (device address, endpoint number, maximum packet size, and so on). It also uses the Page Select (PG) field to index the buffer pointer array, storing the selected buffer pointer and the next sequential buffer pointer. For example, if PG field is a 0, then the host controller will store Page 0 and Page 1.

The host controller constructs a physical data buffer address by concatenating the current buffer pointer (as selected using the current transaction description's PG field) and the transaction description's Transaction Offset field. The host controller uses the endpoint addressing information and I/O-bit to execute a transaction to the appropriate endpoint. When the transaction is complete, the host controller clears the active bit and writes back any additional status information to the Status field in the currently selected transaction description.

The data buffer associated with the iTD must be virtually contiguous memory. Seven page pointers are provided to support eight high-bandwidth transactions regardless of the starting packet's offset alignment into the first page. A starting buffer pointer (physical memory address) is constructed by concatenating the page pointer (example: page 0 pointer) selected by the active transaction descriptions' PG (example value: 0b00) field with the transaction offset field. As the transaction moves data, the host controller must detect when an increment of the current buffer pointer will cross a page boundary. When this occurs the host controller simply replaces the current buffer pointer's page portion with the next page pointer (example: page 1 pointer) and continues to move data. The size of each bus transaction is determined by the value in the Maximum Packet Size field. An iTD supports high-bandwidth pipes via the Mult (multiplier) field. When the Mult field is 1, 2, or 3, the host controller executes the specified number of Maximum Packet sized bus transactions for the endpoint in the current microframe. In other words, the Mult field represents a transaction count for the endpoint in the current microframe. If the Mult field is zero, the operation of the host controller is undefined. The transfer description is used to service all transactions indicated by the Mult field.

For OUT transfers, the value of the Transaction *n* Length field represents the total bytes to be sent during the microframe. The Mult field must be set by software to be consistent with Transaction *n* Length and Maximum Packet Size. The host controller will send the bytes in Maximum Packet Sized portions. After each transaction, the host controller decrements it's local copy of Transaction *n* Length by Maximum Packet Size. The number of bytes the host controller sends is always Maximum Packet Size or Transaction *n* Length, whichever is less. The host controller advances the transfer state in the transfer description, updates the appropriate record in the iTD and moves to the next schedule data structure. The maximum sized transaction supported is 3 x 1024 bytes.

For IN transfers, the host controller issues Mult transactions. It is assumed that software has properly initialized the iTD to accommodate all possible data. During each IN transaction, the host controller must use Maximum Packet Size to detect packet babble errors. The host controller keeps the sum of bytes received in the Transaction *n* Length field.

After all transactions for the endpoint have completed for the microframe, Transaction *n* Length contains the total bytes received. The following actions can occur:

- If the final value of Transaction *n* Length is less than the value of Maximum Packet Size, less data was received than was allowed for from the associated endpoint. This short packet condition does not set USBSTS[UI] (USB interrupt). The host controller does not detect this condition.
- If the device sends more than Transaction *n* Length or Maximum Packet Size bytes (whichever is less), the host controller sets the Babble Detected bit and clears the

Active bit. Note, that the host controller does not update the iTD field Transaction n Length in this error scenario.

- If the Mult field is greater than one, the host controller automatically executes the value of Mult transactions. The host controller does not execute all Mult transactions in the following cases:
 - The endpoint is an OUT and Transaction n Length goes to zero before all the Mult transactions have executed (ran out of data).
 - The endpoint is an IN and the endpoint delivers a short packet, or an error occurs on a transaction before Mult transactions have been executed.

The end of microframe may occur before all of the transaction opportunities have been executed. When this happens, the transfer state of the transfer description is advanced to reflect the progress that was made; the result is written back to the iTD; and the host controller proceeds to processing the next microframe.

18.6.8.2 Software operational model for iTDs

A client buffer request to an isochronous endpoint may span 1 to N microframes.

When N is larger than one, system software may have to use multiple iTDs to read or write data with the buffer (if N is larger than eight, it must use more than one iTD).

The figure below illustrates the simple model of how a client buffer is mapped by system software to the periodic schedule (that is, the periodic frame list and a set of iTDs).

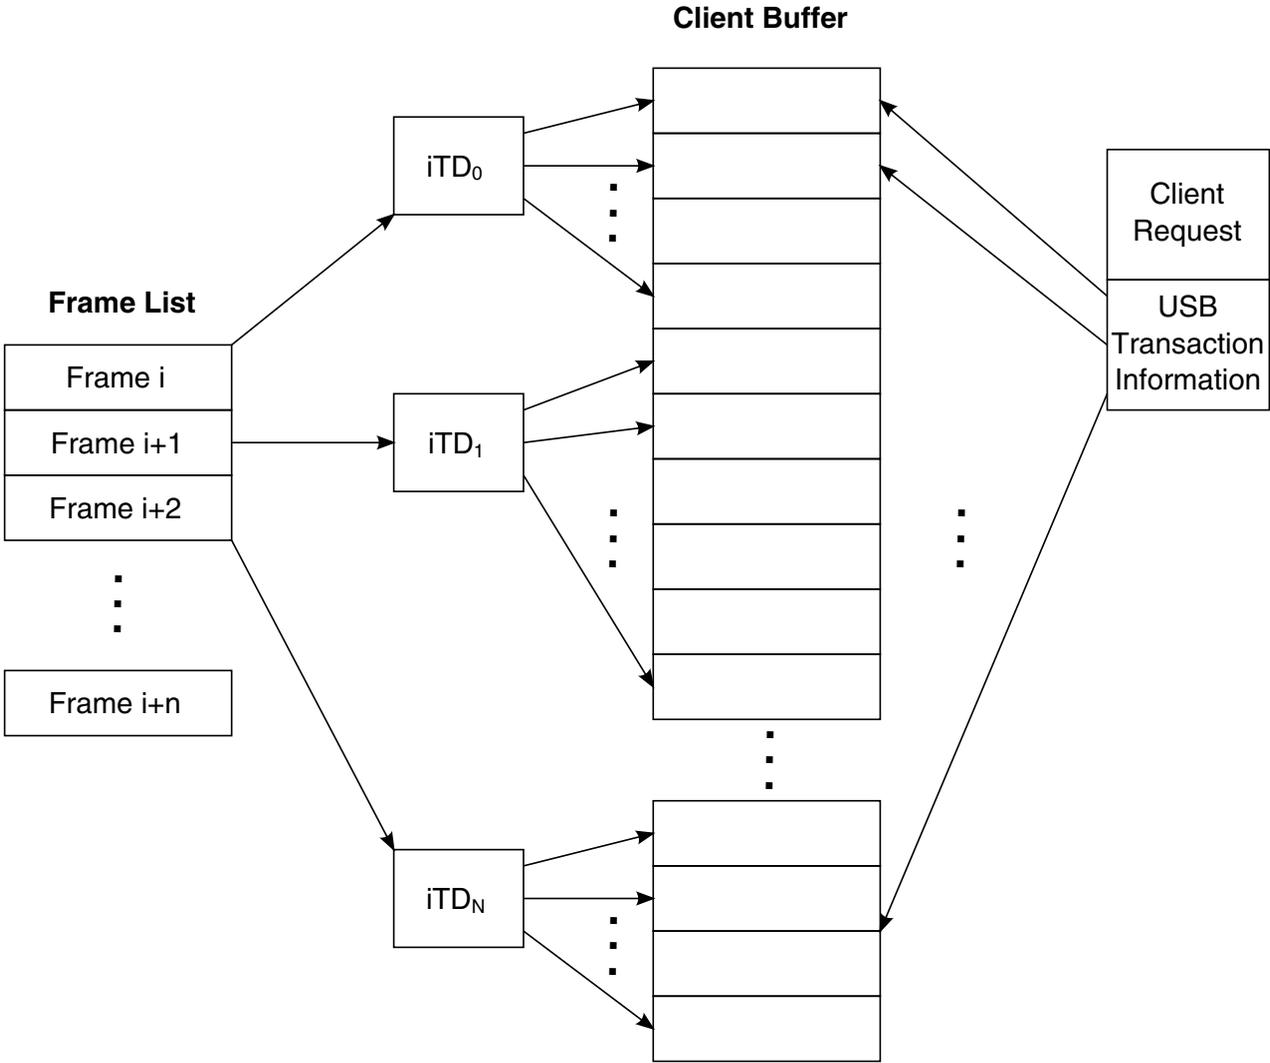


Figure 18-49. Example Association of iTDs to Client Request Buffer

On the right is the client description of its request. The description includes a buffer base address plus additional annotations to identify which portions of the buffer should be used with each bus transaction. In the middle is the iTD data structures used by the system software to service the client request. Each iTD can be initialized to service up to 24 transactions, organized into eight groups of up to three transactions each. Each group maps to one microframe's worth of transactions. The EHCI controller does not provide per transaction results within a microframe. It treats the per microframe transactions as a single logical transfer.

On the left is the host controller's frame list. System software establishes references from the appropriate locations in the frame list to each of the appropriate iTDs. If the buffer is large, system software can use a small set of iTDs to service the entire buffer. System software can activate the transaction description records (contained in each iTD) in any pattern required for the particular data stream.

As noted above, the client request includes a pointer to the base of the buffer and offsets into the buffer to annotate which buffer sections are to be used on each bus transaction that occurs on this endpoint. System software must initialize each transaction description in an iTD to ensure it uses the correct portion of the client buffer. For example, for each transaction description, the PG field is set to index the correct physical buffer page pointer and the Transaction Offset field is set relative to the correct buffer pointer page (for example, the same one referenced by the PG field). When the host controller executes a transaction it selects a transaction description record based on FRINDEX[2-0]. It then uses the current Page Buffer Pointer (as selected by the PG field) and concatenates to the transaction offset field. The result is a starting buffer address for the transaction. As the host controller moves data for the transaction, it must watch for a page wrap condition and properly advance to the next available Page Buffer Pointer. System software must not use the Page 6 buffer pointer in a transaction description where the length of the transfer will wrap a page boundary. Doing so yields undefined behavior. The host controller hardware is not required to alias the page selector to page zero. USB 2.0 isochronous endpoints can specify a period greater than one. Software can achieve the appropriate scheduling by linking iTDs into the appropriate frames (relative to the frame list) and by setting appropriate transaction description elements active bits to a one.

18.6.8.2.1 Periodic scheduling threshold

The isochronous scheduling threshold field in the HCCPARAMS capability register is an indicator to system software as to how the host controller pre-fetches and effectively caches schedule data structures.

It is used by system software when adding isochronous work items to the periodic schedule. The value of this field indicates to system software the minimum distance it can update isochronous data (relative to the current location of the host controller execution in the periodic list) and still have the host controller process them.

The iTD and siTD data structures each describe 8 microframes worth of transactions. The host controller is allowed to cache one (or more) of these data structures in order to reduce memory traffic. There are three basic caching models that account for the fact the isochronous data structures span 8 microframes. The three caching models are: no caching, microframe caching and frame caching.

When software is adding new isochronous transactions to the schedule, it always performs a read of the FRINDEX register to determine the current frame and microframe the host controller is currently executing. Of course, there is no information about where in the microframe the host controller is, so a constant uncertainty factor of one microframe has to be assumed. Combining the knowledge of where the host controller is executing with the knowledge of the caching model allows the definition of simple algorithms for how closely software can reliably work to the executing host controller.

No caching is indicated with a value of zero in the isochronous scheduling threshold field. The host controller may pre-fetch data structures during a periodic schedule traversal (per microframe) but will always dump any accumulated schedule state at the end of the microframe. At the appropriate time relative to the beginning of every microframe, the host controller always begins schedule traversal from the frame list. Software can use the value of the FRINDEX register (plus the constant 1 uncertainty-factor) to determine the approximate position of the executing host controller. When no caching is selected, software can add an isochronous transaction as near as 2 microframes in front of the current executing position of the host controller.

Frame caching is indicated with a non-zero value in bit [7] of the isochronous scheduling threshold field. In the frame-caching model, system software assumes that the host controller caches one (or more) isochronous data structures for an entire frame (8 microframes). Software uses the value of the FRINDEX register (plus the constant 1 uncertainty) to determine the current microframe/frame (assume modulo 8 arithmetic in adding the constant 1 to the microframe number). For any current frame N , if the current microframe is 0 to 6, then software can safely add isochronous transactions to Frame $N + 1$. If the current microframe is 7, then software can add isochronous transactions to Frame $N + 2$.

Microframe caching is indicated with a non-zero value in the least-significant 3 bits of the isochronous scheduling threshold field. System software assumes the host controller caches one or more periodic data structures for the number of microframes indicated in the isochronous scheduling threshold field. For example, if the count value were 2, then the host controller keeps a window of two microframes worth of state (current microframe, plus the next) on chip. On each microframe boundary, the host controller releases the current microframe state and begins accumulating the next microframe state.

18.6.9 Asynchronous schedule

The asynchronous schedule traversal is enabled or disabled through USBCMD[ASE] (asynchronous schedule enable).

If USBCMD[ASE] is cleared, then the host controller simply does not try to access the asynchronous schedule via the ASYNCLISTADDR register. Likewise, if USBCMD[ASE] is set, the host controller does use the ASYNCLISTADDR register to traverse the asynchronous schedule. Modifications to USBCMD[ASE] are not necessarily immediate. Rather the new value of the bit will only be taken into consideration the next time the host controller needs to use the value of the ASYNCLISTADDR register to get the next queue head.

USBSTS[AS] indicates status of the asynchronous schedule. System software enables (or disables) the asynchronous schedule by writing a one (or zero) to USBCMD[ASE]. Software then can poll USBSTS[AS] to determine when the asynchronous schedule has made the desired transition. Software must not modify USBCMD[ASE] unless the value of USBCMD[ASE] equals that of the USBSTS[AS] (asynchronous schedule status).

The asynchronous schedule is used to manage all Control and Bulk transfers. Control and Bulk transfers are managed using queue head data structures. The asynchronous schedule is based at the ASYNCLISTADDR register. The default value of the ASYNCLISTADDR register after reset is undefined and the schedule is disabled when USBCMD[ASE] is cleared.

Software may only write this register with defined results when the schedule is disabled, for example, USBCMD[ASE] and the USBSTS[AS] are cleared. System software enables execution from the asynchronous schedule by writing a valid memory address (of a queue head) into this register. Then software enables the asynchronous schedule by setting USBCMD[ASE]. The asynchronous schedule is actually enabled when USBSTS[AS] is set.

When the host controller begins servicing the asynchronous schedule, it begins by using the value of the ASYNCLISTADDR register. It reads the first referenced data structure and begins executing transactions and traversing the linked list as appropriate. When the host controller completes processing the asynchronous schedule, it retains the value of the last accessed queue head's horizontal pointer in the ASYNCLISTADDR register. Next time the asynchronous schedule is accessed, this is the first data structure that is serviced. This provides round-robin fairness for processing the asynchronous schedule.

A host controller completes processing the asynchronous schedule when one of the following events occur:

- The end of a microframe occurs.
- The host controller detects an empty list condition
- The schedule has been disabled through USBCMD[ASE].

The queue heads in the asynchronous list are linked into a simple circular list as shown in [Figure 18-45](#). Queue head data structures are the only valid data structures that may be linked into the asynchronous schedule. An isochronous transfer descriptor (iTID or siTD) in the asynchronous schedule yields undefined results.

The maximum packet size field in a queue head is sized to accommodate the use of this data structure for all non-isochronous transfer types. The USB Specification, Revision 2.0 specifies the maximum packet sizes for all transfer types and transfer speeds. System software should always parameterize the queue head data structures according to the core specification requirements.

18.6.9.1 Adding queue heads to asynchronous schedule

This is a software requirement section. There are two independent events for adding queue heads to the asynchronous schedule.

The first is the initial activation of the asynchronous list. The second is inserting a new queue head into an activated asynchronous list.

Activation of the list is simple. System software writes the physical memory address of a queue head into the ASYNCLISTADDR register, then enables the list by setting USBCMD[ASE] to a one.

When inserting a queue head into an active list, software must ensure that the schedule is always coherent from the host controllers' point of view. This means that the system software must ensure that all queue head pointer fields are valid. For example qTD pointers have T-Bits set or reference valid qTDs and the Horizontal Pointer references a valid queue head data structure. The following algorithm represents the functional requirements:

```

InsertQueueHead (pQHeadCurrent, pQueueHeadNew)
--
-- Requirement: all inputs must be properly initialized.
--
-- pQHeadCurrent is a pointer to a queue head that is
-- already in the active list
-- pQHeadNew is a pointer to the queue head to be added
--
-- This algorithm links a new queue head into a existing
-- list
--
pQueueHeadNew.HorizontalPointer = pQueueHeadCurrent.HorizontalPointer
pQueueHeadCurrent.HorizontalPointer = physicalAddressOf(pQueueHeadNew)
End InsertQueueHead

```

18.6.9.2 Removing queue heads from asynchronous schedule

This is a software requirement section.

There are two independent events for removing queue heads from the asynchronous schedule.

host operations

The first is shutting down (deactivating) the asynchronous list. The second is extracting a single queue head from an activated list. Software deactivates the asynchronous schedule by setting USBCMD[ASE] to a zero. Software can determine when the list is idle when USBSTS[AS] is cleared. The normal mode of operation is that software removes queue heads from the asynchronous schedule without shutting it down. Software must not remove an active queue head from the schedule. Software should first deactivate all active qTDs, wait for the queue head to go inactive, then remove the queue head from the asynchronous list. Software removes a queue head from the asynchronous list using the following algorithm. Software merely must ensure all of the link pointers reachable by the host controller are kept consistent.

```

UnlinkQueueHead (pQHeadPrevious, pQueueHeadToUnlink, pQHeadNext)
--
-- Requirement: all inputs must be properly initialized.
--
-- pQHeadPrevious is a pointer to a queue head that
-- references the queue head to remove
-- pQHeadToUnlink is a pointer to the queue head to be
-- removed
-- pQHeadNext is a pointer to a queue head still in the
-- schedule. Software provides this pointer with the
-- following strict rules:
-- if the host software is one queue head, then
-- pQHeadNext must be the same as
-- QueueheadToUnlink.HorizontalPointer. If the host
-- software is unlinking a consecutive series of
-- queue heads, QHeadNext must be set by software to
-- the queue head remaining in the schedule.
--
-- This algorithm unlinks a queue head from a circular list
--
pQueueHeadPrevious.HorizontalPointer = pQueueHeadToUnlink.HorizontalPointer
pQueueHeadToUnlink.HorizontalPointer = pQHeadNext
End UnlinkQueueHead

```

If software removes the queue head with the H-bit set, it must select another queue head still linked into the schedule and set its H-bit. This should be completed before removing the queue head. The requirement is that software keep one queue head in the asynchronous schedule, with its H-bit set. At the point software has removed one or more queue heads from the asynchronous schedule, it is unknown whether the host controller

has a cached pointer to them. Similarly, it is unknown how long the host controller might retain the cached information, as it is implementation dependent and may be affected by the actual dynamics of the schedule load. Therefore, once software has removed a queue head from the asynchronous list, it must retain the coherency of the queue head (link pointers). It cannot disturb the removed queue heads until it knows that the host controller does not have a local copy of a pointer to any of the removed data structures.

The method software uses to determine when it is safe to modify a removed queue head is to handshake with the host controller. The handshake mechanism allows software to remove items from the asynchronous schedule, then execute a simple, lightweight handshake that is used by software as a key that it can free (or reuse) the memory associated the data structures it has removed from the asynchronous schedule.

The handshake is implemented with three bits in the host controller. The first bit is a command bit (USBCMD[IAA]-interrupt on async advance doorbell) that allows software to inform the host controller that something has been removed from its asynchronous schedule. The second bit is a status bit (USBSTS[AAI]-interrupt on async advance) that the host controller sets after it has released all on-chip state that may potentially reference one of the data structures just removed. When the host controller sets this status bit, it also clears the command bit. The third bit is an interrupt enable (USBINTR[AAE]-interrupt on async advance enable) that is matched with the status bit. If the status bit is set and the interrupt enable bit is set, then the host controller asserts a hardware interrupt.

The figure below illustrates a general example where consecutive queue heads (B and C) are unlinked from the schedule using the algorithm above. Before the unlink operation, the host controller has a copy of queue head A.

The unlink algorithm requires that as software unlinks each queue head, the unlinked queue head is loaded with the address of a queue head that will remain in the asynchronous schedule.

When the host controller observes that doorbell bit being set, it makes a note of the local reachable schedule information. In this example, the local reachable schedule information includes both queue heads (A & B). It is sufficient that the host controller can set the status bit (and clear the doorbell bit) as soon as it has traversed beyond current reachable schedule information (that is, traversed beyond queue head (B) in this example).

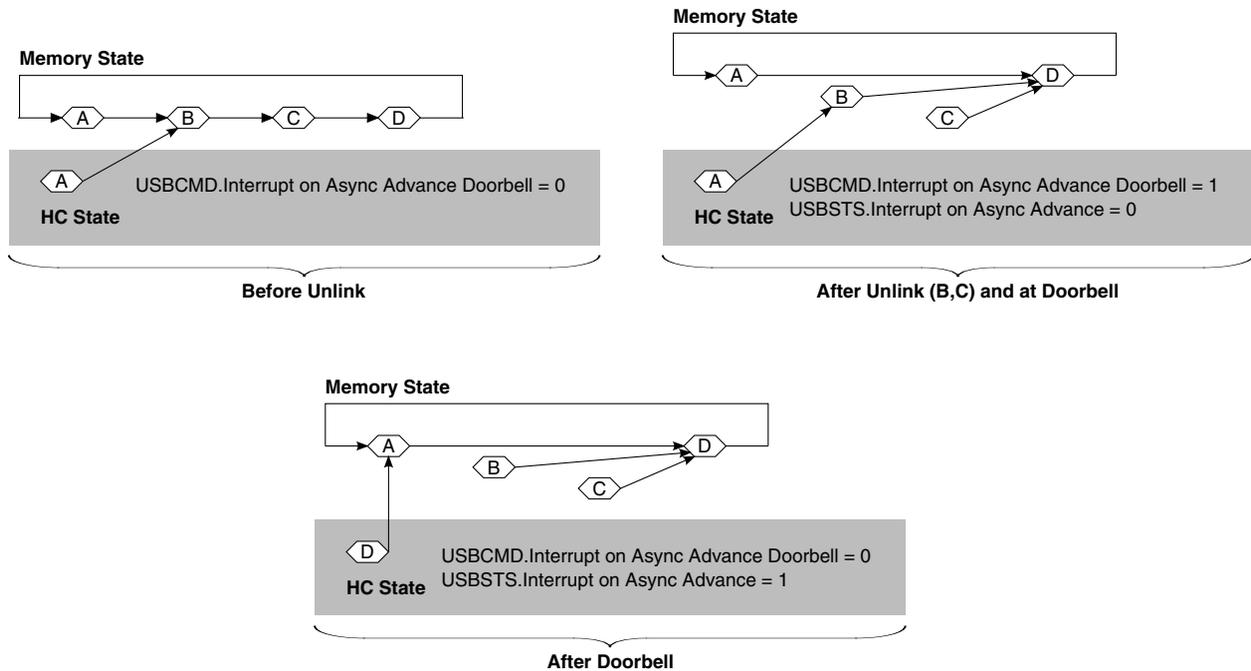


Figure 18-50. Generic queue head unlink scenario

Alternatively, a host controller implementation is allowed to traverse the entire asynchronous schedule list (for example, observed the head of the queue (twice)) before setting USBSTS[AAI].

Software may re-use the memory associated with the removed queue heads after it observes USBSTS[AAI] is set, following assertion of the doorbell. Software should acknowledge the interrupt on async advance status as indicated in the USBSTS register, before using the doorbell handshake again

18.6.9.3 Empty asynchronous schedule detection

EHCI uses two bits to detect when the asynchronous schedule is empty.

The queue head data structure (see [Table 18-66](#)) defines an H-bit in the queue head, which allows software to mark a queue head as being the head of the reclaim list. Host controller also keeps a 1-bit flag in the USBSTS register (Reclamation) that is cleared when the host controller observes a queue head with the H-bit set. The reclamation flag in the status register is set when any USB transaction from the asynchronous schedule is executed (or whenever the asynchronous schedule starts, see [Asynchronous schedule traversal: Start event](#)).

If the controller ever encounters an H-bit of one and a Reclamation bit of zero, the controller simply stops traversal of the asynchronous schedule.

The figure below shows an example illustrating the H-bit in a schedule.

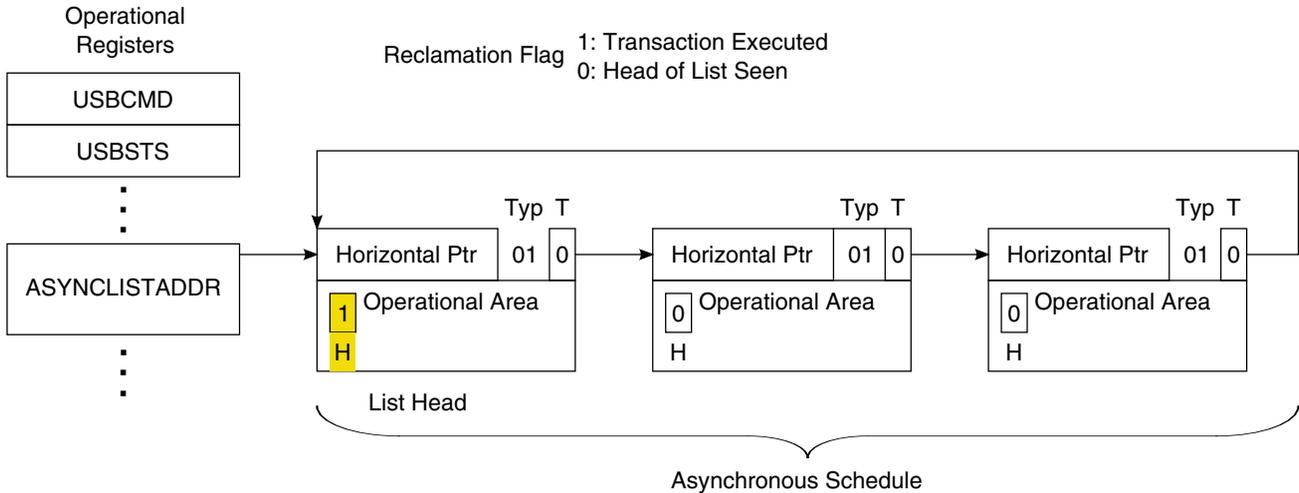


Figure 18-51. Asynchronous schedule list with annotation to mark head of list

18.6.9.4 Asynchronous schedule traversal: Start event

Once the host controller has idled itself using the empty schedule detection, it naturally activates and begins processing from the Periodic Schedule at the beginning of each microframe.

In addition, it may have idled itself early in a microframe. When this occurs (idles early in the microframe) the host controller must occasionally reactivate during the microframe and traverse the asynchronous schedule to determine whether any progress can be made. Asynchronous schedule Start Events are defined to be:

- Whenever the host controller transitions from the periodic schedule to the asynchronous schedule. If the periodic schedule is disabled and the asynchronous schedule is enabled, then the beginning of the microframe is equivalent to the transition from the periodic schedule, or
- The asynchronous schedule traversal restarts from a sleeping state.

18.6.9.5 Reclamation status bit (USBSTS Register)

The operation of the empty asynchronous schedule detection feature depends on the proper management of the Reclamation bit (RCL) in the USBSTS register.

The host controller tests for an empty schedule just after it fetches a new queue head while traversing the asynchronous schedule. The host controller sets USBSTS[RCL] whenever an asynchronous schedule traversal Start Event occurs. USBSTS[RCL] is also set whenever the host controller executes a transaction while traversing the asynchronous schedule. The host controller clears USBSTS[RCL] whenever it finds a queue head with its H-bit set. Software should only set a queue head's H-bit if the queue head is in the asynchronous schedule. If software sets the H-bit in an interrupt queue head, the resulting behavior is undefined. The host controller may clear USBSTS[RCL] when executing from the periodic schedule.

18.6.10 Managing control/bulk/interrupt transfers via queue heads

This section presents an overview of how the host controller interacts with queuing data structures.

Queue heads use the Queue Element Transfer Descriptor (qTD) structure defined in [Queue element transfer descriptor \(qTD\)](#).

One queue head is used to manage the data stream for one endpoint. The queue head structure contains static endpoint characteristics and capabilities. It also contains a working area from where individual bus transactions for an endpoint are executed. Each qTD represents one or more bus transactions, which is defined in the context of the EHCI specification as a transfer.

The general processing model for the host controller's use of a queue head is simple:

- Read a queue head,
- Execute a transaction from the overlay area,
- Write back the results of the transaction to the overlay area
- Move to the next queue head.

If the host controller encounters errors during a transaction, the host controller will set one of the error reporting bits in the queue head's Status field. The Status field accumulates all errors encountered during the execution of a qTD (that is, the error bits in the queue head Status field are sticky until the transfer (qTD) has completed). This state is always written back to the source qTD when the transfer is complete. On transfer (for example, buffer or halt conditions) boundaries, the host controller must auto-advance (without software intervention) to the next qTD. Additionally, the hardware must be able to halt the queue so no additional bus transactions will occur for the endpoint and the host controller will not advance the queue.

18.6.10.1 Buffer pointer list use for data streaming with qTDs

A qTD has an array of buffer pointers, which is used to reference the data buffer for a transfer.

The EHCI specification requires that the buffer associated with the transfer be virtually contiguous. This means that if the buffer spans more than one physical page, it must obey the following rules:

- The first portion of the buffer must begin at some offset in a page and extend through the end of the page.
- The remaining buffer cannot be allocated in small chunks scattered around memory. For each 4K chunk beyond the first page, each buffer portion matches to a full 4K page. The final portion, which may only be large enough to occupy a portion of a page, must start at the top of the page and be contiguous within that page.

The figure below illustrates these requirements.

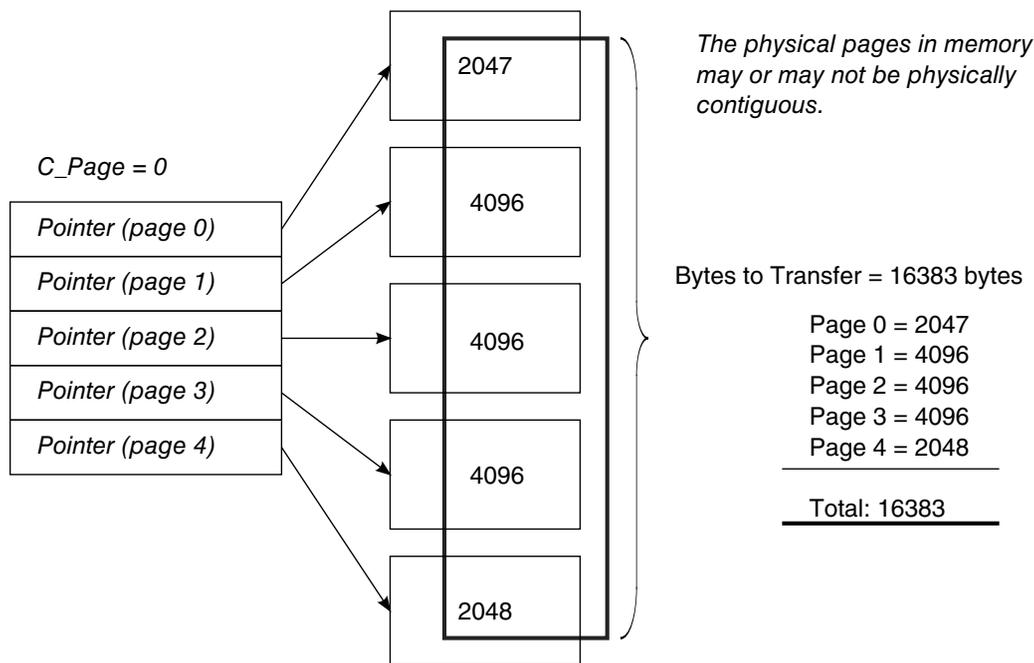


Figure 18-52. Example mapping of qTD buffer pointers to buffer pages

The buffer pointer list in the qTD is long enough to support a maximum transfer size of 20K bytes. This case occurs when all five buffer pointers are used and the first offset is zero. A qTD handles a 16Kbyte buffer with any starting buffer alignment.

The host controller uses the `C_Page` field as an index value to determine which buffer pointer in the list should be used to start the current transaction. The host controller uses a different buffer pointer for each physical page of the buffer. This is always true, even if the buffer is physically contiguous.

The host controller must detect when the current transaction spans a page boundary and automatically move to the next available buffer pointer in the page pointer list. The next available pointer is reached by incrementing `C_Page` and pulling the next page pointer from the list. Software must ensure there are sufficient buffer pointers to move the amount of data specified in the Bytes to Transfer field.

The figure above illustrates a nominal example of how System software would initialize the buffer pointers list and the `C_Page` field for a transfer size of 16383 bytes. `C_Page` is cleared. The upper 20-bits of Page 0 references the start of the physical page. Current Offset (the lower 12-bits of queue head Dword 7) holds the offset in the page for example, 2049 (for example, 4096-2047). The remaining page pointers are set to reference the beginning of each subsequent 4K page.

For the first transaction on the qTD (assuming a 512-byte transaction), the host controller uses the first buffer pointer (page 0 because `C_Page` is cleared) and concatenates the Current Offset field. The 512 bytes are moved during the transaction, the Current Offset and Total Bytes to Transfer are adjusted by 512 and written back to the queue head working area.

During the 4th transaction, the host controller needs 511 bytes in page 0 and one byte in page 1. The host controller will increment `C_Page` (to 1) and use the page 1 pointer to move the final byte of the transaction. After the 4th transaction, the active page pointer is the page 1 pointer and Current Offset has rolled to one, and both are written back to the overlay area. The transactions continue for the rest of the buffer, with the host controller automatically moving to the next page pointer (that is, `C_Page`) when necessary. There are three conditions for how the host controller handles `C_Page`.

- The current transaction does not span a page boundary. The value of `C_Page` is not adjusted by the host controller.
- The current transaction does span a page boundary. The host controller must detect the page cross condition and advance to the next buffer while streaming data to/from the USB.
- The current transaction completes on a page boundary (that is, the last byte moved for the current transaction is the last byte in the page for the current page pointer). The host controller must increment `C_Page` before writing back status for the transaction.

Note that the only valid adjustment the host controller may make to `C_Page` is to increment by one.

18.6.10.2 Adding interrupt queue heads to the periodic schedule

The link path(s) from the periodic frame list to a queue head establishes in which frames a transaction can be executed for the queue head.

Queue heads are linked into the periodic schedule so they are polled at the appropriate rate. System software sets a bit in a queue head's S-Mask to indicate which microframe within a 1 millisecond period a transaction should be executed for the queue head.

Software must ensure that all queue heads in the periodic schedule have S-Mask set to a non-zero value. An S-mask with a zero value in the context of the periodic schedule yields undefined results.

If the desired poll rate is greater than one frame, system software can use a combination of queue head linking and S-Mask values to spread interrupts of equal poll rates through the schedule so that the periodic bandwidth is allocated and managed in the most efficient manner possible. Some examples are illustrated in the table below.

Table 18-77. Example periodic reference patterns for interrupt transfers

Frame # reference sequence	Description
0, 2, 4, 6, 8, S-Mask = 0x01	A queue head for the blnterval of 2 milliseconds (16 microframes) is linked into the periodic schedule so that it is reachable from the periodic frame list locations indicated in the previous column. In addition, the S-Mask field in the queue head is set to 0x01, indicating that the transaction for the endpoint should be executed on the bus during microframe 0 of the frame.
0, 2, 4, 6, 8, ... S-Mask = 0x02	Another example of a queue head with a blnterval of 2 milliseconds is linked into the periodic frame list at exactly the same interval as the previous example. However, the S-Mask is set to 0x02 indicating that the transaction for the endpoint should be executed on the bus during microframe 1 of the frame.

18.6.10.3 Managing transfer complete interrupts from queue heads

The host controller sets an interrupt to be signaled at the next interrupt threshold when the completed transfer (qTD) has an Interrupt on Complete (IOC) bit set, or whenever a transfer (qTD) completes with a short packet.

If system software needs multiple qTDs to complete a client request (that is, like a control transfer) the intermediate qTDs do not require interrupts. System software may only need a single interrupt to notify it that the complete buffer has been transferred. System software may set IOC's to occur more frequently. A motivation for this may be that it wants early notification so that interface data structures can be re-used in a timely manner.

18.6.11 Ping control

USB 2.0 defines an addition to the protocol for high-speed devices called Ping.

Ping is required for all USB 2.0 High-speed bulk and control endpoints. Ping is not allowed for a split-transaction stream. This extension to the protocol eliminates the bad side-effects of Naking OUT endpoints. The Status field has a Ping State bit, which the host controller uses to determine the next actual PID it will use in the next transaction to the endpoint (see [Table 18-64](#)). The Ping State bit is only managed by the host controller for queue heads that meet all of the following criteria:

- The queue head is not an interrupt
- The EPS field equals High-Speed
- The PIDCode field equals OUT

The table below illustrates the state transition table for the host controller's responsibility for maintaining the PING protocol. Refer to Chapter 8 in the *USB Specification, Revision 2.0* for detailed description on the Ping protocol.

Table 18-78. Ping control state transition table

Current	Event		Next
	Host	Device	
Do Ping	PING	Nak	Do Ping
Do Ping	PING	Ack	Do OUT
Do Ping	PING	XactErr ¹	Do Ping
Do Ping	PING	Stall	N/C ²
Do OUT	OUT	Nak	Do Ping
Do OUT	OUT	Nyet	Do Ping ³
Do OUT	OUT	Ack	Do OUT
Do OUT	OUT	XactErr ¹	Do Ping
Do OUT	OUT	Stall	N/C ²

1. Transaction Error (XactErr) is any time the host misses the handshake.
2. No transition change required for the Ping State bit. The Stall handshake results in the endpoint being halted (for example, Active cleared and Halt set). Software intervention is required to restart queue.
3. A Nyet response to an OUT means that the device has accepted the data, but cannot receive any more at this time. Host must advance the transfer state and additionally, transition the Ping State bit to Do Ping.

The Ping State bit is described in [Table 18-64](#). The defined ping protocol allows the host to be imprecise on the initialization of the ping protocol (that is, start in Do OUT when there is uncertainty about the space in the device). The host controller manages the Ping State bit. System software sets the initial value in the queue head when it initializes a

queue head. The host controller preserves the Ping State bit across all queue advancements. This means that when a new qTD is written into the queue head overlay area, the previous value of the Ping State bit is preserved.

18.6.12 Split transactions

USB 2.0 defines extensions to the bus protocol for managing USB 1.x data streams through USB 2.0 hubs. This section describes how the host controller uses the interface data structures to manage data streams with full- and low-speed devices, connected below a USB 2.0 hub, utilizing the split transaction protocol. Refer to the USB 2.0 Specification for the complete definition of the split transaction protocol.

Full- and low-speed devices are enumerated identically as high-speed devices, but the transactions to the full- and low-speed endpoints use the split-transaction protocol on the high-speed bus.

The split transaction protocol is an encapsulation of (or wrapper around) the full- or low-speed transaction. The high-speed wrapper portion of the protocol is addressed to the USB 2.0 hub and transaction translator below which the full- or low-speed device is attached.

EHCI uses dedicated data structures for managing full-speed isochronous data streams. Control, Bulk and Interrupt are managed using the queuing data structures. The interface data structures need to be programmed with the device address and the transaction translator number of the USB 2.0 hub operating as the low-/full-speed host controller for this link. The following sections describe the details of how the host controller processes and manages the split transaction protocol.

18.6.12.1 Split transactions for asynchronous transfers

A queue head in the asynchronous schedule with an EPS field indicating a full-or low-speed device indicates to the host controller that it must use split transactions to stream data for this queue head.

All full-speed bulk and full-, low-speed control are managed via queue heads in the asynchronous schedule.

Software must initialize the queue head with the appropriate device address and port number for the transaction translator that is serving as the full-/low-speed host controller for the links connecting the endpoint. Software must also initialize the split transaction state bit (SplitXState) to Do-Start-Split. Finally, if the endpoint is a control endpoint, then system software must set the Control Transfer Type (C) bit in the queue head to a one. If

this is not a control transfer type endpoint, the C bit must be initialized by software to be a zero. This information is used by the host controller to properly set the Endpoint Type (ET) field in the split transaction bus token. When the C bit is a zero, the split transaction token's ET field is set to indicate a bulk endpoint. When the C bit is a one, the split transaction token's ET field is set to indicate a control endpoint. Refer to Chapter 8 of *USB Specification, Revision 2.0* for details.

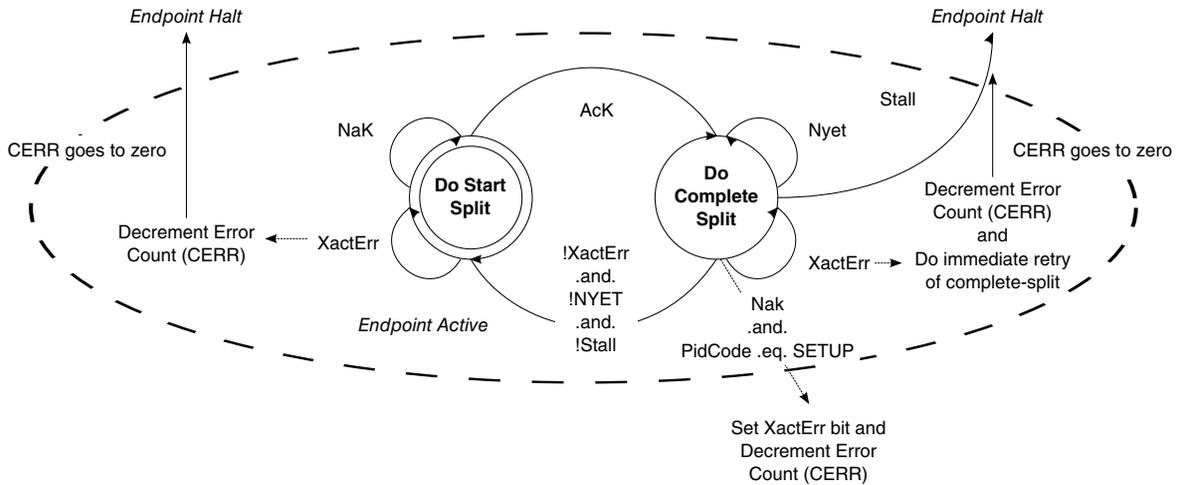


Figure 18-53. Host controller asynchronous schedule split-transaction state machine

18.6.12.1.1 Asynchronous-do-start-split

Do-start-split is the state which software must initialize a full- or low-speed asynchronous queue head.

This state is entered from the Do-Complete-Split state only after a complete-split transaction receives a valid response from the transaction translator that is not a Nyet handshake.

For queue heads in this state, the host controller executes a start-split transaction to the transaction translator. If the bus transaction completes without an error and PID Code indicates an IN or OUT transaction, then the host controller reloads the error counter (Cerr). If it is a successful bus transaction and the PID Code indicates a SETUP, the host controller will not reload the error counter. If the transaction translator responds with a Nak, the queue head is left in this state, and the host controller proceeds to the next queue head in the asynchronous schedule.

If the host controller times out the transaction (no response, or bad response) the host controller decrements Cerr and proceeds to the next queue head in the asynchronous schedule.

18.6.12.1.2 Asynchronous-do-complete-split

This state is entered from the Do-Start-Split state only after a start-split transaction receives an Ack handshake from the transaction translator.

For queue heads in this state, the host controller executes a complete-split transaction to the transaction translator. If the transaction translator responds with a Nyet handshake, the queue head is left in this state, the error counter is reset and the host controller proceeds to the next queue head in the asynchronous schedule. When a Nyet handshake is received for a bus transaction where the queue head's PID Code indicates an IN or OUT, the host controller reloads the error counter (Cerr). When a Nyet handshake is received for a complete-split bus transaction where the queue head's PID Code indicates a SETUP, the host controller must not adjust the value of Cerr.

Independent of PID Code, the following responses have the indicated effects:

- Transaction Error (XactErr). Timeout/data CRC failure. The error counter (Cerr) is decremented by one and the complete split transaction is immediately retried (if possible). If there is not enough time in the microframe to execute the retry, the host controller ensures that the next time the host controller begins executing from the Asynchronous schedule, it must begin executing from this queue head. If another start-split (for some other endpoint) is sent to the transaction translator before the complete-split is really completed, the transaction translator could dump the results (which were never delivered to the host). This is why the core specification states the retries must be immediate. When the host controller returns to the asynchronous schedule in the next microframe, the first transaction from the schedule will be the retry for this endpoint. If Cerr went to zero, the host controller halts the queue.
- NAK. The target endpoint Nak'd the full- or low-speed transaction. The state of the transfer is not advanced and the state is exited. If the PID Code is a SETUP, then the Nak response is a protocol error. The XactErr status bit is set and the Cerr field is decremented.
- STALL. The target endpoint responded with a STALL handshake. The host controller sets the halt bit in the status byte, retires the qTD but does not attempt to advance the queue.

If the PID Code indicates an IN, then any of following responses are expected:

- DATA0/1. On reception of data, the host controller ensures the PID matches the expected data toggle and checks CRC. If the packet is good, the host controller advances the state of the transfer (for example, moves the data pointer by the number of bytes received, decrements the BytesToTransfer field by the number of bytes received, and toggles the dt bit). The host controller then exits this state. The response and advancement of transfer may trigger other processing events, such as retirement of the qTD and advancement of the queue.

If the data sequence PID does not match the expected, the data is ignored, the transfer state is not advanced and this state is exited.

If the PID Code indicates an OUT/SETUP, then any of following responses are expected:

- **ACK.** The target endpoint accepted the data, so the host controller must advance the state of the transfer. The Current Offset field is incremented by Maximum Packet Length or Bytes to Transfer, whichever is less. The Bytes To Transfer field is decremented by the same amount and the data toggle bit (dt) is toggled. The host controller then exits this state.

Advancing the transfer state may cause other processing events such as retirement of the qTD and advancement of the queue.

18.6.12.2 Split transaction interrupt

Split-transaction interrupt-IN/OUT endpoints are managed using the same data structures used for high-speed interrupt endpoints.

They both co-exist in the periodic schedule. Queue heads/qTDs offer the set of features required for reliable data delivery, which is characteristic to interrupt transfer types. The split-transaction protocol is managed completely within this defined functional transfer framework. For example, for a high-speed endpoint, the host controller will visit a queue head, execute a high-speed transaction (if criteria are met) and advance the transfer state (or not) depending on the results of the entire transaction. For low- and full-speed endpoints, the details of the execution phase are different (that is, takes more than one bus transaction to complete), but the remainder of the operational framework is intact.

18.6.12.2.1 Split transaction scheduling mechanisms for interrupt

Full- and low-speed interrupt queue heads have an EPS field indicating full- or low-speed and have a non-zero S-mask field.

The host controller can detect this combination of parameters and assume the endpoint is a periodic endpoint. Low- and full-speed interrupt queue heads require the use of the split transaction protocol. The host controller sets the Endpoint Type (ET) field in the split token to indicate the transaction is an interrupt. These transactions are managed through a transaction translator's periodic pipeline. Software should not set these fields to indicate the queue head is an interrupt unless the queue head is used in the periodic schedule.

System software manages the per/transaction translator periodic pipeline by budgeting and scheduling exactly during which microframes the start-splits and complete-splits for each endpoint will occur. The characteristics of the transaction translator are such that the high-speed transaction protocol must execute during explicit microframes, or the data or response information in the pipeline is lost. The figure below illustrates the general scheduling boundary conditions that are supported by the EHCI periodic schedule and queue head data structure. The S and C_n labels indicate microframes where software can schedule start-splits and complete splits (respectively).

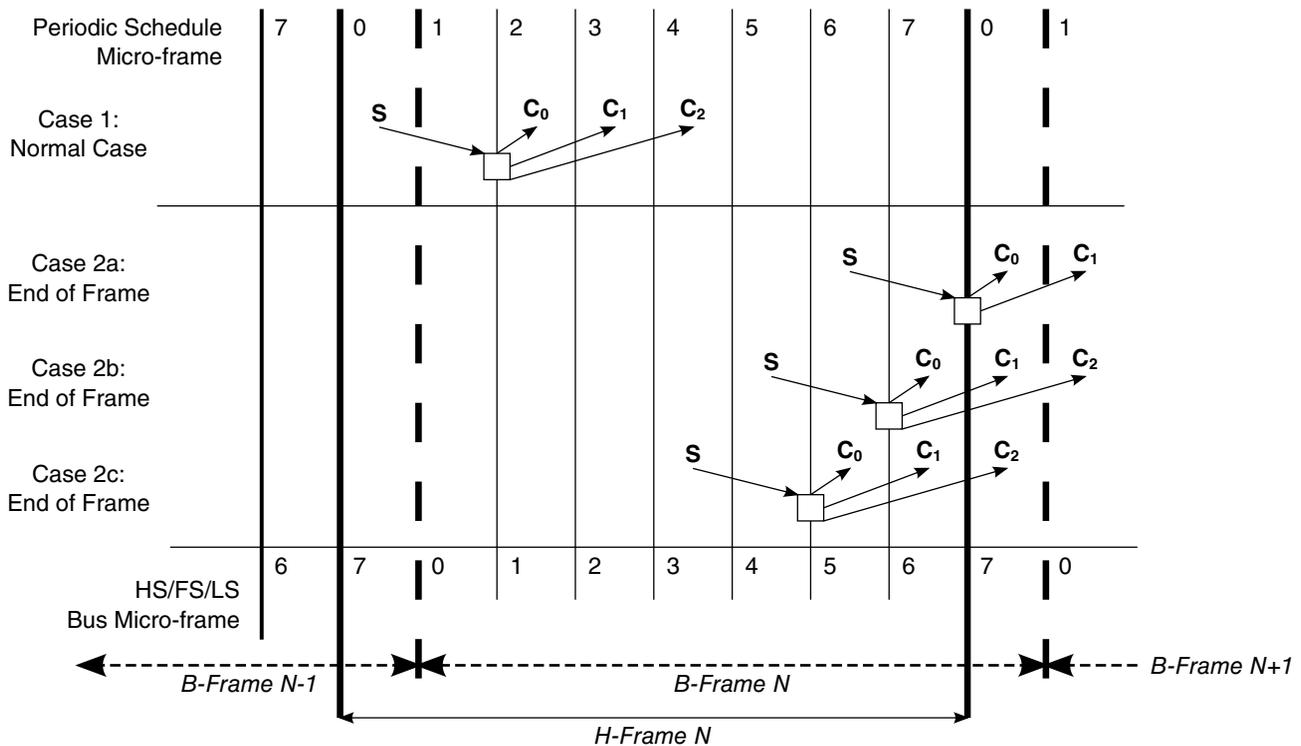


Figure 18-54. Split transaction, interrupt scheduling boundary conditions

The scheduling cases are:

- Case 1: The normal scheduling case is where the entire split transaction is completely bounded by a frame (H-Frame in this case).
- Case 2a through Case 2c: The USB 2.0 hub pipeline rules states clearly, when and how many complete-splits must be scheduled to account for earliest to latest execution on the full/low-speed link. The complete-splits may span the H-Frame boundary when the start-split is in microframe 4 or later. When this occurs, the H-Frame to B-Frame alignment requires that the queue head be reachable from consecutive periodic frame list locations. System software cannot build an efficient schedule that satisfies this requirement unless it uses FSTNs.

The figure below illustrates the general layout of the periodic schedule.

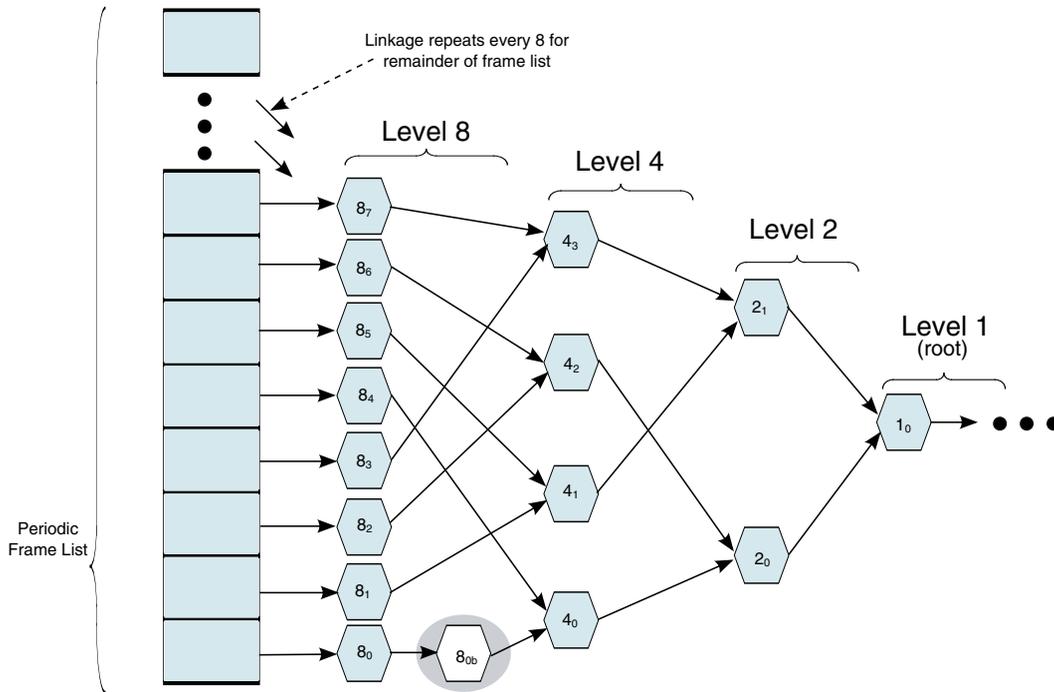


Figure 18-55. General structure of EHCI periodic schedule utilizing interrupt spreading

The periodic frame list is effectively the leaf level a binary tree, which is always traversed leaf to root. Each level in the tree corresponds to a 2^N poll rate. Software can efficiently manage periodic bandwidth on the USB by spreading interrupt queue heads that have the same poll rate requirement across all the available paths from the frame list. For example, system software can schedule eight poll rate 8 queue heads and account for them once in the high-speed bus bandwidth allocation.

When an endpoint is allocated an execution footprint that spans a frame boundary, the queue head for the endpoint must be reachable from consecutive locations in the frame list. An example would be if 8_{0b} were such an endpoint. Without additional support on the interface, to get 8_{0b} reachable at the correct time, software would have to link 8_1 to 8_{0b} . It would then have to move 4_1 and everything linked after into the same path as 4_0 . This upsets the integrity of the binary tree and disallows the use of the spreading technique.

FSTN data structures are used to preserve the integrity of the binary-tree structure and enable the use of the spreading technique. [Periodic frame span traversal node \(FSTN\)](#), defines the hardware and software operational model requirements for using FSTNs.

The following queue head fields are initialized by system software to instruct the host controller when to execute portions of the split-transaction protocol.

- **SplitXState**. This is a single bit residing in the Status field of a queue head ([Table 18-64](#)). This bit is used to track the current state of the split transaction.

- **Frame S-mask.** This is a bit-field where-in system software sets a bit corresponding to the microframe (within an H-Frame) that the host controller should execute a start-split transaction. This is always qualified by the value of the SplitXState bit in the Status field of the queue head. For example, referring to [Figure 18-54](#), case one, the S-mask would have a value of 0b0000_0001 indicating that if the queue head is traversed by the host controller, and the SplitXState indicates Do_Start, and the current microframe as indicated by FRINDEX[2-0] is 0, then execute a start-split transaction.
- **Frame C-mask.** This is a bit-field where system software sets one or more bits corresponding to the microframes (within an H-Frame) that the host controller should execute complete-split transactions. The interpretation of this field is always qualified by the value of the SplitXState bit in the Status field of the queue head. For example, referring to [Figure 18-54](#), case one, the C-mask would have a value of 0b0001_1100 indicating that if the queue head is traversed by the host controller, and the SplitXState indicates Do_Complete, and the current microframe as indicated by FRINDEX[2-0] is 2, 3, or 4, then execute a complete-split transaction. It is software's responsibility to ensure that the translation between H-Frames and B-Frames is correctly performed when setting bits in S-mask and C-mask.

18.6.12.2.2 Host controller operational model for FSTNs

The FSTN data structure is used to manage Low/Full-speed interrupt queue heads that need to be reached from consecutive frame list locations (that is, boundary cases 2a through 2c).

An FSTN is essentially a back pointer, similar in intent to the back pointer field in the siTD data structure.

This feature provides software a simple primitive to save a schedule position, redirect the host controller to traverse the necessary queue heads in the previous frame, then restore the original schedule position and complete normal traversal.

There are four components to the use of FSTNs:

- FSTN data structure, defined in [Periodic frame span traversal node \(FSTN\)](#).
- A Save Place indicator; this is always an FSTN with its Back Path Link Pointer[T] bit cleared.
- A Restore indicator; this is always an FSTN with its Back Path Link Pointer[T] bit set.
- Host controller FSTN traversal rules.

When the host controller encounters an FSTN during microframes 2 through 7 it simply follows the node's Normal Path Link Pointer to access the next schedule data structure. Note that the FSTN's Normal Path Link Pointer[T] bit may set, which the host controller must interpret as the end of periodic list mark.

When the host controller encounters a Save-Place FSTN in microframes 0 or 1, it saves the value of the Normal Path Link Pointer and sets an internal flag indicating that it is executing in Recovery Path mode. Recovery Path mode modifies the host controller's rules for how it traverses the schedule and limits which data structures are considered for execution of bus transactions. The host controller continues executing in Recovery Path mode until it encounters a Restore FSTN or it determines that it has reached the end of the microframe.

The rules for schedule traversal and limited execution while in Recovery Path mode are:

- Always follow the Normal Path Link Pointer when it encounters an FSTN that is a Save-Place indicator. The host controller must not recursively follow Save-Place FSTNs. Therefore, while executing in Recovery Path mode, it must never follow an FSTN's Back Path Link Pointer.
- Do not process an siTD or iTD data structure; simply follow its Next Link Pointer.
- Do not process a QH (Queue Head) whose EPS field indicates a high-speed device; simply follow its Horizontal Link Pointer.
- When a QH's EPS field indicates a Full/Low-speed device, the host controller only considers it for execution if its SplitXState is DoComplete (note: this applies whether the PID Code indicates an IN or an OUT). Refer to the *EHCI Specification* for a complete list of additional conditions that must be met in general for the host controller to issue a bus transaction. Note that the host controller must not execute a Start-split transaction while executing in Recovery Path mode. Refer to the *EHCI Specification* for special handling when in Recovery Path mode.
- Stop traversing the recovery path when it encounters an FSTN that is a Restore indicator. The host controller unconditionally uses the saved value of the Save-Place FSTN's Normal Path Link Pointer when returning to the normal path traversal. The host controller must clear the context of executing a Recovery Path when it restores schedule traversal to the Save-Place FSTN's Normal Path Link Pointer.

If the host controller determines that there is not enough time left in the microframe to complete processing of the periodic schedule, it abandons traversal of the recovery path, and clears the context of executing a recovery path. The result is that at the start of the next consecutive microframe, the host controller starts traversal at the frame list.

An example traversal of a periodic schedule that includes FSTNs is illustrated in the figure below.

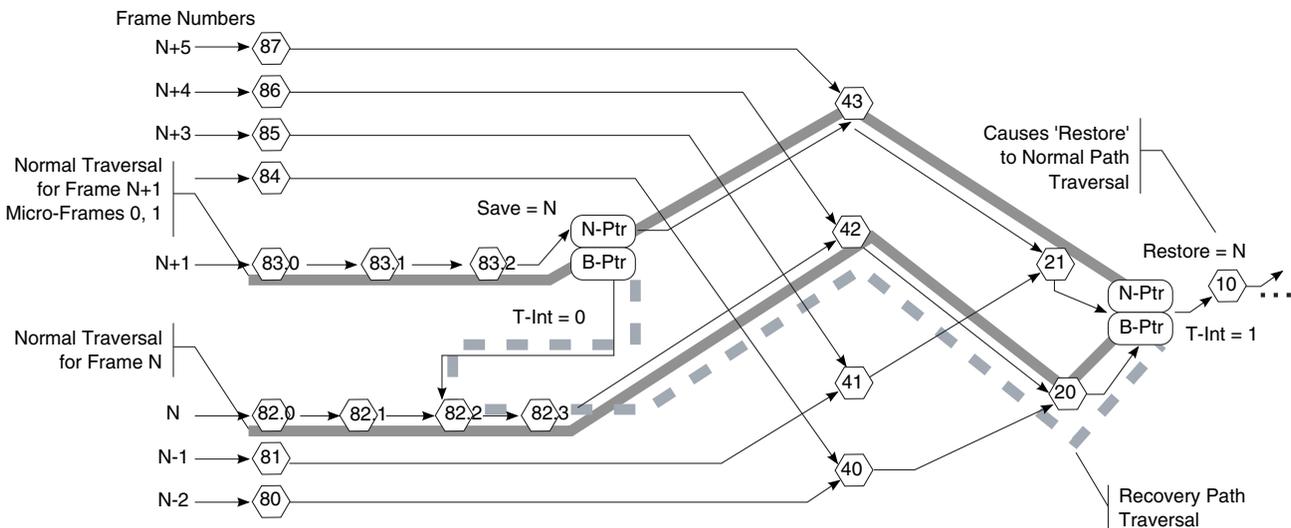


Figure 18-56. Example host controller traversal of recovery path via FSTNs

In frame N (microframes 0-7), for this example, the host controller traverses all of the schedule data structures utilizing the Normal Path Link Pointers in any FSTNs it encounters. This is because the host controller has not yet encountered a Save-Place FSTN so it is not executing in Recovery Path mode. When it encounters the Restore FSTN, (Restore-N), during microframes 0 and 1, it uses Restore-N. Normal Path Link Pointer to traverse to the next data structure (that is, normal schedule traversal). This is because the host controller must use a Restore FSTN's Normal Path Link Pointer when not executing in a Recovery-Path mode. The nodes traversed during frame N include: {82.0, 82.1, 82.2, 82.3, 42, 20, Restore-N, 10 ...}.

In frame N+1 (microframes 0 and 1), when the host controller encounters Save-Path FSTN (Save-N), it observes that Save-N.Back Path Link Pointer.T-bit is zero (definition of a Save-Path indicator). The host controller saves the value of Save-N. Normal Path Link Pointer and follows Save-N.Back Path Link Pointer. At the same time, it sets an internal flag indicating that it is now in Recovery Path mode (the recovery path is annotated in the figure above with a large dashed line). The host controller continues traversing data structures on the recovery path and executing only those bus transactions as noted above, on the recovery path until it reaches Restore FSTN (Restore-N). Restore-N.Back Path Link Pointer.T-bit is set (definition of a Restore indicator), so the host controller exits Recovery Path mode by clearing the internal Recovery Path mode flag and commences (restores) schedule traversal using the saved value of the Save-Place FSTN's Normal Path Link Pointer (for example, Save-N.Normal Path Link Pointer). The nodes traversed during these microframes include: {83.0, 83.1, 83.2, Save-A, 82.2, 82.3, 42, 20, Restore-N, 43, 21, Restore-N, 10 ...}.

In frame N+1 (microframes 2-7), when the host controller encounters Save-Path FSTN Save-N, it unconditionally follows Save-N.Normal Path Link Pointer. The nodes traversed during these microframes include: $\{8_{3,0}, 8_{3,1}, 8_{3,2}, \text{Save-A}, 4_3, 2_1, \text{Restore-N}, 1_0 \dots\}$.

18.6.12.2.3 Software operational model for FSTNs

Software must create a consistent, coherent schedule for the host controller to traverse.

When using FSTNs, system software must adhere to the following rules:

- Each Save-Place indicator requires a matching Restore indicator.

The Save-Place indicator is an FSTN with a valid Back Path Link Pointer and T-bit equal to zero. Note that Back Path Link Pointer[Typ] field must be set to indicate the referenced data structure is a queue head. The Restore indicator is an FSTN with its Back Path Link Pointer[T] bit set.

A Restore FSTN may be matched to one or more Save-Place FSTNs. For example, if the schedule includes a poll-rate 1 level, then system software only needs to place a Restore FSTN at the beginning of this list in order to match all possible Save-Place FSTNs.

- If the schedule does not have elements linked at a poll-rate level of one, and one or more Save-Place FSTNs are used, then System Software must ensure the Restore FSTN's Normal Path Link Pointer's T-bit is set, as this is used to mark the end of the periodic list.
- When the schedule does have elements linked at a poll rate level of one, a Restore FSTN must be the first data structure on the poll rate one list. All traversal paths from the frame list converge on the poll-rate one list. System software must ensure that Recovery Path mode is exited before the host controller is allowed to traverse the poll rate level one list.
- A Save-Place FSTN's Back Path Link Pointer must reference a queue head data structure. The referenced queue head must be reachable from the previous frame list location. In other words, if the Save-Place FSTN is reachable from frame list offset N, then the FSTN's Back Path Link Pointer must reference a queue head that is reachable from frame list offset N-1.

Software should make the schedule as efficient as possible. What this means in this context is that software should have no more than one Save-Place FSTN reachable in any single frame. Note there are times when two (or more, depending on the implementation) could exist as full-/low-speed footprints change with bandwidth adjustments. This could

occur, for example when a bandwidth rebalance causes system software to move the Save-Place FSTN from one poll rate level to another. During the transition, software must preserve the integrity of the previous schedule until the new schedule is in place.

18.6.12.2.4 Tracking split transaction progress for interrupt transfers

To correctly maintain the data stream, the host controller must be able to detect and report errors where data is lost.

For interrupt-IN transfers, data is lost when it makes it into the USB 2.0 hub, but the USB 2.0 host system is unable to get it from the USB 2.0 hub and into the system before it expires from the transaction translator pipeline. When a lost data condition is detected, the queue is halted, thus signaling system software to recover from the error. A data-loss condition exists whenever a start-split is issued, accepted and successfully executed by the USB 2.0 hub, but the complete-splits get unrecoverable errors on the high-speed link, or the complete-splits do not occur at the correct times. One reason complete-splits might not occur at the right time would be due to host-induced system hold-offs that cause the host controller to miss bus transactions because it cannot get timely access to the schedule in system memory.

The same condition can occur for an interrupt-OUT, but the result is not an endpoint halt condition, but rather effects only the progress of the transfer. The queue head has the following fields to track the progress of each split transaction. These fields are used to keep incremental state about which (and when) portions have been executed.

- **C-prog-mask.** This is an eight-bit bit-vector where the host controller keeps track of which complete-splits have been executed. Due to the nature of the transaction translator periodic pipeline, the complete-splits need to be executed in-order. The host controller needs to detect when the complete-splits have not been executed in order. This can only occur due to system hold-offs where the host controller cannot get to the memory-based schedule. C-prog-mask is a simple bit-vector that the host controller sets one of the C-prog-mask bits for each complete-split executed. The bit position is determined by the microframe number in which the complete-split was executed. The host controller always checks C-prog-mask before executing a complete-split transaction. If the previous complete-splits have not been executed then it means one (or more) have been skipped and data has potentially been lost.
- **FrameTag.** This field is used by the host controller during the complete-split portion of the split transaction to tag the queue head with the frame number (H-Frame number) when the next complete split must be executed.
- **S-bytes.** This field can be used to store the number of data payload bytes sent during the start-split (if the transaction was an OUT). The S-bytes field must be used to accumulate the data payload bytes received during the complete-splits (for an IN).

18.6.12.2.5 Split transaction execution state machine for interrupt

In the following section, all references to microframe are in the context of a microframe within an H-Frame.

As with asynchronous full- and low-speed endpoints, a split-transaction state machine is used to manage the split transaction sequence. Aside from the fields defined in the queue head for scheduling and tracking the split transaction, the host controller calculates one internal mechanism that is also used to manage the split transaction. The internal calculated mechanism is:

- **cMicroFrameBit**. This is a single-bit encoding of the current microframe number. It is an eight-bit value calculated by the host controller at the beginning of every microframe. It is calculated from the three least significant bits of the FRINDEX register (that is, $cMicroFrameBit = (1 \text{ shifted-left}(FRINDEX[2-0]))$). The cMicroFrameBit has at most one bit asserted, which always corresponds to the current microframe number. For example, if the current microframe is 0, then cMicroFrameBit will equal 0b0000_0001.

The variable cMicroFrameBit is used to compare against the S-mask and C-mask fields to determine whether the queue head is marked for a start- or complete-split transaction for the current microframe.

The figure below illustrates how a complete interrupt split transaction is managed. There are two phases to each split transaction. The first is a single start-split transaction, which occurs when the SplitXState is at Do_Start and the single bit in cMicroFrameBit has a corresponding bit active in QH[S-mask]. The transaction translator does not acknowledge the receipt of the periodic start-split, so the host controller unconditionally transitions the state to Do_Complete. Due to the available jitter in the transaction translator pipeline, there is more than one complete-split transaction scheduled by software for the Do_Complete state. This translates simply to the fact that there are multiple bits set in the QH[C-mask] field.

The host controller keeps the queue head in the Do_Complete state until the split transaction is complete (see definition below), or an error condition triggers the three-strikes-rule (for example, after the host tries the same transaction three times, and each encounters an error, the host controller stops retrying the bus transaction and halts the endpoint, thus requiring system software to detect the condition and perform system-dependent recovery).

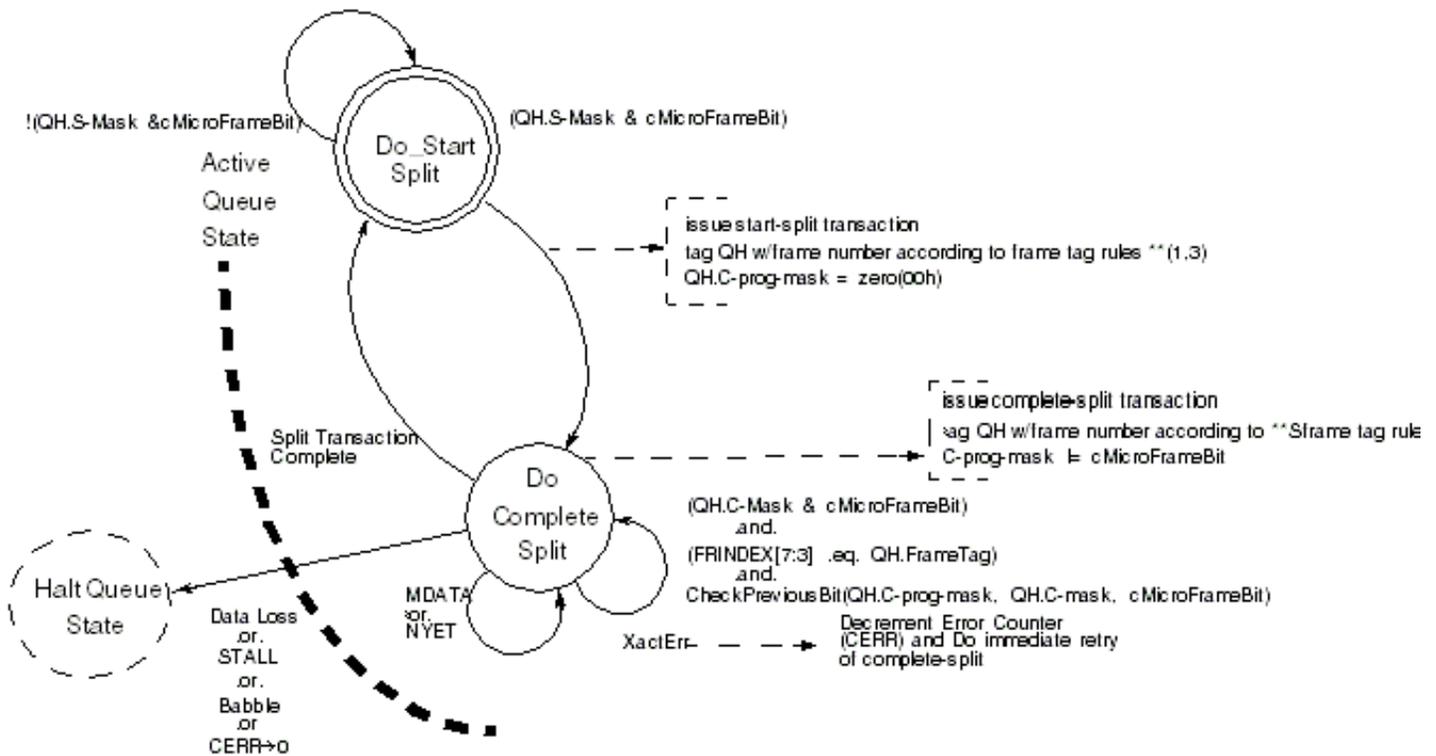


Figure 18-57. Split transaction state machine for interrupt

18.6.12.2.6 Periodic interrupt-do-start-split

This is the state when software must initialize a full- or low-speed interrupt queue head StartXState bit. This state is entered from the Do_Complete Split state only after the split transaction is complete.

This occurs when one of the following events occur: The transaction translator responds to a complete-split transaction with one of the following:

- NAK. A NAK response is a propagation of the full- or low-speed endpoint's NAK response.
- ACK. An ACK response is a propagation of the full- or low-speed endpoint's ACK response. Only occurs on an OUT endpoint.
- DATA 0/1. Only occurs for INs. Indicates that this is the last of the data from the endpoint for this split transaction.
- ERR. The transaction on the low-/full-speed link below the transaction translator had a failure (for example, timeout, bad CRC, and so on).
- NYET (and Last). The host controller issued the last complete-split and the transaction translator responded with a NYET handshake. This means that the start-split was not correctly received by the transaction translator, so it never executed a

transaction to the full- or low-speed endpoint, see [Periodic interrupt-do-complete-split](#) , for the definition of 'Last'.

Each time the host controller visits a queue head in this state (once within the execute transaction state), bit-wise ANDs QH[S-mask] with cMicroFrameBit to determine whether to execute a start-split. If the result is non-zero, then the host controller issues a start-split transaction. If the PID Code field indicates an IN transaction, the host controller must zero-out the QH[S-bytes] field. After the split-transaction has been executed, the host controller sets up state in the queue head to track the progress of the complete-split phase of the split transaction. Specifically, it records the expected frame number into QH[FrameTag] field, sets C-prog-mask to zero (0x00), and exits this state. Note that the host controller must not adjust the value of Cerr as a result of completion of a start-split transaction.

18.6.12.2.7 Periodic interrupt-do-complete-split

This state is entered unconditionally from the Do-Start-Split state after a start-split transaction is executed on the bus.

Each time the host controller visits a queue head in this state (once within the execute transaction state), it checks to determine whether a complete-split transaction should be executed now.

There are four tests to determine whether a complete-split transaction should be executed.

- Test A. cMicroFrameBit is bit-wise ANDed with QH[C-mask] field. A non-zero result indicates that software scheduled a complete-split for this endpoint, during this microframe.
- Test B. QH[FrameTag] is compared with the current contents of FRINDEX[7-3]. An equal indicates a match.
- Test C. The complete-split progress bit vector is checked to determine whether the previous bit is set, indicating that the previous complete-split was appropriately executed. An example algorithm for this test is provided below:

```
Algorithm Boolean CheckPreviousBit(QH.C-prog-mask, QH.C-mask, cMicroFrameBit)
```

```
Begin
```

```
-- Return values:
```

```
-- TRUE - no error
```

```
-- FALSE - error
```

```
--
```

```
Boolean rvalue = TRUE;
```

```
previousBit = cMicroframeBit logical-rotate-right(1)
```

```

-- Bit-wise anding previousBit with C-mask indicates
-- whether there was an intent
-- to send a complete split in the previous microframe. So,
-- if the
-- 'previous bit' is set in C-mask, check C-prog-mask to
-- make sure it
-- happened.
If (previousBit bitAND QH.C-mask)then

If not(previousBit bitAND QH.C-prog-mask) then
    rvalue = FALSE;
    End if
End If

-- If the C-prog-mask already has a one in this bit position,
-- then an aliasing
-- error has occurred. It will probably get caught by the
-- FrameTag Test, but
-- at any rate it is an error condition that as detectable here
-- should not allow
-- a transaction to be executed.
If (cMicroFrameBit bitAND QH.C-prog-mask) then
rvalue = FALSE;
End if
return (rvalue)
End Algorithm
    
```

- **Test D.** Check to see if a start-split should be executed in this microframe. Note this is the same test performed in the Do Start Split state. Whenever it evaluates to TRUE and the controller is NOT processing in the context of a Recovery Path mode, it means a start-split should occur in this microframe. Test D and Test A evaluating to TRUE at the same time is a system software error. Behavior is undefined.

If (Test A and B and C and not(D)) then the host controller will execute a complete-split transaction. When the host controller commits to executing the complete-split transaction, it updates QH[C-prog-mask] by bit-ORing with cMicroFrameBit. On completion of the complete-split transaction, the host controller records the result of the transaction in the queue head and sets QH[FrameTag] to the expected H-Frame number. The effect to the state of the queue head and thus the state of the transfer depends on the response by the

transaction translator to the complete-split transaction. The following responses have the effects (note that any responses that result in decrementing of the Cerr will result in the queue head being halted by the host controller if the result of the decrement is zero):

- NYET (and Last)

On each NYET response, the host controller checks to determine whether this is the last complete-split for this split transaction. Last is defined in this context as the condition where all of the scheduled complete-splits have been executed. If it is the last complete-split (with a NYET response), then the transfer state of the queue head is not advanced (never received any data) and this state exited. The transaction translator must have responded to all the complete-splits with NYETs, meaning that the start-split issued by the host controller was not received. The start-split should be retried at the next poll period.

The test for whether this is the Last complete split can be performed by XOR QH[C-mask] with QH[C-prog-mask]. If the result is all zeros then all complete-splits have been executed. When this condition occurs, the XactErr status bit is set and the Cerr field is decremented.

- NYET (and not Last)

See above description for testing for Last. The complete-split transaction received a NYET response from the transaction translator. Do not update any transfer state (except for C-prog-mask and FrameTag) and stay in this state. The host controller must not adjust Cerr on this response.

- Transaction Error (XactErr). Timeout, data CRC failure, and so on

The Cerr field is decremented and the XactErr bit in the Status field is set. The complete split transaction is immediately retried (if Cerr is non-zero). If there is not enough time in the microframe to complete the retry and the endpoint is an IN, or Cerr is decremented to a zero from a one, the queue is halted. If there is not enough time in the microframe to complete the retry and the endpoint is an OUT and Cerr is not zero, then this state is exited (that is, return to Do Start Split). This results in a retry of the entire OUT split transaction, at the next poll period. Refer to Chapter 11 Hubs (specifically the section on full- and low-speed interrupts) in the *USB Specification Revision 2.0* for detailed requirements on why these errors must be immediately retried.

- ACK

This can only occur if the target endpoint is an OUT. The target endpoint ACK'd the data and this response is a propagation of the endpoint ACK up to the host controller. The host controller must advance the state of the transfer. The Current Offset field is incremented by Maximum Packet Length or Bytes to Transfer, whichever is less. The

field Bytes To Transfer is decremented by the same amount. And the data toggle bit (dt) is toggled. The host controller will then exit this state for this queue head. The host controller must reload Cerr with maximum value on this response. Advancing the transfer state may cause other process events such as retirement of the qTD and advancement of the queue.

- MDATA

This response will only occur for an IN endpoint. The transaction translator responded with zero or more bytes of data and an MDATA PID. The incremental number of bytes received is accumulated in QH[S-bytes]. The host controller must not adjust Cerr on this response.

- DATA0/1

This response may only occur for an IN endpoint. The number of bytes received is added to the accumulated byte count in QH[S-bytes]. The state of the transfer is advanced by the result and the host controller exits this state for this queue head.

Advancing the transfer state may cause other processing events such as retirement of the qTD and advancement of the queue.

If the data sequence PID does not match the expected, the entirety of the data received in this split transaction is ignored, the transfer state is not advanced and this state is exited.

- NAK

The target endpoint Nak'd the full- or low-speed transaction. The state of the transfer is not advanced, and this state is exited. The host controller must reload Cerr with maximum value on this response.

- ERR

There was an error during the full- or low-speed transaction. The ERR status bit is set, Cerr is decremented, the state of the transfer is not advanced, and this state is exited.

- STALL

The queue is halted (an exit condition of the Execute Transaction state). The status field bits: Active bit is cleared and the Halted bit is set and the qTD is retired.

Responses which are not enumerated in the list or which are received out of sequence are illegal and may result in undefined host controller behavior. The other possible combinations of tests A, B, C, and D may indicate that data or response was lost. The table below lists the possible combinations and the appropriate action.

Table 18-79. Interrupt IN/OUT do complete split state execution criteria

Condition	Action	Description
not(A) not(D)	Ignore QHD	Neither a start nor complete-split is scheduled for the current microframe. Host controller should continue walking the schedule.
A not(C)	If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split	Progress bit check failed. These means a complete-split has been missed. There is the possibility of lost data. If PID Code is an IN, then the queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Microframe bit in the status field to a one.
A not(B) C	If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split	QH.FrameTag test failed. This means that exactly one or more H-Frames have been skipped. This means complete-splits and have missed. There is the possibility of lost data. If PID Code is an IN, then the queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Microframe bit in the status field to a one.
A B C not(D)	Execute complete-split	This is the non-error case where the host controller executes a complete-split transaction.
D	If PIDCode = IN Halt QHD If PIDCode = OUT Retry start-split	This is a degenerate case where the start-split was issued, but all of the complete-splits were skipped and all possible intervening opportunities to detect the missed data failed to fire. If PID Code is an IN, then the Queue head must be halted. If PID Code is an OUT, then the transfer state is not advanced and the state exited (for example, start-split is retried). This is a host-induced error and does not effect Cerr. In either case, set the Missed Microframe bit in the status field to a one. Note that when executing in the context of a Recovery Path mode, the host controller is allowed to process the queue head and take the actions indicated above, or it may wait until the queue head is visited in the normal processing mode. Regardless, the host controller must not execute a start-split in the context of a executing in a Recovery Path mode.

18.6.12.2.8 Managing the QH[FrameTag] field

The QH[FrameTag] field in a queue head is completely managed by the host controller.

The rules for setting QH[FrameTag] are simple:

- Rule 1: If transitioning from Do Start Split to Do Complete Split and the current value of FRINDEX[2-0] is 6, QH[FrameTag] is set to $FRINDEX[7-3] + 1$. This accommodates split transactions whose start-split and complete-splits are in different H-Frames (case 2a, see [Figure 18-54](#)).
- Rule 2: If the current value of FRINDEX[2-0] is 7, QH[FrameTag] is set to $FRINDEX[7-3] + 1$. This accommodates staying in Do Complete Split for cases 2a, 2b, and 2c in [Figure 18-54](#).
- Rule 3: If transitioning from Do_Start Split to Do Complete Split and the current value of FRINDEX[2-0] is not 6, or currently in Do Complete Split and the current

value of (FRINDEX[2-0]) is not 7, FrameTag is set to FRINDEX[7-3]. This accommodates all other cases in [Figure 18-54](#).

18.6.12.2.9 Rebalancing the periodic schedule

System software must occasionally adjust a periodic queue head's S-mask and C-mask fields during operation.

This need occurs when adjustments to the periodic schedule create a new bandwidth budget and one or more queue head's are assigned new execution footprints (that is, new S-mask and C-mask values).

It is imperative that system software must not update these masks to new values in the midst of a split transaction. In order to avoid any race conditions with the update, the host controller provides a simple assist to system software. System software sets the Inactivate-on-next-Transaction (I) bit to signal the host controller that it intends to update the S-mask and C-mask on this queue head. System software then waits for the host controller to observe the I-bit is set and transitions the Active bit to a zero. The rules for how and when the host controller clears the Active bit are:

- If the Active bit is cleared, no action is taken. The host controller does not attempt to advance the queue when the I-bit is set.
- If the Active bit is set and the SplitXState is DoStart (regardless of the value of S-mask), the host controller simply clears the Active bit. The host controller is not required to write the transfer state back to the current qTD. Note that if the S-mask indicates that a start-split is scheduled for the current microframe, the host controller must not issue the start-split bus transaction; it must clear the Active bit.

System software must save transfer state before setting the I-bit. This is required so that it can correctly determine what transfer progress (if any) occurred after the I-bit was set and the host controller executed it's final bus-transaction and cleared the Active bit.

After system software has updated the S-mask and C-mask, it must then reactivate the queue head. Since the Active bit and the I-bit cannot be updated with the same write, system software needs to use the following algorithm to coherently re-activate a queue head that has been stopped using the I-bit.

1. Set the Halted bit, then
2. Clear the I-bit, then
3. Set the Active bit and clear the Halted bit in the same write.

Setting the Halted bit inhibits the host controller from attempting to advance the queue between the time the I-bit is cleared and the Active bit is set.

18.6.12.3 Split transaction isochronous

Full-speed isochronous transfers are managed using the split-transaction protocol through a USB 2.0 transaction translator in a USB 2.0 hub.

The host controller utilizes siTD data structure to support the special requirements of isochronous split-transactions. This data structure uses the scheduling model of isochronous TDs (see [Managing isochronous transfers using iTDs](#), for the operational model of iTDs) with the contiguous data feature provided by queue heads. This simple arrangement allows a single isochronous scheduling model and adds the additional feature that all data received from the endpoint (per split transaction) must land into a contiguous buffer.

18.6.12.3.1 Split transaction scheduling mechanisms for isochronous

Full-speed isochronous transactions are managed through a transaction translator's periodic pipeline.

As with full- and low-speed interrupt, system software manages each transaction translator's periodic pipeline by budgeting and scheduling exactly during which microframes the start-splits and complete-splits for each full-speed isochronous endpoint occur. The requirements described in [Split transaction scheduling mechanisms for interrupt](#), apply.

[Figure 18-58](#) illustrates the general scheduling boundary conditions that are supported by the EHCI periodic schedule. The S_n and C_n labels indicate microframes where software can schedule start- and complete-splits (respectively). The H-Frame boundaries are marked with a large, solid bold vertical line. The B-Frame boundaries are marked with a large, bold, dashed line. The bottom of the figure below illustrates the relationship of an siTD to the H-Frame.

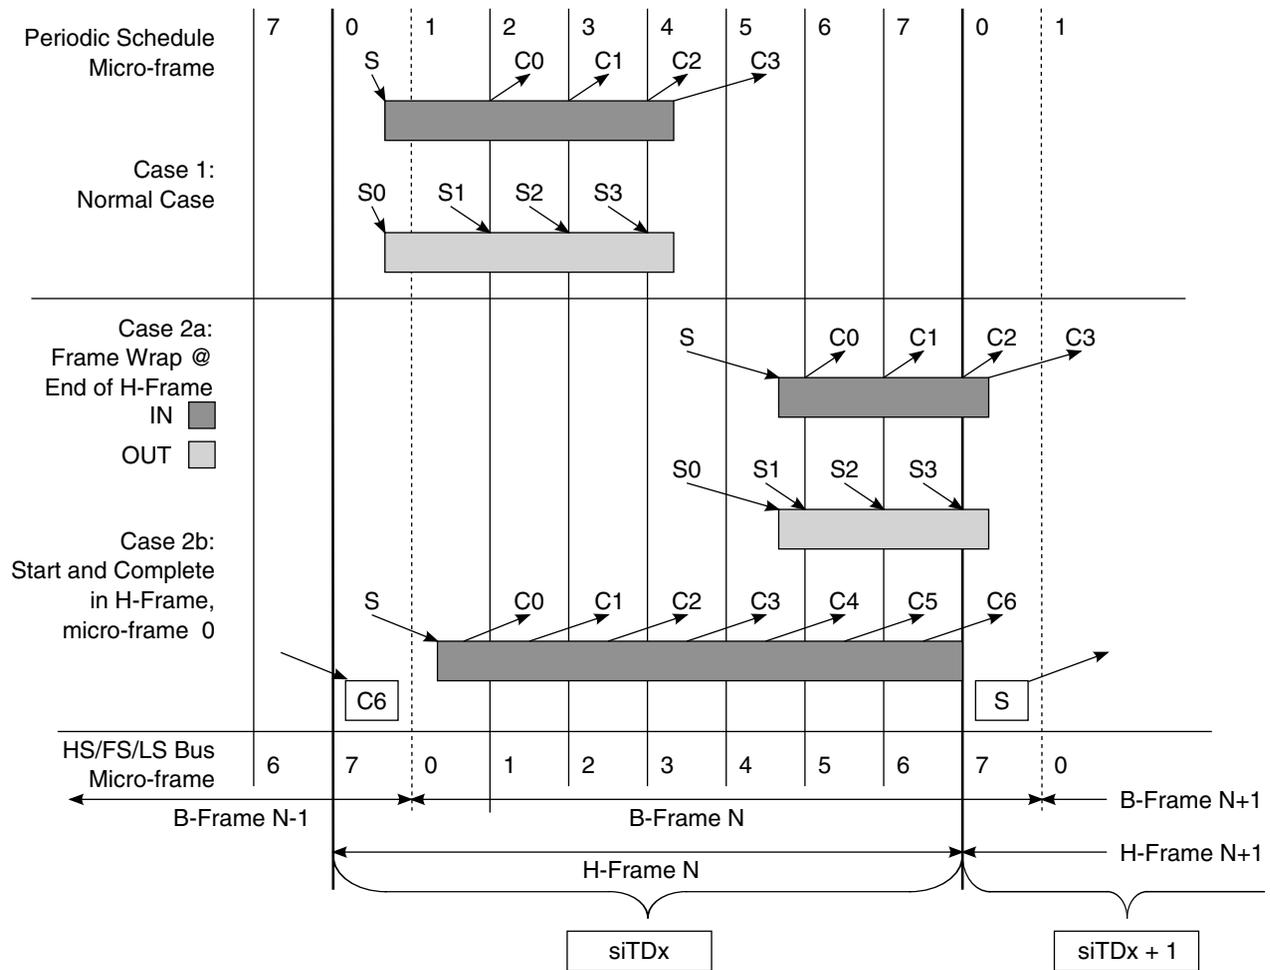


Figure 18-58. Split transaction, isochronous scheduling boundary conditions

When the endpoint is an isochronous OUT, there are only start-splits, and no complete-splits. When the endpoint is an isochronous IN, there is at most one start-split and one to N complete-splits. The scheduling boundary cases are:

- Case 1: The entire split transaction is completely bounded by an H-Frame. For example, the start-splits and complete-splits are all scheduled to occur in the same H-Frame.
- Case 2a: This boundary case is where one or more (at most two) complete-splits of a split transaction IN are scheduled across an H-Frame boundary. This can only occur when the split transaction has the possibility of moving data in B-Frame, microframes 6 or 7 (H-Frame microframe 7 or 0). When an H-Frame boundary wrap condition occurs, the scheduling of the split transaction spans more than one location in the periodic list.(for example, it takes two siTDs in adjacent periodic frame list locations to fully describe the scheduling for the split transaction).

Although the scheduling of the split transaction may take two data structures, all of the complete-splits for each full-speed IN isochronous transaction must use only one data pointer. For this reason, siTDs contain a back pointer.

Software must never schedule full-speed isochronous OUTs across an H-Frame boundary.

- **Case 2b:** This case can only occur for a very large isochronous IN. It is the only allowed scenario where a start-split and complete-split for the same endpoint can occur in the same microframe. Software must enforce this rule by scheduling the large transaction first. Large is defined to be anything larger than 579 byte maximum packet size.

A subset of the same mechanisms employed by full- and low-speed interrupt queue heads are employed in siTDs to schedule and track the portions of isochronous split transactions. The following fields are initialized by system software to instruct the host controller when to execute portions of the split transaction protocol:

- **SplitXState**

This is a single bit residing in the Status field of an siTD (see [Table 18-57](#)). This bit is used to track the current state of the split transaction. The rules for managing this bit are described in [Split transaction execution state machine for isochronous](#).

- **Frame S-mask**

This is a bit-field wherein system software sets a bit corresponding to the microframe (within an H-Frame) that the host controller should execute a start-split transaction. This is always qualified by the value of the SplitXState bit. For example, referring to the IN example in [Figure 18-58](#), case 1, the S-mask would have a value of 0b0000_0001 indicating that if the siTD is traversed by the host controller, and the SplitXState indicates Do Start Split, and the current microframe as indicated by FRINDEX[2-0] is 0, then execute a start-split transaction.

- **Frame C-mask**

This is a bit-field where system software sets one or more bits corresponding to the microframes (within an H-Frame) that the host controller should execute complete-split transactions. The interpretation of this field is always qualified by the value of the SplitXState bit. For example, referring to the IN example in [Figure 18-58](#), case 1, the C-mask would have a value of 0b 0011_1100 indicating that if the siTD is traversed by the host controller, and the SplitXState indicates Do Complete Split, and the current microframe as indicated by FRINDEX[2-0] is 2, 3, 4, or 5, then execute a complete-split transaction.

- **Back Pointer**

This field in a siTD is used to complete an IN split-transaction using the previous H-Frame's siTD. This is only used when the scheduling of the complete-splits span an H-Frame boundary.

There exists a one-to-one relationship between a high-speed isochronous split transaction (including all start- and complete-splits) and one full-speed isochronous transaction. An siTD contains (amongst other things) buffer state and split transaction scheduling information. An siTD's buffer state always maps to one full-speed isochronous data payload. This means that for any full-speed transaction payload, a single siTD's data buffer must be used. This rule applies to both IN and OUTs. An siTD's scheduling information usually also maps to one high-speed isochronous split transaction. The exception to this rule is the H-Frame boundary wrap cases mentioned above.

The siTD data structure describes at most, one frame's worth of high-speed transactions and that description is strictly bounded within a frame boundary. The figure below illustrates some examples. On the top are examples of the full-speed transaction footprints for the boundary scheduling cases described above. In the middle are time-frame references for both the B-Frames (HS/FS/LS Bus) and the H-Frames. On the bottom is illustrated the relationship between the scope of an siTD description and the time references. Each H-Frame corresponds to a single location in the periodic frame list. The implication is that each siTD is reachable from a single periodic frame list location at a time.

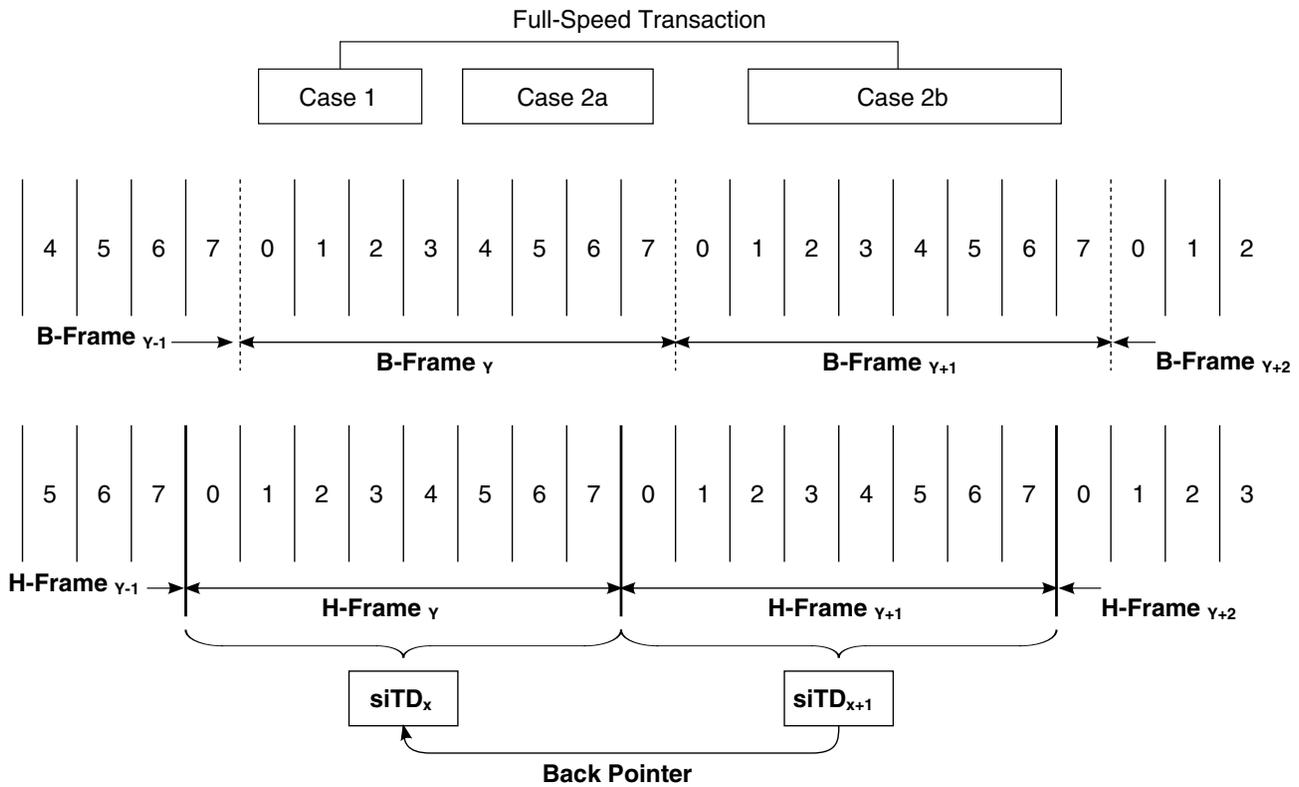


Figure 18-59. siTD scheduling boundary examples

Each case is described below:

- Case 1: One siTD is sufficient to describe and complete the isochronous split transaction because the whole isochronous split transaction is tightly contained within a single H-Frame.
- Case 2a, 2b: Although both INs and OUTs can have these footprints, OUTs always take only one siTD to schedule. However, INs (for these boundary cases) require two siTDs to complete the scheduling of the isochronous split transaction. siTD_x is used to always issue the start-split and the first N complete-splits. The full-speed transaction (for these cases) can deliver data on the full-speed bus segment during microframe 7 of H-Frame_{y+1}, or microframe 0 of H-Frame_{y+2}. The complete splits are scheduled using siTD_{x+2} (not shown). The complete-splits to extract this data must use the buffer pointer from siTD_{x+1}. The only way for the host controller to reach siTD_{x+1} from H-Frame_{y+2} is to use siTD_{x+2}'s back pointer.

Software must apply the following rules when calculating the schedule and linking the schedule data structures into the periodic schedule:

- Software must ensure that an isochronous split-transaction is started so that it will complete before the end of the B-Frame.
- Software must ensure that for a single full-speed isochronous endpoint, there is never a start-split and complete-split in H-Frame, microframe 1. This is mandated as a rule so that case 2a and case 2b can be discriminated. According to the core USB specification, the long isochronous transaction illustrated in Case 2b, could be scheduled so that the start-split was in microframe 1 of H-Frame N and the last complete-split would need to occur in microframe 1 of H-Frame N+1. However, it is impossible to discriminate between cases 2a and case 2b, which has significant impact on the complexity of the host controller.

18.6.12.3.2 Tracking split transaction progress for isochronous transfers

Isochronous endpoints do not employ the concept of a halt on error, however the host controller does identify and report per-packet errors observed in the data stream.

This includes schedule traversal problems (skipped microframes), timeouts and corrupted data received.

In similar kind to interrupt split-transactions, the portions of the split transaction protocol must execute in the microframes they are scheduled. The queue head data structure used to manage full- and low-speed interrupt has several mechanisms for tracking when portions of a transaction have occurred. Isochronous transfers use siTDs for their transfers and the data structures are only reachable using the schedule in the exact microframe in which they are required (so all the mechanism employed for tracking in queue heads is not required for siTDs). Software has the option of reusing siTD several times in the complete periodic schedule. However, it must ensure that the results of split transaction N are consumed and the siTD re-initialized (activated) before the host controller gets back to the siTD (in a future microframe).

Split-transaction isochronous OUTs utilize a low-level protocol to indicate which portions of the split transaction data have arrived. Control over the low-level protocol is exposed in an siTD using the fields Transaction Position (TP) and Transaction Count (T-count). If the entire data payload for the OUT split transaction is larger than 188 bytes, there will be more than one start-split transaction, each of which require proper annotation. If host hold-offs occur, then the sequence of annotations received from the host will not be complete, which is detected and handled by the transaction translator. See [Split transaction scheduling mechanisms for isochronous](#), for a description on how these fields are used during a sequence of start-split transactions.

The fields siTD[T-Count] and siTD[TP] are used by the host controller to drive and sequence the transaction position annotations. It is the responsibility of system software to properly initialize these fields in each siTD. Once the budget for a split-transaction

isochronous endpoint is established, S-mask, T-Count, and TP initialization values for all the siTD associated with the endpoint are constant. They remain constant until the budget for the endpoint is recalculated by software and the periodic schedule adjusted.

For IN-endpoints, the transaction translator simply annotates the response data packets with enough information to allow the host controller to identify the last data. As with split transaction Interrupt, it is the host controller's responsibility to detect when it has missed an opportunity to execute a complete-split. The following field in the siTD is used to track and detect errors in the execution of a split transaction for an IN isochronous endpoint.

- **C-prog-mask.** This is an eight-bit bit-vector where the host controller keeps track of which complete-splits have been executed. Due to the nature of the transaction translator periodic pipeline, the complete-splits need to be executed in-order. The host controller needs to detect when the complete-splits have not been executed in order. This can only occur due to system hold-offs where the host controller cannot get to the memory-based schedule. C-prog-mask is a simple bit-vector that the host controller sets a bit for each complete-split executed. The bit position is determined by the microframe (FRINDEX[2-0]) number in which the complete-split was executed. The host controller always checks C-prog-mask before executing a complete-split transaction. If the previous complete-splits have not been executed, then it means one (or more) have been skipped and data has potentially been lost. System software is required to initialize this field to zero before setting an siTD's Active bit to a one.

If a transaction translator returns with the final data before all of the complete-splits have been executed, the state of the transfer is advanced so that the remaining complete-splits are not executed. It is important to note that an IN siTD is retired based solely on the responses from the transaction translator to the complete-split transactions. This means, for example, that it is possible for a transaction translator to respond to a complete-split with an MDATA PID. The number of bytes in the MDATA's data payload could cause the siTD[Total Bytes to Transfer] field to decrement to zero. This response can occur, before all of the scheduled complete-splits have been executed. In other interface, data structures (for example, high-speed data streams through queue heads), the transition of Total Bytes to Transfer to zero signals the end of the transfer and results in clearing the Active bit. However, in this case, the result has not been delivered by the transaction translator and the host must continue with the next complete-split transaction to extract the residual transaction state. This scenario occurs because of the pipeline rules for a transaction translator. In summary, the periodic pipeline rules require that on a microframe boundary, the transaction translator holds the final two bytes received (if it has not seen an End Of Packet (EOP)) in the full-speed bus pipe stage and gives the remaining bytes to the high-speed pipeline stage. At the microframe boundary, the transaction translator could have received the entire packet (including both CRC bytes)

but not received the packet EOP. In the next microframe, the transaction translator responds with an MDATA and sends all of the data bytes (with the two CRC bytes being held in the full-speed pipeline stage). This could cause the siTD to decrement its Total Bytes to Transfer field to zero, indicating it has received all expected data. The host must still execute one more (scheduled) complete-split transaction in order to extract the results of the full-speed transaction from the transaction translator (for example, the transaction translator may have detected a CRC failure, and this result must be forwarded to the host).

If the host experiences hold-offs that cause the host controller to skip one or more (but not all) scheduled split transactions for an isochronous OUT, then the protocol to the transaction translator is not consistent and the transaction translator detects and reacts to the problem. Likewise, for host hold-offs that cause the host controller to skip one or more (but not all) scheduled split transactions for an isochronous IN, the C-prog-mask is used by the host controller to detect errors. However, if the host experiences a hold-off that causes it to skip all of an siTD, or an siTD expires during a host hold off (for example, a hold-off occurs and the siTD is no longer reachable by the host controller in order for it to report the hold-off event), then system software must detect that the siTDs have not been processed by the host controller (for example, state not advanced) and report the appropriate error to the client driver.

18.6.12.3.3 Split transaction execution state machine for isochronous

In this section, all references to microframe are in the context of a microframe within an H-Frame.

If the active bit in the status byte is a zero, the host controller ignores the siTD and continues traversing the periodic schedule. Otherwise the host controller processes the siTD as specified below. A split transaction state machine is used to manage the split-transaction protocol sequence. The host controller uses the fields defined in [Tracking split transaction progress for isochronous transfers](#), plus the variable `cMicroFrameBit` defined in [Split transaction execution state machine for interrupt](#), to track the progress of an isochronous split transaction. The figure below illustrates the state machine for managing an siTD through an isochronous split transaction. Bold, dotted circles denote the state of the Active bit in the Status field of a siTD. The Bold, dotted arcs denote the transitions between these states. Solid circles denote the states of the split transaction state machine and the solid arcs denote the transitions between these states. Dotted arcs and boxes reference actions that take place either as a result of a transition or from being in a state.

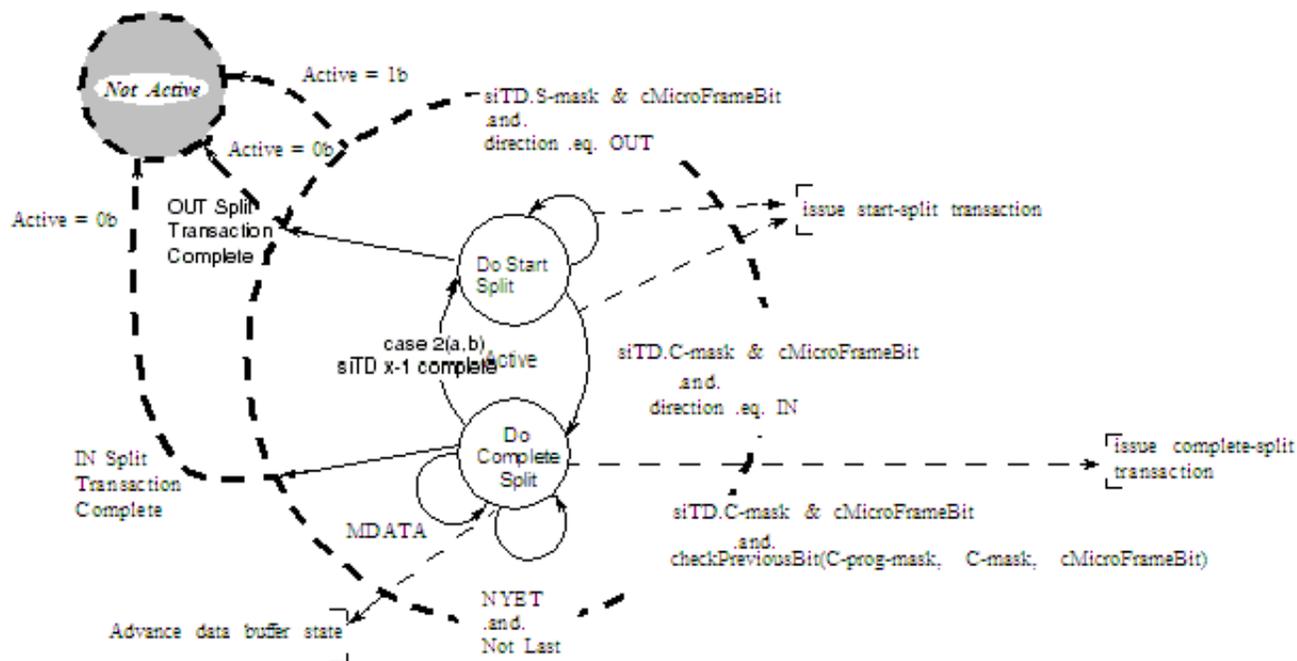


Figure 18-60. Split transaction state machine for isochronous

18.6.12.3.4 Periodic isochronous-do-start-split

Isochronous split transaction OUTs use only this state.

An siTD for a split-transaction isochronous IN is either initialized to this state, or the siTD transitions to this state from Do Complete Split when a case 2a (IN) or 2b scheduling boundary isochronous split-transaction completes.

Each time the host controller reaches an active siTD in this state, it checks the siTD[S-mask] against cMicroFrameBit. If there is a one in the appropriate position, the siTD executes a start-split transaction. By definition, the host controller cannot reach an siTD at the wrong time. If the I/O field indicates an IN, then the start-split transaction includes only the extended token plus the full-speed token. Software must initialize the siTD[Total Bytes To Transfer] field to the number of bytes expected. This is usually the maximum packet size for the full-speed endpoint. The host controller exits this state when the start-split transaction is complete.

The remainder of this section is specific to an isochronous OUT endpoint (that is, the I/O field indicates an OUT). When the host controller executes a start-split transaction for an isochronous OUT it includes a data payload in the start-split transaction. The memory buffer address for the data payload is constructed by concatenating siTD[Current Offset]

with the page pointer indicated by the page select field (siTD[P]). A zero in this field selects Page 0 and a 1 selects Page 1. During the start-split for an OUT, if the data transfer crosses a page boundary during the transaction, the host controller must detect the page cross, update the siTD[P] bit from a zero to a one, and begin using the siTD Page 1 with siTD[Current Offset] as the memory address pointer. The field siTD[TP] is used to annotate each start-split transaction with the indication of which part of the split-transaction data the current payload represents (ALL, BEGIN, MID, END). In all cases, the host controller simply uses the value in siTD[TP] to mark the start-split with the correct transaction position code.

T-count is always initialized to the number of start-splits for the current frame. TP is always initialized to the first required transaction position identifier. The scheduling boundary case (see [Figure 18-59](#)) is used to determine the initial value of TP. The initial cases are summarized in the table below.

Table 18-80. Initial conditions for OUT siTD TP and T-count fields

Case	T-Count	TP	Description
1, 2a	=1	ALL	When the OUT data payload is less than (or equal to) 188 bytes, only one start-split is required to move the data. The one start-split must be marked with an ALL.
1, 2a	!=1	BEGIN	When the OUT data payload is greater than 188 bytes more than one start-split must be used to move the data. The initial start-split must be marked with a BEGIN.

After each start-split transaction is complete, the host controller updates T-count and TP appropriately so that the next start-split is correctly annotated. The table below illustrates all of the TP and T-count transitions, which must be accomplished by the host controller.

Table 18-81. Transaction position (TP)/transaction count (T-count) transition table

TP	T-Count Next	TP Next	Description
ALL	0	N/A	Transition from ALL, to done.
BEGIN	1	END	Transition from BEGIN to END. Occurs when T-count starts at 2.
BEGIN	!=1	MID	Transition from BEGIN to MID. Occurs when T-count starts at greater than 2.
MID	!=1	MID	TP stays at MID while T-count is not equal to 1 (for example, greater than 1). This case can occur for any of the scheduling boundary cases where the T-count starts greater than 3.
MID	1	END	Transition from MID to END. This case can occur for any of the scheduling boundary cases where the T-count starts greater than 2.

The start-split transactions do not receive a handshake from the transaction translator, so the host controller always advances the transfer state in the siTD after the bus transaction is complete. To advance the transfer state the following operations take place:

- The siTD[Total Bytes To Transfer] and the siTD[Current Offset] fields are adjusted to reflect the number of bytes transferred.
- The siTD[P] (page select) bit is updated appropriately.
- The siTD[TP] and siTD[T-count] fields are updated appropriately as defined in [Table 18-81](#).

These fields are then written back to the memory based siTD. The S-mask is fixed for the life of the current budget. As mentioned above, TP and T-count are set specifically in each siTD to reflect the data to be sent from this siTD. Therefore, regardless of the value of S-mask, the actual number of start-split transactions depends on T-count (or equivalently, Total Bytes to Transfer). The host controller must clear the Active bit when it detects that all of the schedule data has been sent to the bus. Setting the Active bit to zero depends on siTD.TP being 00 or 11, and siTD.Total Bytes decrements to 0. Software must ensure that TP, T-count and Total Bytes to Transfer are set to deliver the appropriate number of bus transactions from each siTD. An inconsistent combination yields undefined behavior.

If the host experiences hold-offs that cause the host controller to skip start-split transactions for an OUT transfer, the state of the transfer does not progress appropriately. The transaction translator observes protocol violations in the arrival of the start-splits for the OUT endpoint (that is, the transaction position annotation is incorrect as received by the transaction translator).

Example scenarios are described in [Split transaction for isochronous-processing example](#).

The host controller can optionally track the progress of an OUT split transaction by setting appropriate bits in the siTD[C-prog-mask] as it executes each scheduled start-split. The checkPreviousBit() algorithm defined in [Periodic isochronous-do complete split](#), can be used prior to executing each start-split to determine whether start-splits were skipped. The host controller can use this mechanism to detect missed microframes. It can then clear the siTD's Active bit and stop execution of this siTD. This saves on both memory and high-speed bus bandwidth.

18.6.12.3.5 Periodic isochronous-do complete split

This state is only used by a split-transaction isochronous IN endpoint. This state is entered unconditionally from the Do Start State after a start-split transaction is executed for an IN endpoint.

Each time the host controller visits an siTD in this state, it conducts a number of tests to determine whether it should execute a complete-split transaction. The sequence in which they are applied depends on which microframe the host controller is currently executing, which means that the tests might not be applied until after the siTD referenced from the back pointer has been fetched. The individual tests are as follows.

- Test A

cMicroFrameBit is bit-wise ANDed with the siTD[C-mask] field. A non-zero result indicates that software scheduled a complete-split for this endpoint, during this microframe. This test is always applied to a newly fetched siTD that is in this state.

- Test B

The siTD[C-prog-mask] bit vector is checked to determine whether the previous complete splits have been executed. An example algorithm is given below (this is slightly different than the algorithm used in [Periodic interrupt-do-complete-split](#)). The sequence in which this test is applied depends on the current value of FRINDEX[2-0]. If FRINDEX[2-0] is 0 or 1, it is not applied until the back pointer has been used. Otherwise it is applied immediately.

```
Algorithm Boolean CheckPreviousBit(siTD.C-prog-mask, siTD.C-mask, cMicroFrameBit)
```

```
Begin
```

```
    Boolean rvalue = TRUE;
```

```
    previousBit = cMicroFrameBit rotate-right(1)
```

```
    -- Bit-wise anding previousBit with C-mask indicates whether there
```

```
    -- was an intent to send a complete split in the previous micro-
```

```
    -- frame. So, if the 'previous bit' is set in C-mask, check
```

```
    -- C-prog-mask to make sure it happened.
```

```
    if previousBit bitAND siTD.C-mask then
```

```
        if not (previousBit bitAND siTD.C-prog-mask) then
```

```
rvalue = FALSE
```

```
        End if
```

```
    End if
```

```
    Return rvalue
```

```
End Algorithm
```

If Test A is true and FRINDEX[2-0] is zero or one, this is a case 2a or 2b scheduling boundary (see [Figure 18-58](#)). See [Complete-split for scheduling boundary cases 2a, 2b](#), for details in handling this condition.

If Test A and Test B evaluate to true, the host controller executes a complete-split transaction using the transfer state of the current siTD. When the host controller commits to executing the complete-split transaction, it updates QH[C-prog-mask] by bit-ORing with cMicroFrameBit. The transfer state is advanced based on the completion status of the complete-split transaction. To advance the transfer state of an IN siTD, the host controller must perform the following actions:

1. Decrement the number of bytes received from siTD[Total Bytes To Transfer]
2. Adjust siTD[Current Offset] by the number of bytes received
3. Adjust the siTD[P] (page select) field if the transfer caused the host controller to use the next page pointer
4. Set any appropriate bits in the siTD[Status] field, depending on the results of the transaction.

Note that if the host controller encounters a condition where siTD[Total Bytes To Transfer] is zero, and it receives more data, the host controller must not write the additional data to memory. The siTD[Status-Active] bit must be cleared and the siTD[Status-Babble Detected] bit must be set. The fields siTD[Total Bytes To Transfer], siTD[Current Offset], and siTD[P] are not required to be updated as a result of this transaction attempt.

The host controller accepts (assuming good data packet CRC and sufficient room in the buffer as indicated by the value of siTD[Total Bytes To Transfer]) MDATA and DATA0/1 data payloads up to and including 192 bytes. The host controller may optionally clear siTD[Status-Active] and set siTD[Status-Babble Detected] when it receives MDATA or DATA0/1 with a data payload of more than 192 bytes. The following responses have the noted effects:

- ERR

The full-speed transaction completed with a time-out or bad CRC and this is a reflection of that error to the host. The host controller sets the ERR bit in the siTD[Status] field and clears the Active bit.

- Transaction Error (XactErr)

The complete-split transaction encounters a Timeout, CRC16 failure, and so on. The siTD[Status] field XactErr field is set and the complete-split transaction must be retried immediately. The host controller must use an internal error counter to count the number of retries as a counter field is not provided in the siTD data structure. The host controller will not retry more than two times. If the host controller exhausts the retries or the end of the microframe occurs, the Active bit is cleared.

- DATAx (0 or 1)

This response signals that the final data for the split transaction has arrived. The transfer state of the siTD is advanced and the Active bit is cleared. If the Bytes To Transfer field has not decremented to zero (including the reception of the data payload in the DATAx response), then less data than was expected, or allowed for was actually received. This short packet event does not set the USB interrupt status bit (USBSTS[UI]) to a one. The host controller will not detect this condition.

- NYET (and Last)

On each NYET response, the host controller also checks to determine whether this is the last complete-split for this split transaction. Last was defined in [Periodic interrupt-do-complete-split](#). If it is the last complete-split (with a NYET response), then the transfer state of the siTD is not advanced (never received any data) and the Active bit is cleared. No bits are set in the Status field because this is essentially a skipped transaction. The transaction translator must have responded to all the scheduled complete-splits with NYETs, meaning that the start-split issued by the host controller was not received. This result should be interpreted by system software as if the transaction was completely skipped. The test for whether this is the last complete split can be performed by XORing C-mask with C-prog-mask. A zero result indicates that all complete-splits have been executed.

- MDATA (and Last)

See above description for testing for Last. This can only occur when there is an error condition. Either there has been a babble condition on the full-speed link, which delayed the completion of the full-speed transaction, or software set up the S-mask and/or C-masks incorrectly. The host controller must set the XactErr bit and clear the Active bit.

- NYET (and not Last)

See above description for testing for Last. The complete-split transaction received a NYET response from the transaction translator. Do not update any transfer state (except for C-prog-mask) and stay in this state.

- MDATA (and not Last)

The transaction translator responds with an MDATA when it has partial data for the split transaction. For example, the full-speed transaction data payload spans from microframe X to X+1 and during microframe X, the transaction translator responds with an MDATA and the data accumulated up to the end of microframe X. The host controller advances the transfer state to reflect the number of bytes received.

If Test A succeeds, but Test B fails, it means that one or more of the complete-splits have been skipped. The host controller sets the Missed Micro-Frame status bit and clears the Active bit.

18.6.12.3.6 Complete-split for scheduling boundary cases 2a, 2b

Boundary cases 2a and 2b (INs only) (see [Figure 18-58](#)) require that the host controller use the transaction state context of the previous siTD to finish the split transaction. The table below enumerates the transaction state fields.

Table 18-82. Summary siTD split transaction state

Buffer state	Status	Execution progress
Total bytes To transfer	All bits in the status field	C-prog-mask
P (page select)		
Current Offset		
TP (transaction position)		
T-count (transaction count)		

NOTE

TP and T-count are used only for Host to Device (OUT) endpoints.

If software has budgeted the schedule of this data stream with a frame wrap case, then it must initialize the siTD[Back Pointer] field to reference a valid siTD and have the T bit in the siTD[Back Pointer] field cleared. Otherwise, software must set the T bit in siTD[Back Pointer]. The host controller's rules for interpreting when to use the siTD[Back Pointer] field are listed below. These rules apply only when the siTD's Active bit is a one and the SplitXState is Do Complete Split.

- When cMicroFrameBit is a 0x1 and the siTD_X[Back Pointer] T-bit is zero, or
- If cMicroFrameBit is a 0x2 and siTD_X[S-mask[0]] is zero

When either of these conditions apply, then the host controller must use the transaction state from siTD_{X-1}.

In order to access siTD_{X-1}, the host controller reads on-chip the siTD referenced from siTD_X[Back Pointer].

The host controller must save the entire state from siTD_X while processing siTD_{X-1}. This is to accommodate for case 2b processing. The host controller must not recursively walk the list of siTD[Back Pointers].

If siTD_{X-1} is active (Active bit is set and SplitXStat is Do Complete Split), then both Test A and Test B are applied as described above. If these criteria to execute a complete-split are met, the host controller executes the complete split and evaluates the results as described above. The transaction state (see Table 18-82) of siTD_{X-1} is appropriately advanced based on the results and written back to memory. If the resultant state of siTD_{X-1}'s Active bit is a one, then the host controller returns to the context of siTD_X, and follows its next pointer to the next schedule item. No updates to siTD_X are necessary.

If siTD_{X-1} is active (Active bit is set and SplitXStat is Do Start Split), then the host controller must clear the Active bit and set the Missed Micro-Frame status bit and the resultant status is written back to memory.

If $siTD_{X-1}$'s Active bit is cleared, (because it was cleared when the host controller first visited $siTD_{X-1}$ via $siTD_X$'s back pointer, it transitioned to zero as a result of a detected error, or the results of $siTD_{X-1}$'s complete-split transaction cleared it), then the host controller returns to the context of $siTD_X$ and transitions its SplitXState to Do Start Split. The host controller then determines whether the case 2b start split boundary condition exists (that is, if $cMicroframeBit$ is 1 and $siTD_X[S-mask[0]]$ is 1). If this criterion is met the host controller immediately executes a start-split transaction and appropriately advances the transaction state of $siTD_X$, then follows $siTD_X[Next Pointer]$ to the next schedule item. If the criterion is not met, the host controller simply follows $siTD_X[Next Pointer]$ to the next schedule item. Note that in the case of a 2b boundary case, the split-transaction of $siTD_{X-1}$ will have its Active bit cleared when the host controller returns to the context of $siTD_X$. Also, note that software should not initialize an $siTD$ with C-mask bits 0 and 1 set and an S-mask with bit 0 set. This scheduling combination is not supported and the behavior of the host controller is undefined.

18.6.12.3.7 Split transaction for isochronous-processing example

There is an important difference between how the hardware/software manages the isochronous split transaction state machine and how it manages the asynchronous and interrupt split transaction state machines.

The asynchronous and interrupt split transaction state machines are encapsulated within a single queue head. The progress of the data stream depends on the progress of each split transaction. In some respects, the split-transaction state machine is sequenced using the Execute Transaction queue head traversal state machine.

Isochronous is a pure time-oriented transaction/data stream. The interface data structures are optimized to efficiently describe transactions that need to occur at specific times. The isochronous split-transaction state machine must be managed across these time-oriented data structures. This means that system software must correctly describe the scheduling of split-transactions across more than one data structure.

Then the host controller must make the appropriate state transitions at the appropriate times, in the correct data structures.

For example, the table below illustrates a few frames worth of scheduling required to schedule a case 2a full-speed isochronous data stream.

Table 18-83. Example case 2a-software scheduling $siTD$ s for an IN endpoint

$siTD_x$		Micro-frames								InitialSplitXState
#	Masks	0	1	2	3	4	5	6	7	

Table continues on the next page...

Table 18-83. Example case 2a-software scheduling siTDs for an IN endpoint (continued)

siTD _x		Micro-frames								InitialSplitXState
X	S-mask					1				Do start split
	C-mask	1	1					1	1	
X+1	S-mask					1				Do complete split
	C-mask	1	1					1	1	
X+2	S-mask					1				Do complete split
	C-mask	1	1					1	1	
X+3	S-mask	Repeats previous pattern								Do complete split
	C-mask									

This example shows the first three siTDs for the transaction stream. Since this is the case-2a frame-wrap case, S-masks of all siTDs for this endpoint have a value of 0x10 (a one bit in microframe 4) and C-mask value of 0xC3 (one-bits in microframes 0,1, 6 and 7). Additionally, software ensures that the Back Pointer field of each siTD references the appropriate siTD data structure (and the Back Pointer T-bits are cleared).

The initial SplitXState of the first siTD is Do Start Split. The host controller will visit the first siTD eight times during frame X. The C-mask bits in microframes 0 and 1 are ignored because the state is Do Start Split. During microframe 4, the host controller determines that it can run a start-split (and does) and changes SplitXState to Do Complete Split. During microframes 6 and 7, the host controller executes complete-splits. Notice the siTD for frame X+1 has it's SplitXState initialized to Do Complete Split. As the host controller continues to traverse the schedule during H-Frame X+1, it will visit the second siTD eight times. During microframes 0 and 1 it will detect that it must execute complete-splits.

During H-Frame X+1, microframe 0, the host controller detects that siTD_{X+1}'s Back Pointer[T] bit is a zero, saves the state of siTD_{X+1} and fetches siTD_X. It executes the complete split transaction using the transaction state of siTD_X. If the siTD_X split transaction is complete, siTD's Active bit is cleared and results written back to siTD_X. The host controller retains the fact that siTD_X is retired and transitions the SplitXState in siTD_{X+1} to Do Start Split. At this point, the host controller is prepared to execute the start-split for siTD_{X+1} when it reaches microframe 4. If the split-transaction completes early (transaction-complete is defined in [Periodic isochronous-do complete split](#)), that is, before all the scheduled complete-splits have been executed, the host controller changes siTD_X[SplitXState] to Do Start Split early and naturally skips the remaining scheduled complete-split transactions. For this example, siTD_{X+1} does not receive a DATA0 response until H-Frame X+2, microframe 1.

During H-Frame X+2, microframe 0, the host controller detects that siTD_{X+2}'s Back Pointer[T] bit is zero, saves the state of siTD_{X+2} and fetches siTD_{X+1}. As described above, it executes another split transaction, receives an MDATA response, updates the transfer state, but does not modify the Active bit. The host controller returns to the context of siTD_{X+2}, and traverses its next pointer without any state change updates to siTD_{X+2}.

During H-Frame X+2, microframe 1, the host controller detects siTD_{X+2}'s S-mask[0] bit is zero, saves the state of siTD_{X+2} and fetches siTD_{X+1}. It executes another complete-split transaction, receives a DATA0 response, updates the transfer state and clears the Active bit. It returns to the state of siTD_{X+2} and changes its SplitXState to Do Start Split. At this point, the host controller is prepared to execute start-splits for siTD_{X+2} when it reaches microframe 4.

18.6.13 Port test modes

EHCI host controllers implement the port test modes Test J_State, Test K_State, Test_Packet, Test Force_Enable, and Test SEO_NAK as described in the *USB Specification Revision 2.0*.

The required, port test sequence, assuming the CF-bit in the CONFIGFLAG register is set, is as follows:

1. Disable the periodic and asynchronous schedules by clearing the USBCMD[ASE] and USBCMD[PSE].
2. Place all enabled root ports into the suspended state by setting the Suspend bit in the PORTSC register (PORTSC[SUSP]).
3. Clear USBCMD[RS] (run/stop) and wait for USBSTS[HCH] to transition to a one. In Device mode, the Test Mode starts only if Run/Stop bit is set to 1. In Host mode, the Test Mode starts regardless of Run/Stop bit.
4. Set the Port Test Control field in the port under test PORTSC register to the value corresponding to the desired test mode. If the selected test is Test_Force_Enable, then USBCMD[RS] must then be transitioned back to one, in order to enable transmission of SOFs out of the port under test.
5. When the test is complete, system software must ensure the host controller is halted (HCH bit is a one) then it terminates and exits test mode by setting USBCMD[RST].

18.6.14 Interrupts

The EHCI host controller hardware provides interrupt capability based on a number of sources.

The following list describes the general groups of interrupt sources:

- Interrupts as a result of executing transactions from the schedule (success and error conditions)
- Host controller events (Port change events, and so on)
- Host controller error events

All transaction-based sources are maskable through the host controller's Interrupt Enable register (USBINTR). Additionally, individual transfer descriptors can be marked to generate an interrupt on completion. This section describes each interrupt source and the processing that occurs in response to the interrupt.

During normal operation, interrupts may be immediate or deferred until the next interrupt threshold occurs. The interrupt threshold is a tunable parameter via the interrupt threshold control field in the USBCMD register. The value of this register controls when the host controller generates an interrupt on behalf of normal transaction execution. When a transaction completes during an interrupt interval period, the interrupt signaling the completion of the transfer will not occur until the interrupt threshold occurs. For example, the default value is eight microframes. This means that the host controller will not generate interrupts any more frequently than once every eight microframes.

[Host system error](#) details effects of a host system error.

If an interrupt has been scheduled to be generated for the current interrupt threshold interval, the interrupt is not signaled until after the status for the last complete transaction in the interval has been written back to system memory. This may sometimes result in the interrupt not being signaled until the next interrupt threshold.

Initial interrupt processing is the same, regardless of the reason for the interrupt. When an interrupt is signaled by the hardware, CPU control is transferred to host controller's USB interrupt handler. The precise mechanism to accomplish the transfer is OS specific. For this discussion it is just assumed that control is received. When the interrupt handler receives control, its first action is to read the USBSTS. It then acknowledges the interrupt by clearing all of the interrupt status bits by writing ones to these bit positions. The handler then determines whether the interrupt is due to schedule processing or some other event. After acknowledging the interrupt, the handler (via an OS-specific mechanism), schedules a deferred procedure call (DPC) which will execute later. The DPC routine processes the results of the schedule execution. The precise mechanisms used are beyond the scope of this document.

NOTE

The only method software should use for acknowledging an interrupt is by transitioning the appropriate status bits in the USBSTS register from a one to a zero.

18.6.14.1 Transfer/transaction based interrupts

These interrupt sources are associated with transfer and transaction progress. They are all dependent on the next interrupt threshold.

18.6.14.1.1 Transaction error

A transaction error is any error that caused the host controller to think that the transfer did not complete successfully.

The table below lists the events/responses that the host can observe as a result of a transaction. The effects of the error counter and interrupt status are summarized in the following paragraphs. Most of these errors set the XactErr status bit in the appropriate interface data structure.

Table 18-84. Summary of transaction errors

Event/ Result	Queue Head/qTD/iTD/siTD Side Effects		USBSTS[USBERRINT]
	Cerr	Status Field	
CRC	-1	XactErr set	1 ¹
Timeout	-1	XactErr set	1 ¹
Bad PID ²	-1	XactErr set	1 ¹
Babble	N/A	See Serial bus babble	1
Buffer Error	N/A	See Data buffer error	

¹ If occurs in a queue head, then USBERRINT is asserted only when Cerr counts down from a one to a zero. In addition the queue is halted.

² The host controller received a response from the device, but it could not recognize the PID as a valid PID.

There is a small set of protocol errors that relate only when executing a queue head and fit under the umbrella of a WRONG PID error that are significant to explicitly identify. When these errors occur, the XactErr status bit in the queue head is set and the Cerr field is decremented. When the PID Code indicates a SETUP, the following responses are protocol errors and result in XactErr bit being set and the Cerr field being decremented.

- EPS field indicates a high-speed device and it returns a Nak handshake to a SETUP.
- EPS field indicates a high-speed device and it returns a Nyet handshake to a SETUP.
- EPS field indicates a low- or full-speed device and the complete-split receives a Nak handshake.

18.6.14.1.2 Serial bus babble

When a device transmits more data on the USB than the host controller is expecting for this transaction, it is defined to be babbling.

In general, this is called a packet babble. When a device sends more data than the maximum length number of bytes, the host controller sets the babble detected bit to a one and halts the endpoint if it is using a queue head. Maximum length is defined as the minimum of total bytes to transfer and maximum packet size. The Cerr field is not decremented for a packet babble condition (only applies to queue heads).

A babble condition also exists if IN transaction is in progress at High-speed EOF2 point. This is called a frame babble. A frame babble condition is recorded into the appropriate schedule data structure. In addition, the host controller must disable the port to which the frame babble is detected.

USBSTS[UEI] (USB error interrupt) is set and if the USBINTR[UEE] (USB error interrupt enable) is set, then a hardware interrupt is signaled to the system at the next interrupt threshold. The host controller will never start an OUT transaction that babbles across a microframe EOF.

NOTE

When a host controller detects a data PID mismatch, it must either: disable the packet babble checking for the duration of the bus transaction or do packet babble checking based solely on Maximum Packet Size. The USB core specification defines the requirements on a data receiver when it receives a data PID mismatch (for example, expects a DATA0 and gets a DATA1 or visa-versa). In summary, it must ignore the received data and respond with an ACK handshake, in order to advance the transmitter's data sequence. The EHCI interface allows system software to provide buffers for a Control, Bulk or Interrupt IN endpoint that are not an even multiple of the maximum packet size specified by the device. Whenever a device misses an ACK for an IN endpoint, the host and device are out of synchronization with respect to the progress of the data transfer. The host controller may have advanced the transfer to a buffer that is less than maximum packet size. The device re-sends its maximum packet size data packet, with the original data PID, in response to the next IN token. In order to properly manage the bus protocol, the host controller must disable the packet babble check when it observes the data PID mismatch.

18.6.14.1.3 Data buffer error

This event indicates that an overrun of incoming data or a underrun of outgoing data has occurred for this transaction.

This would generally be caused by the host controller not being able to access required data buffers in memory within necessary latency requirements. These conditions are not considered transaction errors, and do not effect the error count in the queue head. When these errors do occur, the host controller records the fact the error occurred by setting the Data Buffer Error bit in the queue head, iTD or siTD.

If the data buffer error occurs on a non-isochronous IN, the host controller will not issue a handshake to the endpoint. This forces the endpoint to resend the same data (and data toggle) in response to the next IN to the endpoint.

If the data buffer error occurs on an OUT, the host controller must corrupt the end of the packet so that it cannot be interpreted by the device as a good data packet. Simply truncating the packet is not considered acceptable. An acceptable implementation option is to 1's complement the CRC bytes and send them. There are other options suggested in the transaction translator section of the *USB Specification Revision 2.0*.

18.6.14.1.4 USB interrupt (interrupt on completion (IOC))

Transfer descriptors (iTDS, siTDs, and queue heads (qTDs)) contain a bit that can be set to cause an interrupt on their completion.

The completion of the transfer associated with that schedule item causes USBSTS[UI] (USB interrupt) to be set. In addition, if a short packet is encountered on an IN transaction associated with a queue head, then this event also causes USBINT to be set. If USBINTR[UE] (USB interrupt enable) is set, a hardware interrupt is signaled to the system at the next interrupt threshold. If the completion is because of errors, USBSTS[UEI] (USB error interrupt) is also set.

18.6.14.1.5 Short packet

Reception of a data packet that is less than the endpoint's Max Packet size during Control, Bulk or Interrupt transfers signals the completion of the transfer.

Whenever a short packet completion occurs during a queue head execution, USBSTS[UI] (USB interrupt bit) is set. If the USB interrupt enable bit is set (USBINTR[UE]), a hardware interrupt is signaled to the system at the next interrupt threshold.

18.6.14.2 Host controller event interrupts

These interrupt sources are independent of the interrupt threshold (with the one exception being the Interrupt on Async Advance).

18.6.14.2.1 Port change events

Port registers contain status and status change bits.

When the status change bits are set, the host controller sets the USBSTS[PCI]. If the port change interrupt enable bit (PCE) in the USBINTR register is set, the host controller issues a hardware interrupt. The port status change bits in PORTSC include:

- Connect change status (CSC)
- Port enable/disable change (PEC)
- Over-current change (OCC)
- Force port resume (FPR)

18.6.14.2.2 Frame list rollover

This event indicates that the host controller has wrapped the frame list.

The current programmed size of the frame list effects how often this interrupt occurs. If the frame list size is 1024, then the interrupt occurs every 1024 milliseconds, if it is 512, then it occurs every 512 milliseconds, and so on. When a frame list rollover is detected, the host controller sets the frame list rollover bit, USBSTS[FRI]. If USBINTR[FRE] is set (frame list rollover enable), the host controller issues a hardware interrupt. This interrupt is not delayed to the next interrupt threshold.

18.6.14.2.3 Interrupt on async advance

This event is used for deterministic removal of queue heads from the asynchronous schedule.

Whenever the host controller advances the on-chip context of the asynchronous schedule, it evaluates the value of USBCMD[IAA]. If it is set, it sets USBSTS[AAI]. If USBINTR[AAE] is set, the host controller issues a hardware interrupt at the next interrupt threshold. A detailed explanation of this feature is described in [Removing queue heads from asynchronous schedule](#).

18.6.14.2.4 Host system error

The host controller is a bus master and any interaction between the host controller and the system may experience errors.

The type of host error may be catastrophic to the host controller making it impossible for the host controller to continue in a coherent fashion. Behavior for these types of errors is to halt the host controller. Host-based error must result in the following actions:

- USBCMD[RS] is cleared.
- USBSTS[SEI] and USBSTS[HCH] register are set
- If the host system error enable bit, USBINTR[SEE] is set, the host controller issues a hardware interrupt. This interrupt is not delayed to the next interrupt threshold.

The table below summarizes the required actions taken on the various host errors.

Table 18-85. Summary behavior on host system errors

Cycle type	Master abort	Target abort	Data phase parity
Frame list pointer fetch (read)	Fatal	Fatal	Fatal
siTD fetch (read)	Fatal	Fatal	Fatal
siTD status write-back (write)	Fatal	Fatal	Fatal
iTD fetch (read)	Fatal	Fatal	Fatal
iTD status write-back (write)	Fatal	Fatal	Fatal
qTD fetch (read)	Fatal	Fatal	Fatal
qHD status write-back (write)	Fatal	Fatal	Fatal
Data write	Fatal	Fatal	Fatal
Data read	Fatal	Fatal	Fatal

NOTE

After a host system error, software must reset the host controller using USBCMD[RST] before re-initializing and restarting the host controller.

18.7 Device data structures

This section defines the interface data structures used to communicate control, status, and data between device controller driver (DCD) software and the device controller. The data structure definitions in this chapter support a 32-bit memory buffer address space. The interface consists of device queue heads and transfer descriptors.

NOTE

Software must ensure that no interface data structure reachable by the device controller spans a 4K-page boundary.

The data structures defined in the section are (from the device controller's perspective) a mix of read-only and read/writable fields. The device controller must preserve the read-only fields on all data structure writes.

The USB DR module includes DCD software called the USB 2.0 Device API. The device API provides an easy to use Application Program Interface for developing device (peripheral) applications. The device API incorporates and abstracts for the application developer all of the elements of the program interface.

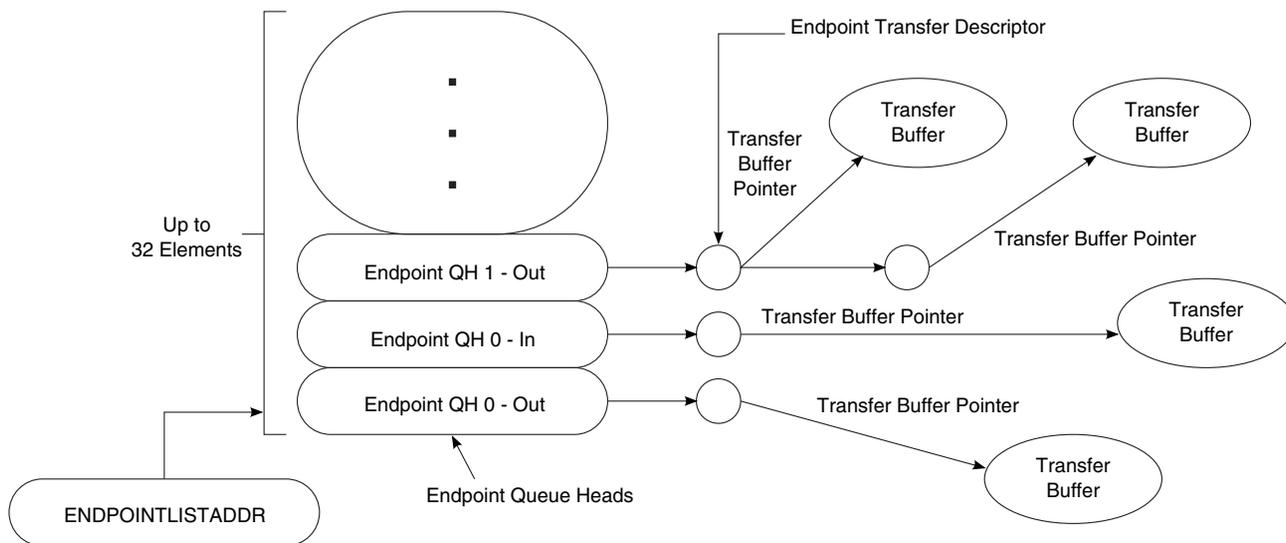


Figure 18-61. Endpoint queue head organization

18.7.1 Endpoint queue head

The device endpoint queue head (dQH) is where all transfers are managed.

The dQH is a 48-byte data structure, but must be aligned on 64-byte boundaries. During priming of an endpoint, the dTD (device transfer descriptor) is copied into the overlay area of the dQH, which starts at the nextTD pointer DWord and continues through the end of the buffer pointers DWords. After a transfer is complete, the dTD status DWord is updated in the dTD pointed to by the currentTD pointer. While a packet is in progress, the overlay area of the dQH is used as a staging area for the dTD so that the Device Controller can access needed information with little minimal latency.

The figure below shows the endpoint queue head structure.

Table 18-86. Endpoint queue head layout

3	3	2	2	2	2	2	2	2	2	2	2	1	1	1	1	15	14	1	1	11	10	9	8	7	6	5	4	3	2	1	0	offset
1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6			3	2													
Mult	zlt	00	Maximum Packet Length													ios	000_0000_0000_0000										0x00					
Current dTD Pointer ¹																									0_0000		0x04					
Next dTD Pointer ¹																									0000		T ¹	0x08 ²				

Table continues on the next page...

Table 18-86. Endpoint queue head layout (continued)

0	Total Bytes ¹	ioc ¹ 000	MultO ¹ 00	Status ¹	0x0C ²
	Buffer Pointer (Page 0) ¹			Current Offset ¹	0x10 ²
	Buffer Pointer (Page 1) ¹			Reserved	0x14 ²
	Buffer Pointer (Page 2) ¹			Reserved	0x18 ²
	Buffer Pointer (Page 3) ¹			Reserved	0x1C ²
	Buffer Pointer (Page 4) ¹			Reserved	0x20 ²
	Reserved				0x24
	Setup Buffer Bytes 3-0 ¹				0x28
	Setup Buffer Bytes 7-4 ¹				0x2C

1. Device controller read/write; all others read-only.
2. Offsets 0x08 through 0x20 contain the transfer overlay.

18.7.1.1 Endpoint capabilities/characteristics

This DWord specifies static information about the endpoint, in other words, this information does not change over the lifetime of the endpoint.

Device controller software should not attempt to modify this information while the corresponding endpoint is enabled.

The table below describes the endpoint capabilities and characteristics fields.

Table 18-87. Endpoint capabilities/characteristics

Bits	Name	Description
31-30	Mult	Mult. This field is used to indicate the number of packets executed per transaction description as given by the following: 00 - Execute N Transactions as demonstrated by the USB variable length packet protocol where N is computed using the Maximum Packet Length (dQH) and the Total Bytes field (dTD) 01 Execute 1 Transaction. 10 Execute 2 Transactions. 11 Execute 3 Transactions. NOTE: Non-ISO endpoints must set Mult = 00. NOTE: ISO endpoints must set Mult = 01, 10, or 11 as needed.
29	zlt	Zero length termination select. This bit is used to indicate when a zero length packet is used to terminate transfers where the total transfer length is a multiple. This bit is not relevant for Isochronous transfers. 0 Enable zero length packet to terminate transfers equal to a multiple of the Maximum Packet Length. (default). 1 Disable the zero length packet on transfers that are equal in length to a multiple Maximum Packet Length.
28-27	-	Reserved, should be cleared. These bit reserved for future use and should be cleared.

Table continues on the next page...

Table 18-87. Endpoint capabilities/characteristics (continued)

Bits	Name	Description
26-16	Maximum Packet Length	Maximum packet length. This directly corresponds to the maximum packet size of the associated endpoint (wMaxPacketSize). The maximum value this field may contain is 0x400 (1024).
15	ios	Interrupt on setup (IOS). This bit is used on control type endpoints to indicate if USBINT is set in response to a setup being received.
14-0	-	Reserved, should be cleared. Bits reserved for future use and should be cleared.

18.7.1.2 Transfer overlay

The seven DWords in the overlay area represent a transaction working space for the device controller.

The general operational model is that the device controller can detect whether the overlay area contains a description of an active transfer. If it does not contain an active transfer, then it will not read the associated endpoint.

After an endpoint is read, the dTD is copied into this queue head overlay area by the device controller. Until a transfer is expired, software must not write the queue head overlay area or the associated transfer descriptor. When the transfer is complete, the device controller will write the results back to the original transfer descriptor and advance the queue.

See dTD for a description of the overlay fields.

18.7.1.3 Current dTD pointer

The current dTD pointer is used by the device controller to locate the transfer in progress.

This word is for USB_DR controller (hardware) use only and should not be modified by DCD software.

The table below describes the current dTD pointer fields.

Table 18-88. Current dTD pointer

Bits	Description
31-5	Current dtd. This field is a pointer to the dTD that is represented in the transfer overlay area. This field is modified by the Device Controller to next dTD pointer during endpoint priming or queue advance.
4-0	Reserved, should be cleared. Bit reserved for future use and should be cleared.

18.7.1.4 Setup buffer

The setup buffer is dedicated storage for the 8-byte data that follows a setup PID.

NOTE

Each endpoint has a TX and an RX dQH associated with it, and only the RX queue head is used for receiving setup data packets.

The table below describes the multiple mode control fields.

Table 18-89. Multiple mode control

DWord	Bits	Description
1	31-0	Setup Buffer 0. This buffer contains bytes 3 to 0 of an incoming setup buffer packet and is written by the device controller to be read by software.
2	31-0	Setup Buffer 1. This buffer contains bytes 7 to 4 of an incoming setup buffer packet and is written by the device controller to be read by software.

18.7.2 Endpoint transfer descriptor (dTD)

The dTD describes the location and quantity of data to be sent/received for given transfer to the device controller.

The DCD should not attempt to modify any field in an active dTD except the Next Link Pointer, which should only be modified as described in [Managing transfers with transfer descriptors](#).

The figure below shows the endpoint transfer descriptor.

Table 18-90. Endpoint transfer descriptor (dTD)

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	offset
Next Link Pointer																											0000	T	0x00			
0	Total Bytes ¹														ioc	000		MultO	00		Status ¹					0x04						
Buffer Pointer (Page 0)											Current Offset ¹											0x08										
Buffer Pointer (Page 1)											0	Frame Number ¹											0x0C									
Buffer Pointer (Page 2)											0000_0000_0000											0x10										
Buffer Pointer (Page 3)											0000_0000_0000											0x14										
Buffer Pointer (Page 4)											0000_0000_0000											0x18										

1. Device controller read/write; all others read-only.

The table below describes the next dTD pointer fields.

Table 18-91. Next dTD pointer

Bits	Description
31-5	Next transfer element pointer. This field contains the physical memory address of the next dTD to be processed. The field corresponds to memory address signals [31:5], respectively.
4-1	Reserved, should be cleared. Bits reserved for future use and should be cleared.
0	Terminate (T). 1=pointer is invalid. 0=Pointer is valid (points to a valid Transfer Element Descriptor). This bit indicates to the Device Controller that there are no more valid entries in the queue.

The table below describes the next dTD token fields.

Table 18-92. dTD token

Bits	Description
31	Reserved, should be cleared. Bit reserved for future use and should be cleared.
30-16	<p>Total Bytes. This field specifies the total number of bytes to be moved with this transfer descriptor. This field is decremented by the number of bytes actually moved during the transaction and only on the successful completion of the transaction.</p> <p>The maximum value software may store in the field is 5*4K(5000H). This is the maximum number of bytes 5 page pointers can access. Although it is possible to create a transfer up to 20K this assumes the 1st offset into the first page is 0. When the offset cannot be predetermined, crossing past the 5th page can be guaranteed by limiting the total bytes to 16K. Therefore, the maximum recommended transfer is 16K(4000H).</p> <p>If the value of the field is zero when the controller fetches this transfer descriptor (and the active bit is set), the device controller executes a zero-length transaction and retires the transfer descriptor.</p> <p>It is not a requirement for IN transfers that Total Bytes To Transfer be a multiple of Maximum Packet Length. If software builds such a transfer descriptor for an IN transfer, the last transaction will always be less that Maximum Packet Length.</p>
15	Interrupt On Complete (IOC). This bit is used to indicate if USBINT is to be set in response to device controller being finished with this dTD.
14-12	Reserved, should be cleared. Bits reserved for future use and should be cleared.
11-10	<p>Multiplier Override (MultiO). This field can be used for transmit ISO's (that is, ISO-IN) to override the multiplier in the QH. This field must be zero for all packet types that are not transmit-ISO.</p> <p>Example:</p> <p>if QH.multiplier = 3; Maximum packet size = 8; Total bytes = 15; MultiO = 0 [default] Three packets are sent: {Data2(8); Data1(7); Data0(0)}</p> <p>if QH.multiplier = 3; Maximum packet size = 8; Total bytes = 15; MultiO = 2 Two packets are sent: {Data1(8); Data0(7)}</p> <p>For maximal efficiency, software should compute MultiO = greatest integer of (Total Bytes/Max. Packet Size) except for the case when Total bytes = 0; then MultiO should be 1.</p> <p>NOTE: Non-ISO and non-TX endpoints must set MultiO = 00.</p>
9-8	Reserved, should be cleared. Bits reserved for future use and should be cleared.

Table continues on the next page...

Table 18-92. dTD token (continued)

Bits	Description
7-0	Status. This field is used by the Device Controller to communicate individual command execution states back to the Device Controller software. This field contains the status of the last transaction performed on this qTD. The bit encodings are: Bit Status Field Description 7 Active 6 Halted 5 Data Buffer Error 3 Transaction Error 4,2,0 Reserved, should be cleared

The following tables describe the buffer pointer page *n* fields.

Table 18-93. Buffer pointer page 0

Bits	Description
31-12	Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers.
11-0	Current Offset. Offset into the 4kb buffer where the packet is to begin.

Table 18-94. Buffer pointer page 1

Bits	Description
31-12	Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers.
11	Reserved
10-0	Frame Number. Written by the device controller to indicate the frame number in which a packet finishes. This is typically be used to correlate relative completion times of packets on an ISO endpoint.

Table 18-95. Buffer pointer pages 2-4

Bits	Description
31-12	Buffer Pointer. Selects the page offset in memory for the packet buffer. Non virtual memory systems will typically set the buffer pointers to a series of incrementing integers.
11-0	Reserved

18.8 Device operational model

The function of the device operation is to transfer a request in the memory image to and from the Universal Serial Bus. Using a set of linked list transfer descriptors, pointed to by a queue head, the device controller will perform the data transfers. The following sections explain the use of the device controller from the device controller driver (DCD) point-of-view and further describe how specific USB bus events relate to status changes in the device controller programmer's interface.

18.8.1 Device controller initialization

After hardware reset, the USB DR module is disabled until the run/stop bit (USBCMD[RS]) is set to a '1'.

In the disabled state, the pull-up on the USB D+ is not active which prevents an attach event from occurring. At a minimum, it is necessary to have the queue heads setup for endpoint zero before the device attach occurs. Shortly after the device is enabled, a USB reset will occur followed by setup packet arriving at endpoint 0. A queue head must be prepared so that the device controller can store the incoming setup packet.

To configure the external ULPI PHY the following initialization sequence is required.

1. After power-on reset the UTMI PHY will be in disabled state and the PLL will be held reset. The UTMI PHY should remain disabled if the ULPI is being used.
2. Set the CONTROL[PHY_CLK_SEL] bits to select the ULPI PHY as the source of USB controller PHY clock.
3. Wait for PHY clock to become valid. This can be determined by polling the CONTROL[PHY_CLK_VALID] status bit. Note that this bit is not valid once the CONTROL[USB_EN] bit is set.

Once the PHY clock is valid the user can proceed to the device controller initialization phase.

In order to initialize a device, the software should perform the following steps:

1. Set the controller mode to device mode. Optionally set USBMODE[SDIS] (streaming disable).

NOTE

Transitioning from host mode to device mode requires a device controller reset before modifying USBMODE.

2. Optionally modify the BURSTSIZE register.
3. Program PORTSC[PTS] if using a non-ULPI PHY.

4. Set CONTROL[USB_EN]
5. Allocate and initialize device queue heads in system memory Minimum: Initialize device queue heads 0 Tx and 0 Rx.

NOTE

All device queue heads must be initialized for control endpoints before the endpoint is enabled. Device queue heads for non-control endpoints must be initialized before the endpoint can be used.

For information on device queue heads, refer to [Device data structures](#).

6. Configure the ENDPOINTLISTADDR pointer.

For additional information on ENDPOINTLISTADDR, refer to the register table.

7. Enable the microprocessor interrupt associated with the USB DR module and optionally change setting of USBCMD[ITC].

Recommended: enable all device interrupts including: USBINT, USBERRINT, Port Change Detect, USB Reset Received, DCSuspend.

For a list of available interrupts refer to the USBINTR and the USBSTS register tables.

8. Set USBCMD[RS] to run mode.

After the run bit is set, a device reset will occur. The DCD must monitor the reset event and set the DEVICEADDR register, set the ENDPTCTRLx registers, and adjust the software state as described in [Bus reset](#).

NOTE

Endpoint 0 is designed as a control endpoint only and does not need to be configured using ENDPTCTRL0 register.

It is also not necessary to initially prime Endpoint 0 because the first packet received will always be a setup packet. The contents of the first setup packet will require a response in accordance with USB device framework command set.

18.8.2 Port state and control

From a chip or system reset, the USB_DR controller enters the powered state.

A transition from the powered state to the attach state occurs when the run/stop bit (USBCMD[RS]) is set to a '1'. After receiving a reset on the bus, the port will enter the defaultFS or defaultHS state in accordance with the protocol reset described in Appendix C.2 of the *USB Specification Rev. 2.0*. The figure below depicts the state of a USB 2.0 device.

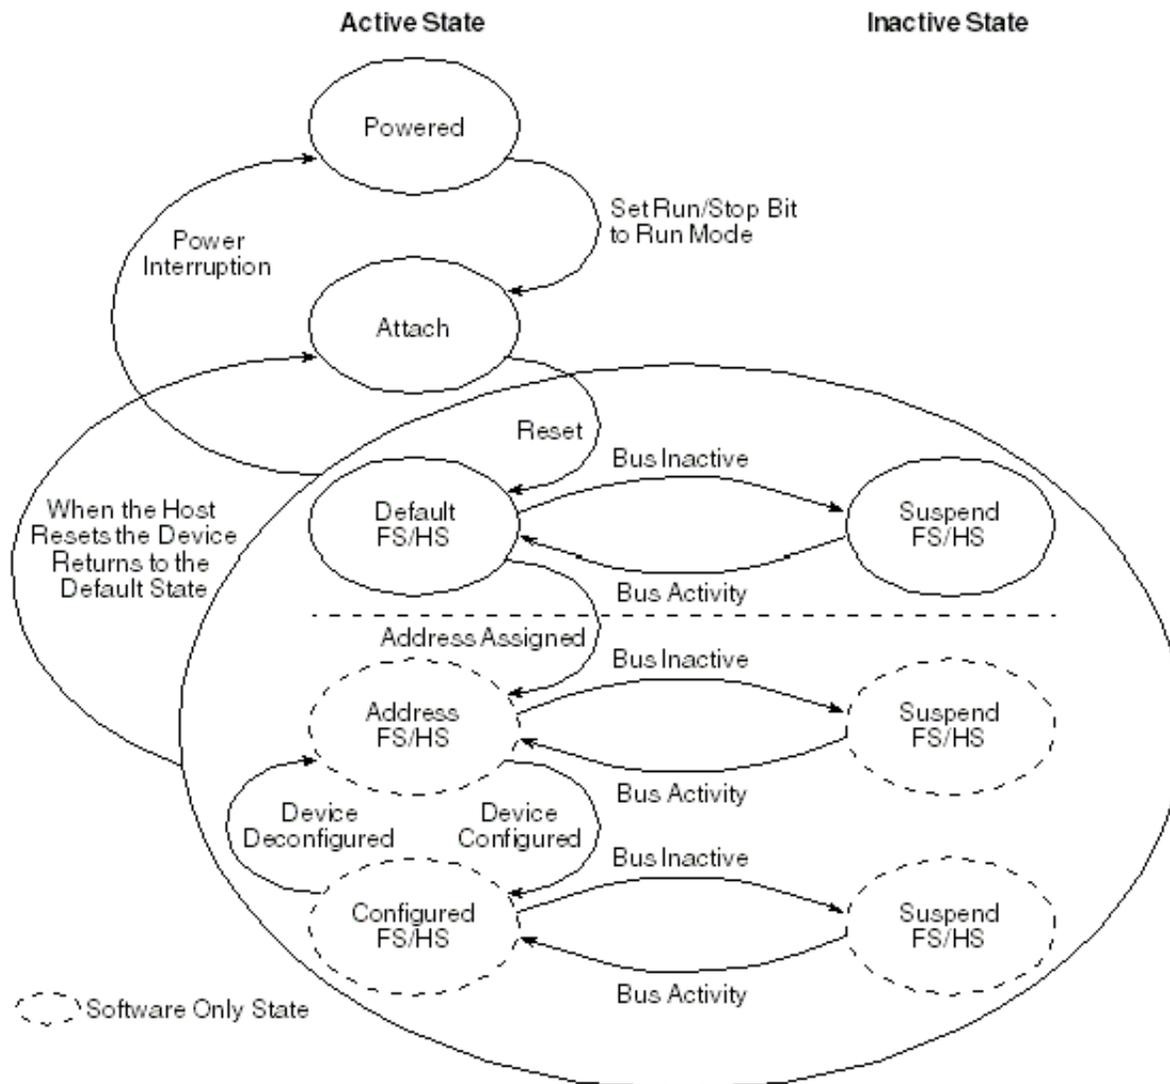


Figure 18-62. USB 2.0 device states

States Powered, Attach, DefaultFS/HS, SuspendFS/HS are implemented in the USB_DR controller and are communicated to the DCD using status bits, as shown in the table below:

Table 18-96. Device controller state information bits

Bits	Register
DCSuspend (SLI)	USBSTS

Table continues on the next page...

Table 18-96. Device controller state information bits (continued)

Bits	Register
USB Reset Received (URI)	USBSTS
Port Change Detect (PCI)	USBSTS
High-Speed Port	PORTSC

It is the responsibility of the DCD to maintain a state variable to differentiate between the defaultFS/HS state and the address/configured states. Change of state from default to address and the configured states is part of the enumeration process described in the device framework section of the USB 2.0 Specification.

As a result of entering the address state, the device address register (DEVICEADDR) must be programmed by the DCD.

Entry into the configured indicates that all endpoints to be used in the operation of the device have been properly initialized by programming the `ENDPTCTRLn` registers and initializing the associated queue heads.

18.8.2.1 Bus reset

A bus reset is used by the host to initialize downstream devices.

When a bus reset is detected, the USB_DR controller will renegotiate its attachment speed, reset the device address to 0, and notify the DCD by interrupt (assuming the USB reset interrupt enable bit, `USBINTR[URE]`, is set). After a reset is received, all endpoints (except endpoint 0) are disabled and any primed transactions are canceled by the device controller. The concept of priming is clarified below, but the DCD must perform the following tasks when a reset is received:

1. Clear all setup token semaphores by reading the `ENDPTSETUPSTAT` register and writing the same value back to the `ENDPTSETUPSTAT` register.
2. Clear all the endpoint complete status bits by reading the `ENDPTCOMPLETE` register and writing the same value back to the `ENDPTCOMPLETE` register.
3. Cancel all primed status by waiting until all bits in the `ENDPTPRIME` are 0 and then writing `0xFFFF_FFFF` to `ENDPTFLUSH`.
4. Read the reset bit in the `PORTSC` register (`PORTSC[PR]`) and make sure that it is still active.
 - A USB reset occurs for a minimum of 3 ms, and the DCD must reach this point in the reset cleanup before end of the reset occurs.
 - If it does not, a hardware reset of the device controller is recommended. A hardware reset can be performed by writing a one to the USB_DR controller reset bit in (`USBCMD[RST]`). Note that a hardware reset will cause the device to

detach from the bus by clearing USBCMD[RS] bit. Thus, the DCD must completely re-initialize the USB_DR controller after a hardware reset.

5. Free all allocated dTDs because they will no longer be executed by the device controller. If this is the first time the DCD is processing a USB reset event, then it is likely that no dTDs have been allocated.

At this time, the DCD may release control back to the OS because no further changes to the device controller are permitted until a Port Change Detect is indicated.

After a Port Change Detect, the device has reached the default state and the DCD can read the PORTSC to determine if the device is operating in FS or HS mode. At this time, the device controller has reached normal operating mode and DCD can begin enumeration according to the USB Chapter 9, Device Framework.

NOTE

The device DCD may use the FS/HS mode information to determine the bandwidth mode of the device.

In some applications, it may not be possible to enable one or more pipes while in FS mode. Beyond the data rate issue, there is no difference in DCD operation between FS and HS modes.

18.8.2.2 Suspend/resume

This section discusses the suspend and resume functions.

18.8.2.2.1 Suspend description

In order to conserve power, USB_DR controller automatically enters the suspended state when no bus traffic has been observed for a specified period.

When suspended, the USB_DR controller maintains any internal status, including its address and configuration. Attached devices must be prepared to suspend at any time they are powered, regardless of if they have been assigned a non-default address, are configured, or neither. Bus activity may cease due to the host entering a suspend mode of its own. In addition, a USB device shall also enter the suspended state when the hub port it is attached to is disabled.

The USB_DR controller exits suspend mode when there is bus activity. It may also request the host to exit suspend mode or selective suspend by using electrical signaling to indicate remote wake-up. The ability of a device to signal remote wake-up is optional. The USB_DR controller is capable of remote wake-up signaling. When the USB_DR controller is reset, remote wake-up signaling must be disabled.

18.8.2.2.2 Suspend operational model

The USB_DR controller moves into the suspend state when suspend signaling is detected or activity is missing on the upstream port for more than a specific period.

After the device controller enters the suspend state, the DCD is notified by an interrupt (assuming DC Suspend Interrupt is enabled). When the USBSTS[SLI] (device controller suspend) is set, the device controller is suspended.

DCD response when the device controller is suspended is application specific and may involve switching to low power operation. Information on the bus power limits in suspend state can be found in USB 2.0 specification.

18.8.2.2.3 Resume

If the USB_DR controller is suspended, its operation is resumed when any non-idle signaling is received on its upstream facing port.

In addition, the USB_DR controller can signal the system to resume operation by forcing resume signaling to the upstream port. Resume signaling is sent upstream by writing a '1' to the PORTSC[FPR] (resume bit) while the device is in suspend state. Sending resume signal to an upstream port should cause the host to issue resume signaling and bring the suspended bus segment (one or more devices) back to the active condition.

NOTE

Before resume signaling can be used, the host must enable it by using the Set Feature command defined in device framework (Chapter 9) of the USB 2.0 Specification.

18.8.3 Managing endpoints

The USB 2.0 specification defines an endpoint, also called a device endpoint or an address endpoint as a uniquely addressable portion of a USB device that can source or sink data in a communications channel between the host and the device.

The endpoint address is specified by the combination of the endpoint number and the endpoint direction.

The channel between the host and an endpoint at a specific device represents a data pipe. Endpoint 0 for a device is always a control type data channel used for device discovery and enumeration. Other types of endpoints support by USB include bulk, interrupt, and

isochronous. Each endpoint type has specific behavior related to packet response and error handling. More detail on endpoint operation can be found in the USB 2.0 specification.

The USB_DR controller supports up to six endpoint specified numbers. The DCD can enable, disable, and configure each endpoint.

Each endpoint direction is essentially independent and can be configured with differing behavior in each direction. For example, the DCD can configure endpoint 1-IN to be a bulk endpoint and endpoint 1-OUT to be an isochronous endpoint. This helps to conserve the total number of endpoints required for device operation. The only exception is that control endpoints must use both directions on a single endpoint number to function as a control endpoint. Endpoint 0, for example, is always a control endpoint and uses the pair of directions.

Each endpoint direction requires a queue head allocated in memory. If the maximum of 6 endpoint numbers, one for each endpoint direction are being used by the device controller, then 12 queue heads are required. The operation of an endpoint and use of queue heads are described later in this document.

18.8.3.1 Endpoint initialization

After hardware reset, all endpoints except endpoint zero are uninitialized and disabled.

The DCD must configure and enable each endpoint by writing to configuration bit in the `ENDPTCTRLn` register. Each 32-bit `ENDPTCTRLn` is split into an upper and lower half. The lower half of `ENDPTCTRLn` is used to configure the receive or OUT endpoint and the upper half is likewise used to configure the corresponding transmit or IN endpoint. Control endpoints must be configured the same in both the upper and lower half of the `ENDPTCTRLn` register otherwise the behavior is undefined. The table below shows how to construct a configuration word for endpoint initialization.

Table 18-97. Device controller endpoint initialization

Field	Value
Data Toggle Reset	1
Data Toggle Inhibit	0
Endpoint Type	00 Control 01 Isochronous 10 Bulk 11 Interrupt
Endpoint Stall	0

18.8.3.1.1 Stalling

There are two occasions where the USB_DR controller may need to return to the host a STALL.

The first occasion is the functional stall, which is a condition set by the DCD as described in the USB 2.0 device framework (Chapter 9). A functional stall is only used on non-control endpoints and can be enabled in the device controller by setting the endpoint stall bit in the `ENDPTCTRLn` register associated with the given endpoint and the given direction. In a functional stall condition, the device controller will continue to return STALL responses to all transactions occurring on the respective endpoint and direction until the endpoint stall bit is cleared by the DCD.

A protocol stall, unlike a function stall, is used on control endpoints is automatically cleared by the device controller at the start of a new control transaction (setup phase). When enabling a protocol stall, the DCD should enable the stall bits (both directions) as a pair. A single write to the `ENDPTCTRLn` register can ensure that both stall bits are set at the same instant.

NOTE

Any write to the `ENDPTCTRLn` register during operational mode must preserve the endpoint type field (that is, perform a read-modify-write).

The table below describes the device controller stall response matrix.

Table 18-98. Device controller stall response matrix

USB Packet	Endpoint Stall Bit.	Effect on STALL Bit.	USB Response
SETUP packet received by a non-control endpoint	N/A	None	STALL
IN/OUT/PING packet received by a non-control endpoint	'1	None	STALL
IN/OUT/PING packet received by a non-control endpoint	'0	None	ACK/NAK/NYET
SETUP packet received by a control endpoint	N/A	Cleared	ACK
IN/OUT/PING packet received by a control endpoint	'1	None	STALL
IN/OUT/PING packet received by a control endpoint	'0	None	ACK/NAK/NYET

18.8.3.2 Data toggle

Data toggle is a mechanism to maintain data coherency between host and device for any given data pipe.

For more information on data toggle, refer to the *Universal Serial Bus Revision 2.0 Specification*.

18.8.3.2.1 Data toggle reset

The DCD may reset the data toggle state bit and cause the data toggle sequence to reset in the device controller by writing a '1' to the data toggle reset bit in the `ENDPTCTRLn` register.

This should only be necessary when configuring/initializing an endpoint or returning from a STALL condition.

18.8.3.2.2 Data toggle inhibit

This feature is for test purposes only and should never be used during normal device controller operation.

Setting the data toggle Inhibit bit active ('1') causes the USB_DR controller to ignore the data toggle pattern that is normally sent and accept all incoming data packets regardless of the data toggle state.

In normal operation, the USB_DR controller checks the DATA0/DATA1 bit against the data toggle state bit to determine if the packet is valid. If Data PID does not match the data toggle state bit maintained by the device controller for that endpoint, the Data toggle is considered not valid. If the data toggle is not valid, the device controller assumes the packet was already received and discards the packet (not reporting it to the DCD). To prevent the USB_DR controller from re-sending the same packet, the device controller will respond to the error packet by acknowledging it with either an ACK or NYET response.

18.8.3.3 Device operational model for packet transfers

All transactions on the USB bus are initiated by the host and in turn, the device must respond to any request from the host within the turnaround time stated in the *Universal Serial Bus Revision 2.0 Specification*.

A USB host will send requests to the USB_DR controller in an order that can not be precisely predicted as a single pipeline, so it is not possible to prepare a single packet for the device controller to execute. However, the order of packet requests is predictable when the endpoint number and direction is considered. For example, if endpoint 2 (transmit direction) is configured as a bulk pipe, then the host sends IN requests to that endpoint. This USB_DR controller prepares packets for each endpoint/direction in anticipation of the host request. The process of preparing the device controller to send or

receive data in response to host initiated transaction on the bus is referred to as 'priming' the endpoint. This term is used throughout the following documentation to describe the USB_DR controller operation so the DCD can be architected properly use priming. Further, note that the term 'flushing' is used to describe the action of clearing a packet that was queued for execution.

18.8.3.3.1 Priming transmit endpoints

Priming a transmit endpoint will cause the device controller to fetch the device transfer descriptor (dTD) for the transaction pointed to by the device queue head (dQH).

After the dTD is fetched, it is stored in the dQH until the device controller completes the transfer described by the dTD. Storing the dTD in the dQH allows the device controller to fetch the operating context needed to handle a request from the host without the need to follow the linked list, starting at the dQH when the host request is received.

After the device has loaded the dTD, the leading data in the packet is stored in a FIFO in the device controller. This FIFO is split into virtual channels so that the leading data can be stored for any endpoint.

After a priming request is complete, an endpoint state of primed is indicated in the ENDPTSTATUS register. For a primed transmit endpoint, the device controller can respond to an IN request from the host and meet the stringent bus turnaround time of High Speed USB.

Since only the leading data is stored in the device controller FIFO, it is necessary for the device controller to begin filling in behind leading data after the transaction starts. The FIFO has been sized to account for the maximum latency that can be incurred by the system memory bus.

18.8.3.3.2 Priming receive endpoints

Priming receive endpoints is identical to priming of transmit endpoints from the point of view of the DCD.

At the device controller the major difference in the operational model is that there is no data movement of the leading packet data simply because the data is to be received from the host.

Note as part of the architecture, the FIFO for the receive endpoints is not partitioned into multiple channels like the transmit FIFO. Thus, the size of the RX FIFO does not scale with the number of endpoints.

18.8.3.4 Interrupt/bulk endpoint operational model

The behaviors of the device controller for interrupt and bulk endpoints are identical.

All valid IN and OUT transactions to bulk pipes will handshake with a NAK unless the endpoint had been primed. Once the endpoint has been primed, data delivery will commence.

A dTD is retired by the device controller when the packets described in the transfer descriptor have been completed. Each dTD describes N packets to be transferred according to the USB Variable Length transfer protocol. The formula and the following tables describe how the device controller computes the number and length of the packets to be sent/received by the USB vary according to the total number of bytes and maximum packet length.

With Zero Length Termination (ZLT) = 0

$$N = \text{INT}(\text{number of bytes}/\text{max. packet length}) + 1$$

With Zero Length Termination (ZLT) = 1

$$N = \text{MAXINT}(\text{number of bytes}/\text{max. packet length})$$

Table 18-99. Variable length transfer protocol example (ZLT = 0)

Bytes (dTD)	Max. packet length (dQH)	N	P1	P2	P3
511	256	2	256	255	-
512	256	3	256	256	0
512	512	2	512	0	-

Table 18-100. Variable length transfer protocol example (ZLT = 1)

Bytes (dTD)	Max. packet length (dQH)	N	P1	P2	P3
511	256	2	256	255	-
512	256	2	256	256	-
512	512	1	512	-	-

NOTE

The MULT field in the dQH must be set to '00' for bulk, interrupt, and control endpoints.

TX-dTD is complete when:

- All packets described dTD were successfully transmitted. *** Total bytes in dTD will equal zero when this occurs.

RX-dTD is complete when:

- All packets described in dTD were successfully received. *** Total bytes in dTD will equal zero when this occurs.
- A short packet (number of bytes < maximum packet length) was received. *** This is a successful transfer completion; DCD must check Total Bytes in dTD to determine the number of bytes that are remaining. From the total bytes remaining in the dTD, the DCD can compute the actual bytes received.
- A long packet was received (number of bytes > maximum packet size) OR (total bytes received > total bytes specified). *** This is an error condition. The device controller will discard the remaining packet, and set the Buffer Error bit in the dTD. In addition, the endpoint is flushed and the USBERR interrupt will become active.

On the successful completion of the packet(s) described by the dTD, the active bit in the dTD is cleared and the next pointer will be followed when the Terminate bit is clear. When the Terminate bit is set, the USB_DR controller will flush the endpoint/direction and cease operations for that endpoint/direction.

On the unsuccessful completion of a packet (see long packet above), the dQH is left pointing to the dTD that was in error. In order to recover from this error condition, the DCD must properly re-initialize the dQH by clearing the active bit and update the nextTD pointer before attempting to re-prime the endpoint.

NOTE

All packet level errors such as a missing handshake or CRC error will be retried automatically by the device controller.

There is no required interaction with the DCD for handling such errors.

18.8.3.4.1 Interrupt/bulk endpoint bus response matrix

The table below shows the interrupt/bulk endpoint bus response matrix.

Table 18-101. Interrupt/bulk endpoint bus response matrix

	Stall	Not Primed	Primed	Underflow	Overflow
Setup	Ignore	Ignore	Ignore	N/A	N/A
In	STALL	NAK	Transmit	BS Error ¹	N/A
Out	STALL	NAK	Receive + NYET/ACK ²	N/A	NAK
Ping	STALL	NAK	ACK	N/A	N/A
Invalid	Ignore	Ignore	Ignore	Ignore	Ignore

1. Force Bit Stuff Error.

2. NYET/ACK-NYET unless the Transfer Descriptor has packets remaining according to the USB variable length protocol then ACK.

SYSEERR-System error should never occur when the latency FIFOs are correctly sized and the DCD is responsive.

18.8.3.5 Control endpoint operation model

This section discusses the control endpoint operation model.

18.8.3.5.1 Setup phase

All requests to a control endpoint begin with a setup phase followed by an optional data phase and a required status phase.

The USB_DR controller will always accept the setup phase unless the setup lockout is engaged.

The setup lockout will engage so that future setup packets are ignored. Lockout of setup packets ensures that while software is reading the setup packet stored in the queue head, that data is not written as it is being read potentially causing an invalid setup packet.

The setup lockout mechanism can be disabled and a tripwire type semaphore will ensure that the setup packet payload is extracted from the queue head without being corrupted by an incoming setup packet. This is the preferred behavior because ignoring repeated setup packets due to long software interrupt latency would be a compliance issue.

Setup Packet Handling

- Disable Setup Lockout by writing '1' to Setup Lockout Mode (SLOM) in USBMODE (once at initialization). Setup lockout is not necessary when using the tripwire as described below.

NOTE

Leaving the Setup Lockout Mode as '0' will result in a potential compliance issue.

- After receiving an interrupt and inspecting ENDPTSETUPSTAT to determine that a setup packet was received on a particular pipe:
 - Write '1' to clear corresponding bit ENDPTSETUPSTAT.
 - Write '1' to Setup Tripwire (SUTW) in USBCMD register.
 - Duplicate contents of dQH.SetupBuffer into local software byte array.
 - Read Setup TripWire (SUTW) in USBCMD register. (if set-continue; if cleared-goto 2)
 - Write '0' to clear Setup Tripwire (SUTW) in USBCMD register.
 - Process setup packet using local software byte array copy and execute status/handshake phases.

NOTE

After receiving a new setup packet the status and/or handshake phases may still be pending from a previous control sequence. These should be flushed and de-allocated before linking a new status and/or handshake dTD for the most recent setup packet.

18.8.3.5.2 Data phase

If the control transfer requires a data stage following the setup phase, the DCD must create a device transfer descriptor for the data phase and prime the transfer.

After priming the packet, the DCD must verify a new setup packet has not been received by reading the ENDPTSETUPSTAT register immediately verifying that the prime had completed. A prime will complete when the associated bit in the ENDPTPRIME register is zero and the associated bit in the ENDPTSTATUS register is a one. If a prime fails, that is, The ENDPTPRIME bit goes to zero and the ENDPTSTATUS bit is not set, then the prime has failed. This can only be due to improper setup of the dQH, dTD or a setup arriving during the prime operation. If a new setup packet is indicated after the ENDPTPRIME bit is cleared, then the transfer descriptor can be freed and the DCD must reinterpret the setup packet.

Should a setup arrive after the data stage is primed, the device controller will automatically clear the prime status (ENDPTSTATUS) to enforce data coherency with the setup packet.

NOTE

The MULT field in the dQH must be set to '00' for bulk, interrupt, and control endpoints.

NOTE

Error handling of data phase packets is the same as bulk packets described previously.

18.8.3.5.3 Status phase

Similar to the data phase, the DCD must create a transfer descriptor (with byte length equal zero) and prime the endpoint for the status phase.

The DCD must also perform the same checks of the ENDPTSETUPSTAT as described above in the data phase.

NOTE

The MULT field in the dQH must be set to '00' for bulk, interrupt, and control endpoints.

NOTE

Error handling of data phase packets is the same as bulk packets described previously.

18.8.3.5.4 Control endpoint bus response matrix

The table below shows the device controller response to packets on a control endpoint, according to the device controller state.

Table 18-102. Control endpoint bus response matrix

Token type	Endpoint state					Setup lockout
	Stall	Not primed	Primed	Underflow	Overflow	
Setup	ACK	ACK	ACK	N/A	SYSERR ¹	
In	STALL	NAK	Transmit	BS Error ²	N/A	N/A
Out	STALL	NAK	Receive + NYET/ ACK ³	N/A	NAK	N/A
Ping	STALL	NAK	ACK	N/A	N/A	N/A
Invalid	Ignore	Ignore	Ignore	Ignore	Ignore	Ignore

1. SYSERR-System error should never occur when the latency FIFOs are correctly sized and the DCD is responsive.
2. Force Bit Stuff Error.
3. NYET/ACK-NYET unless the Transfer Descriptor has packets remaining according to the USB variable length protocol then ACK.

18.8.3.6 Isochronous endpoint operational model

Isochronous endpoints are used for real-time scheduled delivery of data and their operational model is significantly different than the host throttled Bulk, Interrupt, and Control data pipes.

Real time delivery by the USB_DR controller is accomplished by the following:

- Exactly MULT Packets per (micro)Frame are transmitted/received. Note that MULT is a two-bit field in the device Queue Head. The variable length packet protocol is not used on isochronous endpoints.
- NAK responses are not used. Instead, zero length packets and sent in response to an IN request to an unprimed endpoints. For unprimed RX endpoints, the response to an OUT transaction is to ignore the packet within the device controller.
- Prime requests always schedule the transfer described in the dTD for the next (micro)frame. If the ISO-dTD is still active after that frame, then the ISO-dTD is held ready until executed or canceled by the DCD.

The USB_DR controller in host mode uses the periodic frame list to schedule data exchanges to Isochronous endpoints. The operational model for device mode does not use such a data structure. Instead, the same dTD used for Control/Bulk/Interrupt endpoints is also used for isochronous endpoints. The difference is in the handling of the dTD.

The first difference between bulk and ISO-endpoints is that priming an ISO-endpoint is a delayed operation such that an endpoint will become primed only after a SOF is received. After the DCD writes the prime bit, the prime bit is cleared as usual to indicate to software that the device controller completed a priming the dTD for transfer. Internal to the design, the device controller hardware masks that prime start until the next frame boundary. This behavior is hidden from the DCD but occurs so that the device controller can match the dTD to a specific (micro)frame.

Another difference with isochronous endpoints is that the transaction must wholly complete in a (micro)frame. Once an ISO transaction is started in a (micro)frame it will retire the corresponding dTD when MULT transactions occur or the device controller finds a fulfillment condition.

The transaction error bit set in the status field indicates a fulfillment error condition. When a fulfillment error occurs, the frame after the transfer failed to complete wholly, the device controller will force retire the ISO-dTD and move to the next ISO-dTD.

It is important to note that fulfillment errors are only caused due to partially completed packets. If no activity occurs to a primed ISO-dTD, the transaction will stay primed indefinitely. This means it is up to software discard transmit ISO-dTDs that pile up from a failure of the host to move the data.

Finally, the last difference with ISO packets is in the data level error handling. When a CRC error occurs on a received packet, the packet is not retried similar to bulk and control endpoints. Instead, the CRC is noted by setting the Transaction Error bit and the data is stored as usual for the application software to sort out.

- TX Packet Retired
 - MULT counter reaches zero.
 - Fulfillment Error [Transaction Error bit is set]
 - #Packets Occurred > 0 AND # Packets Occurred < MULT

NOTE

For TX-ISO, MULT Counter can be loaded with a lesser value in the dTD Multiplier Override field. If the Multiplier Override is zero, the MULT Counter is initialized to the Multiplier in the QH.

- RX Packet Retired:
 - MULT counter reaches zero.

- Non-MDATA Data PID is received
- Overflow Error:
- Packet received is > maximum packet length. [Buffer Error bit is set]
- Packet received exceeds total bytes allocated in dTD. [Buffer Error bit is set]
- Fulfillment Error [Transaction Error bit is set]
- # Packets Occurred > 0 AND # Packets Occurred < MULT
- CRC Error [Transaction Error bit is set]

NOTE

For ISO, when a dTD is retired, the next dTD is primed for the next frame. For continuous (micro)frame to (micro)frame operation the DCD should ensure that the dTD linked-list is out ahead of the device controller by at least two (micro)frames.

18.8.3.6.1 Isochronous pipe synchronization

When it is necessary to synchronize an isochronous data pipe to the host, the (micro)frame number (FRINDEX register) can be used as a marker.

To cause a packet transfer to occur at a specific (micro)frame number [N], the DCD should interrupt on SOF during frame N - 1. When the FRINDEX = N - 1, the DCD must write the prime bit. The USB_DR controller will prime the isochronous endpoint in (micro)frame N - 1 so that the device controller will execute delivery during (micro)frame N.

NOTE

Priming an endpoint towards the end of (micro)frame N - 1 will not guarantee delivery in (micro)frame N. The delivery may actually occur in (micro)frame N + 1 if device controller does not have enough time to complete the prime before the SOF for packet N is received.

18.8.3.6.2 Isochronous endpoint bus response matrix

The table below shows the isochronous endpoint bus response matrix.

Table 18-103. Isochronous endpoint bus response matrix

	Stall	Not Primed	Primed	Underflow	Overflow
Setup	STALL	STALL	STALL	N/A	N/A
In	NULL ¹ Packet	NULL Packet	Transmit	BS Error ²	N/A
Out	Ignore	Ignore	Receive	N/A	Drop Packet
Ping	Ignore	Ignore	Ignore	Ignore	Ignore
Invalid	Ignore	Ignore	Ignore	Ignore	Ignore

1. Zero Length Packet.
2. Force Bit Stuff Error.

18.8.4 Managing queue heads

The figure below shows the endpoint queue head diagram.

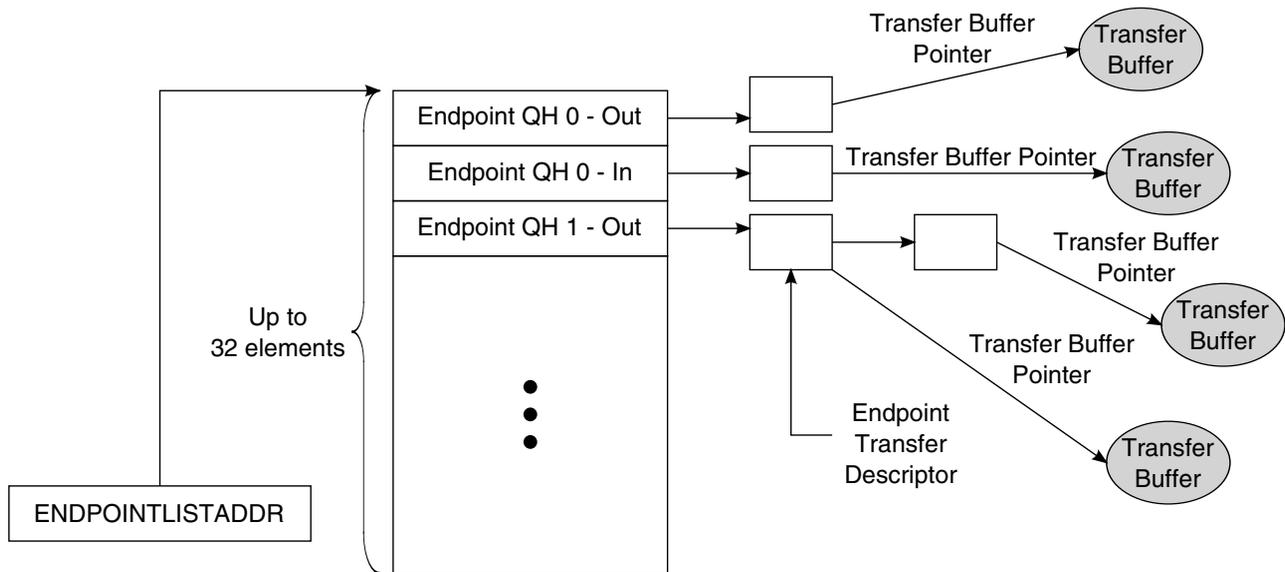


Figure 18-63. Endpoint queue head diagram

The device queue head (dQH) points to the linked list of transfer tasks, each depicted by the device Transfer Descriptor (dTD). An area of memory pointed to by ENDPOINTLISTADDR contains a group of all dQH's in a sequential list as shown in Figure 18-63. The even elements in the list of dQH's are used for receive endpoints (OUT/SETUP) and the odd elements are used for transmit endpoints (IN/INTERRUPT). Device transfer descriptors are linked head to tail starting at the queue head and ending at a terminate bit. Once the dTD has been retired, it will no longer be part of the linked list from the queue head. Therefore, software is required to track all transfer descriptors since pointers will no longer exist within the queue head once the dTD is retired (see [Software link pointers](#)).

In addition to the current and next pointers and the dTD overlay, the dQH also contains the following parameters for the associated endpoint: Multiplier, Maximum Packet Length, Interrupt On Setup. The complete initialization of the dQH including these fields is demonstrated in the next section.

18.8.4.1 Queue head initialization

One pair of device queue heads must be initialized for each active endpoint.

To initialize a device queue head:

- Write the MaxPacketSize field as required by the USB Chapter 9 or application specific protocol.
- Write the multiplier field to 0 for control, bulk, and interrupt endpoints. For ISO endpoints, set the multiplier to 1, 2, or 3 as required bandwidth an in conjunction with the USB Chapter 9 protocol. Note that in FS mode, the multiplier field can only be 1 for ISO endpoints.
- Write the next dTD Terminate bit field to '1.'
- Write the Active bit in the status field to '0.'
- Write the Halt bit in the status field to '0.'

NOTE

The DCD must only modify dQH if the associated endpoint is not primed and there are no outstanding dTDs.

18.8.4.2 Operational model for setup transfers

As discussed in [Control endpoint operation model](#), setup transfer requires special treatment by the DCD.

A setup transfer does not use a dTD but instead stores the incoming data from a setup packet in an 8-byte buffer within the dQH.

Upon receiving notification of the setup packet, the DCD should handle the setup transfer as demonstrated here:

1. Copy setup buffer contents from dQH - RX to software buffer.
2. Acknowledge setup backup by writing a '1' to the corresponding bit in ENDPTSETUPSTAT.

NOTE

The acknowledge must occur before continuing to process the setup packet.

NOTE

After the acknowledge has occurred, the DCD must not attempt to access the setup buffer in the dQH - RX. Only the local software copy should be examined.

3. Check for pending data or status dTD's from previous control transfers and flush if any exist as discussed in [Flushing/depriming an endpoint](#).

NOTE

It is possible for the device controller to receive setup packets before previous control transfers complete. Existing control packets in progress must be flushed and the new control packet completed.

4. Decode setup packet and prepare data phase [optional] and status phase transfer as required by the USB Chapter 9 or application specific protocol.

18.8.5 Managing transfers with transfer descriptors

This section describes software link pointers, transfer descriptors, transfer completion and the device error matrix.

18.8.5.1 Software link pointers

It is necessary for the DCD software to maintain head and tail pointers for the linked list of dTDs for each respective queue head.

This is necessary because the dQH only maintains pointers to the current working dTD and the next dTD to be executed. The operations described in next section for managing dTD will assume the DCD can use reference the head and tail of the dTD linked list.

NOTE

To conserve memory, the reserved fields at the end of the dQH can be used to store the Head and Tail pointers but it still remains the responsibility of the DCD to maintain the pointers.

The figure below shows the software link pointers.

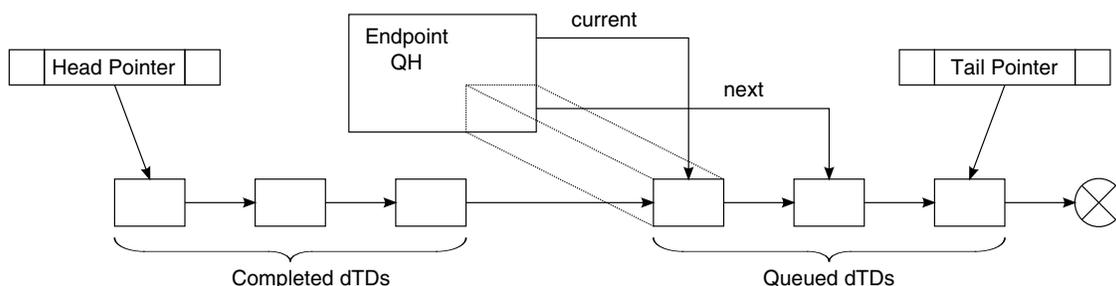


Figure 18-64. Software link pointers

18.8.5.2 Building a transfer descriptor

Before a transfer can be executed from the linked list, a dTD must be built to describe the transfer.

Use the following procedure for building dTDs.

Allocate 8-DWord dTD block of memory aligned to 8-DWord boundaries. Example: bit address 4-0 would be equal to '00000'.

Write the following fields:

1. Initialize first seven DWords to '0'.
2. Set the terminate bit to '1'.
3. Fill in total bytes with transfer size.
4. Set the interrupt on complete if desired.
5. Initialize the status field with the active bit set to '1' and all remaining status bits set to '0'.
6. Fill in buffer pointer page 0 and the current offset to point to the start of the data buffer.
7. Initialize buffer pointer page 1 through page 4 to be one greater than each of the previous buffer pointer.

18.8.5.3 Executing a transfer descriptor

To safely add a dTD, the DCD must account for the event in which the device controller reaches the end of the dTD list at the same time a new dTD is being added to the end of the list.

First, determine whether the link list is empty by checking the DCD driver to see if the pipe is empty (internal representation of linked-list should indicate if any packets are outstanding). Then follow the sequence of actions in the following list as appropriate, depending on whether the link list is empty or not empty.

Case 1: Link list is empty

1. Write dQH next pointer AND dQH terminate bit to '0' as a single DWord operation.
2. Clear active and halt bit in dQH (in case set from a previous error).
3. Prime endpoint by writing '1' to correct bit position in ENDPTPRIME.

Case 2: Link list is not empty

1. Add dTD to end of linked list.

2. Read correct prime bit in ENDPTPRIME-if '1' DONE.
3. Set ATDTW bit in USBCMD register to '1'.
4. Read correct status bit in ENDPTSTATUS. (store in tmp. variable for later).
5. Read ATDTW bit in USBCMD register.
If '0' goto 3.
If '1' continue to 6.
6. Write ATDTW bit in USBCMD register to '0'.
7. If status bit read in (4) is '1' DONE.
8. If status bit read in (4) is '0' then Goto Case 1: Step 1.

18.8.5.4 Transfer completion

After a dTD has been initialized and the associated endpoint primed, the device controller will execute the transfer upon the host-initiated request.

The DCD is notified with a USB interrupt if the Interrupt On Complete bit was set or alternately, the DCD can poll the endpoint complete register to find when the dTD had been executed. After a dTD has been executed, DCD can check the status bits to determine success or failure.

NOTE

Multiple dTD can be completed in a single endpoint complete notification. After clearing the notification, DCD must search the dTD linked list and retire all dTDs that have finished (Active bit cleared).

By reading the status fields of the completed dTDs, the DCD can determine if the transfers completed successfully. Success is determined with the following combination of status bits:

- Active = 0
- Halted = 0
- Transaction Error = 0
- Data Buffer Error = 0

Should any combination other than the one shown above exist, the DCD must take proper action. Transfer failure mechanisms are indicated in the Device Error Matrix.

In addition to checking the status bit the DCD must read the Transfer Bytes field to determine the actual bytes transferred. When a transfer is complete, the Total Bytes transferred is by decremented by the actual bytes transferred. For Transmit packets, a

packet is only complete after the actual bytes reaches zero, but for receive packets, the host may send fewer bytes in the transfer according the USB variable length packet protocol.

18.8.5.5 Flushing/depriming an endpoint

It is necessary for the DCD to flush to deprime one or more endpoints on a USB device reset or during a broken control transfer.

There may also be application specific requirements to stop transfers in progress. The following procedure can be used by the DCD to stop a transfer in progress:

1. Write a '1' to the corresponding bit(s) in ENDPTFLUSH.
2. Wait until all bits in ENDPTFLUSH are '0'.
3. Software note: this operation may take a large amount of time depending on the USB bus activity. It is not desirable to have this wait loop within an interrupt service routine.
4. Read ENDPTSTATUS to ensure that for all endpoints commanded to be flushed, that the corresponding bits are now '0' If the corresponding bits are '1' after step #2 has finished, then the flush failed as described in the following:

Explanation: In very rare cases, a packet is in progress to the particular endpoint when commanded flush using ENDPTFLUSH. A safeguard is in place to refuse the flush to ensure that the packet in progress completes successfully. The DCD may need to repeatedly flush any endpoints that fail to flush by repeating steps 1-3 until each endpoint is successfully flushed.

18.8.5.6 Device error matrix

The table below summarizes packet errors that are not automatically handled by the USB controller.

Table 18-104. Device error matrix

Error	Direction	Packet Type	Data Buffer Error Bit	Transaction Error Bit
Overflow **	RX	Any	1	0
ISO Packet Error	RX	ISO	0	1
ISO Fulfillment Error	Both	ISO	0	1

Notice that the device controller handles all errors on Bulk/Control/Interrupt Endpoints except for a data buffer overflow. However, for ISO endpoints, errors packets are not retried and errors are tagged as indicated. The table below provides the error descriptions.

Table 18-105. Error descriptions

Overflow	Number of bytes received exceeded max. packet size or total buffer length. NOTE: This error also sets the Halt bit in the dQH. If there are dTDs remaining in the linked list for the endpoint, they will not be executed.
ISO Packet Error	CRC Error on received ISO packet. Contents not guaranteed to be correct.
ISO Fulfillment Error	Host failed to complete the number of packets defined in the dQH mult field within the given (micro)frame. For scheduled data delivery the DCD may need to readjust the data queue because a fulfillment error will cause Device Controller to cease data transfers on the pipe for one (micro)frame. During the dead (micro)frame, the Device Controller reports error on the pipe and primes for the following frame.

18.8.6 Servicing interrupts

The interrupt service routine must consider that there are high-frequency, low-frequency operations, and error operations and order accordingly.

18.8.6.1 High-frequency interrupts

High frequency interrupts in particular should be handed in the order shown in the table below. The most important of these is listed first because the DCD must acknowledge a setup buffer in the timeliest manner possible.

Table 18-106. Interrupt handling order

Execution order	Interrupt	Action
1a	USB Interrupt ¹ ENDPTSETUPSTATUS	Copy contents of setup buffer and acknowledge setup packet (as indicated in Managing queue heads). Process setup packet according to USB 2.0 Chapter 9 or application specific protocol.
1b	USB Interrupt ENDPTCOMPLETE	Handle completion of dTD as indicated in Managing queue heads .
2	SOF Interrupt	Action as deemed necessary by application. This interrupt may not have a use in all applications.

1. It is likely that multiple interrupts to stack up on any call to the Interrupt Service Routine AND during the Interrupt Service Routine.

18.8.6.2 Low-frequency interrupts

The low frequency events include the interrupts shown in the table below.

These interrupts can be handled in any order because they do not occur often in comparison to the high-frequency interrupts.

Table 18-107. Low frequency interrupt events

Interrupt	Action
Port Change	Change software state information.
Sleep enable (Suspend)	Change software state information. Low power handling as necessary.
Reset Received	Change software state information. Abort pending transfers.

18.8.6.3 Error interrupts

Error interrupts are the least frequently occurring events. They should be placed last in the interrupt service routine.

The table below shows the error interrupt events.

Table 18-108. Error interrupt events

Interrupt	Action
USB Error Interrupt	This error is redundant because it combines USB Interrupt and an error status in the dTD. The DCD will more aptly handle packet-level errors by checking dTD status field upon receipt of USB Interrupt (w/ ENDPTCOMPLETE).
System Error	Unrecoverable error. Immediate Reset of core; free transfers buffers in progress and restart the DCD.

18.9 Deviations from the EHCI specifications

The host mode operation of the USB DR module is nearly EHCI-compatible with few minor differences. For the most part, the module conforms to the data structures and operations described in Section 3, "Data Structures," and Section 4, "Operational Model," in the EHCI specification. The particulars of the deviations occur in the following areas:

- Embedded transaction translator-Allows direct attachment of FS and LS devices in host mode without the need for a companion controller.

- Device operation-In host mode, the device operational registers are generally disabled and thus device mode is mostly transparent when in host mode. However, there are a couple exceptions documented in the following sections.
- Embedded design interface-The module does not have a PCI interface and therefore the PCI configuration registers described in the EHCI specification are not applicable.

18.9.1 Embedded transaction translator function

The DR module supports directly connected full and low speed devices without requiring a companion controller by including the capabilities of a USB 2.0 high speed hub transaction translator.

Although there is no separate transaction translator block in the system, the transaction translator function normally associated with a high speed hub has been implemented within the DMA and Protocol engine blocks. The embedded transaction translator function is an extension to EHCI interface, but makes use of the standard data structures and operational models that exist in the EHCI specification to support full and low speed devices.

18.9.1.1 Capability registers

The following additions have been added to the capability registers to support the embedded transaction translator function:

- N_TT added to HSCPARAMS-Host Controller Structural Parameters
- N_PTT added to HSCPARAMS-Host Controller Structural Parameters

See [Host controller structural parameters \(USB_HCSPARAMS\)](#), for usage information.

18.9.1.2 Operational registers

The following additions have been added to the operational registers to support the embedded TT:

- ASYNCTTSTS is a new register.
- Addition of two-bit Port Speed (PSPD) to the PORTSC register.

18.9.1.3 Discovery

In a standard EHCI controller design, the EHCI host controller driver detects a full speed (FS) or low speed (LS) device by noting if the port enable bit is set after the port reset operation.

The port enable will only be set in a standard EHCI controller implementation after the port reset operation and when the host and device negotiate a High-Speed connection (that is, Chirp completes successfully).

The module always sets the port enable after the port reset operation regardless of the result of the host device chirp result. The resulting port speed is indicated by the PSPD field in PORTSC. Therefore, the standard EHCI host controller driver requires an alteration to handle directly connected full- and low-speed devices or hubs. The table below summarizes the functional differences between EHCI and EHCI with embedded TT.

Table 18-109. Functional differences between EHCI and EHCI with embedded TT

Standard EHCI	EHCI with embedded transaction translator
After port enable bit is set following a connection and reset sequence, the device/hub is assumed to be HS.	After port enable bit is set following a connection and reset sequence, the device/hub speed is noted from PORTSC.
FS and LS devices are assumed to be downstream from a HS hub thus, all port-level control is performed through the Hub Class to the nearest Hub.	FS and LS device can be either downstream from a HS hub or directly attached. When the FS/LS device is downstream from a HS hub, then port-level control is done using the Hub Class through the nearest Hub. When a FS/LS device is directly attached, then port-level control is accomplished using PORTSC.
FS and LS devices are assumed to be downstream from a HS hub with HubAddr=X. [where HubAddr > 0 and HubAddr is the address of the Hub where the bus transitions from HS to FS/LS (that is, Split target hub)]	FS and LS device can be either downstream from a HS hub with HubAddr = X [HubAddr > 0] or directly attached [where HubAddr = 0 and HubAddr is the address of the Root Hub where the bus transitions from HS to FS/LS (that is, Split target hub is the root hub)]

18.9.1.4 Data structures

The same data structures used for FS/LS transactions through a HS hub are also used for transactions through the root hub.

The following list demonstrates how the hub address and endpoint speed fields should be set for directly attached FS/LS devices and hubs:

1. QH (for direct attach FS/LS)-Async. (Bulk/Control Endpoints) Periodic (Interrupt)
 - Hub Address = 0
 - Transactions to direct attached device/hub.

- QH.EPS = Port Speed
- Transactions to a device downstream from direct attached FS hub.
 - QH.EPS = Downstream Device Speed

NOTE

When QH.EPS = 01 (LS) and PORTSC[PSPD] = 00 (FS), a LS-pre-pid is sent before the transmitting LS traffic.

Maximum Packet Size must be less than or equal 64 or undefined behavior may result.

2. siTD (for direct attach FS)-Periodic (ISO Endpoint)

- All FS ISO transactions:
 - Hub Address = 0
 - siTD.EPS = 00 (full speed)

Maximum packet size must less than or equal to 1023 or undefined behavior may result.

18.9.1.5 Operational model

The operational models are well defined for the behavior of the transaction translator (see *Universal Serial Bus Revision 2.0 Specification*) and for the EHCI controller moving packets between system memory and a USB-HS hub.

Since the embedded transaction translator exists within the DR module there is no physical bus between EHCI host controller driver and the USB FS/LS bus. These sections will briefly discuss the operational model for how the EHCI and transaction translator operational models are combined without the physical bus between. The following sections assume the reader is familiar with both the EHCI and USB 2.0 transaction translator operational models.

18.9.1.5.1 Microframe pipeline

The EHCI operational model uses the concept of H-frames and B-frames to describe the pipeline between the host (H) and the bus (B).

The embedded transaction translator shall use the same pipeline algorithms specified in the *Universal Serial Bus Revision 2.0 Specification* for a Hub-based transaction translator.

All periodic transfers always begin at B-frame 0 (after SOF) and continue until the stored periodic transfers are complete. As an example of the microframe pipeline implemented in the embedded transaction translator, all periodic transfers that are tagged in EHCI to execute in H-frame 0 will be ready to execute on the bus in B-frame 0.

It is important to note that when programming the S-mask and C-masks in the EHCI data structures to schedule periodic transfers for the embedded transaction translator, the EHCI host controller driver must follow the same rules specified in EHCI for programming the S-mask and C-mask for downstream Hub-based transaction translators.

Once periodic transfers are exhausted, any stored asynchronous transfer are moved. Asynchronous transfers are opportunistic in that they shall execute whenever possible and their operation is not tied to H-frame and B-frame boundaries with the exception that an asynchronous transfer can not babble through the SOF (start of B-frame 0).

18.9.1.5.2 Split state machines

The start and complete split operational model differs from EHCI slightly because there is no bus medium between the EHCI controller and the embedded transaction translator.

Where a start or complete-split operation would occur by requesting the split to the HS hub, the start/complete split operation is simple an internal operation to the embedded transaction translator. The table below summarizes the conditions where handshakes are emulated from internal state instead of actual handshakes to HS split bus traffic.

Table 18-110. Emulated handshakes

Condition	Emulate TT response
Start-Split: All asynchronous buffers full	NAK
Start-Split: All periodic buffers full	ERR
Start-Split: Success for start of Async. Transaction	ACK
Start-Split: Start Periodic Transaction	No handshake (Ok)
Complete-Split: Failed to find transaction in queue	Bus Time Out
Complete-Split: Transaction in Queue is Busy	NYET
Complete-Split: Transaction in Queue is Complete	[Actual handshake from FS/LS device]

18.9.1.5.3 Asynchronous transaction scheduling and buffer management

The following *Universal Serial Bus Revision 2.0 Specification* items are implemented in the embedded transaction translator:

- USB 2.0-11.17.3
 - Sequencing is provided and a packet length estimator ensures no full-speed/low-speed packet babbles into SOF time.

- USB 2.0-11.17.4
 - Transaction tracking for 2 data pipes.
- USB 2.0-11.17.5
 - Clear_TT_Buffer capability provided

18.9.1.5.4 Periodic transaction scheduling and buffer management

The following *Universal Serial Bus Revision 2.0 Specification* items are implemented in the embedded transaction translator:

- USB 2.0-11.18.6.[1-2]
 - Abort of pending start-splits
 - EOF (and not started in microframes 6)
 - Idle for more than 4 microframes
 - Abort of pending complete-splits
 - EOF
 - Idle for more than 4 microframes

NOTE

There is no data schedule mechanism for these transactions other than the microframe pipeline. The embedded TT assumes the number of packets scheduled in a frame does not exceed the frame duration (1 msec) or else undefined behavior may result.

18.9.1.5.5 Multiple transaction translators

The maximum number of embedded transaction translators that is currently supported is one as indicated by the N_TT field in the HCSPARAMS register.

See [Host controller structural parameters \(USB_HCSPARAMS\)](#), for more information.

18.9.2 Device operation

The co-existence of a device operational controller within the DR module has little effect on EHCI compatibility for host operation except as noted in this section.

18.9.3 Non-zero fields the register file

Some of the reserved fields and reserved addresses in the capability registers and operational registers have use in device mode, the following must be adhered to:

- Write operations to all EHCI reserved fields (some of which are device fields in the DR module) in the operation registers should always be written to zero. This is an EHCI requirement of the device controller driver that must be adhered to.
- Read operations by the module must properly mask EHCI reserved fields (some of which are device fields in the DR module registers).

18.9.4 SOF interrupt

The SOF interrupt is a free running 125 μ sec interrupt for host mode.

EHCI does not specify this interrupt, but it has been added for convenience and as a potential software time base. Note that the free running interrupt is shared with the device-mode start-of-frame interrupt. See [USB status \(USB_USBSTS\)](#), and [USB interrupt enable \(USB_USBINTR\)](#), for more information.

18.9.5 Embedded design

This is an Embedded USB Host Controller as defined by the EHCI specification and thus does not implement the PCI configuration registers.

18.9.5.1 Frame adjust register

Given that the optional PCI configuration registers are not included in this implementation, there is no corresponding bit level timing adjustments like those provided by the Frame Adjust register in the PCI configuration registers.

Starts of microframes are timed precisely to 125 μ sec using the transceiver clock as a reference clock. That is, 60-MHz transceiver clock for 8-bit physical interfaces and full-speed serial interfaces or 30-MHz transceiver clock.

18.9.6 Miscellaneous variations from EHCI

The modules support multiple physical interfaces which can operate in different modes when the module is configured with the software programmable Physical Interface Modes.

The control bits for selecting the PHY operating mode have been added to the PORTSC register providing a capability that is not defined by the EHCI specification.

18.9.6.1 Discovery

This section discusses port reset and port speed detection.

18.9.6.1.1 Port reset

The port connect methods specified by EHCI require setting the port reset bit in the register for a duration of 10 msec.

Due to the complexity required to support the attachment of devices that are not high speed there are counter already present in the design that can count the 10 msec reset pulse to alleviate the requirement of the software to measure this duration. Therefore, the basic connection is then summarized as the following:

- [Port Change Interrupt] Port connect change occurs to notify the host controller driver that a device has attached.
- Software shall write a '1' to the reset the device.
- Software shall write a '0' to the reset the device after 10 msec.
 - This step, which is necessary in a standard EHCI design, may be omitted with this implementation. Should the EHCI host controller driver attempt to write a '0' to the reset bit while a reset is in progress the write will simple be ignored and the reset will continue until completion.
- [Port Change Interrupt] Port enable change occurs to notify the host controller that the device is now operational and at this point the port speed has been determined.

18.9.6.1.2 Port speed detection

After the port change interrupt indicates that a port is enabled, the EHCI stack should determine the port speed.

Unlike the EHCI implementation which will re-assign the port owner for any device that does not connect at High-Speed, this host controller supports direct attach of non-HS devices. Therefore, the following differences are important regarding port speed detection:

- Port owner is read-only and always reads 0.
- A 2-bit port speed indicator has been added to PORTSC to provide the current operating speed of the port to the host controller driver.
 - A 1-bit high-speed indicator has been added to PORTSC to signify that the port is in HS vs. FS/LS



Chapter 19

Enhanced Serial Peripheral Interface

The enhanced serial peripheral interface (eSPI) allows the device to exchange data with peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.

19.1 Introduction

The enhanced serial peripheral interface (eSPI) allows the device to exchange data with peripheral devices such as EEPROMs, real-time clocks, A/D converters, and ISDN devices.

The eSPI is a full-duplex, synchronous, character-oriented channel that supports a simple interface (receive, transmit, clock and chip selects). The eSPI consists of transmitter and receiver sections, an independent baud-rate generator, and a control unit. The transmitter and receiver sections use the same clock, which is derived from the eSPI baud rate generator in master mode. During an eSPI transfer, data is sent and received simultaneously.

The eSPI receiver and transmitter each have a FIFO of 32 bytes, as shown below. When the eSPI is disabled in the eSPI mode register ($SPMODE[EN] = 0$), it consumes little power.

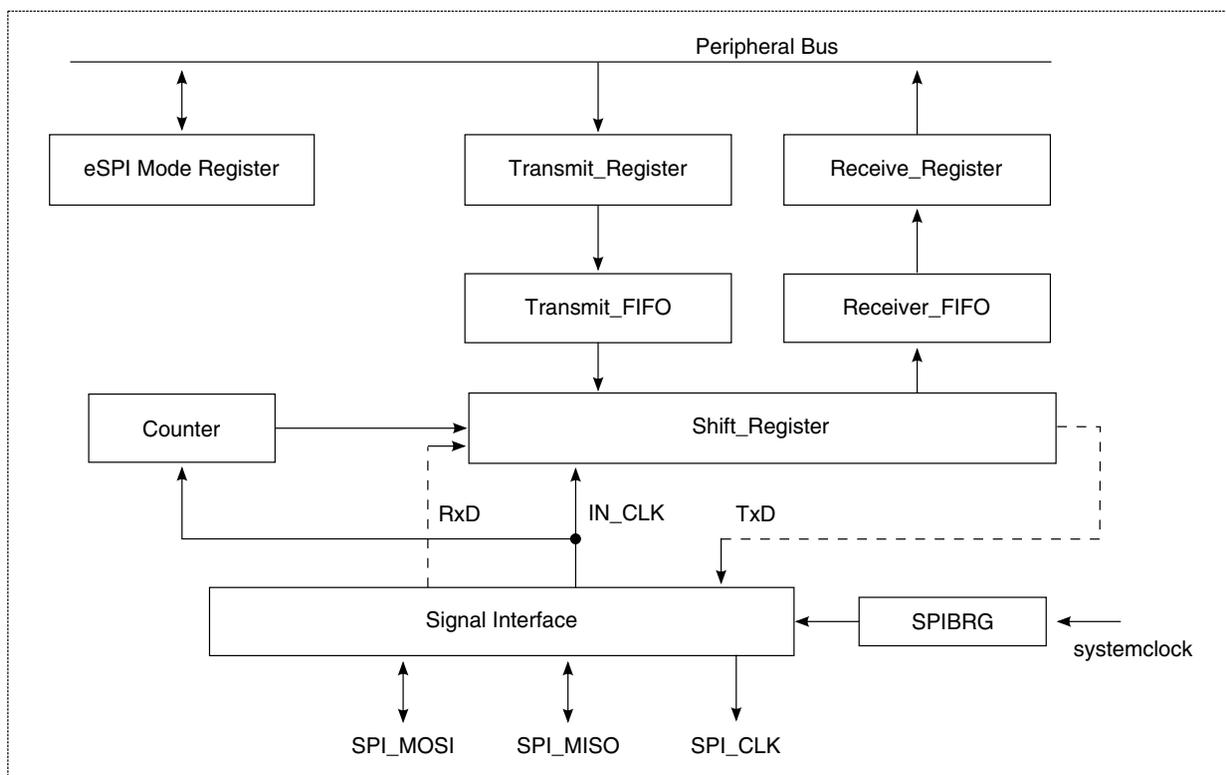


Figure 19-1. eSPI block diagram

19.1.1 Features

The major features of the eSPI are listed as follows:

- Interface contains SPI_MOSI, SPI_MISO, SPI_CLK, and 4 chip selects
- Supports eSPI master
- Supports RapidS™ full clock cycle operation
- Full-duplex or half-duplex master operation
- Supports Winbond dual output read
- Command in transaction level-easier for accessing eSPI devices
- Works with a range from 4-bit to 16-bit data characters
- Supports back-to-back character transmission and reception
- Supports reverse data mode for 8/16 bits data characters
- Supports single-master environment
- Maximum clock rate possible is (platform clock rate/ 4)
- Independent programmable baud rate generator
- Programmable clock phase and polarity.
- Supports four different configurations per chip select

- Local loopback capability for testing
- Supports booting from eSPI interface. See [eSDHC boot](#) for more detailed information.

19.1.2 eSPI transmission and reception process

As the eSPI is a character-oriented communication unit, the core is responsible for packing and unpacking the receive and transmit frames.

A frame consists of all of the characters transmitted or received during a completed eSPI transmission session, from the first character written to the eSPI transmit FIFO access register (SPITF) register to the last character transmitted with the total number as indicated in the command written to the [eSPI command register \(ESPI_SPCOM\)](#) register.

The core receives data by reading the eSPI receive FIFO access register (SPIRF) when the RNE ("not empty") bit in the eSPI event register (SPIE) is set.

The core transmits data by writing it into the SPITF register. After the core writes the final character to SPITF it waits for DON bit in the SPIE register to be set indicating frame was fully transmitted. It might then write a new command to SPCOM register.

The eSPI sets the TNF ("not full") bit in SPIE register whenever its transmit FIFO is not full.

The eSPI core handshake protocol can be implemented by using a polling or interrupt mechanism. When using a polling mechanism, the core reads the SPIE in a predefined frequency and acts according to the value of the SPIE bits. The polling frequency depends on the eSPI serial channel frequency. When using the interrupt mechanism, setting either the TNF (not full) or RNE (not empty) bits of the SPIE causes an interrupt to the core. The core then reads the SPIE and acts appropriately.

19.1.3 Modes of operation

The eSPI can be programmed to work in a single master environment. This section describes eSPI master operation in a single-master configuration.

In master mode, the eSPI sends a message to the slave peripheral, which sends back a simultaneous reply. A single master with multiple slaves uses up to four chip select signals to selectively enable slaves, as shown below.

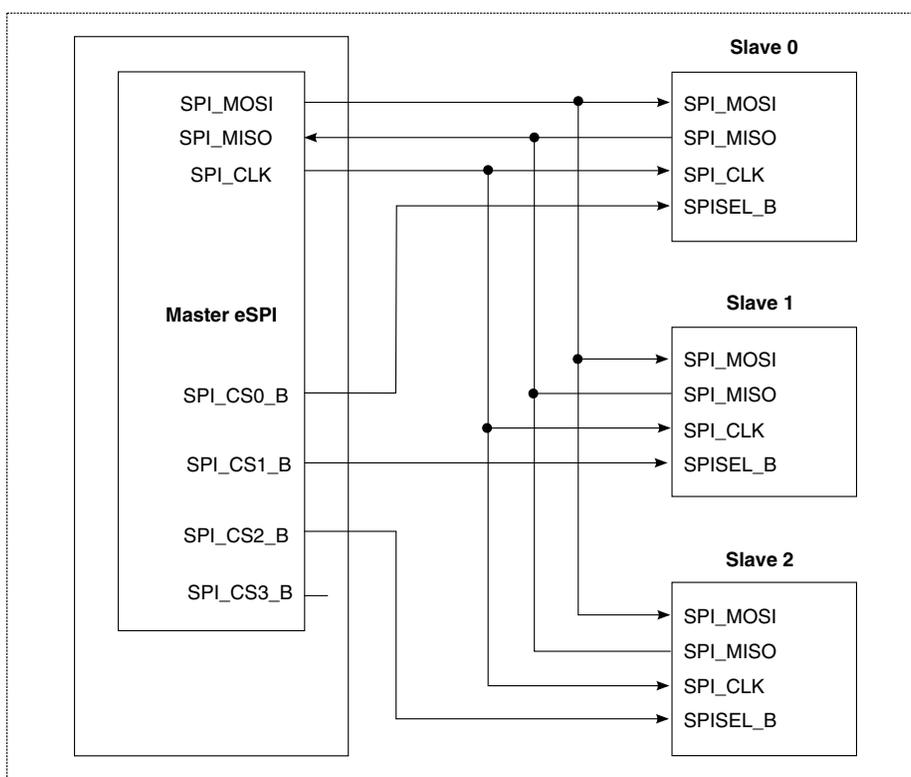


Figure 19-2. Single-master/multi-slave configuration

To start exchanging data, the core writes the data to be sent into the SPITF register. The eSPI then generates programmable clock pulses on SPI_CLK for each character. It shifts Tx data out on the "eSPI master-out slave-in" (SPI_MOSI) and Rx data in on the eSPI "master-in slave-out" (SPI_MISO) simultaneously. During the transmission process the core is responsible for supplying the data whenever the eSPI requests it to ensure smooth operation.

The maximum sustained data rate that the eSPI supports depends on the software latency. However, the eSPI can transfer a single character at very high rates— a maximum (up to system clock / 2) specified by the device hardware specifications (where system clock is defined as platform clock divided by 2). Gaps might be inserted between multiple frames.

19.2 External signal descriptions

The eSPI interface consists of transmit, receive, clock and chip selects.

19.2.1 Overview

Table 19-1. Signal properties

Name	Function
SPI_MISO	Master input slave output
SPI_MOSI	Master output slave input or second master input slave output for Winbond dual output read.
SPI_CLK	Output serial clock connected to the other SPI_CLK.
SPI_CS_B [0:3]	eSPI slave select outputs

19.2.2 ESPI detailed signal descriptions

Table 19-2. ESPI detailed signal descriptions

Signal	I/O	Description
SPI_MISO	I	Master input slave output
		State meaning Asserted-The data that has been received from the eSPI is high Negated-The data that has been received from the eSPI is low
		Timing Assertion-According to the SPI_CLK assertion/negation/in the middle of phase (depends on the SPMODEx configuration register) Negation-According to the SPI_CLK assertion/negation/in the middle of phase (depends on the SPMODEx configuration register)
SPI_MOSI	I/O	Master output slave input or second master input slave output for Winbond dual output read
		State meaning Asserted-The data that has been transmitted from/to the eSPI is high Negated-The data that has been transmitted from/to the eSPI is low
		Timing Assertion-According to the SPI_CLK assertion/negation/in the middle of phase (depends on the SPMODEx configuration register). Negation-According to the SPI_CLK assertion/negation/in the middle of phase (depends on the SPMODEx configuration register).
SPI_CLK	O	Serial clock out
		State meaning Assertion/negation-According to SPMODEx[PM,DIV16] register rate configuration
		Timing Assertion/negation-During frame reception/transmission
SPI_CS_B [0:3]	O	eSPI slave select outputs
		State meaning Asserted- Slave 0, 1, 2, 3 is selected and master controls transmission/reception Negated-idle state
		Timing Assertion-A predefined time before frame starts, during frame transmission/reception, a predefined time after frame ends Negation-When master is idle or controls another slave

The eSPI can be configured as a master in single master environment. The master eSPI generates the transfer clock SPI_CLK using the eSPI baud rate generator (BRG). The eSPI BRG takes as its input the platform clock divided by two.

SPI_CLK is a gated clock, active only during data transfers. Four combinations of SPI_CLK phase and polarity can be configured with the clock invert SPMODEx[CIx] and clock phase SPMODEx[CPx] register bits.

The eSPI master-in slave-out SPI_MISO signal acts as an input for master devices and as an output for slave devices. Conversely, the master-out slave-in SPI_MOSI signal is an output for master devices and an input for slave devices. However, it also acts as a second input for master devices and as a second output for slave devices when using Winbond dual output read.

SPI_CLK is the clock output signal that shifts received data in from SPI_MISO and transmitted data out to SPI_MOSI. eSPI masters must output a slave select signal to enable eSPI slave devices.

19.3 Enhanced serial peripheral interface (eSPI) memory map

The table below shows the memory mapped registers of the eSPI and their offsets. It lists the offset, name, and a cross-reference to the complete description of each register. Note that the full register address comprises CCSRBAR together with the SPI block base address and offset listed in the table below.

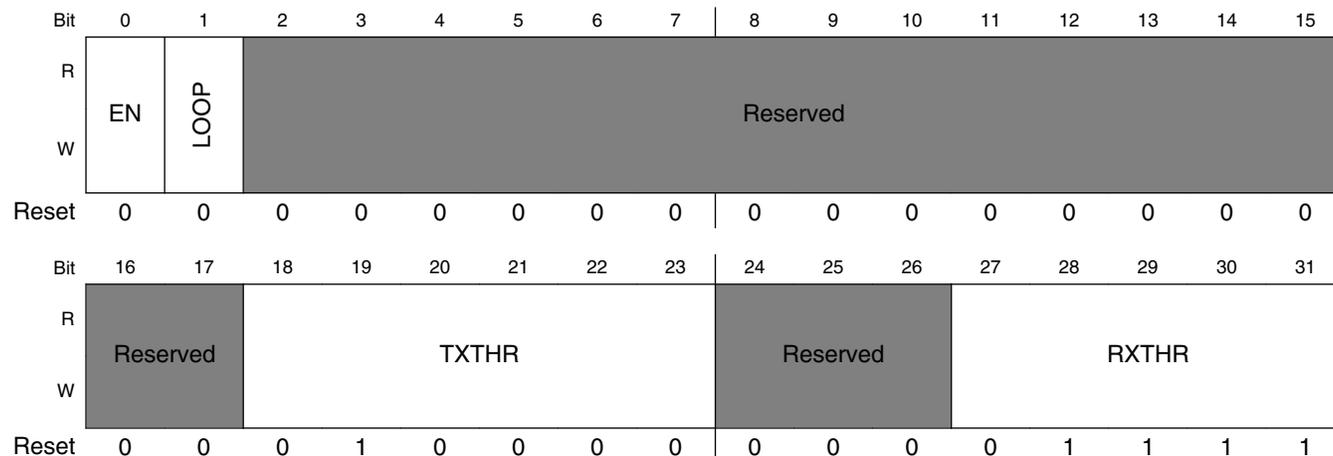
ESPI memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/ page
7000	eSPI mode register (ESPI_SPMODE)	32	R/W	0000_100Fh	19.3.1/1751
7004	eSPI event register (ESPI_SPIE)	32	R/W	0020_0000h	19.3.2/1752
7008	eSPI mask register (ESPI_SPIM)	32	R/W	0000_0000h	19.3.3/1754
700C	eSPI command register (ESPI_SPCOM)	32	W	0000_0000h	19.3.4/1756
7010	eSPI transmit FIFO access register (ESPI_SPITF)	32	W	0000_0000h	19.3.5/1758
7014	eSPI receive FIFO access register (ESPI_SPIRF)	32	R	0000_0000h	19.3.6/1760
7020	eSPI CS0 mode register (ESPI_SPMODE0)	32	R/W	0010_0000h	19.3.7/1761
7024	eSPI CS1 mode register (ESPI_SPMODE1)	32	R/W	0010_0000h	19.3.8/1763
7028	eSPI CS2 mode register (ESPI_SPMODE2)	32	R/W	0010_0000h	19.3.9/1765
702C	eSPI CS3 mode register (ESPI_SPMODE3)	32	R/W	0010_0000h	19.3.10/1767

19.3.1 eSPI mode register (ESPI_SPMODE)

The eSPI mode register (SPMODE) controls eSPI general operation mode.

Address: 7000h base + 0h offset = 7000h



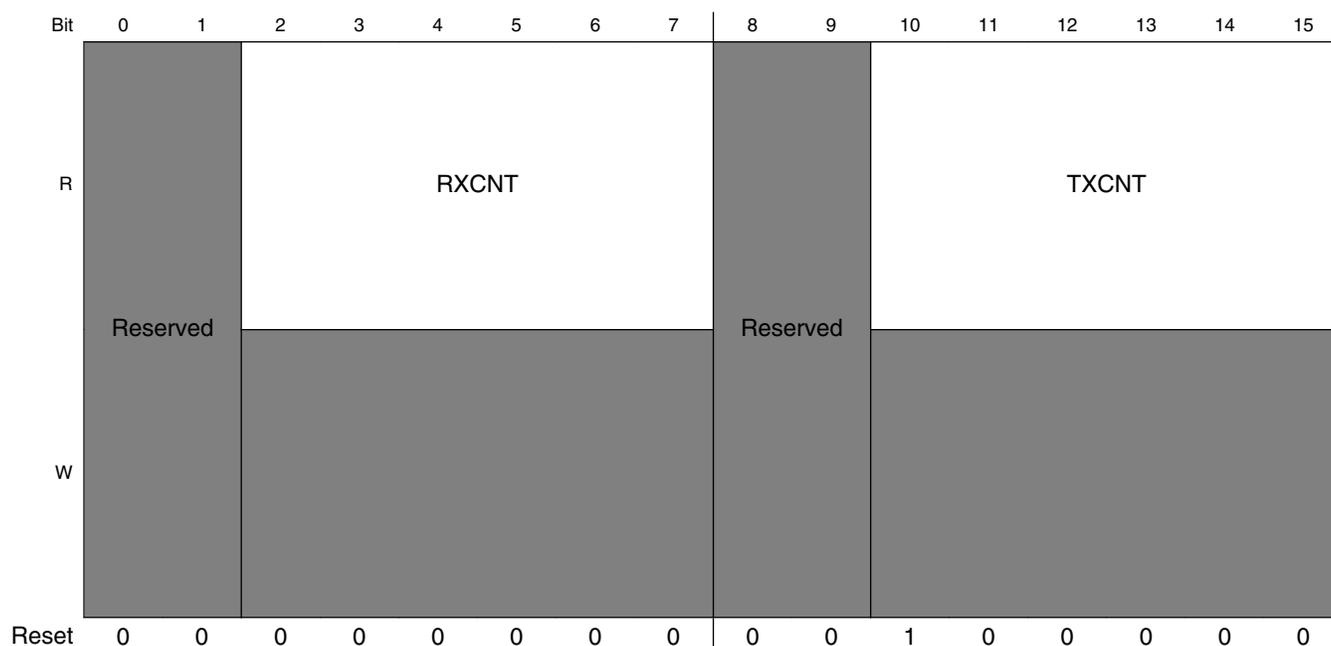
ESPI_SPMODE field descriptions

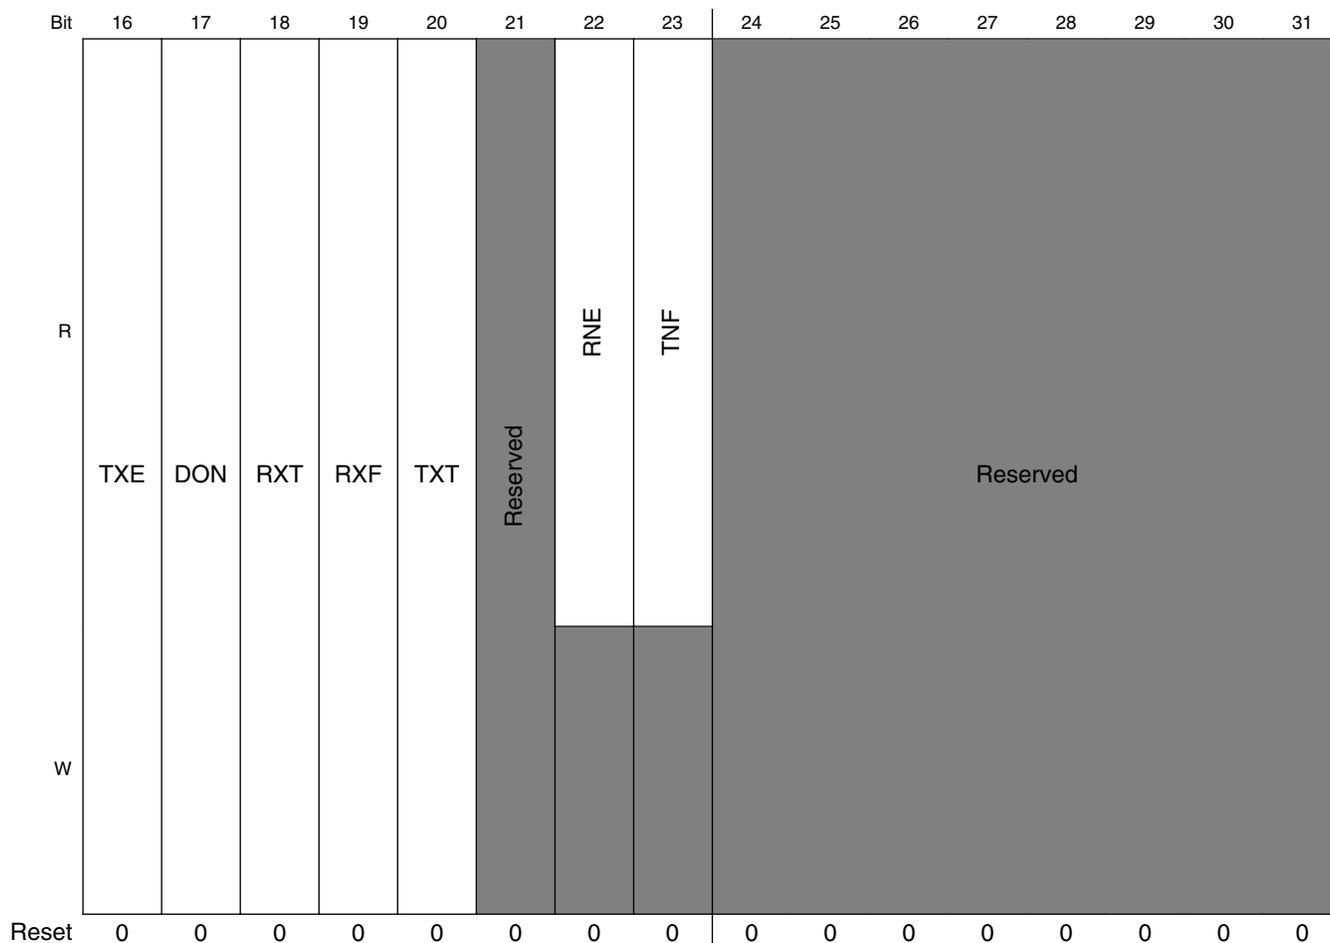
Field	Description
0 EN	Enable eSPI. Any bits in SPMODE must not change when EN is set. 0 The eSPI is disabled. The eSPI is in a idle state and consumes minimal power. The eSPI BRG is not functioning and the input clock is disabled. 1 The eSPI is enabled.
1 LOOP	Loop mode. Enables local loopback operation. 0 Normal operation. 1 Loopback mode. Used to test the eSPI controller internal functionality, the transmitter output is internally connected to the receiver input. The receiver and transmitter operate normally, except that received data is ignored.
2–17 -	This field is reserved. Reserved
18–23 TXTHR	Tx FIFO threshold-if Tx FIFO has less than TXTHR bytes, an interrupt can be issued to the core. Valid values: 1-32
24–26 -	This field is reserved. Reserved
27–31 RXTHR	Rx FIFO threshold-if Rx FIFO has more than RXTHR bytes, an interrupt can be issued to the core. Valid values: 0-31

19.3.2 eSPI event register (ESPI_SPIE)

The eSPI event register (SPIE) generates interrupts and reports events recognized by the eSPI. When an event is recognized, the eSPI sets the corresponding SPIE bit. Clear SPIE bits by writing a 1—writing 0 has no effect. Setting a bit in the eSPI mask register (SPIM) enables interrupt and clearing a bit masks the corresponding interrupt. Unmasked SPIE bits must be cleared before the core clears internal interrupt requests. Bits RNE and TNF are status bits. Fields RXCNT and TXCNT hold Rx and Tx FIFOs' statuses. They are not cleared as a result of writing to SPIE.

Address: 7000h base + 4h offset = 7004h





ESPI_SPIE field descriptions

Field	Description
0-1 -	This field is reserved. Reserved, should be cleared.
2-7 RXCNT	The current number of full Rx FIFO bytes NOTE: For character lengths of 9 to 16 bits-each character occupies 2 bytes in Rx/Tx FIFO.
8-9 -	This field is reserved. Reserved, should be cleared.
10-15 TXCNT	The current number of free Tx FIFO bytes NOTE: For character lengths of 9 to 16 bits-each character occupies 2 bytes in Rx/Tx FIFO
16 TXE	Tx FIFO is empty
17 DON	Last character was transmitted. The last character was transmitted and a new command can be written for the next frame
18 RXT	Rx FIFO has more than RXTHR bytes, that is, at least RXTHR + 1 bytes
19 RXF	Rx FIFO is full

Table continues on the next page...

ESPI_SPIE field descriptions (continued)

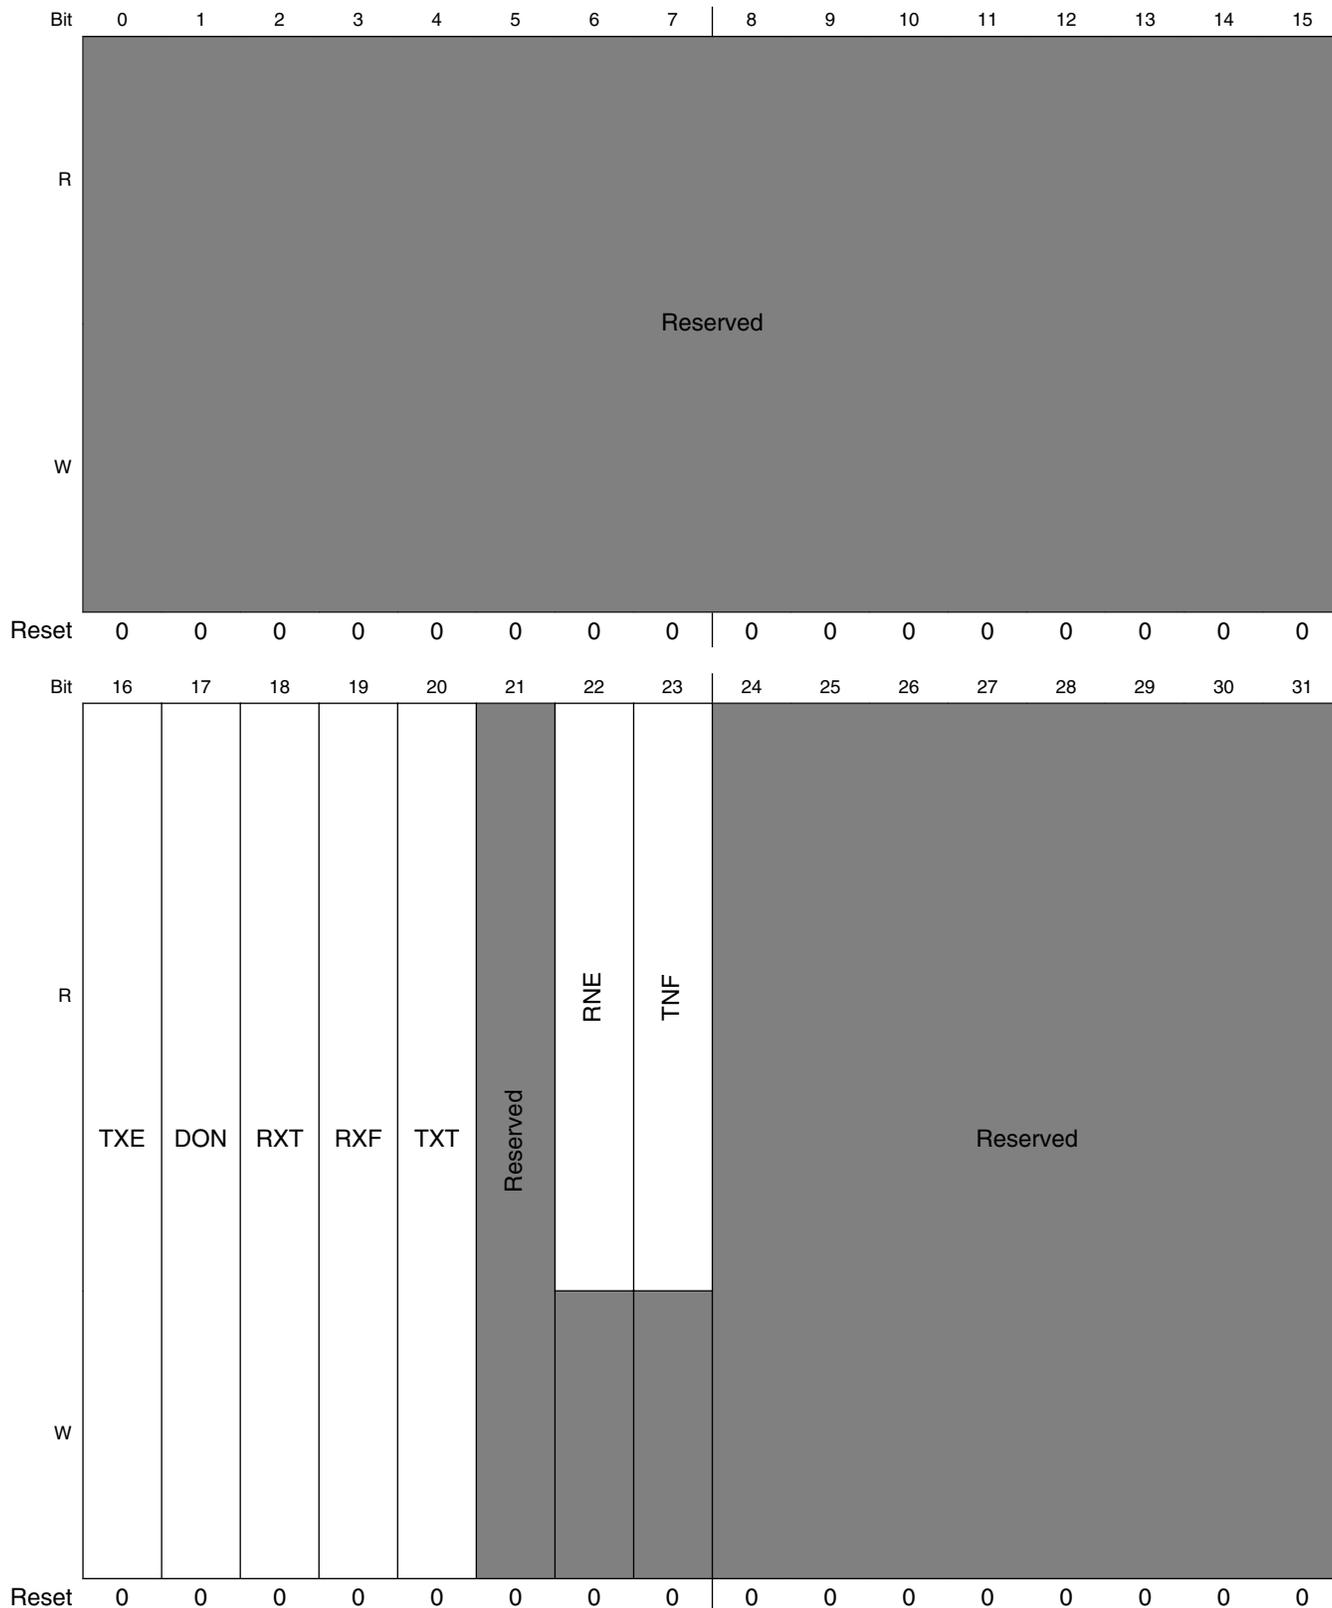
Field	Description
20 TXT	Tx FIFO has less than TXTHR bytes, that is, at most TXTHR - 1 bytes
21 -	This field is reserved. Reserved, should be cleared.
22 RNE	Not empty. Indicates that the Rx FIFO register contains a received character. 0 The Rx FIFO is empty 1 The Rx FIFO has a received character. The core can read the content of Rx FIFO through SPIRF.
23 TNF	Tx FIFO not full. 0 The transmitter FIFO is full. 1 The transmitter FIFO is not full.
24–31 -	This field is reserved. Reserved, should be cleared.

19.3.3 eSPI mask register (ESPI_SPIM)

The eSPI mask register (SPIM) enables/masks interrupts for events recognized by the eSPI. When an event is recognized, the eSPI sets the corresponding SPIE bit. Setting a bit in the eSPI mask register (SPIM) enables and clearing a bit masks the corresponding interrupt. Unmasked SPIE bits must be cleared before the core clears internal interrupt requests.

Bits RNE and TNF in SPIM are status bits. They are not cleared as a result of writing to SPIM.

Address: 7000h base + 8h offset = 7008h



ESPI_SPIM field descriptions

Field	Description
0–15 -	This field is reserved. Reserved, should be cleared.
16 TXE	Tx FIFO empty interrupt mask 0 TXE event will not cause eSPI Interrupt 1 TXE event will cause eSPI Interrupt
17 DON	Last character transmitted mask 0 DON event will not cause eSPI Interrupt 1 DON event will cause eSPI Interrupt
18 RXT	Rx threshold interrupt mask 0 RXT event will not cause eSPI Interrupt 1 RXT event will cause eSPI Interrupt
19 RXF	Rx FIFO full interrupt mask 0 RXF event will not cause eSPI Interrupt 1 RXF event will cause eSPI Interrupt
20 TXT	Tx threshold interrupt mask 0 TXT event will not cause eSPI Interrupt 1 TXT event will cause eSPI Interrupt
21 -	This field is reserved. Reserved, should be cleared.
22 RNE	Rx not empty interrupt mask 0 Not Empty event will not cause eSPI Interrupt 1 Not Empty event will cause eSPI Interrupt
23 TNF	Tx not full interrupt mask 0 Not full event will not cause eSPI Interrupt 1 Not full event will cause eSPI Interrupt
24–31 -	This field is reserved. Reserved, should be cleared.

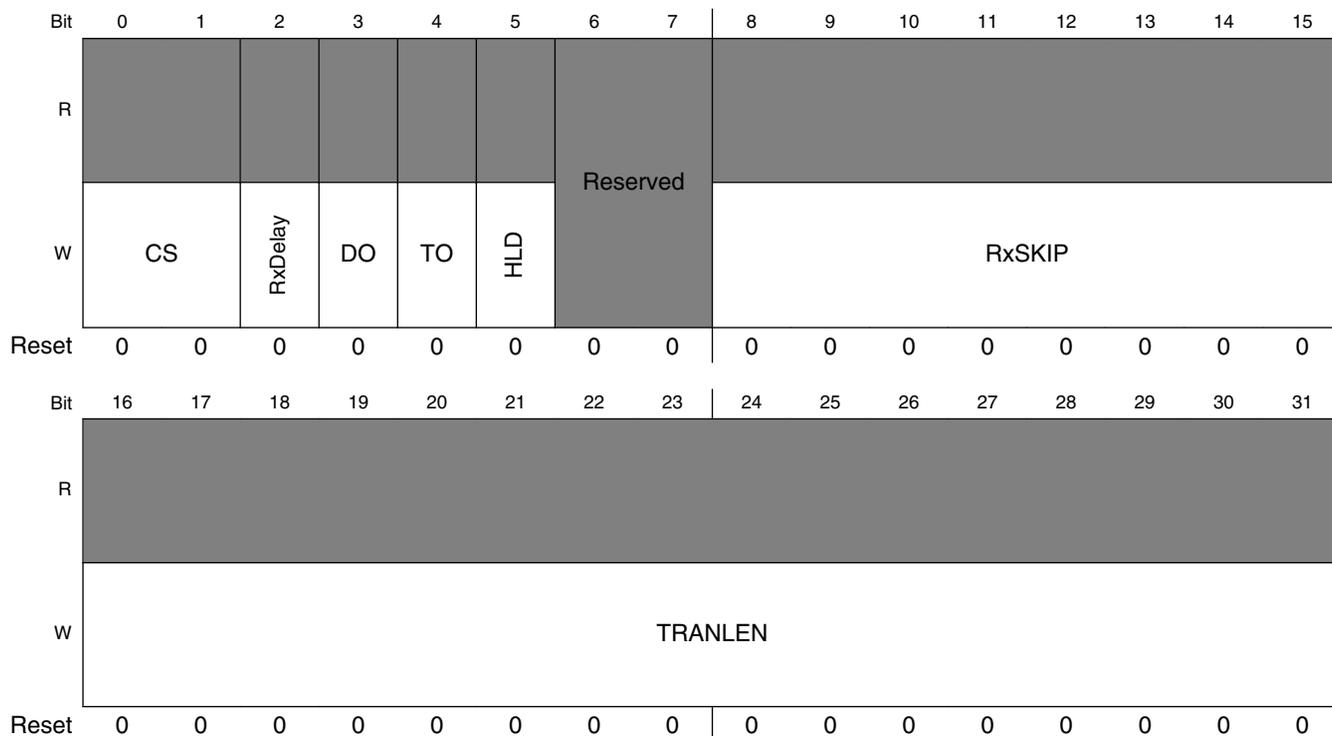
19.3.4 eSPI command register (ESPI_SPCOM)

The eSPI command register (SPCOM) is used by the host to supply information on the new frame.

After SPCOM has been written to initiate the first transaction after startup, commands can be executed only after SPIE[DON] is set. Otherwise they are ignored.

A transaction can be full duplex (regular eSPI) or half duplex. Half duplex can be used for example for write accesses to a flash (only transmit) or for a read access from a flash (first part is transmit without receive, while the second part is receive without transmit).

Address: 7000h base + Ch offset = 700Ch



ESPI_SPCOM field descriptions

Field	Description
0–1 CS	Chip select-chip select for which transaction is destined 00 SPI_CS0_B 01 SPI_CS1_B 10 SPI_CS2_B 11 SPI_CS3_B
2 RxDelay	RxDelay 0 Normal eSPI operation 1 Rx data should be sampled a bit later than regular eSPI (used for full cycle operation such as with Atmel RapidS devices)
3 DO	This mode is useful only for character lengths of 4,6,8. DO and RapidS should not be set simultaneously. 0 Normal eSPI operation 1 Winbond dual output read-when eSPI master reads data 2 data bits are available (on MISO and MOSI)
4 TO	Transmit only 1 No reception is done for the frame (useful for write transactions) 0 Normal operation
5 HLD	HLD 0 Normal operation 1 Mask first generated SPI_CLK. Should be used only for RapidS mode0

Table continues on the next page...

ESPI_SPCOM field descriptions (continued)

Field	Description
6-7 -	This field is reserved. Reserved, should be cleared.
8-15 RxSKIP	If (RXSKIP \neq 0)-Number of characters skipped for reception from frame start. Non-zero values of RxSKIP force the eSPI to half-duplex mode, and therefore this causes TRANLEN-RxSKIP characters to be skipped for transmission. RXSKIP is useful for reads of SPI Flash memories where the first valid read data is received several characters after the transmission begins (after the eSPI has transmitted an instruction opcode and address). NOTE: If TO = 1, RxSKIP must be set to 0. NOTE: If RXSKIP = 0 and TO = 0, the eSPI changes to full duplex mode.
16-31 TRANLEN	Transaction length - (number of characters in the frame - 1)

19.3.5 eSPI transmit FIFO access register (ESPI_SPITF)

The 32-bit write-only eSPI transmit FIFO access register (SPITF) holds the characters to be written to the transmit FIFO. The number of bits in each character is specified by SPMODEx[LENx]. Each time SPIE[TNF] is set, the core can write more data to the SPITF register, if there is no error indication in the SPIE.

For character lengths of 4 to 8 bits, SPITF contains up to 4 characters (unless end of frame). The lsbs are in bits 7, 15, 23, and 31 of SPITF.

For character lengths of 9 to 16 bits, SPITF contains up to 2 characters (unless end of frame). For 16 bits with SPMODEx[REVx] = 1, the lsb is in bits 15 and 31 of SPITF. For other options, lsbs are in bits 7 and 23 while msbs are in bits (23-LENx) and (39-LENx) of SPITF.

For example: REV = 0, LEN = 10 (0xA), SPITF[0-15] = 0xFB05-bitstream is: (lsb first) 1101111101 (msb last).

NOTE

The user must write N bytes of SPITF ($1 \leq N \leq 4$) that do not exceed the number of free bytes in the transmit FIFO. It is valid for the user to write only 1 or 2 bytes of SPITF (at offset 0x010) if the user wishes to write fewer characters than the maximum supported by SPITF for the particular character length in use.

The following figures show examples of the contents of SPITF with various parameters set.

SPITF Example - SPMODE_x[REV_x]=0, SPMODE_x[LEN_x]=3, LSB Sent First:

	0	3	4	5	6	7	8	11	12	13	14	15	16	19	20	21	22	23	24	27	28	29	30	31
R																								
W	-	MSB 0	Data 0	LSB0	-	MSB 1	Data 1	LSB1	-	MSB 2	Data 2	LSB2	-	MSB 3	Data 3	LSB3								

SPITF Example - SPMODE_x[REV_x]=x, SPMODE_x[LEN_x]=7:

	0	1	6	7	8	9	14	15	16	17	22	23	24	30	31
R															
W	MSB0	Data 0	LSB 0	MSB1	Data 1	LSB 1	MSB2	Data 2	LSB 2	MSB3	Data 3	LSB 3			

SPITF Example - SPMODE_x[REV_x]=0, SPMODE_x[LEN_x]=10, LSB Sent First:

	0	6	7	8	12	13	14	15	16	22	23	24	28	29	30	31
R																
W	Data 0 LS Byte	LSB0	-	MSB 0	Data 0 MS Byte	Data 1 LS Byte	LSB1	-	MSB 1	Data 1 MS Byte						

SPITF Example - SPMODE_x[REV_x]=1, SPMODE_x[LEN_x]=15, MSB Sent First:

	0	1	7	8	14	15	16	17	23	24	30	31
R												
W	MSB0	Data 0 MS Byte	Data 0 LS Byte	LS B0	MSB1	Data 1 MS Byte	Data 1 LS Byte	LS B1				

Address: 7000h base + 10h offset = 7010h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R																																	
W	DATA																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ESPI_SPITF field descriptions

Field	Description
0–31 DATA	Varies as parameters set

19.3.6 eSPI receive FIFO access register (ESPI_SPIRF)

The 32-bit read-only eSPI receive data register (SPIRF) is used to hold characters read from the receive FIFO. Each time SPIE[RNE] is set, the core can read the SPIRF register.

For character lengths of 4 to 8 bits, SPIRF contains up to 4 characters. The msbs are in bits 0, 8, 16, and 24. For character lengths of 9 to 16 bits, SPIRF contains up to 2 characters. The msbs are in bits 0 and 16. SPMODEx[REVx] does not affect the msb or lsb bit positions when reading the SPIRF register.

The user must read N bytes of SPIRF ($1 \leq N \leq 4$) that do not exceed the amount of data in the receive FIFO. The user can read less bytes than the amount of data in the receive FIFO. For example, a 1-byte read of SPIRF when configured for 8-bit characters with 4 characters of data in the receive FIFO results in the 3 unread characters shuffling down to the lower 24 bits of SPIRF in preparation for the following SPIRF read.

SPIRF Example - SPMODEx[LENx]=3

	0	1	2	3	4	7	8	9	10	11	12	15	16	17	18	19	20	23	24	25	26	27	28	31
R	MSB0	Data 0	LS B0	-	MSB1	Data 1	LSB 1	-	MSB2	Data 2	LSB 2	-	MSB3	Data 3	LSB 3	-								
W																								

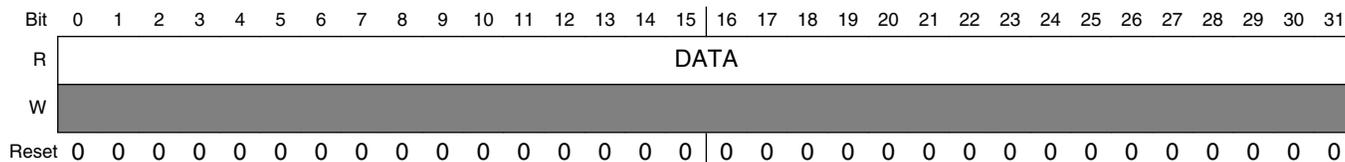
SPIRF Example - SPMODEx[LENx]=10:

	0	1	7	8	9	10	11	15	16	17	23	24	25	26	27	31
R	MSB 0	Data 0 MS Byte	Data 0 LS Byte	LSB0	-	MSB 1	Data 1 MS Byte	Data 1 LS Byte	LSB1	-						
W																

SPIRF Example - SPMODEx[LENx]=15:

	0	1	7	8	14	15	16	17	23	24	30	31
R	MSB0	Data 0 MS Byte	Data 0 LS Byte	LSB 0	MSB1	Data 1 MS Byte	Data 1 LS Byte	LSB 1				
W												

Address: 7000h base + 14h offset = 7014h



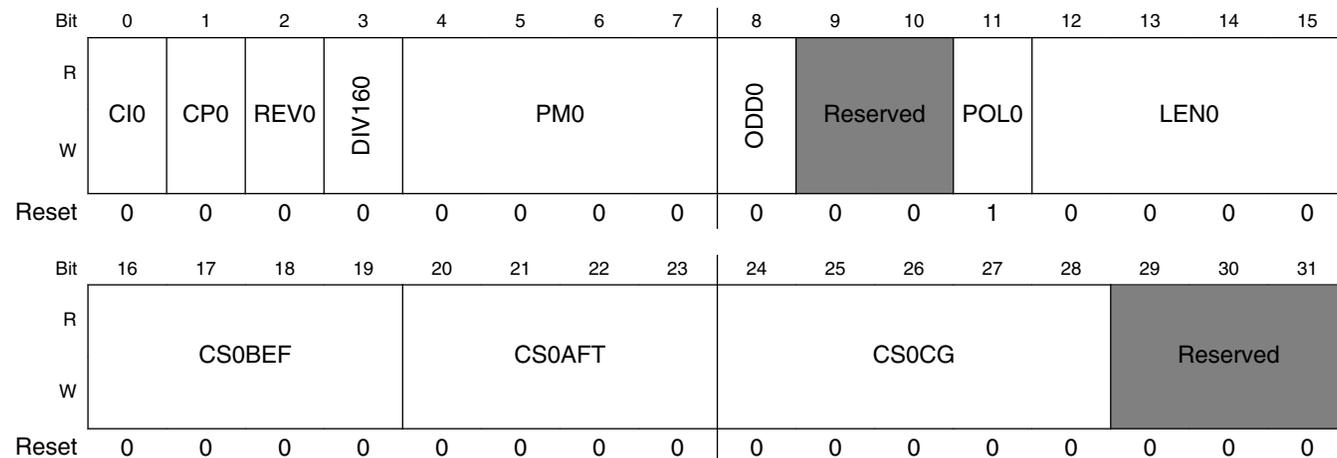
ESPI_SPIRF field descriptions

Field	Description
0–31 DATA	Varies as parameters set

19.3.7 eSPI CS0 mode register (ESPI_SPMODE0)

The eSPI CS0 mode register (SPMODE0) controls eSPI master operation with chip select 0.

Address: 7000h base + 20h offset = 7020h



ESPI_SPMODE0 field descriptions

Field	Description
0 CI0	Clock invert. Inverts eSPI clock polarity. See Figure 19-13 and Figure 19-14 for more information 0 The inactive state of SPI_CLK is low. 1 The inactive state of SPI_CLK is high.
1 CP0	Clock phase. Selects the transfer format. See Figure 19-13 and Figure 19-14 for more information. 0 SPI_CLK starts toggling at the middle of the data transfer. 1 SPI_CLK starts toggling at the beginning of the data transfer.

Table continues on the next page...

ESPI_SPMODE0 field descriptions (continued)

Field	Description
2 REV0	Reverse data mode. Determines the receive and transmit character bit order. 0 lsb of the character sent and received first 1 msb of the character sent and received first-for 8/16 bits data character only
3 DIV160	Divide by 16. Selects the clock source for the eSPI baud rate generator (eSPI BRG) when configured as an eSPI master. NOTE: System clock as used here is defined to be platform clock divided by 2. 0 System clock is the input to the eSPI BRG. 1 System clock/16 is the input to the eSPI BRG.
4–7 PM0	Prescale modulus select. Specifies the divide ratio of the prescale divider in the eSPI clock generator. The eSPI baud rate generator clock source (either system clock or system clock divided by 16, depending on DIV16 bit) is divided by $2 \times ([PM] + 1)$, a range from 2 to 32. For example, if the prescale modulus is set to PM=0x0011 and DIV16 is set, the SPI_CLK/system clock rate will be $16 \times (2 \times (0x0011 + 1)) = 128$ NOTE: System clock as used here is defined to be platform clock divided by 2
8 ODD0	0 Even division - $2 \times (PM + 1) \times (15 \times DIV16 + 1)$ - 50% duty cycle 1 Odd division - $(2 \times PM + 1) \times (15 \times DIV16 + 1)$ (except for PM = 0 where it divides by $2 \times (7 \times DIV16 + 1)$); duty cycle is $(PM + 1) \div (2 \times PM + 1)$ for DIV16 = 0; duty cycle is 50% for DIV16 = 1.
9–10 -	This field is reserved. Reserved
11 POL0	CS0 Polarity. 1 Asserted Low, Negated High 0 Asserted High, Negated Low.
12–15 LEN0	Character length in bits per character. Supports a range from 1-bit to 16-bit data characters.
16–19 CS0BEF	CS assertion time in bits before frame start (that is, before clock toggles) Example: CS0BEF = 0010 inserts 2 bits time gap between CS0 assertion to clock toggle
20–23 CS0AFT	CS assertion time in bits after frame end (that is, after clock finishes toggling) Example: CS0AFT = 0010 inserts 2 bits time gap between clock stop to CS0 negation
24–28 CS0CG	Clock gap insert gaps between transmitted frames according to this size (during this time, chip select is negated). Chip select is negated minimum time of 1 bit time. Example: CS0CG = 00101 inserts $5 + 1 = 6$ bits time gap between every two consecutive frames
29–31 -	This field is reserved. Reserved

19.3.8 eSPI CS1 mode register (ESPI_SPMODE1)

The eSPI CS1 mode register (SPMODE1) controls eSPI master operation with chip select 1.

Address: 7000h base + 24h offset = 7024h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	CI1	CP1	REV1	DIV161	PM1				ODD1	Reserved		POL1	LEN1			
Reset	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R													Reserved			
W	CS1BEF				CS1AFT				CS1CG				Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ESPI_SPMODE1 field descriptions

Field	Description
0 CI1	<p>Clock invert. Inverts eSPI clock polarity. See Figure 19-13 and Figure 19-14 for more information</p> <p>0 The inactive state of SPI_CLK is low. 1 The inactive state of SPI_CLK is high.</p>
1 CP1	<p>Clock phase. Selects the transfer format. See Figure 19-13 and Figure 19-14 for more information.</p> <p>0 SPI_CLK starts toggling at the middle of the data transfer. 1 SPI_CLK starts toggling at the beginning of the data transfer.</p>
2 REV1	<p>Reverse data mode. Determines the receive and transmit character bit order.</p> <p>0 lsb of the character sent and received first 1 msb of the character sent and received first-for 8/16 bits data character only</p>
3 DIV161	<p>Divide by 16. Selects the clock source for the eSPI baud rate generator (eSPI BRG) when configured as an eSPI master.</p> <p>NOTE: System clock as used here is defined to be platform clock divided by 2.</p> <p>0 System clock is the input to the eSPI BRG. 1 System clock/16 is the input to the eSPI BRG.</p>
4–7 PM1	<p>Prescale modulus select. Specifies the divide ratio of the prescale divider in the eSPI clock generator. The eSPI baud rate generator clock source (either system clock or system clock divided by 16, depending on DIV16 bit) is divided by 2 x ([PM] + 1), a range from 2 to 32. For example, if the prescale modulus is set to PM = 0x0011 and DIV16 is set, the SPI_CLK/system clock rate will be 16 x (2 x (0x0011 + 1)) = 128</p>

Table continues on the next page...

ESPI_SPMODE1 field descriptions (continued)

Field	Description
	NOTE: System clock as used here is defined to be platform clock divided by 2.
8 ODD1	0 Even division- $2 \times (PM + 1) \times (15 \times DIV16 + 1)$ - 50% duty cycle 1 Odd division- $(2 \times PM + 1) \times (15 \times DIV16 + 1)$ (except for $PM = 0$ where it divides by $2 \times (7 \times DIV16 + 1)$); duty cycle is $(PM + 1) \div (2 \times PM + 1)$ for $DIV16 = 0$; duty cycle is 50% for $DIV16 = 1$.
9–10 -	This field is reserved. Reserved
11 POL1	CS1 Polarity. 1 Asserted Low, Negated High 0 Asserted High, Negated Low.
12–15 LEN1	Character length in bits per character. Supports a range from 1-bit to 16-bit data characters.
16–19 CS1BEF	CS assertion time in bits before frame start (that is, before clock toggles) Example: CS1BEF = 0010 inserts 2 bits time gap between CS1 assertion to clock toggle.
20–23 CS1AFT	CS assertion time in bits after frame end (that is, after clock finishes toggling) Example: CS1AFT = 0010 inserts 2 bits time gap between clock stop to CS1 negation.
24–28 CS1CG	Clock Gap insert gaps between transmitted frames according to this size (during this time, chip select is negated). Chip select is negated minimum time of 1 bit time. Example: CS1CG = 00101 inserts $5 + 1 = 6$ bits time gap between every two consecutive frames.
29–31 -	This field is reserved. Reserved

19.3.9 eSPI CS2 mode register (ESPI_SPMODE2)

The eSPI CS2 mode register (SPMODE2) controls the eSPI master operation with chip select 2.

Address: 7000h base + 28h offset = 7028h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	CI2	CP2	REV2	DIV162				PM2	ODD2	Reserved		POL2				LEN2
Reset	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																
W																Reserved
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ESPI_SPMODE2 field descriptions

Field	Description
0 CI2	<p>Clock invert. Inverts eSPI clock polarity. See Figure 19-13 and Figure 19-14 for more information</p> <p>0 The inactive state of SPI_CLK is low. 1 The inactive state of SPI_CLK is high.</p>
1 CP2	<p>Clock phase. Selects the transfer format. See Figure 19-13 and Figure 19-14 for more information.</p> <p>0 SPI_CLK starts toggling at the middle of the data transfer. 1 SPI_CLK starts toggling at the beginning of the data transfer.</p>
2 REV2	<p>Reverse data mode. Determines the receive and transmit character bit order.</p> <p>0 lsb of the character sent and received first 1 msb of the character sent and received first-for 8/16 bits data character only</p>
3 DIV162	<p>Divide by 16. Selects the clock source for the eSPI baud rate generator (eSPI BRG) when configured as an eSPI master.</p> <p>NOTE: System clock as used here is defined to be platform clock divided by 2.</p> <p>0 System clock is the input to the eSPI BRG. 1 System clock/16 is the input to the eSPI BRG.</p>
4–7 PM2	<p>Prescale modulus select. Specifies the divide ratio of the prescale divider in the eSPI clock generator. The eSPI baud rate generator clock source (either system clock or system clock divided by 16, depending on DIV16 bit) is divided by 2 x ([PM] + 1), a range from 2 to 32. For example, if the prescale modulus is set to PM = 0x0011 and DIV16 is set, the system clock/SPI_CLK ratio will be 16 x (2 x (0x0011 + 1)) = 128</p>

Table continues on the next page...

ESPI_SPMODE2 field descriptions (continued)

Field	Description
	NOTE: System clock as used here is defined to be platform clock divided by 2.
8 ODD2	0 Even division- $2 \times (PM + 1) \times (15 \times DIV16 + 1)$ - 50% duty cycle 1 Odd division- $(2 \times PM + 1) \times (15 \times DIV16 + 1)$ (except for $PM = 0$ where it divides by $2 \times (7 \times DIV16 + 1)$); duty cycle is $(PM + 1) \div (2 \times PM + 1)$ for $DIV16 = 0$; duty cycle is 50% for $DIV16 = 1$
9–10 -	This field is reserved. Reserved
11 POL2	CS2 Polarity. 1 Asserted Low, Negated High 0 Asserted High, Negated Low.
12–15 LEN2	Character length in bits per character. Supports a range from 1-bit to 16-bit data characters.
16–19 CS2BEF	CS assertion time in bits before frame start (that is, before clock toggles) Example: CS2BEF = 0010 inserts 2 bits time gap between CS2 assertion to clock toggle
20–23 CS2AFT	CS assertion time in bits after frame end (that is, after clock finishes toggling) Example: CS2AFT = 0010 inserts 2 bits time gap between clock stop to CS2 negation
24–28 CS2CG	Clock Gap insert gaps between transmitted frames according to this size (during this time, chip select is negated). Chip select is negated minimum time of 1 bit time. Example: CS1CG = 00101 inserts $5 + 1 = 6$ bits time gap between every two consecutive frames
29–31 -	This field is reserved. Reserved

19.3.10 eSPI CS3 mode register (ESPI_SPMODE3)

The eSPI CS3 mode register (SPMODE3) controls the eSPI master operation with chip select 3.

Address: 7000h base + 2Ch offset = 702Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	CI3	CP3	REV3	DIV163	PM3				ODD3	Reserved		POL3	LEN3			
Reset	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R													Reserved			
W	CS3BEF				CS3AFT				CS3CG				Reserved			
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

ESPI_SPMODE3 field descriptions

Field	Description
0 CI3	<p>Clock invert. Inverts eSPI clock polarity. See Figure 19-13 and Figure 19-14 for more information</p> <p>0 The inactive state of SPI_CLK is low. 1 The inactive state of SPI_CLK is high.</p>
1 CP3	<p>Clock phase. Selects the transfer format. See Figure 19-13 and Figure 19-14 for more information.</p> <p>0 SPI_CLK starts toggling at the middle of the data transfer. 1 SPI_CLK starts toggling at the beginning of the data transfer.</p>
2 REV3	<p>Reverse data mode. Determines the receive and transmit character bit order.</p> <p>0 lsb of the character sent and received first 1 msb of the character sent and received first - for 8/16 bits data character only</p>
3 DIV163	<p>Divide by 16. Selects the clock source for the eSPI baud rate generator (eSPI BRG) when configured as an eSPI master.</p> <p>NOTE: System clock as used here is defined to be platform clock divided by 2</p> <p>0 System clock is the input to the eSPI BRG. 1 System clock/16 is the input to the eSPI BRG.</p>
4–7 PM3	<p>Prescale modulus select. Specifies the divide ratio of the prescale divider in the eSPI clock generator. The eSPI baud rate generator clock source (either system clock or system clock divided by 16, depending on DIV16 bit) is divided by 2 x ([PM] + 1), a range from 2 to 32. For example, if the prescale modulus is set to PM = 0x0011 and DIV16 is set, the SPI_CLK/system clock rate will be 16 x (2 x (0x0011 + 1)) = 128</p>

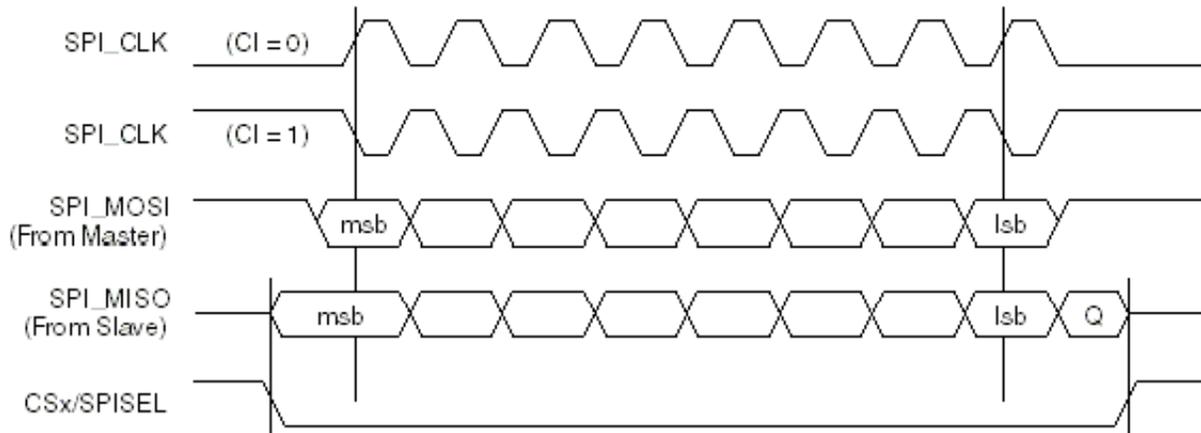
Table continues on the next page...

ESPI_SPMODE3 field descriptions (continued)

Field	Description
	NOTE: System clock as used here is defined to be platform clock divided by 2
8 ODD3	0 Even division- $2 \times (PM + 1) \times (15 \times DIV16 + 1)$ - 50% duty cycle 1 Odd division- $(2 \times PM + 1) \times (15 \times DIV16 + 1)$ (except for $PM = 0$ where it divides by $2 \times (7 \times DIV16 + 1)$); duty cycle is $(PM + 1) \div (2 \times PM + 1)$ for $DIV16 = 0$; duty cycle is 50% for $DIV16 = 1$
9–10 -	This field is reserved. Reserved
11 POL3	CS Polarity. 1 Asserted Low, Negated High 0 Asserted High, Negated Low.
12–15 LEN3	Character length in bits per character. Supports a range from 1-bit to 16-bit data characters.
16–19 CS3BEF	CS assertion time in bits before frame start (that is, before clock toggles) Example: CS3BEF = 0010 inserts 2 bits time gap between CS3 assertion to clock toggle
20–23 CS3AFT	CS assertion time in bits after frame end (that is, after clock finishes toggling) Example: CS3AFT = 0010 inserts 2 bits time gap between clock stop to CS3 negation
24–28 CS3CG	Clock Gap insert gaps between transmitted frames according to this size (during this time, chip select is negated). Chip select is negated minimum time of 1 bit time. Example: CS1CG = 00101 inserts $5 + 1 = 6$ bits time gap between every two consecutive frames
29–31 -	This field is reserved. Reserved

19.4 eSPI transfer formats

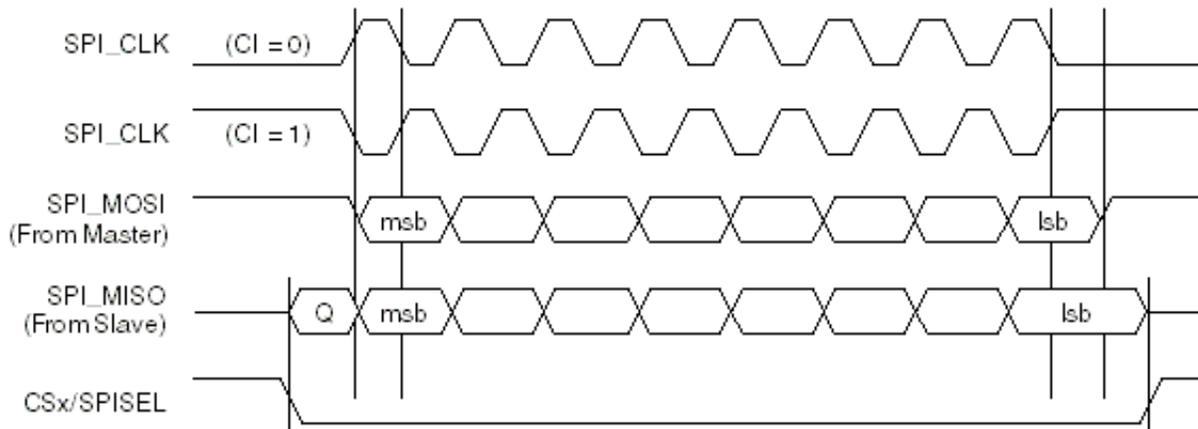
Figure below shows the eSPI transfer format in which SPI_CLK starts toggling in the middle of the transfer ($SPMODEx[CPx] = 0$).



NOTE: Q = Undefined Signal.

Figure 19-13. eSPI transfer format with $SPMODEx[CPx] = 0$

Figure below shows the eSPI transfer format in which SPI_CLK starts toggling at the beginning of the transfer ($SPMODEx[CPx] = 1$).



NOTE: Q = Undefined Signal.

Figure 19-14. eSPI transfer format with $SPMODEx[CPx] = 1$

19.5 CI and CP values for various eSPI devices

1. Regular devices-eSPI mode0-CI = CP = 0
2. Regular devices-eSPI mode3-CI = CP = 1

For Winbond devices DO should also be set for dual output read command.

3. RapidS mode0-CI = 0, CP = 1, HLD = 1
4. RapidS mode3-CI = 1, CP = 0

19.6 eSPI programming examples

This section provides eSPI programming examples for 24-bit address memory and 16-bit address memory.

19.6.1 24-bit address example

The following sequence initializes the eSPI to read 36 bytes from 24-bit address memory, start address = 0x00_0040:

1. Configure a parallel I/O signal to operate as the eSPI CS1 output signal.
2. Write 0xFFFF_FFFF to SPIE to clear any previous events. Configure SPIM to enable all desired eSPI interrupts.
3. Configure SPMODE = 0x8000_100F to enable normal operation, eSPI enabled.
4. Configure SPMODE1 = 0x2417_1108-REV1 = 1, PM1 = 4 (divide eSPI input clock by 10), LEN1 = 7, POL1 = 1, CS1BEF = CS1AFT = CS1CG = 1.
5. Configure SPITF = 0x0300_0040-0x03 is read opcode while 0x00_0040 is the 24-bit start address.
6. Configure SPCOM = 0x0004_0027 so 4 bytes are skipped (1 for opcode and 3 for 24-bit address), TRANLEN = 36 + 4 - 1.

19.6.2 16-bit address example

The following sequence initializes the eSPI to read 36 bytes from 16-bit address memory, start address = 0x0040:

1. Configure a parallel I/O signal to operate as the eSPI CS1 output signal.
2. Write 0xFFFF_FFFF to SPIE to clear any previous events. Configure SPIM to enable all desired eSPI interrupts.
3. Configure SPMODE = 0x8000_100F to enable normal operation, eSPI enabled.
4. Configure SPMODE1 = 0x2417_1108-REV1 = 1, PM1 = 4 (divide eSPI input clock by 10), LEN1 = 7, POL1 = 1, CS1BEF = CS1AFT = CS1CG = 1.
5. Configure SPITF = 0x0300_40xx (xx is don't care)-0x03 is read opcode while 0x0040 is the 16-bit start address.
6. Configure SPCOM = 0x0003_0026 so 3 bytes are skipped (1 for opcode and 2 for 16-bit address), TRANLEN = 36 + 3 - 1.

Chapter 20

General Purpose I/O (GPIO)

This chapter describes the general-purpose I/O module, including pin descriptions, register settings, and interrupt capabilities.

20.1 Introduction

This chapter describes the general-purpose I/O module, including pin descriptions, register settings, and interrupt capabilities. [Figure 20-1](#) shows the block diagram of the GPIO module.

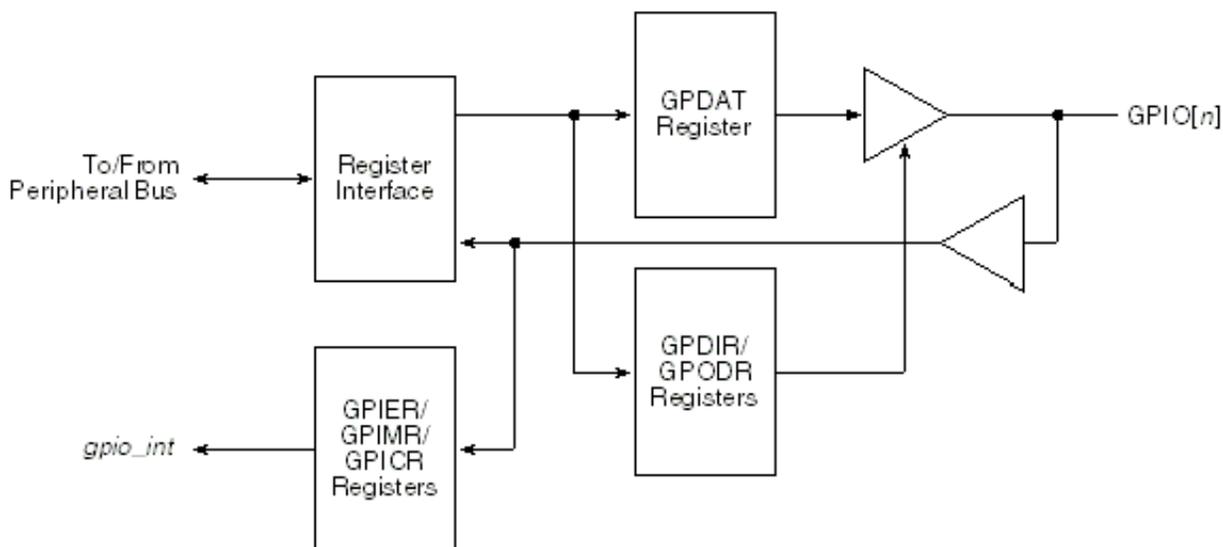


Figure 20-1. GPIO module block diagram

20.1.1 Overview

The GPIO module supports 16 general-purpose I/O ports.

Each port can be configured as an input or as an output. If a port is configured as an input, it can optionally generate an interrupt on detection of a change. If a port is configured as an output, it can be individually configured as an open-drain or a fully active output.

20.1.2 Features

The GPIO unit implements the following features:

- 16 input/output ports
- All signals are configured as inputs when the device comes out of reset and also when HRESET_B is asserted.
- Open-drain capability on all ports
- All input ports can optionally generate an interrupt upon changing their state.

20.2 External signal description

The following section provides information about GPIO signals.

20.2.1 Signals overview

Table 20-1 provides detailed descriptions of the external GPIO signals.

Table 20-1. GPIO external signals-detailed signal descriptions

Signal	I/O	Description
GPIO[0:15]	I/O	General purpose I/O. Each signal can be set individually to act as input or output, according to application needs.
		State Meaning Asserted/Negated-Defined per application.
		Timing Assertion/Negation-Inputs can be asserted completely asynchronously. Outputs are asynchronous to any externally visible clock

20.3 GPIO memory map/register definition

The GPIO has programmable registers that occupy memory-mapped space beginning at offset 0x0_F000 . Note that reading undefined portions of the memory map returns all zeros and writing has no effect.

All GPIO registers are 32 bits wide and are located on 32-bit address boundaries.

GPIO memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
F000	GPIO direction register (GPIO_GPDIR)	32	R/W	0000_0000h	20.3.1/1773
F004	GPIO open drain register (GPIO_GPODR)	32	R/W	0000_0000h	20.3.2/1774
F008	GPIO data register (GPIO_GPDAT)	32	R/W	0000_0000h	20.3.3/1774
F00C	GPIO interrupt event register (GPIO_GPIER)	32	w1c	See section	20.3.4/1775
F010	GPIO interrupt mask register (GPIO_GPIMR)	32	R/W	0000_0000h	20.3.5/1775
F014	GPIO external interrupt control register (GPIO_GPICR)	32	R/W	0000_0000h	20.3.6/1776

20.3.1 GPIO direction register (GPIO_GPDIR)

The GPIO direction registers defines the direction of the individual ports.

Address: F000h base + 0h offset = F000h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W	DRn															Reserved																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

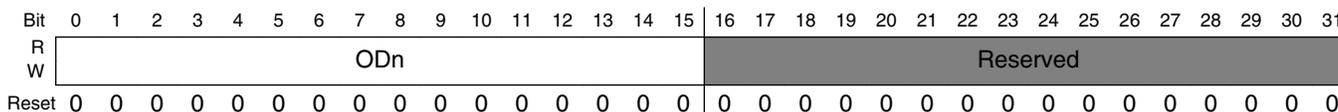
GPIO_GPDIR field descriptions

Field	Description
0–15 DRn	Direction. Indicates whether a signal is used as an input or an output. Bits DR0-DRn correspond to signals GPIO[0:n]. 0 The corresponding signal is an input. 1 The corresponding signal is an output.
16–31 -	This field is reserved.

20.3.2 GPIO open drain register (GPIO_GPODR)

The GPIO open drain register defines the way individual ports drive their output.

Address: F000h base + 4h offset = F004h



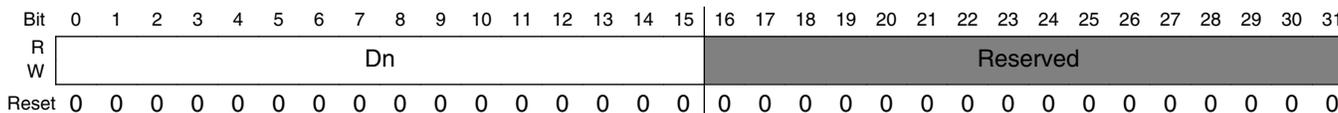
GPIO_GPODR field descriptions

Field	Description
0–15 ODn	Open-drain configuration. Indicates whether a signal is actively driven as an output or an open-drain driver. This register has no effect on signals programmed as inputs in the corresponding GP DIR. Bits OD0-OD <i>n</i> correspond to signals GPIO[0: <i>n</i>]. 0 The I/O signal is actively driven as an output. 1 The I/O signal is an open-drain driver. As an output, the signal is driven active-low, otherwise it is three-stated.
16–31 -	This field is reserved.

20.3.3 GPIO data register (GPIO_GP DAT)

The GPIO data register carries the data in/out for the individual ports.

Address: F000h base + 8h offset = F008h



GPIO_GP DAT field descriptions

Field	Description
0–15 Dn	Data. Write data is latched and presented on external signals if GP DIR has configured the port as an output. Read operation always returns the data at the signal. Bits D0-D <i>n</i> correspond to signals GPIO[0: <i>n</i>].
16–31 -	This field is reserved.

20.3.4 GPIO interrupt event register (GPIO_GPIER)

The GPIO interrupt event register carries information of the events that caused an interrupt. Each bit in GPIER, corresponds to an interrupt source. GPIER bits are cleared by writing ones. However, writing zero has no effect.

The GPIER register does not consider the interrupt mask (which is configured via the GPIMR register). Therefore, a GPIER bit can assert even though the associated interrupt is masked.

Address: F000h base + Ch offset = F00Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	EVn															Reserved																
W	w1c																															
Reset	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*	x*

* Notes:

- The user should write 1s to clear before using. x = Undefined at reset.

GPIO_GPIER field descriptions

Field	Description
0–15 EVn	Interrupt events. Indicates whether an interrupt event occurred on the corresponding GPIO signal. Bits EV0-EVn correspond to signals GPIO[0:n]. 0 No interrupt event occurred on the corresponding GPIO signal. 1 Interrupt event occurred on the corresponding GPIO signal.
16–31 -	This field is reserved.

20.3.5 GPIO interrupt mask register (GPIO_GPIMR)

The GPIO interrupt mask register defines the interrupt masking for the individual ports. When a masked interrupt request occurs, the corresponding GPIER bit is set, regardless of the GPIMR state. When one or more non-masked interrupt events occur, the GPIO module issues an interrupt to the on chip interrupt controller.

Address: F000h base + 10h offset = F010h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	IMn															Reserved																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GPIO_GPIMR field descriptions

Field	Description
0–15 IMn	<p>Interrupt mask. Indicates whether an interrupt event is masked or not masked. Bits IM0-IMn correspond to signals GPIO[0:n].</p> <p>0 The input interrupt signal is masked (disabled). 1 The input interrupt signal is not masked (enabled).</p>
16–31 -	This field is reserved.

20.3.6 GPIO external interrupt control register (GPIO_GPICR)

The GPIO interrupt control register determines whether the corresponding port line asserts an interrupt request on either a high-to-low change or any change on the state of the signal.

Address: F000h base + 14h offset = F014h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	EDn																Reserved															
W	EDn																Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GPIO_GPICR field descriptions

Field	Description
0–15 EDn	<p>Edge detection mode. Bits ED0-EDn correspond to signals GPIO[0:n]. The corresponding port line asserts an interrupt request according to the following:</p> <p>0 Any change on the state of the port generates an interrupt request. 1 High-to-low change on the port generates an interrupt request.</p>
16–31 -	This field is reserved.

Chapter 21

Global Utilities

This chapter describes the global utilities of the chip. It provides signal descriptions, register descriptions, and a functional description of these utilities.

21.1 Overview

This chapter describes the global utilities of the chip. It provides signal descriptions, register descriptions, and a functional description of these utilities.

The global utilities module controls power management, I/O device enabling, power-on-reset (POR) configuration monitoring, alternate function selection for multiplexed signals, and clock control.

21.2 Global utilities features

This section provides an overview of global utilities features.

21.2.1 Power management and block disables

The following features affect the device's overall power consumption:

- Dynamic power management mode
- Software-controlled power management (doze, nap, sleep)
- Externally controlled power management (doze, sleep)
- Static power management (I/O block disables)

21.2.2 Accessing current POR configuration settings

The POR configuration values of all device parameters sampled from pins at reset are available through memory-mapped registers in the global utilities module.

21.2.3 Clock control

The global utilities module also selects the internal clock signal driven on CLK_OUT.

21.3 External signal description

The following subsections provide information about signals that serve as global utilities.

21.3.1 Signals overview

This table summarizes the external signals used by the global utilities module.

Table 21-1. External signal summary

Signal name	I/O	Description	Reference (Section/ page)
ASLEEP	O	Signals that the device has reached a sleep state	Detailed signal descriptions
CKSTP_IN0	I	Checkstop input 0	Detailed signal descriptions
CKSTP_IN1	I	Checkstop input 1	Detailed signal descriptions
CKSTP_OUT0	O	Checkstop output 0	Detailed signal descriptions
CKSTP_OUT1	O	Checkstop output 1	Detailed signal descriptions
CLK_OUT	O	Clock out. Selected by CLKOCR values	Detailed signal descriptions

21.3.2 Detailed signal descriptions

This table describes signals in the global utilities module in detail.

Table 21-2. Detailed signal descriptions

Signal	I/O	Description
ASLEEP	O	Asleep. See Sleep mode . After negation of HRESET_B, ASLEEP is asserted until the device completes its power-on reset sequence and reaches its ready state.
		State meaning Asserted-Indicates that the device is either still in its power-on reset sequence or it has reached a sleep state after a power-down command is issued by software. Negated-The device is not in sleep mode. (It has either awakened from a power-down state, or has completed the POR sequence.)
		Timing Assertion-May occur at any time; may be asserted asynchronously to the input clocks. Negation-Negates synchronously with SYSCLK when leaving power-on sequence; otherwise negation is asynchronous.
CKSTP_IN0_B	I	Checkstop in 0
		State meaning Asserted-Indicates that the e500 core 0 must enter a hard stop condition. All e500 clocks are turned off. CKSTP_OUT0_B is asserted. The rest of chip device logic, including memory controllers, internal memories and registers, and I/O interfaces, remains functional. Negated-Indicates that normal operation should proceed.
		Timing Assertion-May occur at any time; may be asserted asynchronously to the input clocks. Negation-Must remain asserted until the device is reset with assertion of HRESET_B.
CKSTP_IN1_B	I	Checkstop in 1
		State meaning Asserted-Indicates that the e500 core 1 must enter a hard stop condition. All e500 core 1 clocks are turned off. CKSTP_OUT1_B is asserted. The rest of chip device logic, including memory controllers, internal memories and registers, and I/O interfaces, remains functional. Negated-Indicates that normal operation should proceed.
		Timing Assertion-May occur at any time; may be asserted asynchronously to the input clocks. Negation-Must remain asserted until the device is reset with assertion of HRESET_B.
CKSTP_OUT0_B	O	Checkstop out 0
		State Meaning Asserted-Indicates that the e500 core 0 of the chip is in a checkstop state. The rest of the device logic remains functional. Negated-Indicates normal operation. After CKSTP_OUT0_B has been asserted, it is negated after the next assertion (high-to-low transition) of HRESET_B.
		Timing Assertion-May occur at any time; may be asserted asynchronously to the input clocks. Negation-Must remain asserted until the device has been reset with a hard reset.
CKSTP_OUT1_B	O	Checkstop out 1
		State meaning Asserted-Indicates that the e500 core 1 of the chip is in a checkstop state. The rest of the device logic remains functional. Negated-Indicates normal operation. After CKSTP_OUT1_B has been asserted, it is negated after the next assertion (high-to-low transition) of HRESET_B.
		Timing Assertion-May occur at any time; may be asserted asynchronously to the input clocks. Negation-Must remain asserted until the device has been reset with a hard reset.

Table continues on the next page...

Table 21-2. Detailed signal descriptions (continued)

Signal	I/O	Description
CLK_OUT	O	Clock out. Reflects clock signal selected by CLKOCR (see Clock out control register (GUTS_CLKOCR)).
		State meaning Asserted-If CLKOCR[ENB] = 1, clock signal selected by CLKOCR[CLK_SEL] is driven. High impedance-If CLKOCR[ENB] = 0.
		Timing Assertion/negation-Depends on the value of CLKOCR[CLK_SEL].

21.4 GUTS Memory Map/Register Definition

This table summarizes the global utilities registers and their addresses.

GUTS memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
E_0000	POR PLL ratio status register (GUTS_PORPLLSR)	32	R	See section	21.4.1/1782
E_0004	POR boot mode status register (GUTS_PORBMSR)	32	R	See section	21.4.2/1784
E_000C	POR device status register (GUTS_PORDEVSR)	32	R	See section	21.4.3/1786
E_0010	POR debug mode status register (GUTS_PORDBGMSR)	32	R	See section	21.4.4/1789
E_0014	POR device status register 2 (GUTS_PORDEVSR2)	32	R	See section	21.4.5/1791
E_0020	General-purpose POR configuration register (GUTS_GPPORCR)	32	R	See section	21.4.6/1794
E_0060	Alternate function signal multiplex control (GUTS_PMUXCR)	32	R/W	0000_0000h	21.4.7/1794
E_0070	Device disable control register (GUTS_DEVDISR)	32	R/W	See section	21.4.8/1796
E_0080	Power management control and status register (GUTS_POWMGTCR)	32	R/W	0000_0000h	21.4.9/1800
E_008C	Power management clock disable register (GUTS_PMCDR)	32	R/W	0000_00E0h	21.4.10/1803
E_0090	Machine check summary register (GUTS_MCPSUMR)	32	w1c	0000_0000h	21.4.11/1804
E_0094	Reset request status and control register (GUTS_RSTRSCR)	32	R/W	0000_0000h	21.4.12/1806
E_0098	Exception reset control register (GUTS_ECTRSTCR)	32	R/W	0000_0000h	21.4.13/1808
E_009C	Automatic reset status register (GUTS_AUTORSTSR)	32	w1c	0000_0000h	21.4.14/1810
E_00A0	Processor version register (GUTS_PVR)	32	R	8021_1051h	21.4.15/1812

Table continues on the next page...

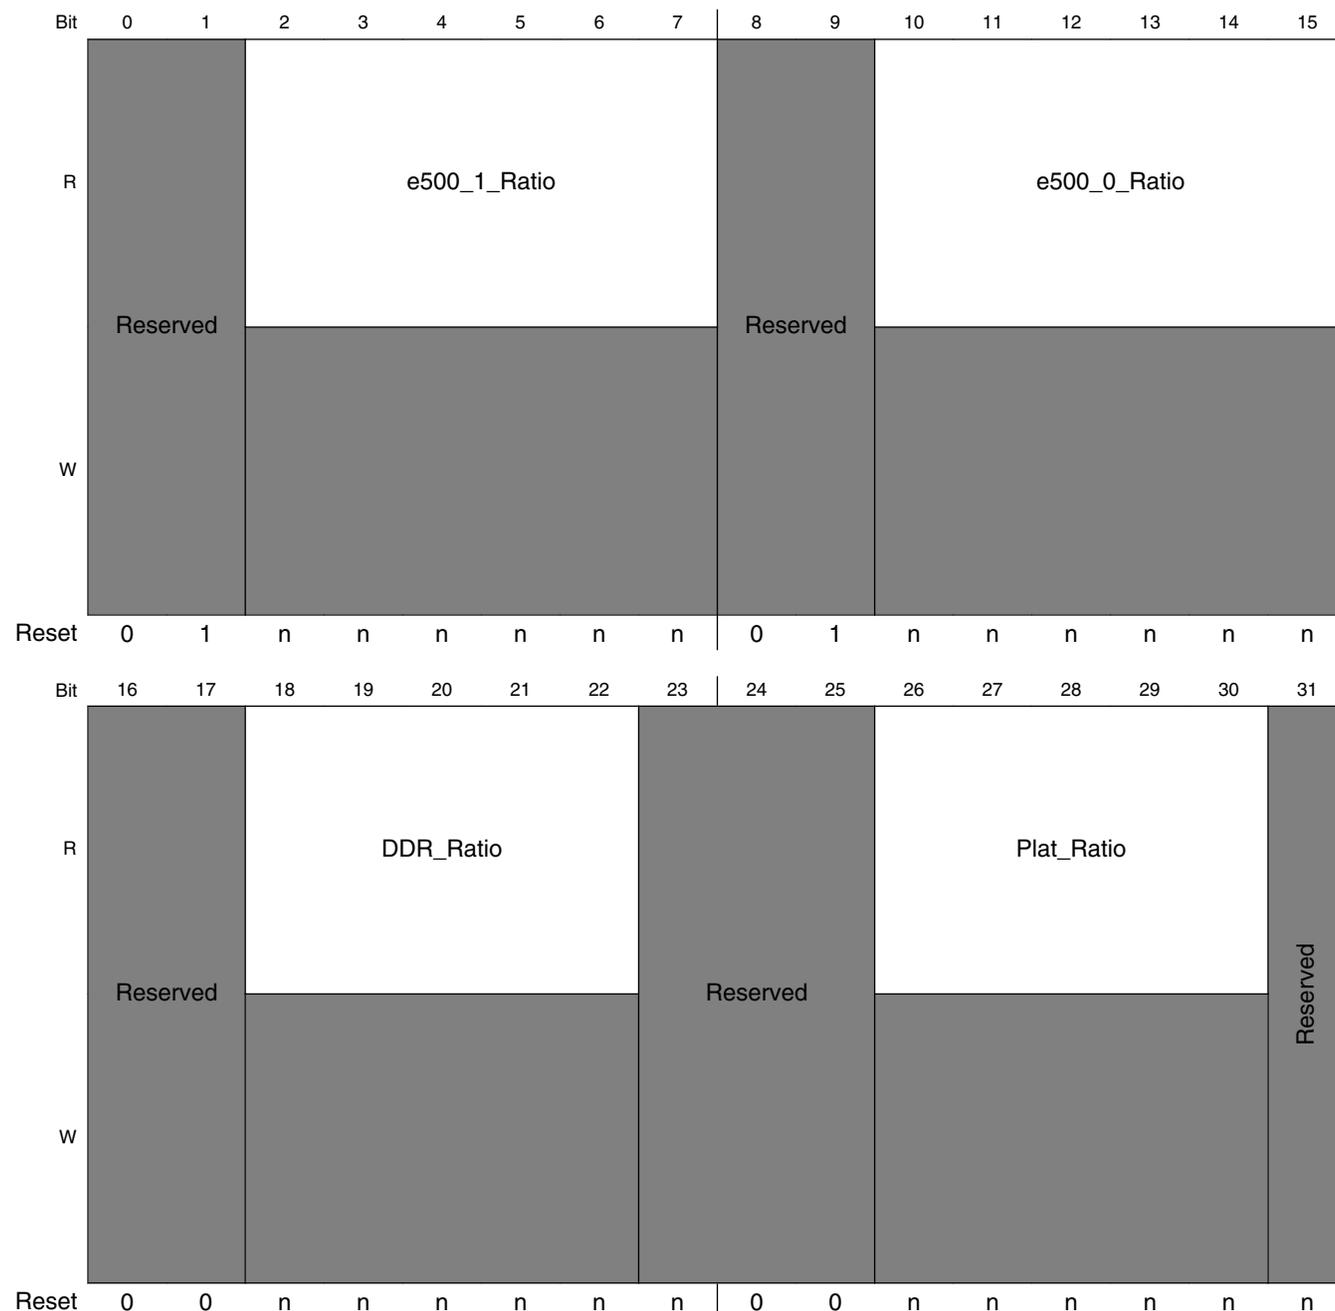
GUTS memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
E_00A4	System version register (GUTS_SVR)	32	R	See section	21.4.16/1812
E_00B0	Reset control register (GUTS_RSTCR)	32	R/W	0000_0000h	21.4.17/1813
E_00C0	IO voltage select status register (GUTS_IOVSELSR)	32	R	0000_0000h	21.4.18/1814
E_0B28	DDR clock disable register (GUTS_DDRCLKDR)	32	R/W	0000_0000h	21.4.19/1815
E_0E00	Clock out control register (GUTS_CLKOCR)	32	R/W	0000_0000h	21.4.20/1816
E_0E20	ECM control register (GUTS_ECMCR)	32	R/W	0000_0000h	21.4.21/1817
E_0F64	SDHC debug control register (GUTS_SDHCDCR)	32	R/W	See section	21.4.22/1818
E_3000	SRDS Control Register 0 (GUTS_SRDSCR0)	32	R/W	See section	21.4.23/1819
E_3004	SRDS Control Register 1 (GUTS_SRDSCR1)	32	R/W	See section	21.4.24/1821
E_3008	SRDS Control Register 2 (GUTS_SRDSCR2)	32	R/W	0080_2040h	21.4.25/1823
E_3010	SRDS Control Register 4 (GUTS_SRDSCR4)	32	R/W	See section	21.4.26/1824
E_3014	SRDS Control Register 5 (GUTS_SRDSCR5)	32	R/W	See section	21.4.27/1826
E_3018	SRDS Control Register 6 (GUTS_SRDSCR6)	32	R/W	See section	21.4.28/1827

21.4.1 POR PLL ratio status register (GUTS_PORPLLSR)

PORPLLSR contains the settings for the PLL ratios as set by the `cfg_sys_pll[0:2]`, `cfg_ddr_pll[0:3]`, `cfg_core0_pll[0:1]`, and `cfg_core1_pll[0:1]` POR configuration pins. See [System PLL ratio](#), [DDR PLL Ratio](#), and [e500 core PLL ratios](#) for more information.

Address: E_0000h base + 0h offset = E_0000h



GUTS_PORPLLSR field descriptions

Field	Description
0–1 -	This field is reserved. Reserved
2–7 e500_1_Ratio	Clock ratio between the e500 core 1 and the CCB clock. The 3 lsbs of this field correspond to the values on <code>cfg_core1_pll[0:2]</code> at the negation of <code>HRESET_B</code> . (See e500 core PLL ratios) Patterns not shown are reserved. 000010 1:1 000011 3:2 000100 2:1 000101 5:2 000110 3:1 000111 7:2 001000 4:1 001001 9:2
8–9 -	This field is reserved. Reserved
10–15 e500_0_Ratio	Clock ratio between the e500 core 0 and the CCB clock. The 3 lsbs of this field correspond to the values on <code>cfg_core0_pll[0:2]</code> at the negation of <code>HRESET_B</code> . (See e500 core PLL ratios) Patterns not shown are reserved. 000010 1:1 000011 3:2 000100 2:1 000101 5:2 000110 3:1 000111 7:2 001000 4:1 001001 9:2
16–17 -	This field is reserved. Reserved
18–22 DDR_Ratio	Clock ratio between the DDR Complex clock and <code>DDRCLK</code> . 000011 3:1 000100 4:1 000110 6:1 001000 8:1 001010 10:1 001100 12:1 001110 Reserved 000111 Synchronous Mode-DDR Complex Clocked by CCB clock
23–25 -	This field is reserved. Reserved
26–30 Plat_Ratio	Clock ratio between the CCB (platform) clock and <code>SYSCLK</code> . 00100 4:1 00101 5:1 00110 6:1 01000 8:1

Table continues on the next page...

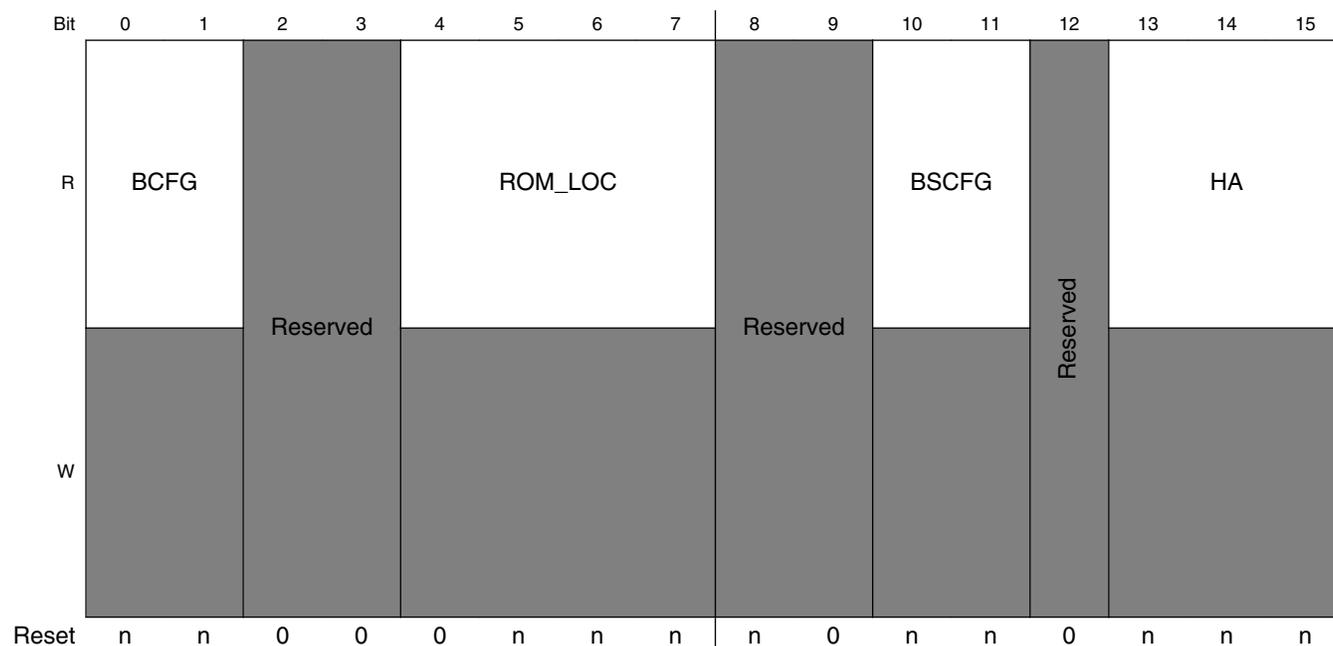
GUTS_PORPLLSR field descriptions (continued)

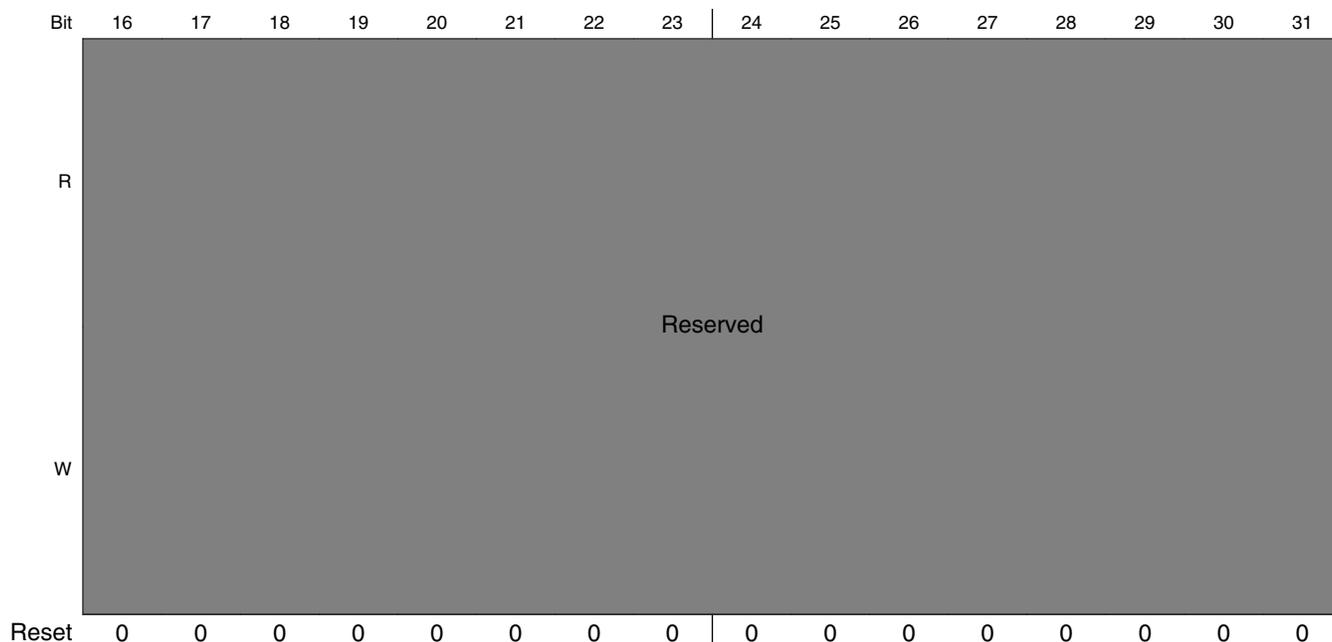
Field	Description
	01010 Reserved 01100 Reserved
31 -	This field is reserved. Reserved

21.4.2 POR boot mode status register (GUTS_PORBMSR)

The PORBMSR reports setting of the POR configuration pins that control the boot mode settings (described in [Boot ROM location](#) , [CPU boot configuration](#) , and [Boot sequencer configuration](#)) and the default settings of PCI Express and serial RapidIO host/agent mode (described in [Host/agent configuration](#)).

Address: E_0000h base + 4h offset = E_0004h





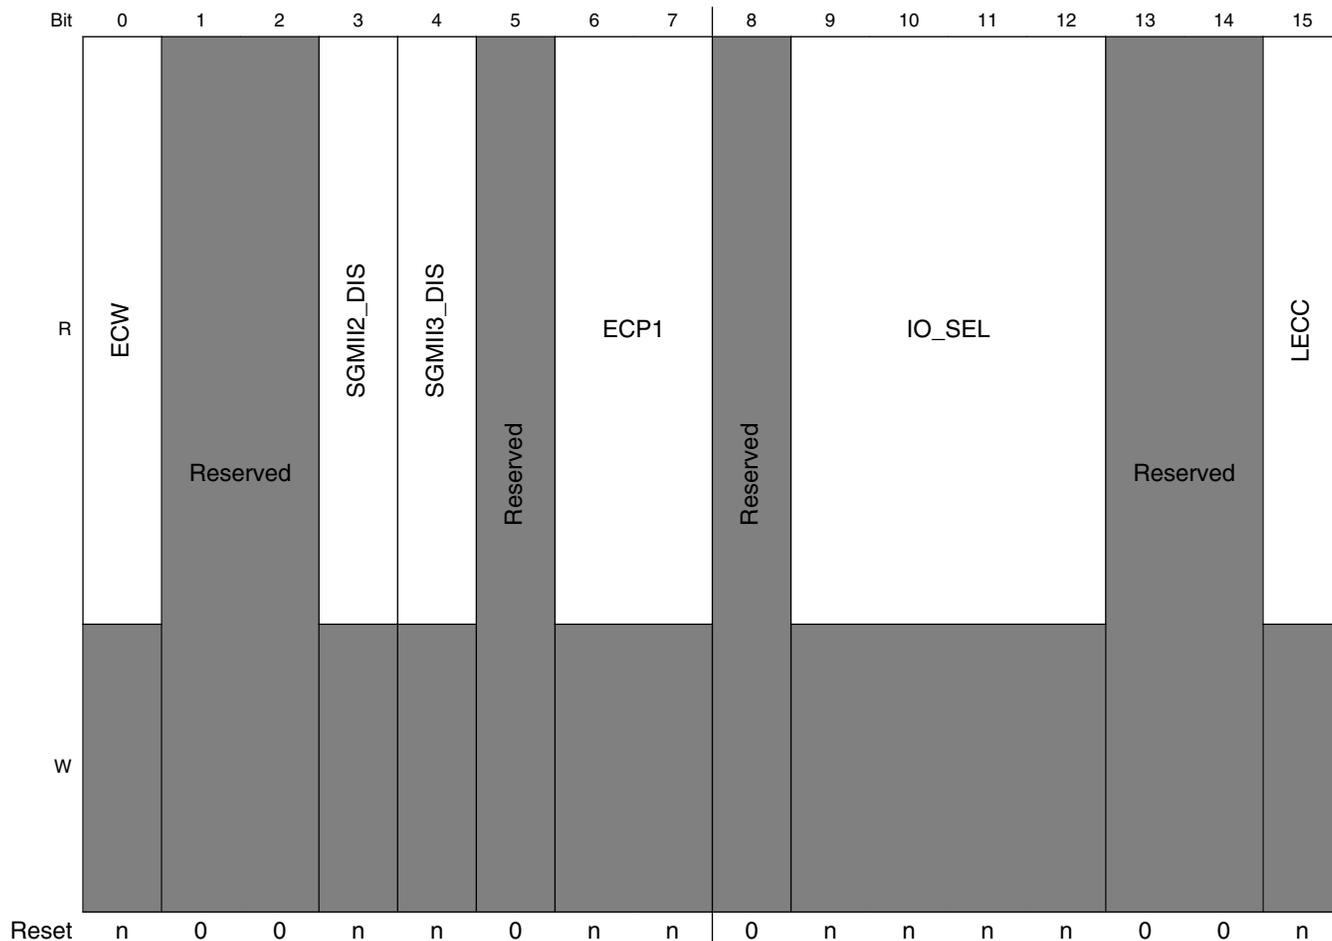
GUTS_PORBMSR field descriptions

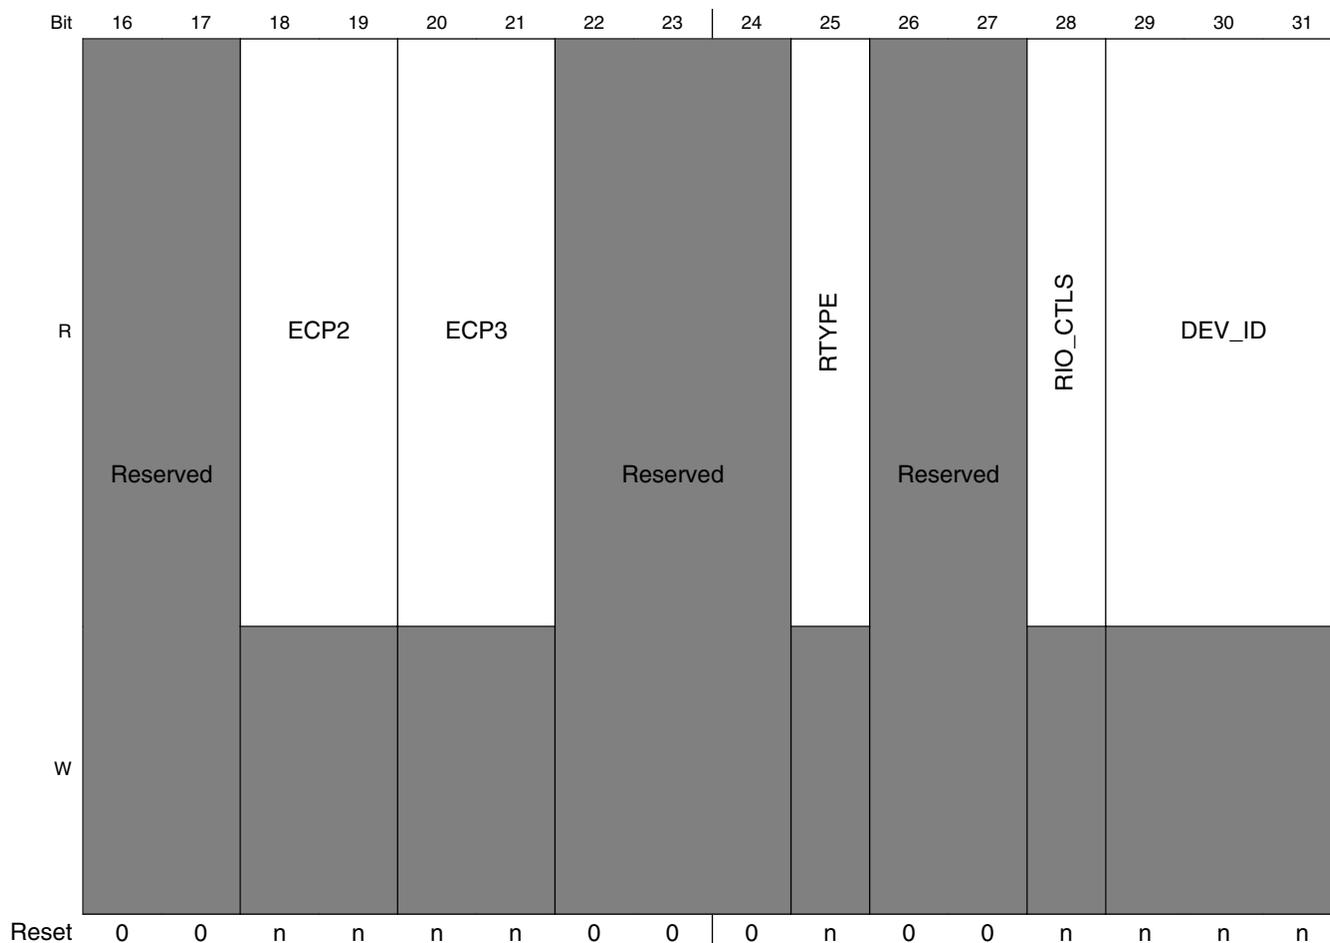
Field	Description
0–1 BCFG	CPU boot configuration (See CPU boot configuration .) For BCFG: cfg_cpu0_boot and cfg_cpu1_boot, see Table 4-15 .
2–3 -	This field is reserved. Reserved
4–7 ROM_LOC	Location of boot ROM. (See Boot ROM location .) For ROM_LOC: cfg_rom_loc[0:3], see Table 4-12 .
8–9 -	This field is reserved. Reserved
10–11 BSCFG	Boot sequencer configuration (Boot sequencer configuration .) For BSCFG: cfg_boot_seq[0:1], Table 4-16 .
12 -	This field is reserved. Reserved
13–15 HA	Host/agent mode configuration. When the chip is an agent on an interface, it is prevented from mastering transactions on that interface until the external host configures the interface appropriately. (Host/agent configuration .) For HA: cfg_host_agt[0:2]], see Table 4-13 .
16–31 -	This field is reserved. Reserved

21.4.3 POR device status register (GUTS_PORDEVSR)

PORDEVSR reports other POR settings for I/O devices as described in Section [SerDes reference clock configuration](#) , [eTSEC1 and eTSEC2 width](#) , and [eTSEC2 protocol](#) .

Address: E_0000h base + Ch offset = E_000Ch





GUTS_PORDEVSR field descriptions

Field	Description
0 ECW	eTSEC1 and eTSEC2 controller width (See eTSEC1 and eTSEC2 width .) For ECW: <code>cfg_tsec_reduce</code> , Table 4-24 .
1–2 -	This field is reserved. Reserved
3 SGMII2_DIS	eTSEC2 in SGMII mode disabled (See eTSEC2 SGMII mode .) For SGMII2_DIS: <code>cfg_sgmi2</code> , see Table 4-22 .
4 SGMII3_DIS	eTSEC3 SGMII mode disabled (See eTSEC3 SGMII mode .) For SGMII3_DIS: <code>cfg_sgmi3</code> , see Table 4-23 .
5 -	This field is reserved. Reserved
6–7 ECP1	eTSEC1 controller protocol (See eTSEC1 protocol .) For ECP1: <code>cfg_tsec1_prctl[0:1]</code> , see Table 4-25 .
8 -	This field is reserved. Reserved
9–12 IO_SEL	I/O port selection mode (See I/O port selection .) For IO_SEL: <code>cfg_IO_ports[0:3]</code> , see Table 4-14 .

Table continues on the next page...

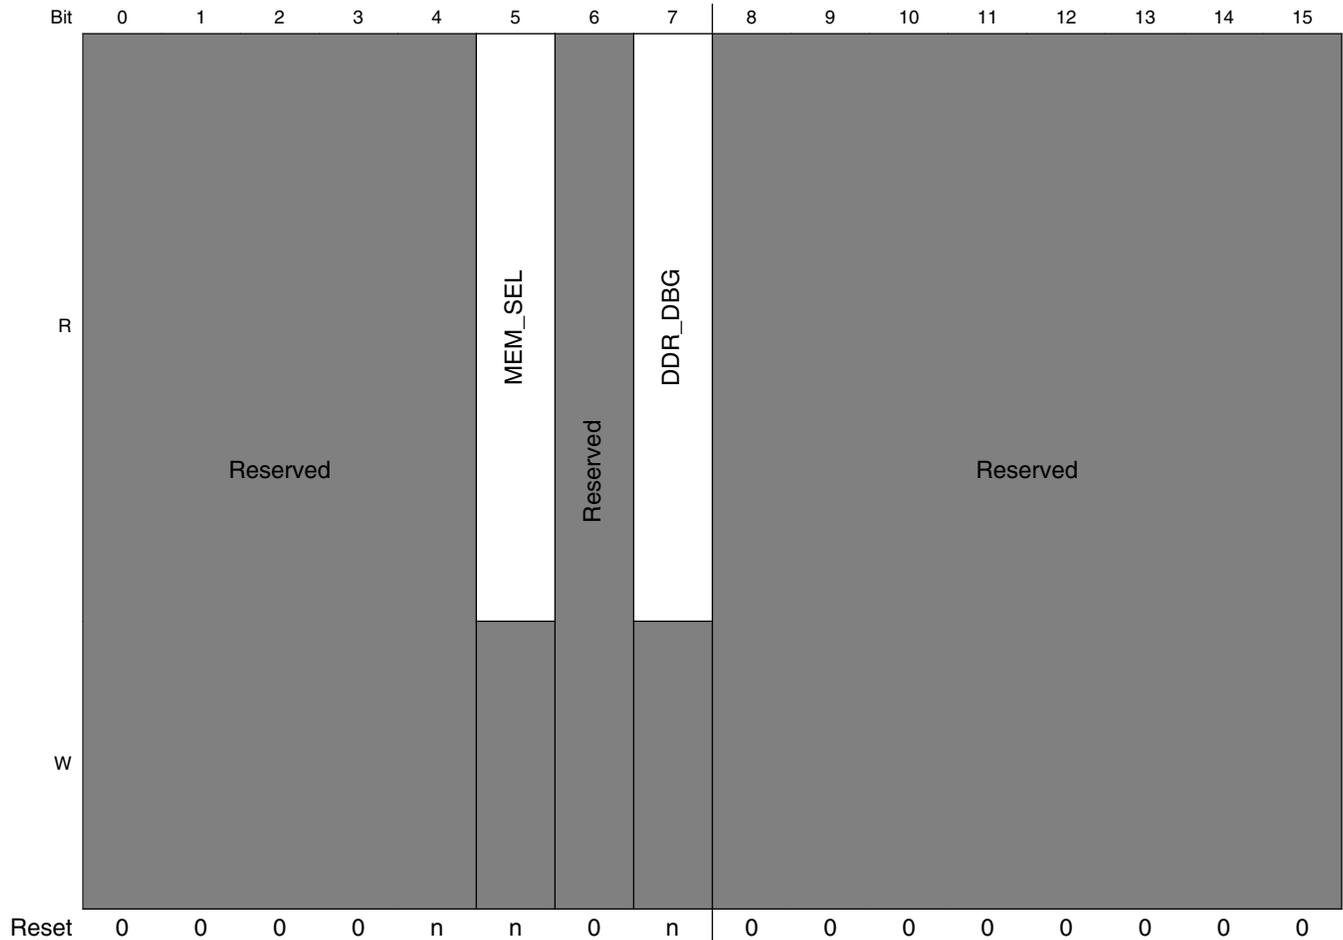
GUTS_PORDEVSR field descriptions (continued)

Field	Description
13–14 -	This field is reserved. Reserved
15 LECC	eLBC ECC enable For LECC: <code>cfg_elbc_ecc</code> , see Table 4-33 .
16–17 -	This field is reserved. Reserved
18–19 ECP2	eTSEC2 controller protocol (See eTSEC2 protocol) For ECP2: <code>cfg_tsec2_prctl[0:1]</code> , see Table 4-26 .
20–21 ECP3	eTSEC3 controller protocol (See eTSEC3 protocol .) For ECP3: <code>cfg_tsec3_prctl[0:1]</code> , see Table 4-27 . NOTE: eTSEC3 parallel interface operation is only available when eTSEC1 and eTSEC2 are configured for reduced controller width (ECW = 0).
22–24 -	This field is reserved. Reserved
25 RTYPE	DRAM Type for DDR Controller (See DDR SDRAM type .) For RTYPE: <code>cfg_dram_type</code> , see Table 4-17 .
26–27 -	This field is reserved. Reserved
28 RIO_CTL5	RapidIO system size (See RapidIO system size .) 0 PE Device does not support common transport large systems (up to 256 devices) 1 PE Device supports common transport large systems (up to 65,536 devices)
29–31 DEV_ID	Serial RapidIO device ID (See RapidIO device ID)

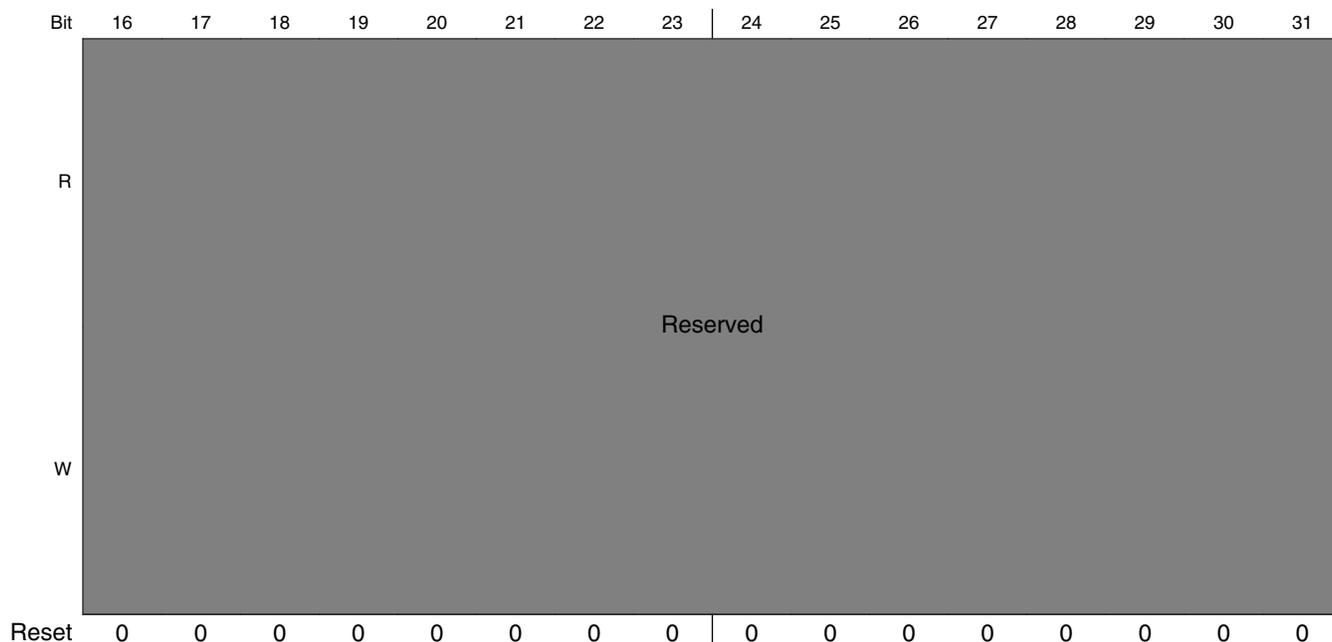
21.4.4 POR debug mode status register (GUTS_PORDBGMSR)

PORDBGMSR holds debug mode settings from the POR configuration pins as described in [Memory debug configuration](#) , and [DDR debug configuration](#) .

Address: E_0000h base + 10h offset = E_0010h



GUTS Memory Map/Register Definition



GUTS_PORDBGMSR field descriptions

Field	Description
0–4 -	This field is reserved. Reserved
5 MEM_SEL	Memory select. Indicates which controller is driving MSRCID[0:4] and MDVAL. (See DDR debug configuration) For MEM_SEL: cfg_mem_debug, see Table 4-30 .
6 -	This field is reserved. Reserved
7 DDR_DBG	DDR controller debug configuration (See DDR debug configuration .) For DDR_DBG: cfg_ddr_debug, see Table 4-31 .
8–31 -	This field is reserved. Reserved

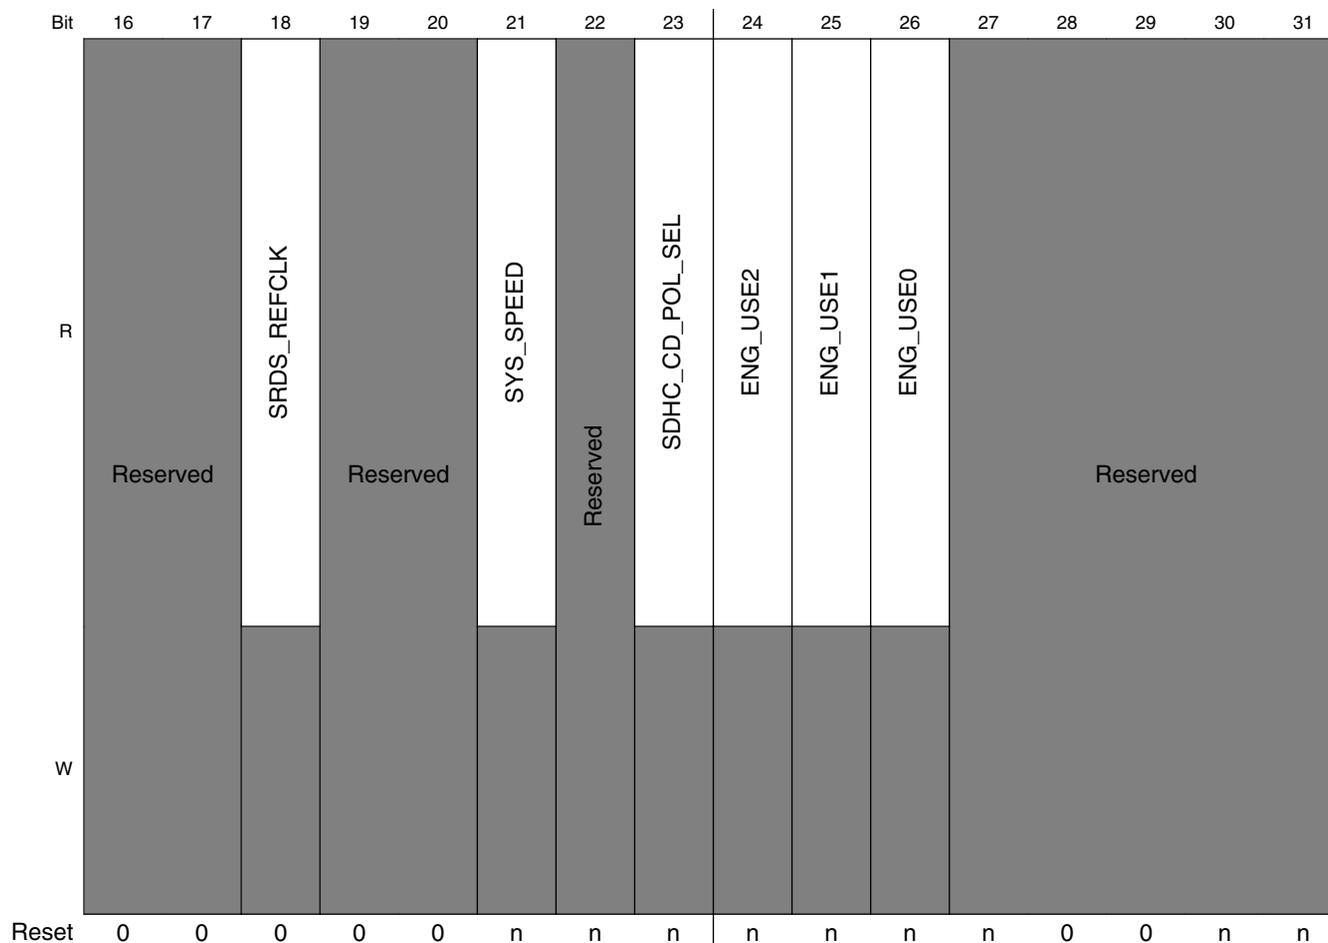
21.4.5 POR device status register 2 (GUTS_PORDEVSR2)

The PORDEVSR2 reports POR settings.

Address: E_0000h base + 14h offset = E_0014h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved							ENG_USE7	ENG_USE6	ENG_USE5	SRDS_PLL_TOE	ENG_USE3	CORE0_SPEED	CORE1_SPEED	PLAT_SPEED	DDR_SPEED
W	n							0	0	0	0	0	n	n	n	n
Reset	n	0	0	0	n	0	n	0	0	0	0	0	n	n	n	n

4013 Memory Map/Register Definition



GUTS_PORDEVS2 field descriptions

Field	Description
0-6 -	This field is reserved. Reserved
7 ENG_USE7	For future engineering use 7 (Driven by POR reset signal). See Table 4-33 for more information.
8 ENG_USE6	For future engineering use 6 (Driven by POR reset signal). See Table 4-33 for more information.
9 ENG_USE5	For future engineering use 5 (Driven by POR reset signal). See Table 4-33 for more information.
10 SRDS_PLL_TOE	SerDes PLL lock time-out counter 0 Enable SerDes PLL lock time-out counter. The power-on-reset sequence waits for the SerDes PLL to lock while the time-out counter has not expired. 1 Disable SerDes PLL lock time-out counter. The power-on-reset sequence waits indefinitely for the SerDes PLL to lock. (default)
11 ENG_USE3	For future engineering use 3 (Driven by POR reset signal)
12 CORE0_SPEED	Core 0 Speed For CORE0_SPEED: <code>cfg_core0_speed</code> , see Table 4-37 .

Table continues on the next page...

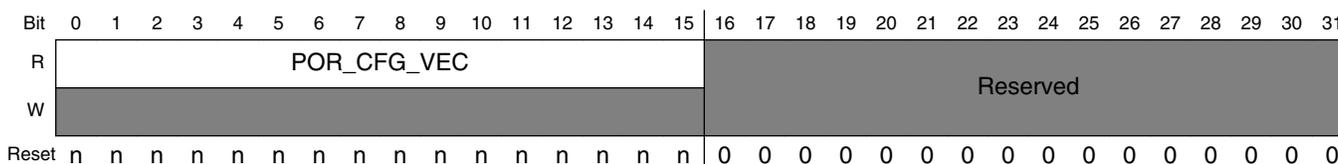
GUTS_PORDEVSR2 field descriptions (continued)

Field	Description
13 CORE1_SPEED	Core 1 Speed For CORE1_SPEED: <code>cfg_core1_speed</code> , see Table 4-38 .
14 PLAT_SPEED	Platform Speed For PLAT_SPEED: <code>cfg_plat_speed</code> , see Table 4-36 .
15 DDR_SPEED	DDR Speed For DDR_SPEED: <code>cfg_ddr_speed</code> , see Table 4-39 .
16–17 -	This field is reserved. Reserved
18 SRDS_REFCLK	SerDes reference clock. For SRDS_REFCLK: <code>cfg_srds_refclk</code> , see Table 4-18 .
19–20 -	This field is reserved. Reserved
21 SYS_SPEED	System Speed For SYS_SPEED: <code>cfg_sys_speed</code> , see Table 4-35 .
22 -	This field is reserved. Reserved
23 SDHC_CD_POL_SEL	eSDHC card-detect polarity select For SDHC_CD_POL_SEL: <code>cfg_sdhc_cd_pol_sel</code> , see Table 4-41 .
24 ENG_USE2	For future engineering use 2 (Driven by POR reset signal). See Table 4-33 , for more information.
25 ENG_USE1	For future engineering use 1 (Driven by POR reset signal). See Table 4-33 , for more information.
26 ENG_USE0	For future engineering use 0 (Driven by POR reset signal). See Table 4-33 , for more information.
27–31 -	This field is reserved. Reserved

21.4.6 General-purpose POR configuration register (GUTS_GPPORCR)

GPPORCR stores the value sampled from the local bus address/data signals, LAD[0:15], during POR, as described in [General-purpose POR configuration](#) , Software can use this value to inform the operating system about initial system configuration. Typical interpretations include circuit board type, board ID number, or a list of available peripherals.

Address: E_0000h base + 20h offset = E_0020h



GUTS_GPPORCR field descriptions

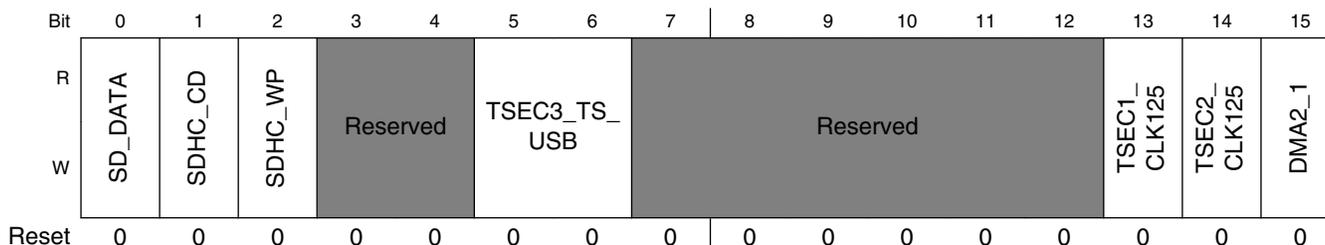
Field	Description
0–15 POR_CFG_VEC	General-purpose POR configuration vector sampled from local bus address/data signals at the negation of HRESET_B . Note that if nothing is driven on these signals during reset, the value of this register is indeterminate.
16–31 -	This field is reserved. Reserved

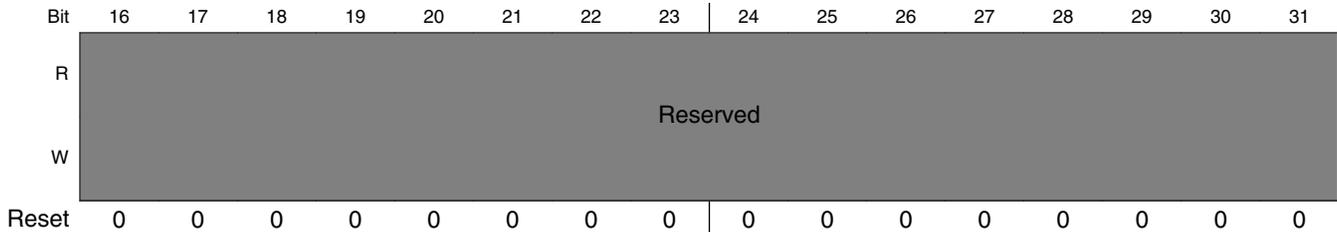
21.4.7 Alternate function signal multiplex control (GUTS_PMUXCR)

PMUXCR contains bits that enable the exposure of an alternate function replacing the pins' primary functions.

Specifically, DMA2, eSDHC, TSEC1, TSEC2, and USB each have features that can be exposed replacing the primary functions of certain SPI and GPIO pins.

Address: E_0000h base + 60h offset = E_0060h





GUTS_PMUXCR field descriptions

Field	Description
0 SD_DATA	Exposes SDHC_DAT[4-7] signals used for MMC 8-bit Mode. 0 SDHC_DAT[4:7] is not exposed to pins. The pins retain their primary function as SPI. 1 SDHC_DAT[4:7] is exposed to pins as follows: SPI_CS[0] functions as SDHC_DAT[4] SPI_CS[1] functions as SDHC_DAT[5] SPI_CS[2] functions as SDHC_DAT[6] SPI_CS[3] functions as SDHC_DAT[7]
1 SDHC_CD	Exposes SDHC_CD_B signal used for Secure Digital card detection. 0 SDHC_CD_B is not exposed to a pin. The pin retains its primary function as GPIO. 1 SDHC_CD_B is exposed to a pin as follows: GPIO[8] functions as SDHC_CD_B
2 SDHC_WP	Exposes SDHC_WP signal used for Secure Digital write protect detection. 0 SDHC_WP is not exposed to a pin. The pin retains its primary function as GPIO. 1 SDHC_WP is exposed to a pin as follows: GPIO[9] functions as SDHC_WP
3-4 -	This field is reserved. Reserved
5-6 TSEC3_TS_USB	Exposes USB_PCTL0 and USB_PCTL1 signals used for host-mode USB port status indication. Patterns not shown are reserved. 00 USB_PCTL0 and USB_PCTL1 are not exposed to pins. The pins retain their primary function as GPIO. 10 USB_PCTL0 and USB_PCTL1 signals are exposed to pins as follows: GPIO[10] functions as USB_PCTL0 GPIO[11] functions as USB_PCTL1
7-12 -	This field is reserved. Reserved
13 TSEC1_CLK125	Selects the 125MHz reference clock source for TSEC1. This configuration field provides the option to select a shared or dedicated clock source for TSEC1. When TSEC1 is configured for RGMII, RTBI, or gigabit GMII, the TSEC1_TX_CLK pin may be used as the 125 MHz reference clock source for TSEC1. Note that IEEE 1588 accuracy can be adversely impacted in systems using multiple unsynchronized gigabit Ethernet ports. To enable IEEE 1588 accuracy in such systems, a dedicated reference clock may be selected.

Table continues on the next page...

GUTS_PMUXCR field descriptions (continued)

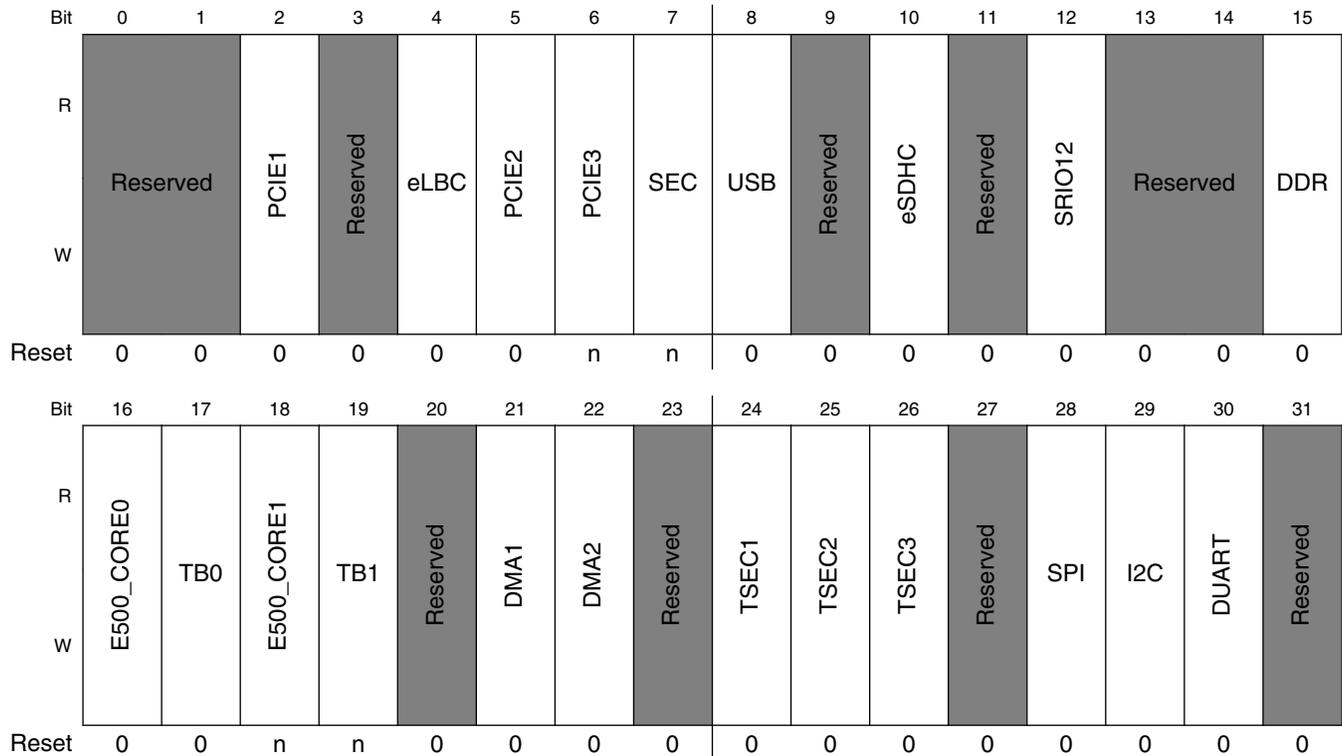
Field	Description
	<p>NOTE: TSEC3 always uses EC_GTX_CLK125 for the 125MHz reference clock.</p> <p>0 Selects EC_GTX_CLK125 as the clock source (shared). 1 Selects TSEC1_TX_CLK as the clock source (dedicated).</p>
14 TSEC2_CLK125	<p>Selects the 125MHz reference clock source for TSEC2.</p> <p>This configuration field provides the option to select a shared or dedicated clock source for TSEC2. When TSEC2 is configured for RGMII, RTBI, or gigabit GMII, the TSEC2_TX_CLK pin may be used as the 125 MHz reference clock source for TSEC2. Note that IEEE 1588 accuracy can be adversely impacted in systems using multiple unsynchronized gigabit Ethernet ports. To enable IEEE 1588 accuracy in such systems, a dedicated reference clock may be selected.</p> <p>NOTE: TSEC3 always uses EC_GTX_CLK125 for the 125MHz reference clock.</p> <p>0 Selects EC_GTX_CLK125 as the clock source (shared). 1 Selects TSEC2_TX_CLK as the clock source (dedicated).</p>
15 DMA2_1	<p>Exposes DMA2 channel 1 signals.</p> <p>0 DMA2 channel 1 is not exposed to pins; the pins retain their primary function as local bus chip selects 1 DMA2 channel 1 is exposed to pins as follows: LCS5_B functions as DMA2_DREQ1_B LCS6_B functions as DMA2_DACK1_B LCS7_B functions as DMA2_DDONE1_B</p>
16–31 -	<p>This field is reserved. Reserved</p>

21.4.8 Device disable control register (GUTS_DEVDISR)

DEVDISR contains disable bits for various functional blocks. Note that bits with a reset value of *n* depend on the state of POR configuration signals at reset.

All functional blocks are enabled after reset; unneeded blocks can be disabled to reduce power consumption. Blocks disabled by DEVDISR must not be re-enabled without a hard reset (except bits TB0 and TB1). [Shutting down unused blocks](#) has more information on the use of DEVDISR.

Address: E_0000h base + 70h offset = E_0070h



GUTS_DEVDISR field descriptions

Field	Description
0-1 -	This field is reserved. Reserved
2 PCIE1	PCI Express 1 controller disable 0 PCI Express 1 controller enable 1 PCI Express 1 controller disable
3 -	This field is reserved. Reserved
4 eLBC	Local bus controller disable 0 Local bus controller enable 1 Local bus controller disable
5 PCIE2	PCI Express 2 controller disable 0 PCI Express 2 controller enable 1 PCI Express 2 controller disable
6 PCIE3	PCI Express 3 controller disable 0 PCI Express 3 controller enable 1 PCI Express 3 controller disable
7 SEC	Security disable 0 Security enable 1 Security disable

Table continues on the next page...

GUTS_DEVDISR field descriptions (continued)

Field	Description
8 USB	USB disable 0 USB enable 1 USB disable
9 -	This field is reserved. Reserved
10 eSDHC	eSDHC disable 0 eSDHC enable 1 eSDHC disable
11 -	This field is reserved. Reserved
12 SRIO12	Serial RapidIO 1 and 2 controller disable 0 Serial RapidIO 1 and 2 controller enable 1 Serial RapidIO 1 and 2 controller disable
13–14 -	This field is reserved. Reserved
15 DDR	DDR controller disable 0 DDR SDRAM controller enable 1 DDR SDRAM controller disable
16 E500_CORE0	e500 core 0 disable 0 e500 0 core enable 1 e500 0 core disable. Places the core in the core_stopped state in which it does not respond to interrupts. Equivalent to nap mode. Instruction fetching is stopped, snooping is disabled, and clocks are shut down to all functional units of the core including the timer facilities. For more information, see Shutting down unused blocks .
17 TB0	Time base (timer facilities) of the e500 core 0 disable 0 Timer facilities enabled 1 Timer facilities disabled.
18 E500_CORE1	e500 core 1 disable 0 e500 1 core enable 1 e500 1 core disable. Places the core in the core_stopped state in which it does not respond to interrupts. Equivalent to nap mode. Instruction fetching is stopped, snooping is disabled, and clocks are shut down to all functional units of the core including the timer facilities. For more information, see Shutting down unused blocks .
19 TB1	Time base (timer facilities) of the e500 core 1 disable 0 Timer facilities enabled 1 Timer facilities disabled.
20 -	This field is reserved. Reserved
21 DMA1	DMA 1 controller disabled

Table continues on the next page...

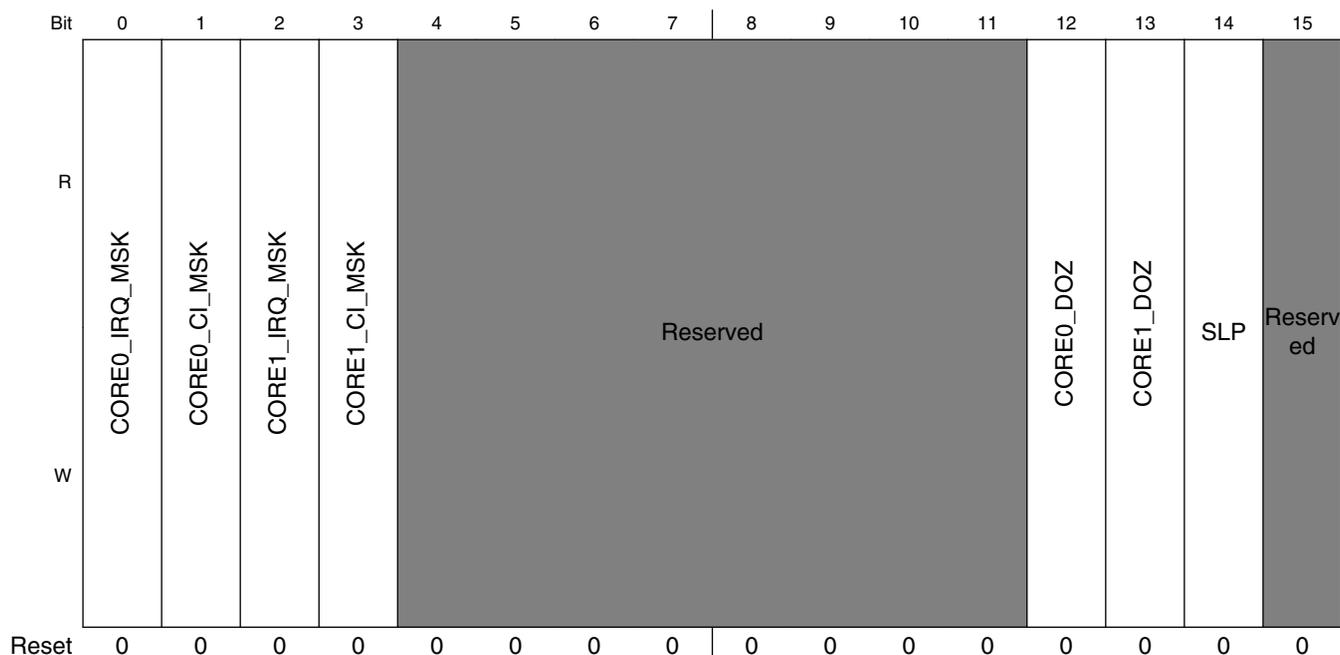
GUTS_DEVDISR field descriptions (continued)

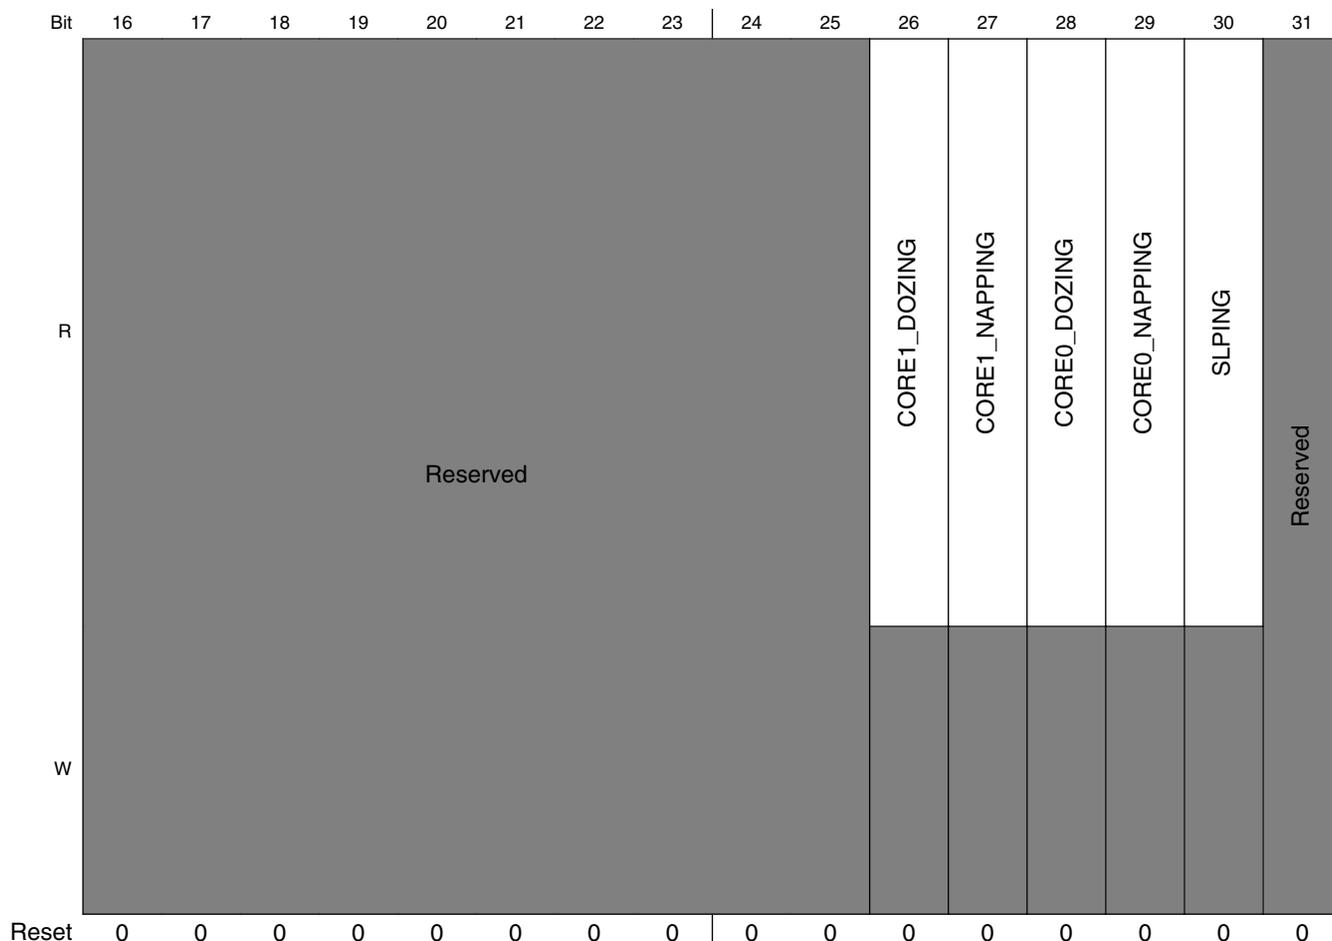
Field	Description
	0 DMA 1 controller enabled 1 DMA 1 controller disabled
22 DMA2	DMA 2 controller disabled 0 DMA 2 controller enabled 1 DMA 2 controller disabled
23 -	This field is reserved. Reserved
24 TSEC1	Three-speed Ethernet controller 1 disable 0 eTSEC1 enabled 1 eTSEC1 disabled
25 TSEC2	Three-speed Ethernet controller 2 disable 0 eTSEC2 enabled 1 eTSEC2 disabled.
26 TSEC3	Three-speed Ethernet controller 3 disable 0 eTSEC3 enabled 1 eTSEC3 disabled
27 -	This field is reserved. Reserved
28 SPI	SPI controller disable 0 SPI enabled 1 SPI disabled
29 I2C	I ² C controllers disabled 0 I ² C controllers enabled 1 I ² C controllers disabled
30 DUART	Dual UART controller disabled 0 DUART enabled 1 DUART disabled
31 -	This field is reserved. Reserved

21.4.9 Power management control and status register (GUTS_POWMGTCR)

POWMGTCSR contains bits for placing the chip into low power states and for controlling when it wakes up. It also contains power management status bits. See [Interrupts and power management controlled by POWMGTCR](#) for more information.

Address: E_0000h base + 80h offset = E_0080h





GUTS_POWMGTCR field descriptions

Field	Description
0 CORE0_IRQ_MSK	Core 0 interrupt input mask 0 Interrupts cause the device to wake up from a low-power state. 1 Interrupts are masked as a wake-up condition. The device remains in a low-power state despite the presence of an interrupt request.
1 CORE0_CI_MSK	Core 0 critical interrupt input mask 0 Critical interrupts cause the device to wake up from a low power state. 1 Critical interrupts are masked as a wake-up condition. The device remains in a low-power state despite the presence of a critical interrupt.
2 CORE1_IRQ_MSK	Core 1 interrupt input mask 0 Interrupts cause the device to wake up from a low-power state. 1 Interrupts are masked as a wake-up condition. The device remains in a low-power state despite the presence of an interrupt request.
3 CORE1_CI_MSK	Core 1 critical interrupt input mask 0 Critical interrupts cause the device to wake up from a low power state. 1 Critical interrupts are masked as a wake-up condition. The device remains in a low-power state despite the presence of a critical interrupt.

Table continues on the next page...

GUTS_POWMGTCR field descriptions (continued)

Field	Description
4–11 -	This field is reserved. Reserved
12 CORE0_DOZ	Core 0 doze mode. 0 No request to put device in doze mode. Note that this bit is automatically cleared on MCP, UDE, SRESET, <i>core_tbit</i> (from the core) and also <i>int_b</i> and <i>cint_b</i> if not masked. 1 Device is to be placed in doze mode. Instruction fetching is halted in the e500 core 0. Note that this bit is logically ORed with HID0[DOZE].
13 CORE1_DOZ	Core 1 doze mode. 0 No request to put device in doze mode. Note that this bit is automatically cleared on MCP, UDE, SRESET, <i>core_tbit</i> (from the core) <i>int_b</i> and <i>cint_b</i> also and if not masked. 1 Device is to be placed in doze mode. Instruction fetching is halted in the e500 core 1. Note that this bit is logically ORed with HID0[DOZE].
14 SLP	Sleep mode 0 No request to put device in sleep mode. 1 Device is to be placed in sleep mode. Instruction fetching is halted, snooping of L1 caches is disabled, and most functional blocks are shut down in both the e500 cores and the system logic.
15–25 -	This field is reserved. Reserved
26 CORE1_DOZING	Core 1 doze status 0 Device is not in doze mode. 1 The chip is in doze mode because POWMGTCR[DOZ] is set or because HID0[DOZE] and MSR[WE] (in the e500 core 1) are set. The core has halted instruction fetching, but all other functional blocks in the core and device are running.
27 CORE1_NAPPING	Core 1 nap status 0 Device is not in nap mode. 1 The chip is in nap mode because HID0[NAP] and MSR[WE] are set. The core has halted instruction fetching, snooping of the L1 caches is disabled, and all of the core's functional units except the timer facilities are shut down. All functional blocks in the device are running.
28 CORE0_DOZING	Core 0 doze status 0 Device is not in doze mode. 1 The chip is in doze mode because POWMGTCR[DOZ] is set or because HID0[DOZE] and MSR[WE] (in the e500 core 0) are set. The core has halted instruction fetching, but all other functional blocks in the core and device are running.
29 CORE0_NAPPING	Core 0 nap status 0 Device is not in nap mode. 1 The chip is in nap mode because HID0[NAP] and MSR[WE] are set. The core has halted instruction fetching, snooping of the L1 caches is disabled, and all of the core's functional units except the timer facilities are shut down. All functional blocks in the device are running.
30 SLPING	Sleep status 0 Device is not attempting to reach sleep mode. 1 The device is attempting to SLEEP because POWMGTCR[SLP] is set or because HID0[SLEEP] and MSR[WE] (in the e500 core) are set. Most functional blocks in the core and device are shut down or are attempting to shut down.

Table continues on the next page...

GUTS_POWMGTCSR field descriptions (continued)

Field	Description
31 -	This field is reserved. Reserved. Should be cleared.

21.4.10 Power management clock disable register (GUTS_PMCDR)

The power management clock disable register (PMCDR) contains disable bits for various functional blocks. The register determine the blocks which will shut down the clock in the sleep power states.

Address: E_0000h base + 8Ch offset = E_008Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								TSEC1	TSEC2	TSEC3	Reserved				
W	Reserved								TSEC1	TSEC2	TSEC3	Reserved				
Reset	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0	0

GUTS_PMCDR field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24 TSEC1	Three-speed Ethernet controller 1 disable clock in sleep mode. 0 eTSEC1 enabled clock 1 eTSEC1 disabled clock
25 TSEC2	Three-speed Ethernet controller 2 disable clock in sleep mode. 0 eTSEC2 enabled clock 1 eTSEC2 disabled clock
26 TSEC3	Three-speed Ethernet controller 3 disable clock in sleep mode.

Table continues on the next page...

GUTS_PMCDR field descriptions (continued)

Field	Description
	0 eTSEC3 enabled clock 1 eTSEC3 disabled clock
27–31 -	This field is reserved. Reserved

21.4.11 Machine check summary register (GUTS_MCPSUMR)

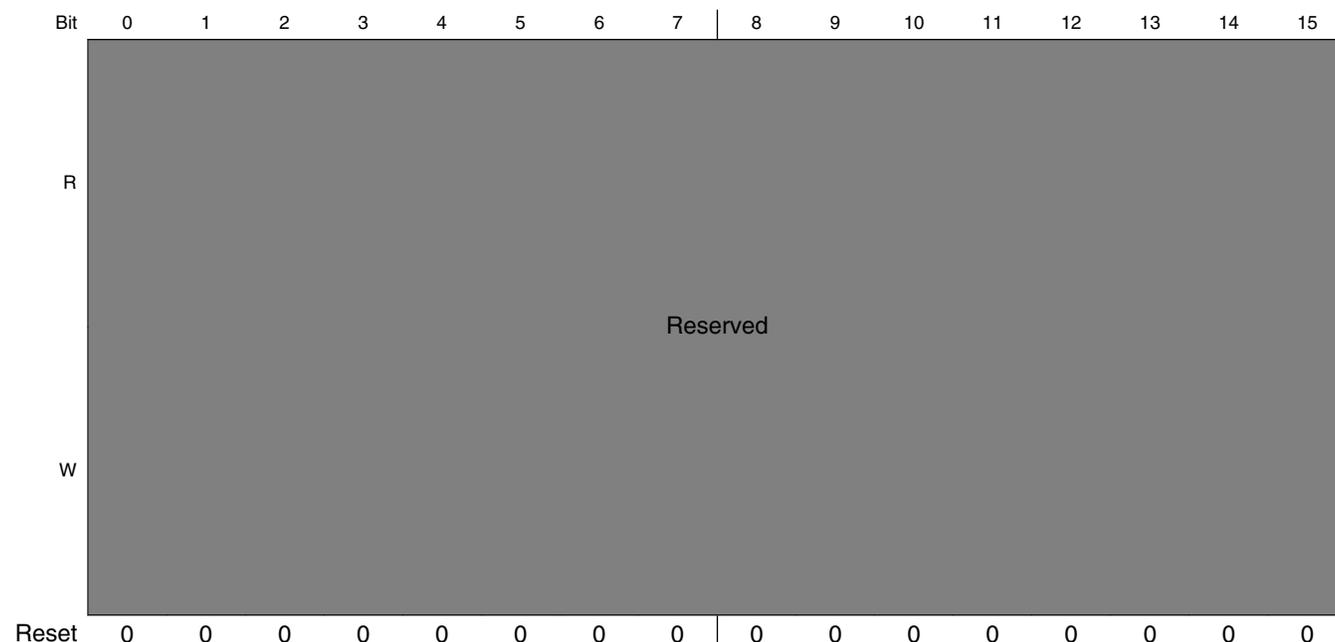
MCPSUMR contains bits summarizing some of the sources of a pending machine check interrupt. All MCPSUMR bits function as write-1-to-clear.

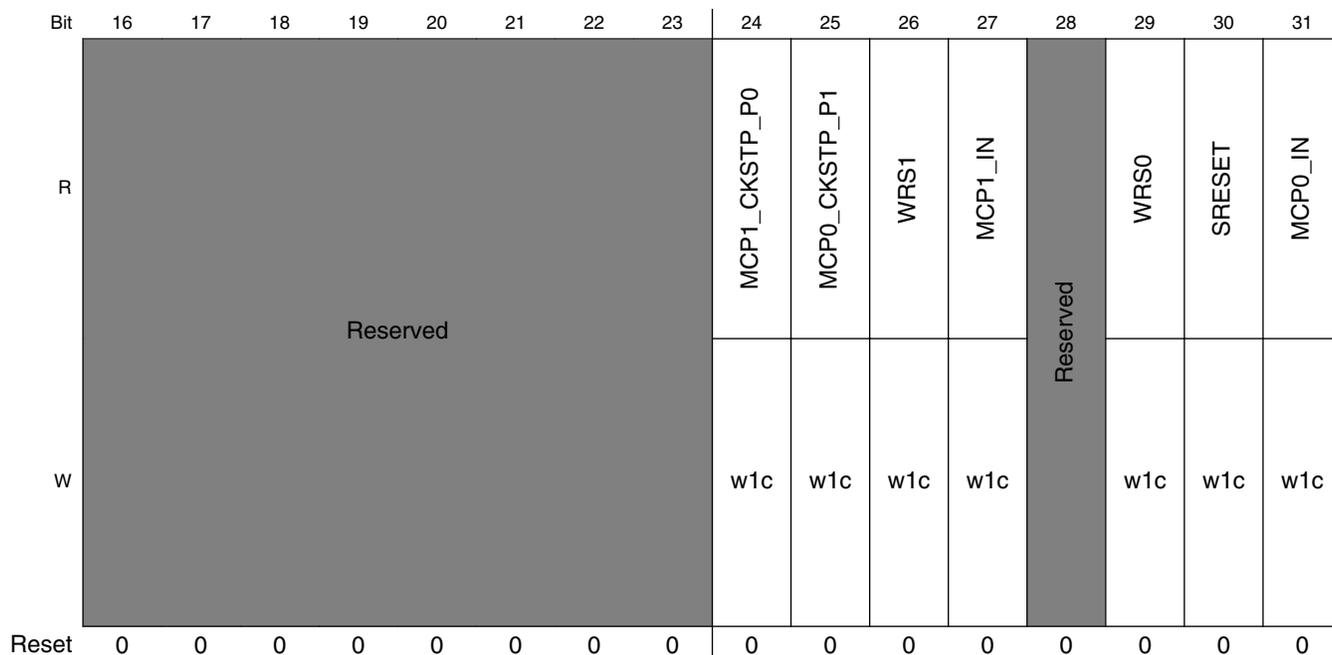
NOTE

Register fields designated as write-1-to-clear are cleared only by writing ones to them. Writing zeros to them has no effect.

Note that other conditions can cause a machine check condition not summarized in MCPSUMR. For example, uncorrectable read errors cause the assertion of *core_fault_in*, which may directly cause a machine check (if $HID1[RFXE] = 1$). If $RFXE = 0$, the assertion of *core_fault_in* does not directly cause a machine check interrupt, but must be handled by the block that generated the error.

Address: E_0000h base + 90h offset = E_0090h





GUTS_MCPSUMR field descriptions

Field	Description
0–23 -	This field is reserved. Reserved
24 MCP1_CKSTP_P0	Machine check to core 1 in response to checkstop from core 0 0 Machine check exception was not caused by checkstop of core 0 1 Machine check exception was caused by checkstop of core 0
25 MCP0_CKSTP_P1	Machine check to core 0 in response to checkstop from core 1 0 Machine check exception was not caused by checkstop of core 1 1 Machine check exception was caused by checkstop of core 1
26 WRS1	Core 1 watchdog timer machine check 0 Machine check exception was not caused by watchdog timer. 1 Machine check was caused by a soft reset condition from the e500_1 watchdog timer as configured in the core1's TSR. Specifically, TSR[WRS] = 01 and a watchdog reset condition occurred.
27 MCP1_IN	MCP1_B signal asserted 0 Machine check exception was not caused by MCP1_B assertion. 1 Machine check exception was caused by the assertion of the MCP1_B input signal.
28 -	This field is reserved. Reserved
29 WRS0	Core 0 watchdog timer machine check 0 Machine check exception was not caused by watchdog timer. 1 Machine check was caused by a soft reset condition from the e500_0 watchdog timer as configured in the core0's TSR. Specifically, TSR[WRS] = 01 and a watchdog reset condition occurred.
30 SRESET	Soft reset machine check

Table continues on the next page...

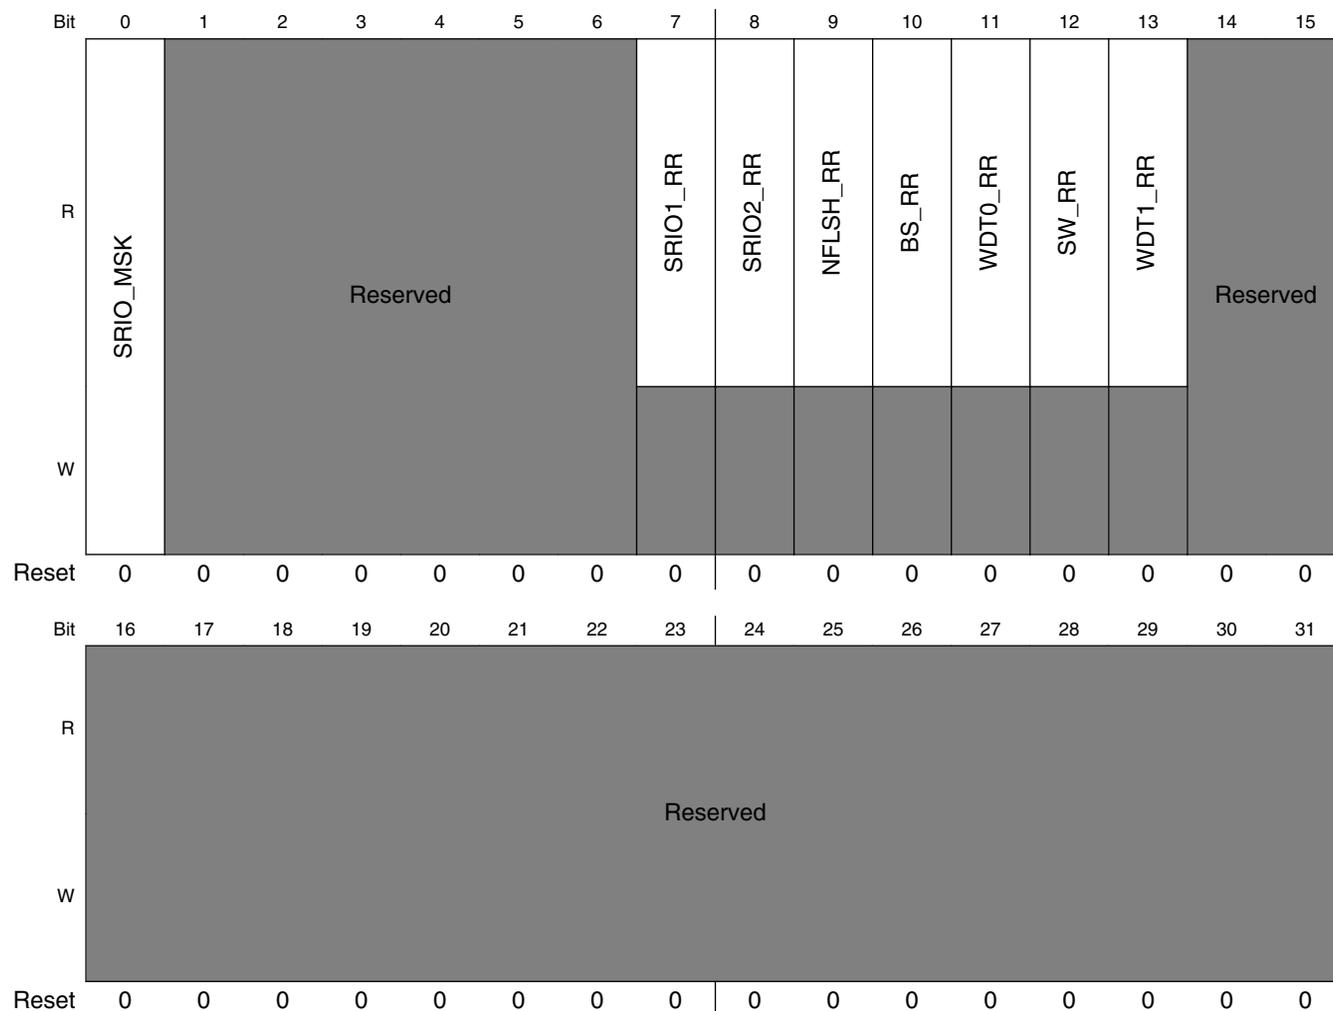
GUTS_MCPSUMR field descriptions (continued)

Field	Description
	0 Machine check exception was not caused by SRESET_B assertion. 1 Machine check exception was caused by the assertion of the SRESET_B input signal.
31 MCP0_IN	MCP0_B signal asserted 0 Machine check exception was not caused by MCP0_B assertion. 1 Machine check exception was caused by the assertion of the MCP0_B input signal.

21.4.12 Reset request status and control register (GUTS_RSTRSCR)

The RSTRSCR contains mask bits for RapidIO and status for boot sequencer, watchdog timer, RapidIO, and a software settable reset request bit.

Address: E_0000h base + 94h offset = E_0094h



GUTS_RSTRSCR field descriptions

Field	Description
0 SRIO_MSK	Serial RapidIO (SRIO) reset request mask This bit masks the reset request for both SRIO1 and SRIO2.
1–6 -	This field is reserved. Reserved
7 SRIO1_RR	Serial RapidIO 1 (SRIO1) reset request
8 SRIO2_RR	Serial RapidIO 2 (SRIO2) reset request
9 NFLSH_RR	NAND Flash ECC error during boot reset request.
10 BS_RR	Boot sequence reset request
11 WDT0_RR	Watchdog timer reset request in core 0. Occurs when TSR[WRS] = 10 for core 0 and a watchdog reset condition is reached.
12 SW_RR	Software settable reset request
13 WDT1_RR	Watchdog timer reset request in core 1. Occurs when TSR[WRS] = 10 for core 1 and a watchdog reset condition is reached.
14–31 -	This field is reserved. Reserved

21.4.13 Exception reset control register (GUTS_ECTRSTCR)

The ECTRSTCR contains control bits for the exception reset of core 0 and core 1 in response to checkstop.

Address: E_0000h base + 98h offset = E_0098h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GUTS_ECTRSTCR field descriptions

Field	Description
0 RST_CKSTP_P0_EN	Enable automatic reset of core 0 in response to checkstop
1 RST_CKSTP_P1_EN	Enable automatic reset of core 1 in response to checkstop
2-3 -	This field is reserved. Reserved
4 CKSTP_OUT0_DIS	Disable assertion of CKSTP_OUT0_B pin
5 CKSTP_OUT1_DIS	Disable assertion of CKSTP_OUT1_B pin
6-7 -	This field is reserved. Reserved
8 MCP0_CKSTP_P1_EN	Enable machine check to core 0 in response to check stop from core 1

Table continues on the next page...

GUTS_ECTRSTCR field descriptions (continued)

Field	Description
9 MCP1_CKSTP_ P0_EN	Enable machine check to core 1 in response to check stop from core 0
10–31 -	This field is reserved. Reserved

21.4.14 Automatic reset status register (GUTS_AUTORSTSR)

The AUTORSTSR contains the automatic reset status bits for core 0 and core 1.

Address: E_0000h base + 9Ch offset = E_009Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	RST_CKSTP_P0	RST_CKSTP_P1	Reserved						RST_MPIC_P0	RST_MPIC_P1	Reserved		RST_CORE_P0	RST_CORE_P1	Reserved	
W	w1c	w1c							w1c	w1c			w1c	w1c		
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	READY_P0	READY_P1	Reserved													
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GUTS_AUTORSTSR field descriptions

Field	Description
0 RST_CKSTP_P0	Core 0 was reset in response to check stop
	0 No reset
	1 Reset occurred.

Table continues on the next page...

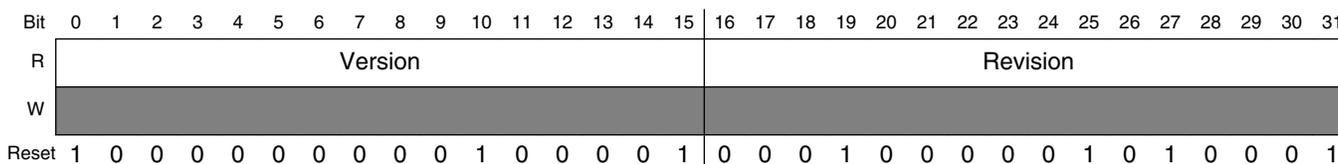
GUTS_AUTORSTSR field descriptions (continued)

Field	Description
1 RST_CKSTP_P1	Core 1 was reset in response to check stop 0 No reset 1 Reset occurred.
2-7 -	This field is reserved. Reserved
8 RST_MPIC_P0	Core 0 was reset in response to MPIC reset request 0 No reset 1 Reset occurred
9 RST_MPIC_P1	Core 1 was reset in response to MPIC reset request 0 No reset 1 Reset occurred
10-11 -	This field is reserved. Reserved
12 RST_CORE_P0	Core 0 was reset in response to internal core request to reset itself by setting bit DBCR0[RST] register 0 No reset 1 Reset occurred
13 RST_CORE_P1	Core 1 was reset in response to internal core request to reset itself by setting bit DBCR0[RST] register 0 No reset 1 Reset occurred
14-15 -	This field is reserved. Reserved
16 READY_P0	Core 0 ready pin. This bit reflects what is driven on the READY_P0 external signal. 0 Core 0 not ready 1 Core 0 ready
17 READY_P1	Core 1 ready pin. This bit reflects what is driven on the READY_P1 external signal. 0 Core 1 not ready 1 Core 1 ready
18-31 -	This field is reserved. Reserved

21.4.15 Processor version register (GUTS_PVR)

The PVR contains the e500 processor version number. It is a memory-mapped copy of the PVR in the e500 core (and is therefore accessible to external devices). [Processor version register \(PVR\) and system version register \(SVR\)](#) lists the complete values for the chip.

Address: E_0000h base + A0h offset = E_00A0h



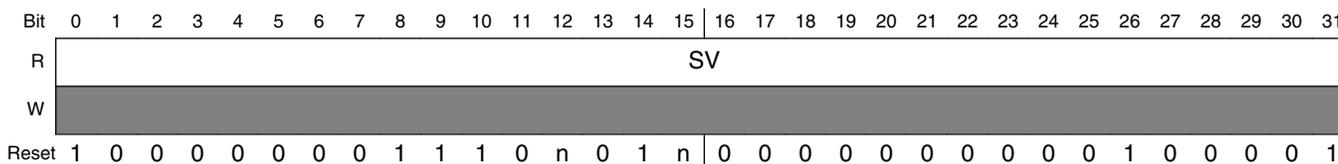
GUTS_PVR field descriptions

Field	Description
0–15 Version	A 16-bit number that identifies the version of the processor. Different version numbers indicate major differences between processors, such as which optional facilities and instructions are supported. (See Processor version register (PVR) and system version register (SVR) for specific values.)
16–31 Revision	A 16-bit number that distinguishes between implementations of the version. Different revision numbers indicate minor differences between processors having the same version number, such as clock rate and engineering change level. (See Processor version register (PVR) and system version register (SVR) for specific values.)

21.4.16 System version register (GUTS_SVR)

The SVR contains the system version number for the chip implementation. This value can also be read through the SVR SPR of the e500 core. [Processor version register \(PVR\) and system version register \(SVR\)](#) lists the complete values for the chip.

Address: E_0000h base + A4h offset = E_00A4h



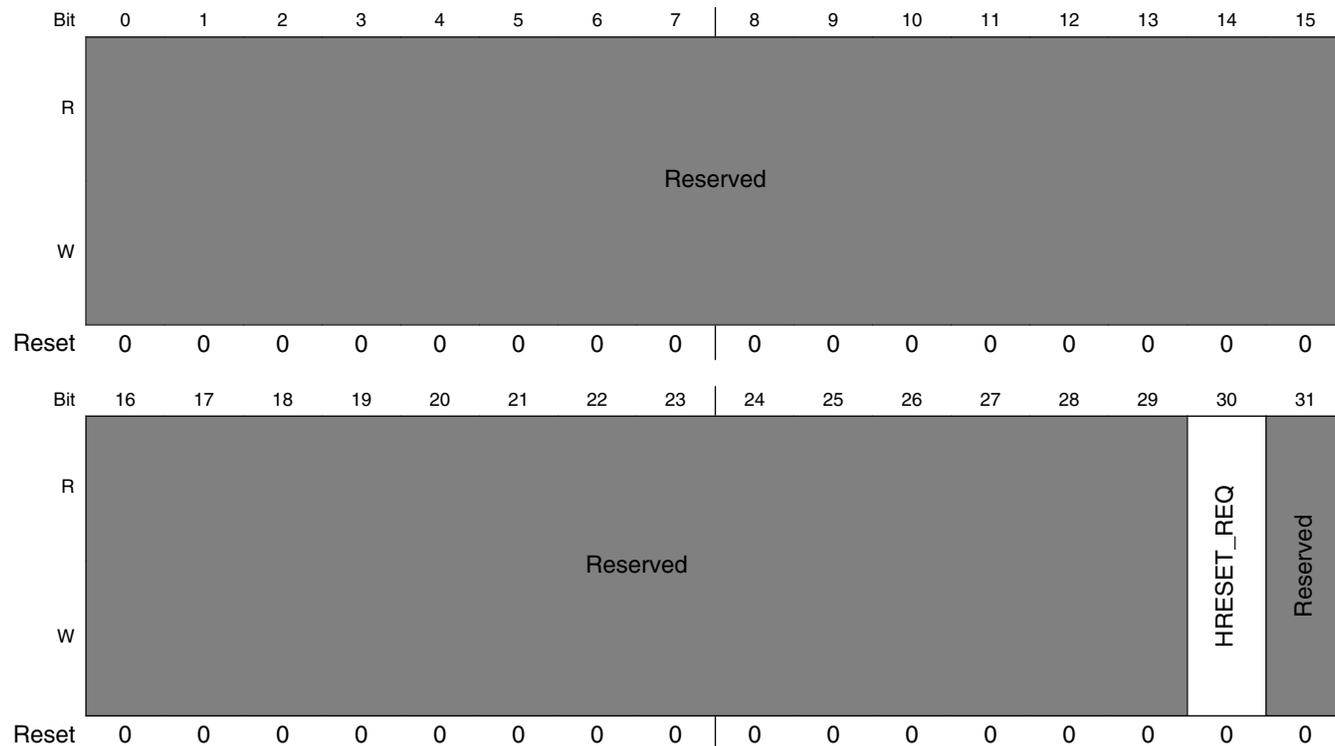
GUTS_SVR field descriptions

Field	Description
0–31 SV	System version numbers for the chip system logic 0x80E2_0021 for P2020 without security 0x80EA_0021 for P2020 with security, 0x80EB_0021 for P2010 with security 0x80E3_0021 for P2010 without security

21.4.17 Reset control register (GUTS_RSTCR)

The RSTCR contains the reset control bits.

Address: E_0000h base + B0h offset = E_00B0h



GUTS_RSTCR field descriptions

Field	Description
0–29 -	This field is reserved. Reserved
30 HRESET_REQ	Hardware reset request

Table continues on the next page...

GUTS_RSTCR field descriptions (continued)

Field	Description
31 -	This field is reserved. Reserved

21.4.18 IO voltage select status register (GUTS_IOVSELSR)

The IOVSELSR contains voltage configuration status bits for the chip driver/receivers on the BVDD (local bus, GPIO[8:15]), CVDD (eSDHC, SPI, USB), and LVDD (eTSEC1, eTSEC2, eTSEC3) power planes. Voltage selection is achieved by tying the associated dedicated configuration inputs: LVDD_VSEL, BVDD_VSEL[0:1], and CVDD_VSEL[0:1] appropriately.

Address: E_0000h base + C0h offset = E_00C0h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	Reserved						CVDDV		Reserved		BVDDV		Reserved		LVDDV		
W	Reserved						CVDDV		Reserved		BVDDV		Reserved		LVDDV		
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

GUTS_IOVSELSR field descriptions

Field	Description
0–21 -	This field is reserved. Reserved
22–23 CVDDV	Indicates the I/O voltage for the Secure Digital Bus, SPI, and USB on the CVDD supply. 00 3.3 V 01 2.5 V 10 1.8 V 11 3.3 V
24–25 -	This field is reserved. Reserved
26–27 BVDDV	Indicates the I/O voltage for the local bus and GPIO[8:15] on the BVDD supply. 00 3.3 V 01 2.5 V 10 1.8 V 11 3.3 V

Table continues on the next page...

GUTS_IOVSELSR field descriptions (continued)

Field	Description
28–29 -	This field is reserved. Reserved
30–31 LVDDV	Indicates the I/O voltage for eTSEC1, eTSEC2, eTSEC3, and Ethernet management interfaces on the LVDD supply. Patterns not shown are reserved. 00 3.3 V 01 2.5 V

21.4.19 DDR clock disable register (GUTS_DDRCLKDR)

The DDRCLKDR contains bits that allow disabling the clocks of the DDR SDRAM controller.

Address: E_0000h base + B28h offset = E_0B28h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										DDR_MCK0_	DDR_MCK1_	DDR_MCK2_	DDR_MCK3_	DDR_MCK4_	DDR_MCK5_
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GUTS_DDRCLKDR field descriptions

Field	Description
0–25 -	This field is reserved. Reserved
26 DDR_MCK0_DIS	DDR controller clock 0 disable 0 MCK0 is enabled. 1 MCK0 is disabled.
27 DDR_MCK1_DIS	DDR controller clock 1 disable 0 MCK1 is enabled. 1 MCK1 is disabled.

Table continues on the next page...

GUTS_DDRCLKDR field descriptions (continued)

Field	Description
28 DDR_MCK2_DIS	DDR controller clock 2 disable 0 MCK2 is enabled. 1 MCK2 is disabled.
29 DDR_MCK3_DIS	DDR controller clock 3 disable 0 MCK3 is enabled. 1 MCK3 is disabled.
30 DDR_MCK4_DIS	DDR controller clock 4 disable 0 MCK4 is enabled. 1 MCK4 is disabled.
31 DDR_MCK5_DIS	DDR controller clock 5 disable 0 MCK5 is enabled. 1 MCK5 is disabled.

21.4.20 Clock out control register (GUTS_CLKOCR)

The CLKOCR contains control bits that select the clock sources to be placed on the clock out (CLK_OUT) signal.

Address: E_0000h base + E00h offset = E_0E00h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved										CLK_SEL					
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GUTS_CLKOCR field descriptions

Field	Description
0 ENB	Clock out enable 0 CLK_OUT signal is three-stated 1 CLK_OUT signal is driven according to CLKOCR[CLK_SEL]
1–25 -	This field is reserved. Reserved
26–31 CLK_SEL	Clock out select. Patterns not shown are reserved. 000000 CCB (platform) clock

Table continues on the next page...

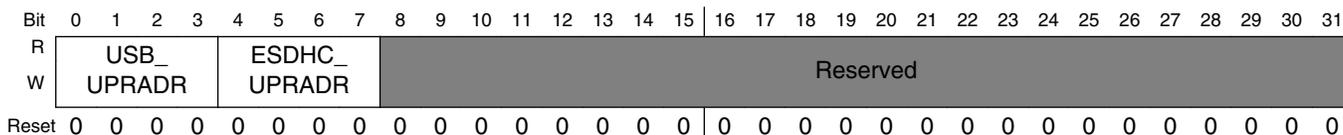
GUTS_CLKOCR field descriptions (continued)

Field	Description
000001	CCB (platform) clock divided by 2
000010	SYSCLK (echoes SYSCLK input)
000011	SYSCLK divided by 2 (demonstrates platform PLL lock)

21.4.21 ECM control register (GUTS_ECMCR)

The ECMCR contains the four uppermost bits of the USB and eSDHC address bus for transactions initiated by these blocks.

Address: E_0000h base + E20h offset = E_0E20h



GUTS_ECMCR field descriptions

Field	Description
0–3 USB_UPRADR	The uppermost bits of the USB address bus for all transactions initiated by the USB.
4–7 ESDHC_UPRADR	The uppermost bits of the ESDHC address bus for all transactions initiated by the ESDHC.
8–31 -	This field is reserved. Reserved

21.4.22 SDHC debug control register (GUTS_SDHCDCR)

The SDHCDCR enables the SDHC card-detect and write detect polarities to be inverted. The reset value *n* denotes that the value is based on a POR configuration.

Address: E_0000h base + F64h offset = E_0F64h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R																	
W	CD_	Reserved															
IN	INV																
Reset	n	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	Reserved																
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

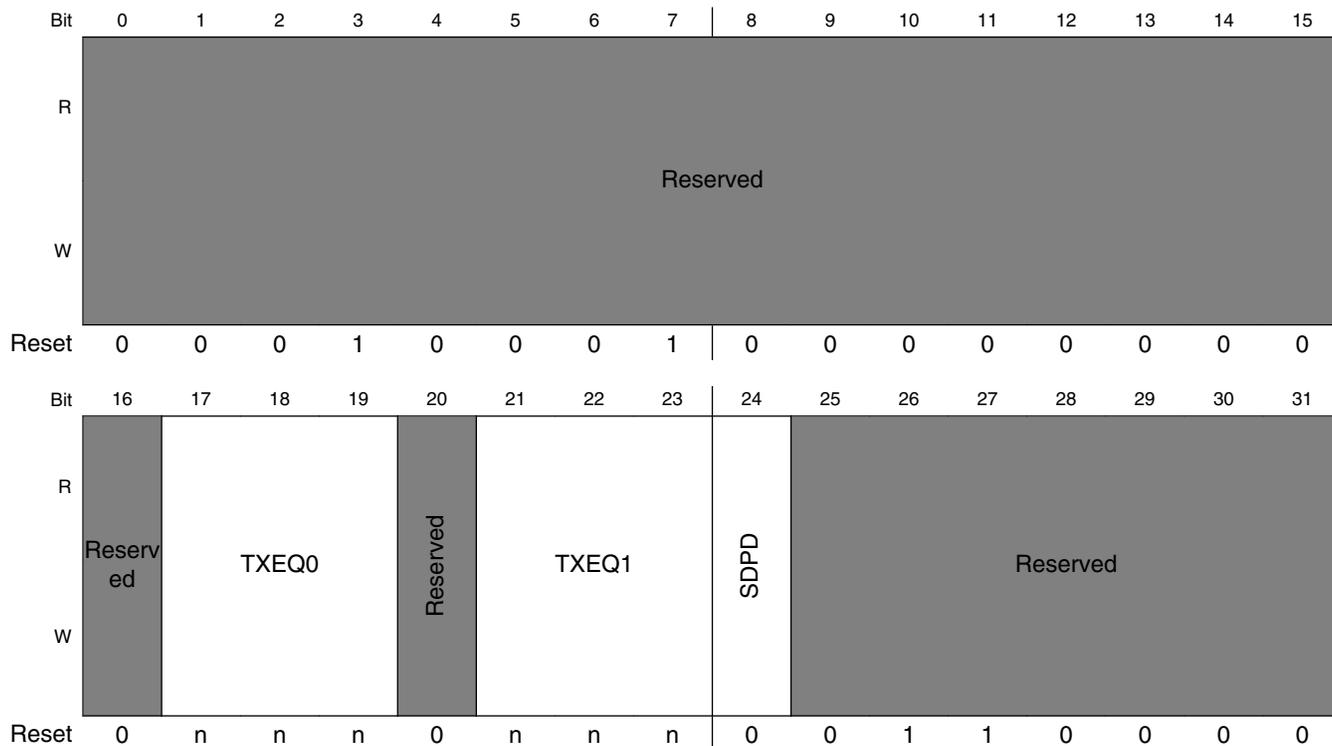
GUTS_SDHCDCR field descriptions

Field	Description
0 CD_INV	Card-detect inversion 0 No inversion of SDHC card-detect 1 Invert SDHC card-detect
1–31 -	This field is reserved. Reserved

21.4.23 SRDS Control Register 0 (GUTS_SRDSR0)

The SRDSR0 contains the functional control bits for the SerDes logic.

Address: E_0000h base + 3000h offset = E_3000h



GUTS_SRDSR0 field descriptions

Field	Description
0–16 -	This field is reserved. Reserved
17–19 TXEQ0	Transmit equalization selection bus for lane 0. Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane 0. If register field SRDSR5[SDTXL0] = 000, then the equalization definitions are: 000 No equalization 001 1.09x relative amplitude 010 1.2x relative amplitude 011 1.33x relative amplitude 100 1.5x relative amplitude 101 1.71x relative amplitude 110 2.0x relative amplitude If register field SRDSR5[SDTXL0] = 011 (0.750 full swing), then the equalization definitions are:

Table continues on the next page...

GUTS_SRDSCR0 field descriptions (continued)

Field	Description
	<p>000 No equalization</p> <p>001 1.13x relative amplitude</p> <p>010 1.29x relative amplitude</p> <p>011 1.5x relative amplitude</p> <p>100 1.8x relative amplitude</p> <p>101 Reserved</p> <p>110 Reserved</p> <p>111 Reserved</p> <p>Recommended setting per protocol:</p> <ul style="list-style-type: none"> • PCI Express: 100 • SRIO: 011
20 -	<p>This field is reserved.</p> <p>Reserved</p>
21–23 TXEQ1	<p>Transmit equalization selection bus for lane 1. Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane 1</p> <p>Patterns not shown below are reserved.</p> <p>If register field SRDSCR5[SCTXL1] = 000, then the equalization definitions are:</p> <p>000 No equalization</p> <p>001 1.09x relative amplitude</p> <p>010 1.2x relative amplitude</p> <p>011 1.33x relative amplitude</p> <p>100 1.5x relative amplitude</p> <p>101 1.71x relative amplitude</p> <p>110 2.0x relative amplitude</p> <p>111 Reserved</p> <p>If register field SRDSCR5[SCTXL1] = 011(0.750 full swing), then the equalization definitions are:</p> <p>000 No equalization</p> <p>001 1.13x relative amplitude</p> <p>010 1.29x relative amplitude</p> <p>011 1.5x relative amplitude</p> <p>100 1.8x relative amplitude</p> <p>Recommended setting per protocol:</p> <ul style="list-style-type: none"> • PCI Express: 100 • SRIO: 011
24 SDPD	<p>SerDes power down. This power down signal shuts down the PLL, all of the receiver amplifiers, all of the samplers and places the transmitters in three-state.</p> <p>Recommended setting: 0</p> <p>0 Application mode</p> <p>1 Block power down</p>

Table continues on the next page...

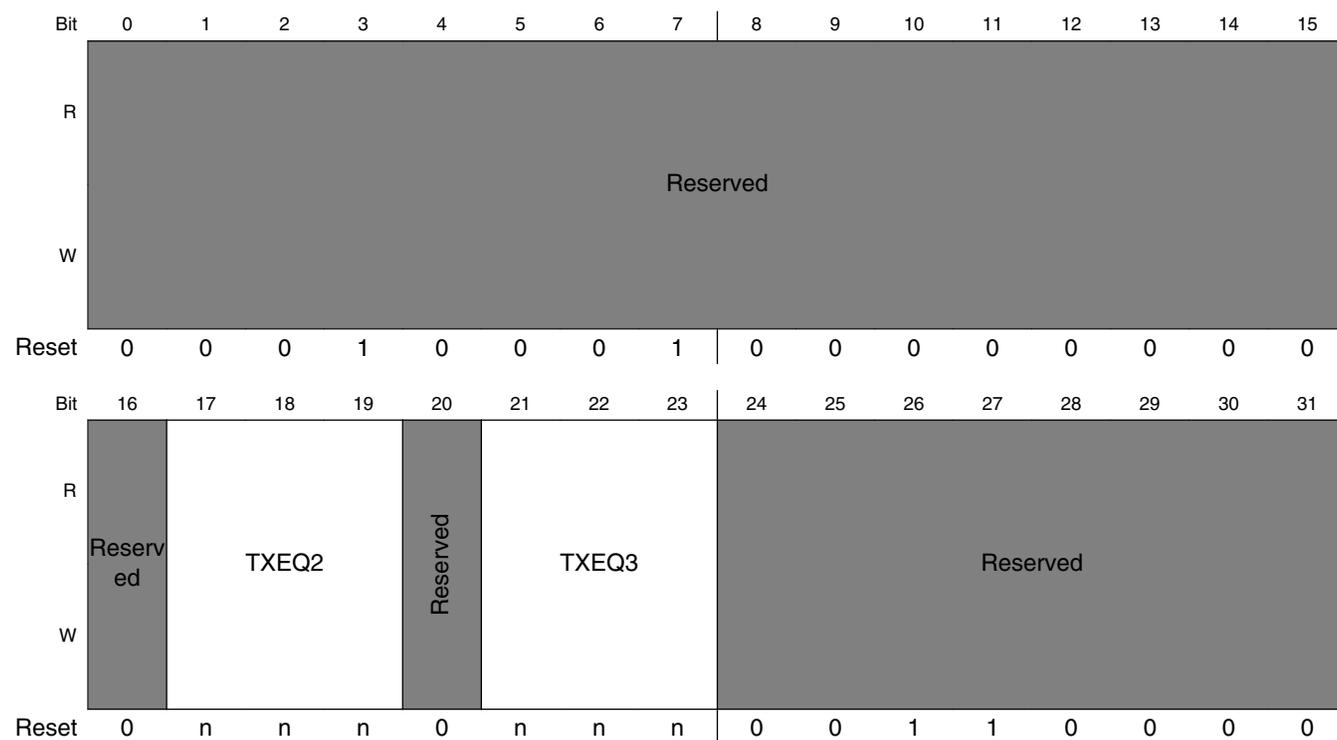
GUTS_SRDSCR0 field descriptions (continued)

Field	Description
25–31 -	This field is reserved. Reserved

21.4.24 SRDS Control Register 1 (GUTS_SRDSCR1)

The SRDSCR1 contains the functional control bits for the SerDes logic.

Address: E_0000h base + 3004h offset = E_3004h



GUTS_SRDSCR1 field descriptions

Field	Description
0–16 -	This field is reserved. Reserved
17–19 TXEQ2	Transmit equalization selection bus for lane 2. Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane 2. If register field SRDSCR6[SDTXL2] = 000, then the equalization definitions are: 000 No equalization 001 1.09x relative amplitude 010 1.2x relative amplitude

Table continues on the next page...

GUTS_SRDSCR1 field descriptions (continued)

Field	Description
	<p>011 1.33x relative amplitude 100 1.5x relative amplitude 101 1.71x relative amplitude 110 2.0x relative amplitude 111 Reserved</p> <p>If register field SRDSCR6[SDTXL2] = 011(0.750 full swing), then the equalization definitions are:</p> <p>000 No equalization 001 1.13x relative amplitude 010 1.29x relative amplitude 011 1.5x relative amplitude 100 1.8x relative amplitude 101 Reserved 110 Reserved 111 Reserved</p> <p>Recommended setting per protocol:</p> <ul style="list-style-type: none"> • PCI Express: 100 • SGMII: 100 • SRIO: 011
20 -	This field is reserved. Reserved
21–23 TXEQ3	<p>Transmit equalization selection bus for lane 3. Sets the peak value for output swing of transmitters and the amount of transmit equalization for lane 3</p> <p>If register field SRDSCR6[SDTXL3] = 000, then the equalization definitions are:</p> <p>000 No equalization 001 1.09x relative amplitude 010 1.2x relative amplitude 011 1.33x relative amplitude 100 1.5x relative amplitude 101 1.71x relative amplitude 110 2.0x relative amplitude 111 Reserved</p> <p>If register field SRDSCR6[SDTXL3] = 011(0.750 full swing), then the equalization definitions are:</p> <p>000 No equalization 001 1.13x relative amplitude 010 1.29x relative amplitude 011 1.5x relative amplitude 100 1.8x relative amplitude 101 Reserved 110 Reserved</p>

Table continues on the next page...

GUTS_SRDSCR1 field descriptions (continued)

Field	Description
	Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 100 • SGMII: 100 • SRIO: 011
24–31 -	This field is reserved. Reserved

21.4.25 SRDS Control Register 2 (GUTS_SRDSCR2)

The SRDSCR2 contains the functional control bits for the SerDes logic.

Individual lanes can be powered down using SRDSCR2[0:7]. It requires the entire SerDes to reset in order to activate a lane from powerdown.

Address: E_0000h base + 3008h offset = E_3008h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W																
	PD0	PD1	Reserved	PD2	PD3	Reserved			X3S0	X3S1	Reserved	X3S2	X3S3			
Reset	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved															
Reset	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	0

GUTS_SRDSCR2 field descriptions

Field	Description
0 PD0	Lane 0 power down 0 Normal 1 Power down Lane 0
1 PD1	Lane 1 power down 0 Normal 1 Power down Lane 1
2–3 -	This field is reserved. Reserved
4 PD2	Lane 2 power down 0 Normal 1 Power down Lane 2
5 PD3	Lane 3 power down 0 Normal 1 Power down Lane 3

Table continues on the next page...

GUTS_SRDSCR2 field descriptions (continued)

Field	Description
6–9 -	This field is reserved. Reserved
10 X3S0	Lane 0 transmitter three-state This is a configuration bit and reading this bit do not reflect the current state of lane. 0 Normal 1 The transmitter output is disabled and place in a three-state condition
11 X3S1	Lane 1 transmitter three-state This is a configuration bit and reading this bit do not reflect the current state of lane. 0 Normal 1 The transmitter output is disabled and place in a three-state condition
12–13 -	This field is reserved. Reserved
14 X3S2	Lane 2 transmitter three-state This is a configuration bit and reading this bit do not reflect the current state of lane. 0 Normal 1 The transmitter output is disabled and place in a three-state condition
15 X3S3	Lane 3 transmitter three-state This is a configuration bit and reading this bit do not reflect the current state of lane. 0 Normal 1 The transmitter output is disabled and place in a three-state condition
16–31 -	This field is reserved. Reserved

21.4.26 SRDS Control Register 4 (GUTS_SRDSCR4)

The SRDSCR4 contains the functional control bits used for the SerDes logic.

Address: E_0000h base + 3010h offset = E_3010h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15	
R	Reserved																	
W	Reserved																	
Reset	0	0	0	0	n	n	0	0		0	0	0	0	0	0	0	0	
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31	
R	Reserved			LSTS2				Reserved				LSTS3			Reserved			
W	Reserved			LSTS2				Reserved				LSTS3			Reserved			
Reset	0	0	0	n	n	n	0	0		0	0	0	0	n	n	n	0	0

GUTS_SRDCR4 field descriptions

Field	Description
0–18 -	This field is reserved. Reserved
19–21 LSTS2	<p>Receiver electrical idle detection control for Lane 2.</p> <p>For PCI Express and SRIO protocols, the recommended setting below should be used. For SGMII protocol, when desired, changing this bit field from the default setting of 001 to 100 increases the SerDes Lane2 receiver's Vrx_diffp-p level on electrical idle detection. Refer to the VLOS parameter (Loss of signal threshold) in the "SGMII DC Receiver Timing Specification" section of the device hardware specifications , for related detail information. Patterns not shown are reserved.</p> <p>000 Loss of signal detect function is disabled.</p> <p>001 Default SGMII levels (Low = 30 mV, High = 100 mV)</p> <p>100 Default PCI Express levels (Low = 65 mV, High = 175mV)</p> <p>Recommended setting per protocol:</p> <p>PCI Express: 100</p> <p>SGMII: 001</p> <p>SRIO: 000</p>
22–26 -	This field is reserved. Reserved
27–29 LSTS3	<p>Receiver electrical idle detection control for Lane 3.</p> <p>For PCI Express and SRIO protocols, the recommended setting below should be used. For SGMII protocol, when desired, changing this bit field from the default setting of 001 to 100 increases the SerDes Lane3 receiver's Vrx_diffp-p level on electrical idle detection. Refer to the VLOS parameter (Loss of signal threshold) in the "SGMII DC Receiver Timing Specification" section of the device hardware specifications , for related detail information. Patterns not shown are reserved.</p> <p>000 Loss of signal detect function is disabled.</p> <p>001 Default SGMII levels (Low = 30 mV, High = 100 mV)</p> <p>100 Default PCI Express (Low = 65 mV, High = 175mV)</p> <p>Recommended Setting per protocol:</p> <p>PCI Express: 100</p> <p>SGMII: 001</p> <p>SRIO: 000</p>
30–31 -	This field is reserved. Reserved

21.4.27 SRDS Control Register 5 (GUTS_SRDSCR5)

The SRDSCR5 contains the functional control bits for the SerDes logic.

Address: E_0000h base + 3014h offset = E_3014h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	Reserved															SDTXL0			Reserved				SDTXL1										
W	Reserved															SDTXL0			Reserved				SDTXL1										
Reset	0	0	0	0	0	0	0	n	n	0	0	0	0	0	0	n	n	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GUTS_SRDSCR5 field descriptions

Field	Description
0–20 -	This field is reserved. Reserved
21–23 SDTXL0	Controls Lane 0 transmitter amplitude levels. Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 000 • SGMII: 000 • SRIO: 000 000 No amplitude reduction (default) 001 Reserved 010 Reserved 011 0.750 full swing 100 Reserved 101 Reserved 110 Reserved 111 Reserved
24–28 -	This field is reserved. Reserved
29–31 SDTXL1	Controls Lane 1 transmitter amplitude levels. Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 000 • SGMII: 000 • SRIO: 000 000 No amplitude reduction (default) 001 Reserved 010 Reserved 011 0.750 full swing 100 Reserved 101 Reserved 110 Reserved 111 Reserved

21.4.28 SRDS Control Register 6 (GUTS_SRDSCR6)

The SRDSCR6 contains the functional control bits for the SerDes logic.

Address: E_0000h base + 3018h offset = E_3018h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
R	Reserved															SDTXL2			Reserved				SDTXL3										
W	Reserved															SDTXL2			Reserved				SDTXL3										
Reset	0	0	0	0	0	0	0	n	n	0	0	0	0	0	0	n	n	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

GUTS_SRDSCR6 field descriptions

Field	Description
0–20 -	This field is reserved. Reserved
21–23 SDTXL2	Controls Lane 2 transmitter amplitude levels. Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 000 • SGMII: 000 • SRIO: 000 000 No amplitude reduction (default) 001 Reserved 010 Reserved 011 0.750 full swing 100 Reserved 101 Reserved 110 Reserved 111 Reserved
24–28 -	This field is reserved. Reserved
29–31 SDTXL3	Controls Lane 3 transmitter amplitude levels. Recommended setting per protocol: <ul style="list-style-type: none"> • PCI Express: 000 • SGMII: 000 • SRIO: 000 000 No amplitude reduction (default) 001 Reserved 010 Reserved 011 0.750 full swing 100 Reserved 101 Reserved 110 Reserved 111 Reserved

21.5 Functional description

This section describes the global utilities from a functional perspective.

21.5.1 Power management

The chip has features to minimize power consumption at several levels.

Dynamic power management locally minimizes power consumption when a block is idle. Software can also shut down clocks to individual blocks when they are not needed through a memory-mapped register (DEVDISR). Additionally, software running on the e500 cores can access the cores's SPRs to put the device into doze or nap power down state. Finally, software can access a memory-mapped register (POWMGTCCR) in the global utilities block to put the device in the doze or sleep states.

Note that the software that writes to either DEVDISR or POWMGTCR can be running either on the e500 cores or on an external master that can write to the chip memory-mapped registers through the PCI Express interfaces.

These features are described in further detail in this section.

21.5.1.1 Relationship between both cores and device power management states

The chip has three low-power states: doze, nap, and sleep.

The mapping of both cores and device power management states shows state transitions from the perspective of the e500 cores.

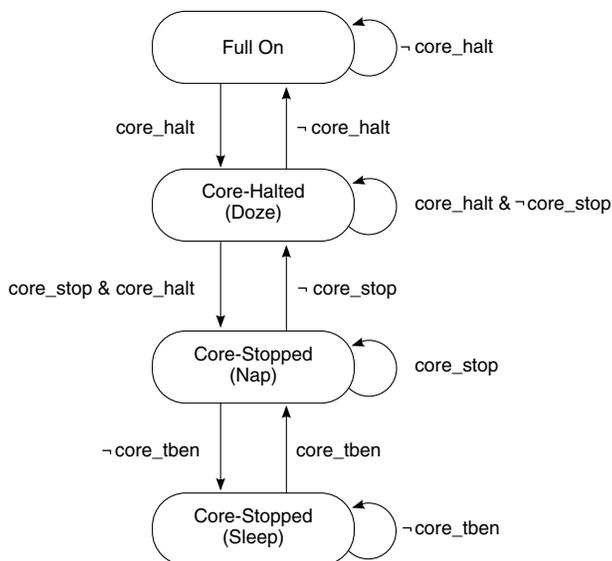


Figure 21-29. e500 cores power management state diagram

For each operating state represented in the diagram, the cores's state is listed first, with the corresponding state of the device shown beneath it in parenthesis. Note that there are many other variables that control the state transitions between device power management states. These additional variables are described in more detail in [Power-down sequence coordination](#).

This table lists basic characteristics of the low-power modes and the full on mode.

Table 21-31. Power management modes-basic description

Mode	Description	Core responds To		Signal states	
		Snoop	Interrupts	READY_P0 or READY_P1	ASLEEP
Full On	All units operating normally.	Yes	Yes	Asserted	Negated
Doze0/1	Core 0 or stops dispatching new instructions (core 0 or 1 is halted)	Yes	Yes	Negated	Negated
Nap0/1	Core 0 or 1 is stopped with clocks off except to time base Should flush data cache before entering	No	Yes	Negated	Negated
Sleep	Core is stopped with clocks off. Clocks powered down to all blocks (including core time base) except to the interrupt controller (PIC) unit	No	Yes	Negated	Asserted

21.5.1.2 CKSTP_IN0_B/CKSTP_IN1_B is not power management

CKSTP_IN0_B/CKSTP_IN1_B are not described here because they are not considered power management signals, although asserting these do stop the cores and a stopped core is technically in a low-power mode.

CKSTP_IN0_B/CKSTP_IN1_B are described in [Detailed signal descriptions](#).

21.5.1.3 Dynamic power management

Many blocks in the chip can dynamically turn off clocks within the block when sections of the block are idle.

This feature is always enabled and occurs automatically.

21.5.1.4 Shutting down unused blocks

As described in [Device disable control register \(GUTS_DEVDISR\)](#), DEVDISR provides a way to shut down certain functional blocks within the chip when they are not needed in a particular system.

DEVDISR can be written by the both e500 cores or by an external master. Powering down a block in this way turns off all clocks to that block.

DEVDISR was designed with the expectation that, once initialized by software, it would be modified only by a hard system reset (HRESET_B). It is recommended that this register be written only during system initialization. Blocks disabled by DEVDISR must not be re-enabled without a hard reset. (Setting DEVDISR[TB0/1] disables the cores's timer facilities, and setting DEVDISR[E500_0/1] places the cores in the core_stopped state in which they do not respond to interrupts.) The results of re-enabling previously disabled blocks (by clearing the corresponding DEVDISR field) without a hard reset are boundedly undefined.

NOTE

Functional blocks disabled using DEVDISR cannot respond to configuration accesses. Any access to configuration, control, and status registers of a disabled block is a programming error.

21.5.1.5 Software-controlled power-down states

e500 software can place the device in doze or nap power-down states by writing to HID0 in the core.

In addition, external masters can write to the memory-mapped POWMGTCR in the chip to cause the device to enter doze or sleep modes.

21.5.1.5.1 Doze mode

In doze mode, the given e500 core suspends instruction execution, significantly reducing the power consumption of the cores.

Snooping of the L1 data cache is still supported and thus the data in the data cache is kept coherent. Interrupts directed to the each core as described in [Interrupts to the e500 processor core](#), are monitored by the device and cause the chip to use the defined handshake mechanism to exit the cores from doze mode to allow the cores to recognize and process the interrupt; however, unless the interrupt subroutine turns off (or masks) the control bits that enabled doze mode (MSR[WE] and HID0[DOZE]), the device re-enters doze mode after the interrupt has been serviced. See [Interrupts and power management](#), for more information.

The given e500 core's timer facilities are still enabled during doze mode, and core time base interrupts can be generated. All device logic external to the core remains fully operational in doze mode.

21.5.1.5.2 Nap mode

In nap mode all clocks internal to the e500 core are turned off except for their timer facilities clock (the core time base).

The L1 caches do not respond to snoops in nap mode, so if coherency with external I/O transactions is required, the L1 cache must be flushed before entering nap mode.

Similar to doze mode, interrupts occurring in nap mode cause the device to wake up the e500 core in order to service the interrupt. However, unless the interrupt service routine changes the control bits that caused the device to enter nap mode (MSR[WE], and HID0[NAP]), the core returns to nap mode after the interrupt is serviced. See [Interrupts and power management](#), for more information.

All device logic external to the e500 core remains fully operational in nap mode.

21.5.1.5.3 Sleep mode

In sleep mode, all clocks internal to both e500 cores are turned off, including the timer facilities clock.

All I/O interfaces in the device logic are also shut down except for the three eTSECs. Only the clocks to the chip PIC and three eTSECs (if enabled in the PMCDR register) are still running so that an external interrupt, ethernet magic packet, or the PIC Global Timer can wake up the device.

After the core and I/O interfaces have shut down, ASLEEP is asserted and READY_P0 and READY_P1 are negated.

NOTE

Internal interrupt sources like the core interval timer or watchdog timer depend on an active clock for their operation and these are disabled in sleep mode.

21.5.1.6 Power management control fields

The e500 cores provide the following fields to signal power management requests to the chip device logic.

- MSR[WE]-Used to qualify the values of HID0[DOZE0/1, NAP] in the generation of the internal *doze* and *nap* signals.
- HID0[DOZEN]-Signals the chip to initiate doze mode for the given core.
- HID0[NAPn]-Signals the chip to initiate nap mode for the given core.

These register fields and their functional relationship are shown in [Figure 21-30](#). The *PowerPC e500 Core Reference Manual* has details on accessing these power management control bits.

An external master can also initiate power management requests by setting the applicable DOZ or SLP bits in the memory-mapped power management control and status register (POWMGTCSR). Because the core responds to snoops while dozing but not while napping, maintaining cache coherency requires significant preparation by the core before entering nap mode. For this reason only the core can initiate a nap during normal operation while other masters can initiate a doze.

21.5.1.7 Power-down sequence coordination

To preserve cache coherency and otherwise avoid loss of system state, a core's transition to low-power modes is coordinated by a set of handshaking signals, shown in the figure below, and protocols with all other chip functional blocks that respond to power-down requests.

The mode-transition protocol is executed automatically under these conditions. It is shown in [Figure 21-29](#) and described in [Table 21-32](#).

The column in [Table 21-32](#) showing the global utilities block as initiating a low-power mode corresponds to the external masters that can write to the POWMGTCR that resides in the global utilities block. For the chip, these are the PCI Express and serial RapidIO interfaces. However, note that the core can also write to POWMGTCR and, in this case, can initiate power management through the global utilities block.

Table 21-32. Power management entry protocol and initiating functional units

Low-power mode	Entry protocol	Initiating functional unit	
		Global utilities	Core
Doze	<ol style="list-style-type: none"> 1. Assert <i>core0/1_halt</i> input to core. 2. Wait for <i>core0/1_halted</i> handshake from core. 	√	√
Nap	<ol style="list-style-type: none"> 1. Follow doze protocol 2. Assert <i>core0/1_stop</i> input to core. 3. Wait for <i>core0/1_stopped</i> handshake from core. 	-	√
Sleep	<ol style="list-style-type: none"> 1. >Follow doze protocol; send stop requests to rest of device. 2. Follow nap protocol. 3. Wait for all interfaces to acknowledge stop requests. 4. Assert ASLEEP, negate READY, power down all clocks except to PIC unit. 	√	-

As shown in the figure below, the e500 cores enter low-power modes only in response to the *core0/1_halt*, *core0/1_stop*, or *core0/1_tben* inputs from the chip power management logic. These inputs may be prompted by the core (by setting the NAP0/1 or DOZE0/1 bits in the HID0 when enabled by setting MSR[WE]) or by an external master (by setting POWMGTCR[DOZ0/1,SLP]).

The following figure shows how all the clocking to the core timer facilities are disabled by clearing HID0[TB0/1EN]. When enabled, (HID0[TB0/1EN] = 1), the clock source is either the CCB clock divided by eight (the default) or a synchronized version of the RTC input.

21.5.1.8 Interrupts and power management

Whether low-power modes are automatically re-enabled after an interrupt is processed differs depending on whether the low power mode was entered due to a write to the core MSR[WE] bit or the low power mode was entered due to a write to POWMGTCR.

21.5.1.8.1 Interrupts and power management controlled by MSR[WE]

When an interrupt is asserted to the CPU, the core complex saves portions of the MSR to MCSRR1, CSRR1, or SRR1 (depending on the type of interrupt), and restores those values on return from the routine.

MSR[WE], which gates the *doze* and *nap* power management outputs (internal device signals) from the core complex, is always among the bits saved and restored; hence these outputs negate to the chip power management logic when the interrupt begins processing in the core. They return to their previous state when the core executes an **rfi**, **rfdi**, or **rfmci** instruction. [Interrupts to the e500 processor core](#) lists interrupts that cause the device to wake up.

NOTE

Returning *doze* and *nap* signals to their original state when MSR[WE] is restored differs from how power management is implemented on earlier PowerPC devices where MSR[POW], which enables power-down requests, is cleared when the processor exits a low-power state and is not automatically restored, as it is in Book E implementations.

21.5.1.8.2 Interrupts and power management controlled by POWMGTCR

The IRQ_MSK and CI_MSK fields of the POWMGTCR register prevent *int_b* interrupts or *cint_b* critical interrupts from waking the device from a low power state.

This is true regardless of the method used to enter the low power state.

Any unmasked interrupt (not masked by the mask bits in the POWMGTCR register) causes the POWMGTCR[DOZ0/1,SLP] fields to be cleared when it occurs. When such an interrupt occurs, the device returns to the normal operating mode and does not automatically attempt to return to a low power state after the interrupt is handled.

Note that interrupts caused by the unconditional debug event (UDE_B) and machine check (MCP_B) signals are not masked by the IRQ_MSK and CI_MSK fields; therefore, when these signals assert, the POWMGTCR[DOZ0/1,SLP] fields are cleared and the

device will return to full power operation. See [Power management control and status register \(GUTS_POWMGTCR\)](#), for detailed information about the bits of POWMGTCR.

Note also that unmasked interrupts that occur while the device is in the process of going into the sleep state (before sleep is completely attained) can also cause the device to clear the POWMGTCR[DOZ0/1,SLP] fields and return the device to full power operation.

21.5.1.9 Snooping in power-down modes

When the chip is in doze mode, the e500 core is in the core-halted state and it snoops its L1 caches and full coherency is maintained.

In deeper power-down modes, however, the e500 core does not respond to snoops.

The device does not perform dynamic bus snooping as described in the *e500 Reference Manual*. That is, when the e500 core is in the core-stopped state (which is the state of the core when the device is in either the nap or sleep state), the core is not awakened to perform snoops on global transactions. Therefore, before entering nap or sleep modes, the L1 caches should be flushed if coherency is required during these power-down modes.

21.5.1.10 Software considerations for power management

Setting MSR[WE] generates a request to the chip logic (external to the core complex) to enter a power saving state.

It is assumed that the desired power-saving state (doze or nap) was set up by setting the appropriate HIDO bit, typically at system start-up time. Setting WE has no direct effect on instruction execution, but is reflected on the internal *doze* and *nap* signals, depending on the HIDO settings. To ensure a clean transition into and out of a power-saving mode, the following program sequence is recommended:

```

sync
mtmsr (WE)
isync
loop:      br loop

```

21.5.1.11 Requirements for reaching and recovering from sleep state

In order to successfully reach the sleep state, I/O traffic to the device must be stopped.

The logic that controls the power down sequence waits for all I/O interfaces to become idle. In some applications this may happen eventually without actively shutting down interfaces, but most likely, software will have to take steps to shut down the eTSEC, serial RapidIO, and PCI Express interfaces-among others-before issuing the command (either the write to the core MSR[WE] as described above or writing to POWMGTCR) to put the device into sleep state.



Chapter 22

Device Performance Monitor

This chapter describes the device performance monitor facility, which can be used to monitor and optimize performance.

22.1 Introduction

This chapter describes the device performance monitor facility, which can be used to monitor and optimize performance.

The e500 core implements a separate performance monitor for strictly core-related behavior, such as instruction timing and L1 cache operations. This is described in the *PowerPC e500 Core Reference Manual (E500CORERM)*.

[Performance monitor events](#) briefly describes the events that can be monitored. Refer to the individual chapters for a better understanding of these events.

The device-level performance monitor facility can be used to monitor and record selected behaviors of the integrated device. Although the performance monitor described here is similar in many respects to the performance monitor facility implemented on the e500 core, it differs in that it is implemented using memory-mapped registers and it counts events outside the e500 core—for example, PCI, DDR, and L2 cache events.

Performance monitor counters (PMC0-PMC11) are used to count events selected by the performance monitor local control registers. PMC0 is a 64-bit counter specifically designated to count cycles. PMC1-PMC11 are 32-bit counters that can monitor 64 counter-specific events in addition to counting 64 reference events.

The benefits of the on-chip performance monitor are numerous. They include the following:

- Because some systems or software environments are not easily characterized by signal traces or benchmarks, the performance monitor can be used to understand the P2020 behavior in any system or software environment.

- The performance monitor facility can be used to aid system developers when bringing up and debugging systems.
- System performance can be increased by monitoring memory hierarchy behavior. This can help to optimize algorithms used to schedule or partition tasks and to refine the data structures and distribution used by each task.

22.1.1 Overview

This figure is a high-level block diagram of the performance monitor, which consists of a global control register (PMGC0), one 64-bit counter (PMC0), eleven 32-bit counters, and two control registers per counter (24 total control registers).

The global control register PMGC0 affects all counters and takes priority over local control registers. The local control registers are divided into two groups, as follows:

- Local control A registers control counter freezing, overflow condition enable, event selection, and burstiness. Local control register PMLCA0, which controls counter PMC0, does not contain event selection because PMC0 counts only cycles.
- Local control B registers control the start and stop triggering, contain the counters' threshold values, and the value of the threshold multiplier. Local control register PMLCB0, which controls PMC0, does not contain threshold information because PMC0 only counts cycles.

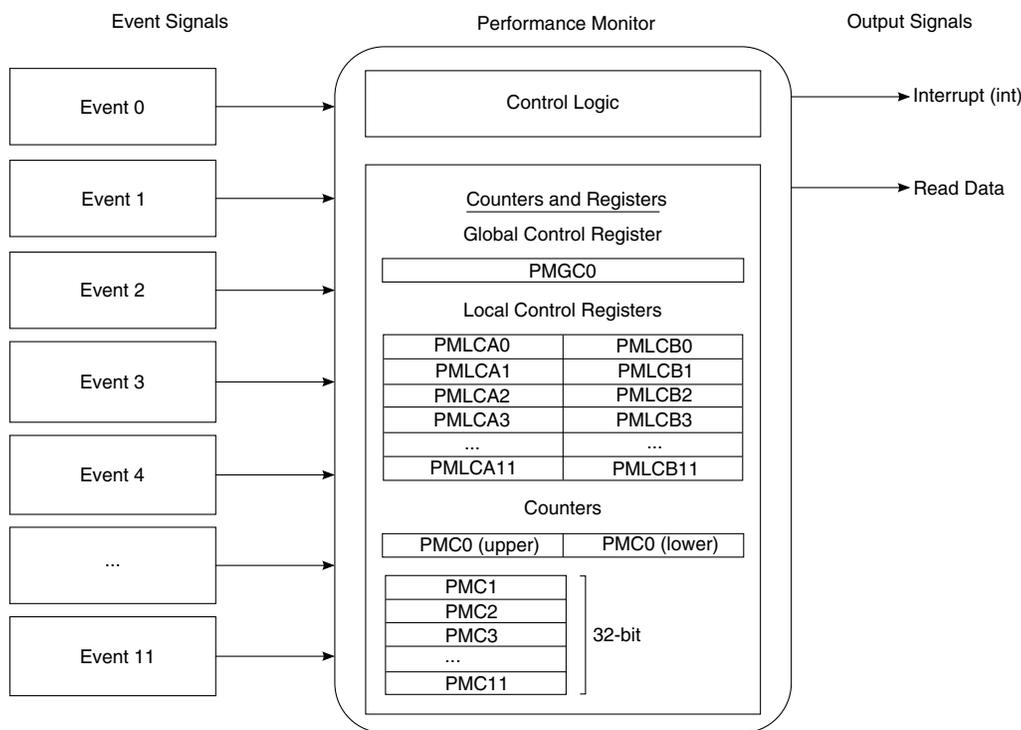


Figure 22-1. Performance monitor block diagram

Performance monitor events are signaled by the functional blocks in the integrated device and are selectively recorded in the PMCs. Sixty-four of these events are referred to as reference events, which can be counted on any of the eleven 32-bit counters. Counter-specific events can be counted only on the counter where the event is defined.

The performance monitor can generate an interrupt on overflow. Several control registers specify how a performance monitor interrupt is signaled. The PMCs can also be programmed to freeze when an interrupt is signaled.

22.1.2 Features

The P2020 performance monitor offers a rich set of features that permits a complete performance characterization of the implementation.

These features include:

- One 64-bit counter exclusively dedicated to counting cycles
- Eleven 32-bit counters that count the occurrence of selected events
- One global control register (affects all counters) and two local control registers per counter
- Ability to count up to 64 reference events that may be counted on any of the eleven 32-bit counters

Signal descriptions

- Ability to count up to 704 counter-specific events
- Triggering and chaining capability
- Duration and quantity threshold counting
- Burstiness feature that permits counting of burst events with a programmable time between bursts
- Ability to generate an interrupt on overflow

22.2 Signal descriptions

The performance monitor does not have any signals that are driven externally (off-chip) but it does assert the internal interrupt (*int*) signal on a performance monitor interrupt condition.

22.3 PERFMON Memory Map/Register Definition

Performance monitor registers reside in the run-time register block starting at offset 0xE_1000. Undefined 4-byte address spaces within offset 0x000-0xFFF are reserved. This section describes the registers implemented to support the performance monitor facilities. This table lists the performance monitor registers. These registers can be read or written only with 32-bit accesses.

The performance monitor uses twelve counter registers and a group of local control registers that are used to specify the method of counting. Two local control registers are associated with each counter in addition to a global control register that applies to all counters.

In addition to these registers, the interrupt control provides four pairs of mask registers that can be used to monitor message, interprocessor control, timer, and external interrupts. See [Performance monitor mask registers \(PMMRs\)](#) .

PERFMON memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
E_1000	Performance monitor global control register (PERFMON_PMGC0)	32	R/W	0000_0000h	22.3.1/1845
E_1010	Performance monitor local control register A0 (PERFMON_PMLCA0)	32	R/W	0000_0000h	22.3.2/1846

Table continues on the next page...

PERFMON memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
E_1014	Performance monitor local control register B0 (PERFMON_PMLCB0)	32	R/W	0000_0000h	22.3.3/1847
E_1018	Performance monitor counter 0 lower (PERFMON_PMC0_lower)	32	R/W	0000_0000h	22.3.4/1848
E_101C	Performance monitor counter 0 upper (PERFMON_PMC0_upper)	32	R/W	0000_0000h	22.3.5/1849
E_1020	Performance monitor local control register An (PERFMON_PMLCA1)	32	R/W	0000_0000h	22.3.6/1849
E_1024	Performance monitor local control register Bn (PERFMON_PMLCB1)	32	R/W	0000_0000h	22.3.7/1850
E_1028	Performance monitor counter n (PERFMON_PMC1)	32	R/W	0000_0000h	22.3.8/1852
E_1030	Performance monitor local control register An (PERFMON_PMLCA2)	32	R/W	0000_0000h	22.3.6/1849
E_1034	Performance monitor local control register Bn (PERFMON_PMLCB2)	32	R/W	0000_0000h	22.3.7/1850
E_1038	Performance monitor counter n (PERFMON_PMC2)	32	R/W	0000_0000h	22.3.8/1852
E_1040	Performance monitor local control register An (PERFMON_PMLCA3)	32	R/W	0000_0000h	22.3.6/1849
E_1044	Performance monitor local control register Bn (PERFMON_PMLCB3)	32	R/W	0000_0000h	22.3.7/1850
E_1048	Performance monitor counter n (PERFMON_PMC3)	32	R/W	0000_0000h	22.3.8/1852
E_1050	Performance monitor local control register An (PERFMON_PMLCA4)	32	R/W	0000_0000h	22.3.6/1849
E_1054	Performance monitor local control register Bn (PERFMON_PMLCB4)	32	R/W	0000_0000h	22.3.7/1850
E_1058	Performance monitor counter n (PERFMON_PMC4)	32	R/W	0000_0000h	22.3.8/1852
E_1060	Performance monitor local control register An (PERFMON_PMLCA5)	32	R/W	0000_0000h	22.3.6/1849
E_1064	Performance monitor local control register Bn (PERFMON_PMLCB5)	32	R/W	0000_0000h	22.3.7/1850
E_1068	Performance monitor counter n (PERFMON_PMC5)	32	R/W	0000_0000h	22.3.8/1852
E_1070	Performance monitor local control register An (PERFMON_PMLCA6)	32	R/W	0000_0000h	22.3.6/1849
E_1074	Performance monitor local control register Bn (PERFMON_PMLCB6)	32	R/W	0000_0000h	22.3.7/1850
E_1078	Performance monitor counter n (PERFMON_PMC6)	32	R/W	0000_0000h	22.3.8/1852
E_1080	Performance monitor local control register An (PERFMON_PMLCA7)	32	R/W	0000_0000h	22.3.6/1849
E_1084	Performance monitor local control register Bn (PERFMON_PMLCB7)	32	R/W	0000_0000h	22.3.7/1850
E_1088	Performance monitor counter n (PERFMON_PMC7)	32	R/W	0000_0000h	22.3.8/1852
E_1090	Performance monitor local control register An (PERFMON_PMLCA8)	32	R/W	0000_0000h	22.3.6/1849

Table continues on the next page...

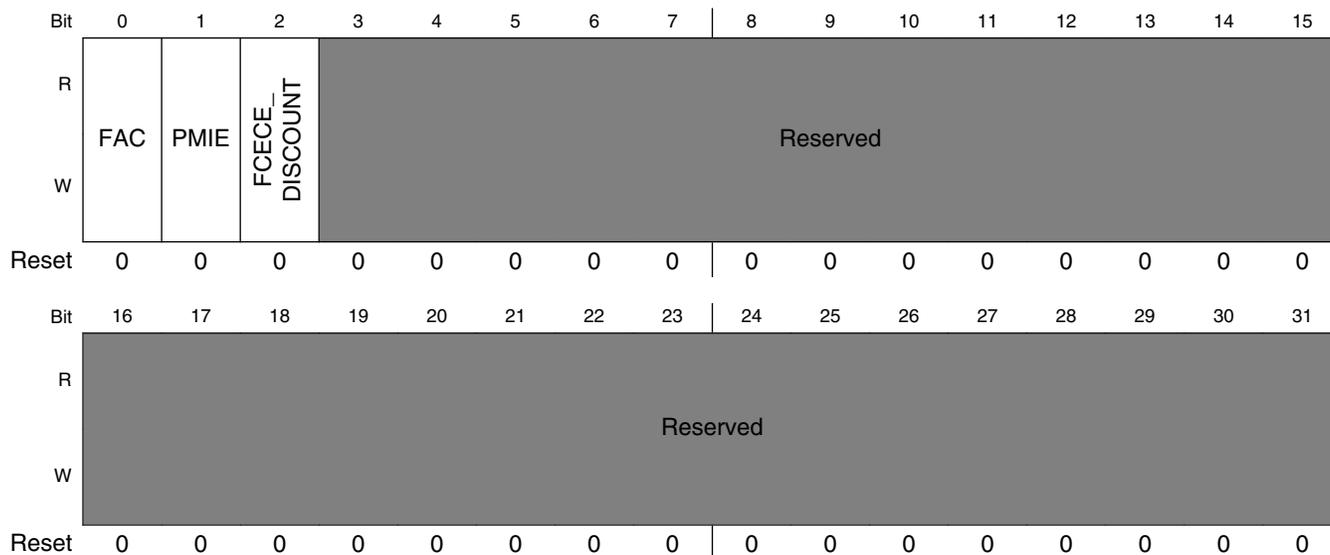
PERFMON memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
E_1094	Performance monitor local control register Bn (PERFMON_PMLCB8)	32	R/W	0000_0000h	22.3.7/1850
E_1098	Performance monitor counter n (PERFMON_PMC8)	32	R/W	0000_0000h	22.3.8/1852
E_10A0	Performance monitor local control register An (PERFMON_PMLCA9)	32	R/W	0000_0000h	22.3.6/1849
E_10A4	Performance monitor local control register Bn (PERFMON_PMLCB9)	32	R/W	0000_0000h	22.3.7/1850
E_10A8	Performance monitor counter n (PERFMON_PMC9)	32	R/W	0000_0000h	22.3.8/1852
E_10B0	Performance monitor local control register An (PERFMON_PMLCA10)	32	R/W	0000_0000h	22.3.6/1849
E_10B4	Performance monitor local control register Bn (PERFMON_PMLCB10)	32	R/W	0000_0000h	22.3.7/1850
E_10B8	Performance monitor counter n (PERFMON_PMC10)	32	R/W	0000_0000h	22.3.8/1852
E_10C0	Performance monitor local control register An (PERFMON_PMLCA11)	32	R/W	0000_0000h	22.3.6/1849
E_10C4	Performance monitor local control register Bn (PERFMON_PMLCB11)	32	R/W	0000_0000h	22.3.7/1850
E_10C8	Performance monitor counter n (PERFMON_PMC11)	32	R/W	0000_0000h	22.3.8/1852

22.3.1 Performance monitor global control register (PERFMON_PMGC0)

The performance monitor global control register (PMGC0) is a 32-bit register used to control all PMCs.

Address: E_1000h base + 0h offset = E_1000h



PERFMON_PMGC0 field descriptions

Field	Description
0 FAC	Freeze all counters. 0 PMCs are incremented (if permitted by other PMGC0/PMLC bits). 1 PMCs are not incremented. Set by hardware when an interrupt is signaled and FCECE =1.
1 PMIE	Performance monitor interrupt enable. Interrupts are caused by PMC overflows. 0 Interrupts are disabled. 1 Interrupts are enabled and occur when an enabled condition or event occurs.
2 FCECE_ DISCOUNT	Freeze counters on enabled condition or event. An enabled condition or event is defined as: The msb = 1 in PMC <i>n</i> and PMLCA <i>n</i> [CE] = 1. The use of the trigger and freeze counter conditions depends on the enabled condition. 0 PMCs can be incremented (if permitted by other control bits). 1 PMCs can be incremented (if permitted by other control bits) only until an enabled condition or event occurs, at which time PMGC0[FAC] is set. It is up to software to clear FAC.
3–31 -	This field is reserved. Reserved

22.3.2 Performance monitor local control register A0 (PERFMON_PMLCA0)

The performance monitor local control registers (PMLCA *n* and PMLCB *n*) are used to control the operation of the PMCs. The performance monitor local control A and B registers are paired, 32-bit control registers that are associated with an individual counter to specify how the counter is used and what event is monitored on that counter.

This figure shows the performance monitor local control A0 register (PMLCA0).

Address: E_1000h base + 10h offset = E_1010h

Bit	0	1	2	3	4	5	6	7		8	9	10	11	12	13	14	15
R																	
W																	
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23		24	25	26	27	28	29	30	31
R	Reserved																
W	Reserved																
Reset	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0

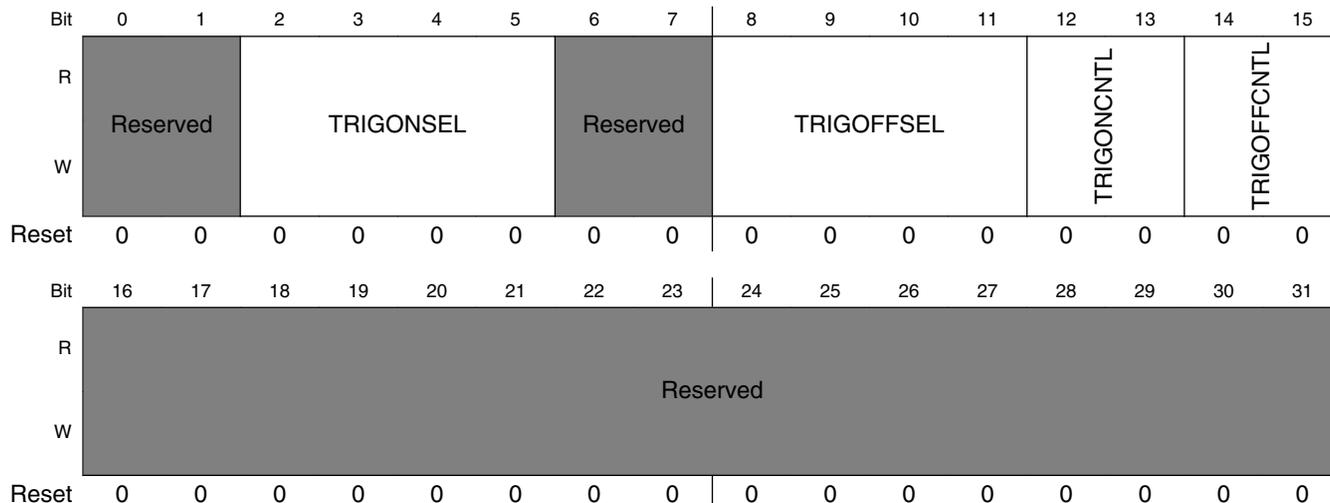
PERFMON_PMLCA0 field descriptions

Field	Description
0 FC	Freeze counter. Basic counter enable. 0 The PMCs are enabled and incremented (if permitted by other SPM control bits). 1 The PMCs are disabled-they do not increment.
1-4 -	This field is reserved. Reserved
5 CE	Condition enable. Controls counter overflow condition. Should be cleared when PMC0 is used as a trigger or is selected for chaining. 0 Overflow conditions for PMC0 cannot occur (PMC0 cannot cause interrupts or freeze counters) 1 Overflow conditions occur when PMC0[msb] is set.
6-31 -	This field is reserved. Reserved

22.3.3 Performance monitor local control register B0 (PERFMON_PMLCB0)

This figure shows the performance monitor local control B0 register (PMLCB0).

Address: E_1000h base + 14h offset = E_1014h



PERFMON_PMLCB0 field descriptions

Field	Description
0–1 -	This field is reserved. Reserved
2–5 TRIGONSEL	Trigger-on select. The number of the counter that starts event counting. When the specified counter's TRIGONCNTL event overflows, the current counter begins counting. No triggering occurs if the value is self-referential, that is, when set to the current counter number.
6–7 -	This field is reserved. Reserved
8–11 TRIGOFFSEL	Trigger-off select. The number of the counter that stops event counting. When the specified counter's TRIGONCNTL event overflows, the current counter stops counting. No triggering occurs if the value is self-referential, that is, when set to the current counter number.
12–13 TRIGONCNTL	Trigger-on control. Indicates the condition under which triggering to start counting occurs 00 Trigger off (no triggering to start) 01 Trigger on change 10 Trigger on overflow 11 Reserved
14–15 TRIGOFFCNTL	Trigger-off control. Indicates the condition under which triggering to stop occurs 00 Trigger off (no triggering to stop) 01 Trigger on change

Table continues on the next page...

PERFMON_PMLCB0 field descriptions (continued)

Field	Description
	10 Trigger on overflow 11 Reserved
16–31 -	This field is reserved. Reserved

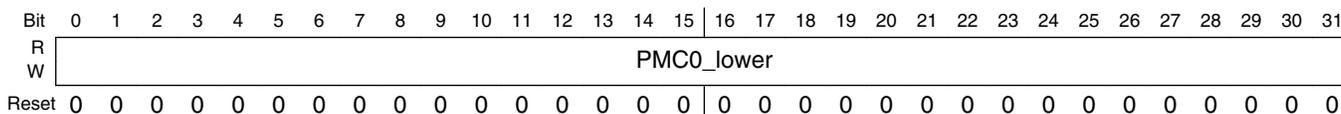
22.3.4 Performance monitor counter 0 lower (PERFMON_PMC0_lower)

PMC0-PMC11 are used to count events selected by the performance monitor local control registers. PMC0_lower, shown in figure below, is associated with two 32-bit registers that form a 64-bit counter designated to count clock cycles. PMC0 upper represents the upper 32 bits of counter 0, and PMC0 lower represents the lower 32 bits.

NOTE

Because accessing a PMC manually has priority over incrementing it due to event counting, writing a PMC while it is counting may affect the count. Likewise, writing a performance monitor control register while its target counter is counting may also affect the count.

Address: E_1000h base + 18h offset = E_1018h



PERFMON_PMC0_lower field descriptions

Field	Description
0–31 PMC0_lower	Event count. Counts only clock cycles

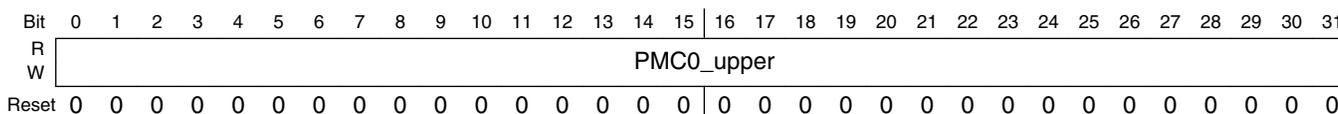
22.3.5 Performance monitor counter 0 upper (PERFMON_PMC0_upper)

PMC0-PMC11 are used to count events selected by the performance monitor local control registers. PMC0_upper, shown in figure below, is associated with two 32-bit registers that form a 64-bit counter designated to count clock cycles. PMC0 upper represents the upper 32 bits of counter 0, and PMC0 lower represents the lower 32 bits.

NOTE

Because accessing a PMC manually has priority over incrementing it due to event counting, writing a PMC while it is counting may affect the count. Likewise, writing a performance monitor control register while its target counter is counting may also affect the count.

Address: E_1000h base + 1Ch offset = E_101Ch



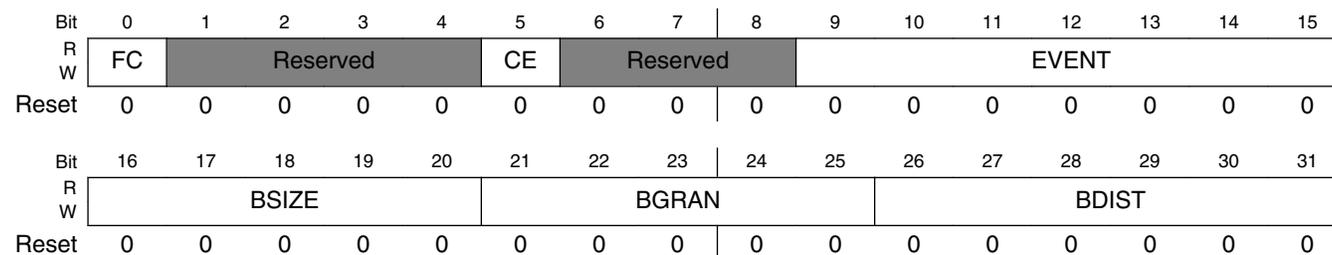
PERFMON_PMC0_upper field descriptions

Field	Description
0–31 PMC0_upper	Event count. Counts only clock cycles

22.3.6 Performance monitor local control register An (PERFMON_PMLCAn)

This figure shows the performance monitor local control A1-A11 registers.

Address: E_1000h base + 20h offset + (16d × i), where i=0d to 10d



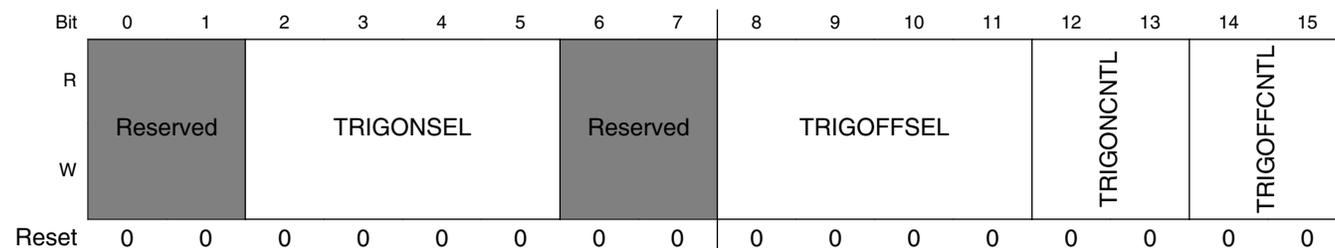
PERFMON_PMLCAn field descriptions

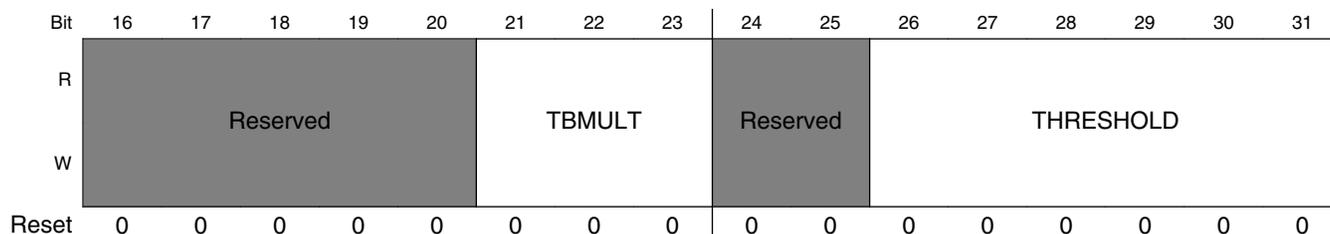
Field	Description
0 FC	Freeze counter 0 The PMCs are incremented (if permitted by other PMC control bits). 1 The PMCs are not incremented (if permitted by other PMC control bits).
1–4 -	This field is reserved. Reserved
5 CE	Condition enable 0 Overflow conditions for PMC <i>n</i> cannot occur (PMC <i>n</i> cannot cause interrupts or freeze counters). Should be cleared when PMC <i>n</i> is used as a trigger or is selected for chaining. 1 Overflow conditions occur when PMC <i>n</i> [msb] is set.
6–8 -	This field is reserved. Reserved
9–15 EVENT	Event selector. Up to 128 events selectable. Note that with counter-specific events, an offset of 64 must be used when programming the field, because counter-specific events occupy the bottom 64 values of the 7-bit event field where events are numbered. For example, to specify counter-specific event 0, the event field must be programmed to 64. See Table 22-43 for definition of events.
16–20 BSIZE	Burst size. Fewest event occurrences that constitute a burst, that is, a rapid sequence of events followed by a relatively long pause. A value less than two implies regular event counting. Any non-threshold, regular event may be counted in a bursty fashion. See Burstiness counting for more information.
21–25 BGRAN	Burst granularity. The maximum number of clock cycles between events that are considered part of a single burst. See Burstiness counting .
26–31 BDIST	Burst distance (used with TBMULT). The number of clock cycles between bursts. Must be set to a value greater than BSIZE for proper burstiness counting behavior. 000000 Regular counting

22.3.7 Performance monitor local control register Bn (PERFMON_PMLCBn)

This figure shows performance monitor local control B1-11 registers.

Address: E_1000h base + 24h offset + (16d × i), where i=0d to 10d





PERFMON_PMLCBn field descriptions

Field	Description
0–1 -	This field is reserved. Reserved
2–5 TRIGONSEL	Trigger-on select. Set this field equal to the number of the counter that should trigger event counting to start. When the specified counter's TRIGONCNTL event overflows, the current counter begins counting. No triggering occurs when TRIGONSEL = current counter.
6–7 -	This field is reserved. Reserved
8–11 TRIGOFFSEL	Trigger-off select. Set this field equal to the number of the counter that should trigger event counting to stop. When the specified counter's TRIGONCNTL event overflows, the current counter stops counting. No triggering occurs when TRIGOFFSEL = current counter.
12–13 TRIGONCNTL	Trigger-on control. Indicates the condition under which triggering to start counting occurs 00 Trigger off (no triggering to start) 01 Trigger on change 10 Trigger on overflow 11 Reserved
14–15 TRIGOFFCNTL	Trigger-off control. Indicates the condition under which triggering to stop occurs 00 Trigger off (no triggering to stop) 01 Trigger on change 10 Trigger on overflow 11 Reserved
16–20 -	This field is reserved. Reserved
21–23 TBMULT	Threshold and burstiness multiplier. Threshold events are counted when the event duration exceeds a specified threshold value. The threshold is scaled based on the TBMULT settings. TBMULT is not used to scale the threshold value for quantity threshold events. The burst distance for burstiness counting is also scaled using the TBMULT settings. For all events that scale the threshold, the threshold field is multiplied by the factors shown below (ranging from 1 to 128). 000 1 001 2 010 4 011 8 100 16 101 32 110 64 111 128
24–25 -	This field is reserved. Reserved

Table continues on the next page...

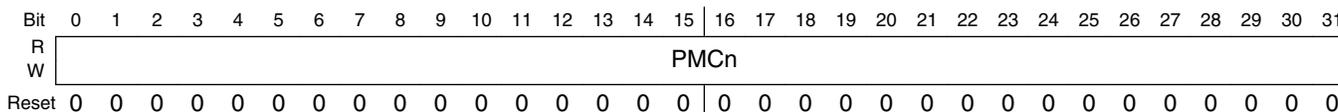
PERFMON_PMLCBn field descriptions (continued)

Field	Description
26–31 THRESHOLD	Threshold. Only events whose (number of) occurrences exceed this value are counted. By varying the threshold value, software can characterize the events subject to the threshold. For example, if PMC2 counts eTSEC BD read latencies for which the duration exceeds the threshold, software can obtain the distribution of eTSEC BD read latencies for a given program by monitoring the program using various threshold values.

22.3.8 Performance monitor counter *n* (PERFMON_PMCn)

PMC1-PMC11 are 32-bit counters that can monitor 64 unique events in addition to the 64 reference events that can be counted on all of these registers.

Address: E_1000h base + 28h offset + (16d × i), where i=0d to 10d



PERFMON_PMCn field descriptions

Field	Description
0–31 PMCn	Event count. An overflow is indicated when the msb = 1. Manually setting the msb can cause an immediate interrupt.

22.4 Functional description

This section describes the use of some features of the performance monitor.

22.4.1 Performance monitor interrupt

PMCs can generate an interrupt on an overflow when the msb of a counter changes from 0 to 1.

For the interrupt to be signaled, the condition enable bit (PMLCAn[CE]) and performance monitor interrupt enable bit (PMGC0[PMIE]) must be set. When an interrupt is signaled and the freeze-counters-on-enabled-condition-or-event bit (PMGC0[FCECE]) is set, PMGC0[FAC] is set by hardware and all of the registers are frozen. Software can clear the interrupt condition by resetting the performance monitor and clearing the most significant bit of the counter that generated the overflow.

22.4.2 Event counting

Using the control registers (see $PMGC0$, $PMLCAn$, and $PMLCBn$), the twelve PMCs can count the occurrences of specific events.

The 64-bit $PMC0$ is designated to count only clock cycles. However, to provide flexibility, a total of 64 reference events can be counted on any of the 32-bit PMCs ($PMC1$ - $PMC11$). Additionally, up to 64 unique events can be counted on each 32-bit counter.

The performance monitor must be reset before event counting sequences. The performance monitor can be reset by first freezing one or more counters and then clearing the freeze condition to allow the counters to count according to the settings in the performance monitor registers. Counters can be frozen individually by setting $PMLCAn[FC]$ bits, or simultaneously by setting $PMGC0[FAC]$. Simply clearing these freeze bits will then allow the performance monitor to begin counting based on the register settings.

Note that using $PMLCAn[FC]$ to reset the performance monitor resets only the specified counter. Performance monitor registers can be configured through reads or writes while the counters are frozen as long as freeze bits are not cleared by the register accesses.

22.4.3 Threshold events

The threshold feature allows characterization of events that can take a variable number of clock cycles to occur.

Threshold events are counted only if the latency is greater than the threshold value specified in $PMLCBn[THRESHOLD]$. There are two types of threshold events.

The first type of threshold events are duration threshold events. For duration threshold event sequences, the PMC increments only when the duration of the event is equal to or greater than the threshold value. The threshold value is scaled by a multiple specified in $PMLCBn[TBMULT]$.

A duration threshold event requires two signals: The first indicates when a threshold event sequence begins, and the second indicates when it ends. An internal counter determines when the threshold count is exceeded and when the PMC can increment. This internal counter decrements during a threshold event sequence until it reaches the value of one. A new sequence cannot begin until the current one completes. Additional threshold start signals are ignored during a sequence until a threshold stop signal occurs. If both a start and stop signal are asserted during the same cycle in a current sequence, the

stop terminates the current sequence and the start signals the beginning of a new one. However, if both signals are asserted during the same cycle while not in a current event sequence, both signals are ignored. This figure is a timing diagram for duration threshold event counting.

An illegal condition exists if the threshold value obtained from $PMLCBn[THRESHOLD]$ and $PMLCBn[TBMULT]$ is less than two. Under these conditions the intent of threshold counting is ambiguous.

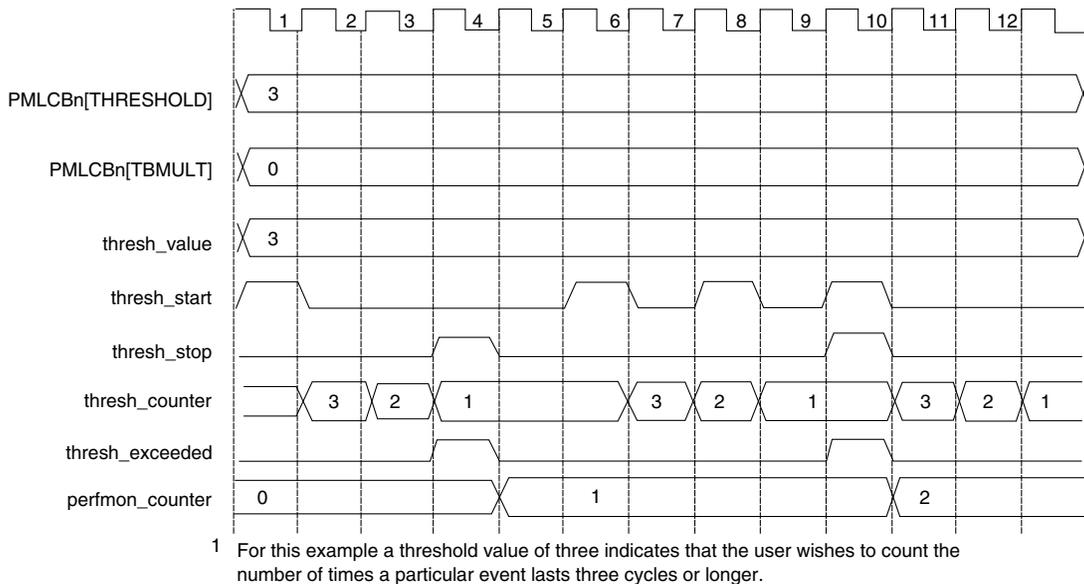


Figure 22-43. Duration threshold event sequence timing diagram

The second type of threshold event is the quantity threshold event. For these types of threshold event sequences the performance monitor counter is only incremented when the specified threshold event exceeds the threshold value. These events do not use the multiplier register field ($PMLCBn[TBMULT]$) like the duration threshold events. This type of threshold event is generally used to monitor the usage of buffers and queues. For example, the usage of a specific queue could be characterized by measuring the amount of time the queue is completely full or partially full. For this example the threshold field would be used to specify how many entries are required to be valid in the queue for that event to be counted.

22.4.4 Chaining

By configuring one counter to increment each time another counter overflows, several counters can be chained together to provide event counts larger than 32 bits.

Each counter in a chain adds 32 bits to the maximum count. The register chaining sequence is not arbitrary and is specified indirectly by selecting the register overflow event to be counted. Selecting an event has the effect of selecting a source register because all available chaining events, as shown in [Table 22-43](#), are dedicated to specific registers.

Note that the chaining overflow event occurs when the counter reaches its maximum value and wraps, not when the register's msb is set. For this overflow to occur, `PMLCAn[CE]` should be cleared to avoid signaling an interrupt when the counter's most significant bit is set. Note that several cycles may be required for the chained counters to reflect the true count because of the internal delay between when an overflow occurs and a counter increments.

22.4.5 Triggering

Triggering allows one counter to start or stop counting on the change of another counter or on the overflow of another counter.

More specifically, if `PMC1` is set to start or stop counting as a result of a change or overflow in counter `PMC2`, then counter `PMC2` must be identified in the local control register of counter `PMC1`. This is done by appropriately setting the trigger-on select bit or trigger-off select bit (`PMLCB1[TRIGOFFSEL]` or `PMLCB1[TRIGONSEL]`).

Additionally, the condition that triggers the counter must be selected by configuring the corresponding control bits (`PMLCB1[TRIGONCNTL]` or `PMLCB1[TRIGOFFCNTL]`). Assuming the counter is enabled by other control register settings, the counter increments (or freezes) when its specified event occurs after the trigger-on (or off) condition occurs.

When trigger on and trigger off are both selected, the trigger-off condition is ignored until the trigger-on condition has occurred. Furthermore, when a trigger-off condition occurs, the counter state is preserved; it is not restarted by subsequent trigger-on conditions.

Triggering is disabled when the counter's trigger-select bits specify itself as the trigger source. Similarly, triggering is disabled when the trigger control bits are cleared.

22.4.6 Burstiness counting

The burstiness counting feature makes it easier to characterize events that occur in rapid succession followed by a relatively long pause.

As shown in this table, event bursts are defined by size, granularity, and distance.

Table 22-42. Burst definition

Parameter	Description	RegisterField
Size	The minimum number of events constituting a burst	PMLCAn [BSIZE]
Granularity	The maximum time between individual events counted as members of the same burst	PMLCAn [BGRAN]
Distance	The minimum time between bursts	PMLCAn [BDIST] x PMLCBn[TBMULT]

This figure shows the relationships between size, granularity, and distance. Burstiness counting can be performed for all events except threshold events.

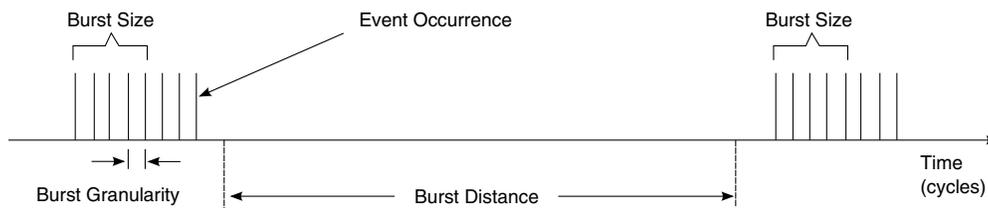


Figure 22-44. Burst size, distance, granularity, and burstiness counting

The burstiness size field (PMLCAn[BSIZE]) specifies the minimum number of event occurrences that constitute a burst. A burst is identified when the number of event occurrences equals or exceeds PMLCAn[BSIZE]. Furthermore, these individual event occurrences must be separated by no more clock cycles than the value in the burstiness granularity field (PMLCAn[BGRAN]). Note that, although a burst is identified when the minimum number of events occurs, it is not counted until the burst sequence has ended. A burst sequence ends when the specified burstiness granularity is exceeded, at which point the last valid event has occurred for that sequence.

PMLCAn[BGRAN] specifies the maximum number of cycles between individual events for them to qualify as members of the same burst sequence.

The burstiness distance field (PMLCAn[BDIST]) and threshold/burstiness multiplier field (PMLCBn[TBMULT]) specify the acceptable number of cycles between the end of a burst sequence and the beginning of a new sequence for a group of event occurrences to be counted as an individual burst. The product of the burstiness distance field and the threshold/burstiness multiplier field determine the burstiness distance value used to determine when another burst sequence can begin. Note that the burst distance count begins when a new burst sequence ends and the PMC is incremented. No new burst sequence may begin until the burst distance count has reached zero. After the burst distance count reaches zero, it holds the zero value indicating that a new burst sequence can be counted. The burst distance count begins again when a new burst sequence is identified and counted.

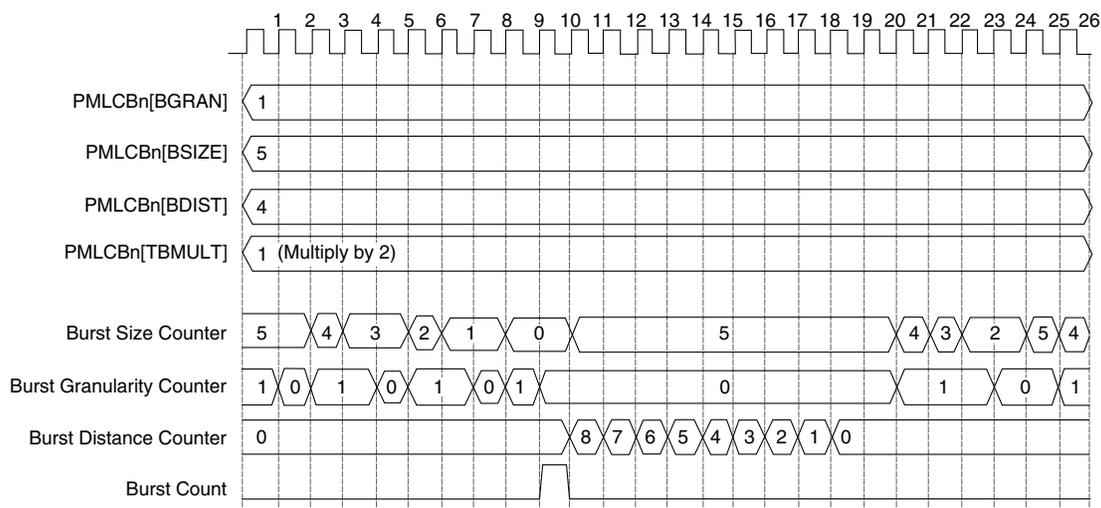
Burstiness counting is disabled when the definition of a burst is ambiguous, that is, when the burst size field is less than two, or the burst distance is zero. When burstiness counting is disabled, regular counting is allowed.

Figure 22-44 shows that the burst distance is measured from the end of one burst sequence and that a new burst sequence may not begin until the burst distance count expires.

Three internal counters track the different values required for burstiness counting.

- Burstiness size is monitored by a counter. It is loaded with the value specified in the local control register when the burst granularity counter and the burst distance counters reach zero, and no new event is occurring. It always decrements when the following conditions occur: its value is not already zero, an event occurs, and the burst distance count equals zero.
- Burstiness granularity is monitored by a counter that is loaded with the specified value in the local control register on the rising edge of an event occurrence whenever the burst distance count equals zero. The granularity counter is decremented (if it has not already reached zero) when an event is not occurring and burst distance count equals zero.
- Burstiness distance is measured by a counter that is loaded with the product of $PMLCB_n[BDIST]$ and $PMLCB_n[TBMULT]$ when a burst sequence has been identified and counted. This counter is decremented when burstiness counting is enabled (and the counter has not already reached zero).

A burst is counted at the end of a burst sequence when the three burst parameter counters are all equal to zero. This figure shows a burstiness counting example.



¹ For this example: count bursts of 5 event occurrences, burst granularity of 1, and acceptable distance between bursts of 8 (product of TBMULT and BDIST).

Figure 22-45. Burstiness counting timing diagram

22.4.7 Performance monitor events

This table lists performance monitor events specified in PMLCA1-PMLC11.

The event assignment column indicates the event's type and number, using the following formats:

- Ref:#-Reference events are shared across counters PMC1-PMC11. The number indicates the event. For example, Ref:6 means that PMC1-PMC11 share reference event 6.
- C[0-11]:#-Counter-specific events. C8 indicates an event assigned to PMC8. Thus C8-62 means PMC8 is assigned event 62 (PIC interrupt wait cycles).

Counter events not specified in this table are reserved.

Note that some functional blocks may operate at a frequency other than that of the performance monitor. In these cases, the count represented by the performance monitor counters may be inaccurate.

Table 22-43. Performance monitor events

Event Counted	Number	Description of Event Counted
General Events		
Nothing	Ref-0	Register counter holds current value
System cycles	C0 and Ref-63	CCB (platform) clock cycles

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
DDR Memory Controller Events		
Cycles a read is returning data from DRAM	Ref-19	Each data beat returned to the memory controller on the DRAM interface
Cycles a read or write transfers data from (or to) DRAM	Ref-27	Each data beat transferred to or from the DRAM
Pipelined read misses in the row open table	C1-121	Row open table read misses issued while a read is outstanding
Pipelined read or write misses in the row open table	C2-64	Row open table read or write misses issued while a read or write is outstanding
Non-pipelined read misses in the row open table	C3-124	Row open table read misses issued when no reads are outstanding
Non-pipelined read or write misses in the row open table	C4-64	Row open table read or write misses issued when no reads or writes are outstanding
Pipelined read hits in the row open table	C5-120	Row open table read hits issued when a read is outstanding
Pipelined read or write hits in the row open table	C6-64	Row open table read or write hits issued when a read or write is outstanding
Non-pipelined read hits in the row open table	C7-121	Row open table read hits issued when no reads are outstanding
Non-pipelined read or write hits in the row open table	C8-64	Row open table read or write hits issued when no reads or writes are outstanding
Forced page closings not caused by a refresh	C1-64	Precharges issued to the DRAM for any reason except refresh. The possibilities are as follows: <ul style="list-style-type: none"> • A new transaction must be issued to an already active bank and sub-bank that has a different row open. • A new transaction must be issued, but the row open table is full and there is no bank/sub-bank match between the current transaction and the row open table. • The BSTOPRE interval expired for an open row.
Row open table misses	C2-65 and C11-66	Transactions that miss in the row open table
Row open table hits	C3-64 and C10-68	Transaction that hit in the row open table
Force page closings	C4-65	Forced page closings including those due to refreshes
Read-modify-write transactions due to ECC	C5-64	If ECC is enabled and a transaction requires byte enables, a read-modify-write sequence is issued on the DRAM interface.
Forced page closings due to collision with bank and sub-bank	C9-64	Increments if a new transaction must be issued to an active bank and sub-bank that has a different row open
Reads or writes from core 0-1	Ref-13	-
Reads or writes from eTSEC 1-3	C3-65	-
Reads or writes from high speed interfaces (PCI Express 1-3, and serial RapidIO 1-2)	C4-67	-
Reads or writes from DMA 1-2	C5-66	-
Reads or writes from Security	C6-69	-
Row open table hits for reads or writes from core 0-1	Ref-14	-

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
Row open table hits for reads or writes from eTSEC 1-4	C6-65	-
Row open table hits for reads or writes from high speed interfaces (PCI Express 1-3, and serial RapidIO 1-2)	C7-65	-
Row open table hits for reads or writes from DMA 1-2	C8-66	-
Row open table hits for reads or writes from Security	C7-68 and C10-65	-
Memory Target Queue 1 Events		
DMA Controller 1 Events		
Channel 0 read request	C1-66	DMA channel 0 read request active in the system
Channel 1 read request	C2-69	DMA channel 1 read request active in the system
Channel 2 read request	C3-68	DMA channel 2 read request active in the system
Channel 3 read request	C4-70	DMA channel 3 read request active in the system
Channel 0 write request	C1-67	DMA channel 0 write request active in the system
Channel 1 write request	C2-70	DMA channel 1 write request active in the system
Channel 2 write request	C3-69	DMA channel 2 write request active in the system
Channel 3 write request	C4-71	DMA channel 3 write request active in the system
Channel 0 descriptor request	C5-105	DMA channel 0 descriptor request active in the system
Channel 1 descriptor request	C6-108	DMA channel 1 descriptor request active in the system
Channel 2 descriptor request	C7-105	DMA channel 2 descriptor request active in the system
Channel 3 descriptor request	C8-105	DMA channel 3 descriptor request active in the system
Channel 0 read DW or less	C1-68 and C5-117	DMA channel 0 read double word valid
Channel 1 read DW or less	C2-71 and C6-122	DMA channel 1 read double word valid
Channel 2 read DW or less	C3-70 and C7-118	DMA channel 2 read double word valid
Channel 3 read DW or less	C4-72 and C8-116	DMA channel 3 read double word valid
Channel 0 write DW or less	C1-69	DMA channel 0 write double word valid
Channel 1 write DW or less	C2-72	DMA channel 1 write double word valid
Channel 2 write DW or less	C3-71	DMA channel 2 write double word valid
Channel 3 write DW or less	C4-73	DMA channel 3 write double word valid
DMA Controller 2 Events		
Channel 0 read request	C8-86	DMA channel 0 read request active in the system
Channel 1 read request	C5-94	DMA channel 1 read request active in the system
Channel 2 read request	C10-75	DMA channel 2 read request active in the system
Channel 3 read request	C11-73	DMA channel 3 read request active in the system
Channel 0 write request	C8-90	DMA channel 0 write request active in the system
Channel 1 write request	C5-104	DMA channel 1 write request active in the system

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
Channel 2 write request	C10-76	DMA channel 2 write request active in the system
Channel 3 write request	C11-74	DMA channel 3 write request active in the system
Channel 0 descriptor request	C4-101	DMA channel 0 descriptor request active in the system
Channel 1 descriptor request	C5-85	DMA channel 1 descriptor request active in the system
Channel 2 descriptor request	C6-79	DMA channel 2 descriptor request active in the system
Channel 3 descriptor request	C7-91	DMA channel 3 descriptor request active in the system
Channel 0 read DW or less	C1-78 and C8-91	DMA channel 0 read double word valid
Channel 1 read DW or less	C2-80 and C9-108	DMA channel 1 read double word valid
Channel 2 read DW or less	C3-82 and C10-77	DMA channel 2 read double word valid
Channel 3 read DW or less	C4-107 and C11-75	DMA channel 3 read double word valid
Channel 0 write DW or less	C5-95	DMA channel 0 write double word valid
Channel 1 write DW or less	C6-92	DMA channel 1 write double word valid
Channel 2 write DW or less	C7-92	DMA channel 2 write double word valid
Channel 3 write DW or less	C8-92	DMA channel 3 write double word valid
e500 Coherency Module (ECM) Events		
ECM request wait core 0	C8-77	Asserted for every cycle core 0 request occurs
ECM request wait core 1	C11-98	Asserted for every cycle core 1 request occurs
ECM request wait I ² C/Security/Test Port	C7-77	Asserted for every cycle I ² C/Security/Test Port request occurs
ECM request wait PCI Express 1-3/ DMA1-2/SRIO1-2	C5-80	Asserted for every cycle PCI Express 1-3/DMA1-2/SRIO1-2 request occurs
ECM request wait eTSEC1-3	C6-80	Asserted for every cycle eTSEC1-3 request occurs
ECM request wait SAP/USB/eSDHC	C4-84	Asserted for every cycle SAP/USB/eSDHC request occurs
ECM dispatch	Ref-15	ECM dispatch (includes address only's) Note: all ECM dispatch events are for committed dispatches
ECM dispatch from core 0	C1-80	ECM dispatch from core 0 (includes address only's)
ECM dispatch from core 1	C10-95	ECM dispatch from core 1 (includes address only's)
ECM dispatch from eTSEC1	C4-85	-
ECM dispatch from eTSEC2	C8-80	-
ECM dispatch from eTSEC3	C9-80	-
ECM dispatch from SRIO1	C10-96	-
ECM dispatch from SRIO2	C11-65	-
ECM dispatch from PCI Express 1	C7-78	-
ECM dispatch from PCI Express 2	C10-98	-
ECM dispatch from PCI Express 3	C8-81	-
ECM dispatch from DMA1	C9-82	-
ECM dispatch from DMA2	C10-99	-
ECM dispatch from Security	C10-100	-

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
ECM dispatch from RIO Message Unit, Door Bell, or Port Write	C8-78	-
ECM dispatch from USB Host	C10-66	-
ECM dispatch from USB Dual Role	C9-68	-
ECM dispatch to DDR	C7-79	ECM dispatch to DDR (excludes address only)
ECM dispatch to L2	C10-102	-
ECM dispatch to SRAM	C8-79	-
ECM dispatch to LBC	C9-83	-
ECM dispatch to SRIO1	C1-85	-
ECM dispatch to SRIO2	C11-67	-
ECM dispatch to PCI Express 1	C2-85	-
ECM dispatch to PCI Express 2	C1-81	-
ECM dispatch to PCI Express 3	C2-84	-
ECM dispatch to CCSR	C9-87	-
ECM dispatch write	C7-81	-
ECM dispatch write allocate	C1-82	-
ECM dispatch read	C2-86	-
ECM dispatch read atomic	C3-86	-
ECM data bus grant DDR	C1-83	-
ECM data bus grant I ² C/Security/Test Port	C3-87	-
ECM data bus grant PCI Express 1-3/ DMA1-2/SRIO1-2	C4-89	-
ECM data bus grant LBC	Ref-16	-
ECM data bus grant eTSEC1-3	Ref-17	-
ECM data bus grant SAP/USB/eSDHC	Ref-18	-
ECM global data bus beat	C5-86	-
I2C/Security Events		
I2C/Security requests	C1-119	Number of requests (total)
I2C/Security read requests	C2-76	Number of read requests
I2C/Security data beats	C3-75	Number of data beats (total)
I2C/Security read data beats	C4-79	Number of read data beats
I2C/Security request less than 32 bytes	C7-76	Number of requests less than 32 bytes
USB/eSDHC Events		
USB/eSDHC requests	C1-99	Number of requests (total)
USB/eSDHC read requests	C10-78	Number of read requests
USB/eSDHC data beats	C11-76	Number of data beats (total)
USB/eSDHC read data beats	C5-96	Number of read data beats
USB/eSDHC request less than 32 bytes	C3-101	Number of requests less than 32 bytes
Interrupt Controller (PIC) Events		

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
PIC total interrupt count	Ref-26	Total number of interrupts serviced
PIC interrupt wait cycles	C8-126	Counts cycles when an interrupt waits to be acknowledge
PIC interrupt service cycles	C2-83	Number of cycles there is an interrupt currently being serviced.
PIC interrupt select 0 (Duration threshold)	C1-120	THRESHOLD: select 0-3: interrupt count over threshold. (Note: only unmasked, nonzero priority requests are acknowledged). The four interrupts are selected through register pairs, PM0MRn-PM3MRn. See Performance monitor mask registers (PMMRs) .
PIC interrupt select 1 (Duration threshold)	C3-123	
PIC interrupt select 2 (Duration threshold)	C5-119	
PIC interrupt select 3 (Duration threshold)	C6-124	
eTSEC 1 Events		
DMA write data beats	C3-109	DMA write data beats
DMA read data beats	C4-110	DMA read data beats
DMA write request	C5-106	DMA Write Request
DMA read request	C6-109	DMA Read Request
Number of dropped frames	C9-88	Number of dropped frames
TxBD read lifetime (Duration Threshold)	Ref-34	TxBD read lifetime
RxBD read lifetime (Duration Threshold)	Ref-38	RxBD read lifetime
TxBD write lifetime (Duration Threshold)	Ref-42	TxBD write lifetime
RxBD write lifetime (Duration Threshold)	Ref-46	RxBD write lifetime
Read data lifetime (Duration Threshold)	Ref-50	Read data lifetime
Rx IP packets checked for checksum	C9-92	Rx IP packets checked for checksum
TX IP packet with checksum	C1-105	Tx IP packet with checksum
TX TCP/UDP packet with checksum	C2-113	Tx TCP/UDP packet with checksum
TCP/UDP packets checked for c.s.	C3-114	TCP/UDP packets checked for c.s.
IP or TCP/UDP Rx checksum error	C4-115	IP or TCP/UDP Rx checksum error
Number of rejected frames by filer	C5-111	Number of rejected frames by filer
Number of rejected frames due to filer error	C6-114	Number of rejected frames due to filer error
Number of cycles Rx FIFO > 1/4 full	C5-110	Number of cycles Rx FIFO > 1/4 full
Number of cycles Rx FIFO > 1/2 full	C6-113	Number of cycles Rx FIFO > 1/2 full
Number of cycles Rx FIFO > 3/4 full	C7-110	Number of cycles Rx FIFO > 3/4 full
Number of cycles Rx FIFO = full	C8-110	Number of cycles Rx FIFO = full
Number of accepted frames match	C9-110	-
Number of accepted frames to station address	C10-80	-
Number of accepted unicaset frames via hash	C11-78	-
Number of accepted group frames via hash	C6-96	-

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
Number of accepted frames via exact match	C7-100	-
Number of rejected frames at layer 2	C9-111	-
Number of RX interrupts signaled	C8-94	-
Number of TX interrupts signaled	C9-112	-
RX data write lifetime (Duration Threshold)	C10-81	-
Number of RX packets received while RX FIFO is full	C11-79	-
eTSEC 2 Events		
DMA write data beats	C5-107	DMA write data beats
DMA read data beats	C6-110	DMA read data beats
DMA write request	C7-107	DMA Write Request
DMA read request	C8-107	DMA Read Request
Number of dropped frames	C2-110	Number of dropped frames
TxBD read lifetime (Duration Threshold)	Ref-35	TxBD read lifetime
RxBD read lifetime (Duration Threshold)	Ref-39	RxBD read lifetime
TxBD write lifetime (Duration Threshold)	Ref-43	TxBD write lifetime
RxBD write lifetime (Duration Threshold)	Ref-47	RxBD write lifetime
Read data lifetime (Duration Threshold)	Ref-51	Read data lifetime
Rx IP packets checked for checksum	C3-115	Rx IP packets checked for checksum
TX IP packet with checksum	C4-116	Tx IP packet with checksum
TX TCP/UDP packet with checksum	C5-101	Tx TCP/UDP packet with checksum
TCP/UDP packets checked for c.s.	C6-115	TCP/UDP packets checked for c.s.
IP or TCP/UDP Rx checksum error	C7-111	IP or TCP/UDP Rx checksum error
Number of rejected frames by filer	C8-111	Number of rejected frames by filer
Number of rejected frames due to filer error	C9-93	Number of rejected frames due to filer error
Number of cycles Rx FIFO > 1/4 full	C7-113	Number of cycles Rx FIFO > 1/4 full
Number of cycles Rx FIFO > 1/2 full	C8-113	Number of cycles Rx FIFO > 1/2 full
Number of cycles Rx FIFO > 3/4 full	C9-96	Number of cycles Rx FIFO > 3/4 full
Number of cycles Rx FIFO = full	C1-108	Number of cycles Rx FIFO = full
Number of accepted frames match	C9-113	-
Number of accepted frams to station address	C10-82	-
Number of accepted unicaset frames via hash	C11-80	-
Number of accepted group frames via hash	C6-97	-
Number of accepted frams via exact match	C7-101	-
Number of rejected frames at layer 2	C9-114	-

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
Number of RX interrupts signaled	C10-83	-
Number of TX interrupts signaled	C11-81	-
RX data write lifetime (Duration Threshold)	C8-96	-
Number of RX packets received while RX FIFO is full	C9-115	-
eTSEC 3 Events		
DMA write data beats	C7-108	DMA write data beats
DMA read data beats	C8-108	DMA read data beats
DMA write request	C9-90	DMA Write Request
DMA read request	C1-103	DMA Read Request
Number of dropped frames	C4-112	Number of dropped frames
TxBD read lifetime (Duration Threshold)	Ref-36	TxBD read lifetime
RxBD read lifetime (Duration Threshold)	Ref-40	RxBD read lifetime
TxBD write lifetime (Duration Threshold)	Ref-44	TxBD write lifetime
RxBD write lifetime (Duration Threshold)	Ref-48	RxBD write lifetime
Read data lifetime (Duration Threshold)	Ref-52	Read data lifetime
Rx IP packets checked for checksum	C6-116	Rx IP packets checked for checksum
TX IP packet with checksum	C7-112	Tx IP packet with checksum
TX TCP/UDP packet with checksum	C8-112	Tx TCP/UDP packet with checksum
TCP/UDP packets checked for c.s.	C9-94	TCP/UDP packets checked for c.s.
IP or TCP/UDP Rx checksum error	C1-106	IP or TCP/UDP Rx checksum error
Number of rejected frames by filer	C2-114	Number of rejected frames by filer
Number of rejected frames due to filer error	C3-116	Number of rejected frames due to filer error
Number of cycles Rx FIFO > 1/4 full	C9-97	Number of cycles Rx FIFO > 1/4 full
Number of cycles Rx FIFO > 1/2 full	C1-109	Number of cycles Rx FIFO > 1/2 full
Number of cycles Rx FIFO > 3/4 full	C2-116	Number of cycles Rx FIFO > 3/4 full
Number of cycles Rx FIFO = full	C3-118	Number of cycles Rx FIFO = full
Number of accepted frames match	C9-116	-
Number of accepted frams to station address	C10-84	-
Number of accepted unicaset frames via hash	C11-82	-
Number of accepted group frames via hash	C6-98	-
Number of accepted frams via exact match	C7-102	-
Number of rejected frames at layer 2	C9-117	-
Number of RX interrupts signaled	C8-97	-
Number of TX interrupts signaled	C9-118	-

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
RX data write lifetime (Duration Threshold)	C10-85	-
Number of RX packets received while RX FIFO is full	C11-83	-
PCI Express 1 Events		
Inbound PCI Express read	C8-119	A single pulse to indicate an inbound PCI Express read has occurred.
Inbound PCI Express write	C9-101	A single pulse to indicate an inbound PCI Express write has occurred.
Inbound PCI Express data	C5-124	A level signal to indicate the amount of data transferred if any for inbound G2PI request. Active for every beat of G2PI data.
Outbound PCI Express read	C6-126	A single pulse to indicate an outbound PCI Express read has occurred.
Outbound PCI Express write	C7-125	A single pulse to indicate an outbound PCI Express write has occurred.
Outbound PCI Express data	C8-124	A level signal to indicate the amount of data transferred if any for outbound PCI Express request. Active for every beat of PCI Express data.
Inbound Static Queue 0 start (Duration Threshold)	Ref-54	Lifetime of ISQ entry 0 or 6.
Outbound Static Queue 0 start (Duration Threshold)	Ref-55	Lifetime of OSQ entry 0.
PCI Express 2 Events		
Inbound PCI Express read	C3-102	A single pulse to indicate an inbound PCI Express read has occurred.
Inbound PCI Express write	C4-103	A single pulse to indicate an inbound PCI Express write has occurred.
Inbound PCI Express data	C5-97	A level signal to indicate the amount of data transferred if any for inbound PCI Express request. Active for every beat of PCI Express data.
Outbound PCI Express read	C1-95	A single pulse to indicate an outbound PCI Express read has occurred.
Outbound PCI Express write	C2-101	A single pulse to indicate an outbound PCI Express write has occurred.
Outbound PCI Express data	C11-87	A level signal to indicate the amount of data transferred if any for outbound PCI Express request. Active for every beat of PCI Express data.
Inbound Static Queue 0 start (Duration Threshold)	Ref-32	Lifetime of ISQ entry 0 or 6.
Outbound Static Queue 0 start (Duration Threshold)	Ref-33	Lifetime of OSQ entry 0.
PCI Express 3 Events		
Inbound PCI Express read	C9-124	A single pulse to indicate an inbound PCI Express read has occurred.
Inbound PCI Express write	C10-90	A single pulse to indicate an inbound PCI Express write has occurred.

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
Inbound PCI Express data	C11-89	A level signal to indicate the amount of data transferred if any for inbound PCI Express request. Active for every beat of PCI Express data.
Outbound PCI Express read	C6-102	A single pulse to indicate an outbound PCI Express read has occurred.
Outbound PCI Express write	C7-98	A single pulse to indicate an outbound PCI Express write has occurred.
Outbound PCI Express data	C8-127	A level signal to indicate the amount of data transferred if any for outbound PCI Express request. Active for every beat of PCI Express data.
Inbound Static Queue 0 start (Duration Threshold)	Ref-61	Lifetime of ISQ entry 0 or 6.
Outbound Static Queue 0 start (Duration Threshold)	Ref-62	Lifetime of OSQ entry 0.
Serial RapidIO 1 Events		
Inbound packet accepted	C1-124	Packet accepted on RIO
Inbound priority 0 packet accepted	C2-126	Packet accepted from RIO of priority 0
Inbound non-idles received	C4-126	Non idles received. This can be used to determine the RIO link utilization. This is actually 1/2 of the actual count.
Inbound packet retry occurred	C5-83	Packet retry occurred, any cause
Inbound buffer is full	C6-85	Clock cycle occurred in which the inbound buffer is full to any priority (measured from when EOP received to EOP transferred on OCN). Event asserted for as many clock cycles as this is true.
Inbound buffer is full to priority 0	C7-123	Clock cycle occurred in which the inbound buffer is full to priority 0 (measured from when EOP received to EOP transferred on OCN). Event asserted for as many clock cycles as this is true.
Outbound non-idles transmitted	Ref-59	Non idles transmitted. This can be used to determine the RIO link utilization. This is actually 1/2 of the actual count.
Outbound packet sent to RapidIO interface	C2-127	Outbound packet sent to RapidIO interface
Outbound packet was misaligned	C4-76	Outbound packet was misaligned
Outbound packet was retried	C5-84	Outbound packet was retried
Outbound packet was reordered	C6-73	Outbound packet was reordered
Outbound packet sent to RIO of priority 0	C7-124	Outbound packet sent to RIO of priority 0
Outbound buffer is full	C8-123	Clock cycle occurred in which the outbound buffer is full to any priority. Event asserted for as many clock cycles as this is true.
Outbound buffer is full to priority 0	C9-103	Clock cycle occurred in which the outbound buffer is full to priority 0. Event asserted for as many clock cycles as this is true.
Serial RapidIO 2 Events		
Inbound packet accepted	C3-77	Packet accepted on RIO
Inbound priority 0 packet accepted	C10-64	Packet accepted from RIO of priority 0
Inbound non-idles received	C11-64	Non idles received. This can be used to determine the RIO link utilization. This is actually 1/2 of the actual count.
Inbound packet retry occurred	C8-70	Packet retry occurred, any cause

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
Inbound buffer is full	C7-71	Clock cycle occurred in which the inbound buffer is full to any priority (measured from when EOP received to EOP transferred on OCN). Event asserted for as many clock cycles as this is true.
Inbound buffer is full to priority 0	C9-65	Clock cycle occurred in which the inbound buffer is full to priority 0 (measured from when EOP received to EOP transferred on OCN). Event asserted for as many clock cycles as this is true.
Outbound non-idles transmitted	Ref-41	Non idles transmitted. This can be used to determine the RIO link utilization. This is actually 1/2 of the actual count.
Outbound packet sent to RapidIO interface	C1-72	Outbound packet sent to RapidIO interface
Outbound packet was misaligned	C5-65	Outbound packet was misaligned
Outbound packet was retried	C11-126	Outbound packet was retried
Outbound packet was reordered	C7-109	Outbound packet was reordered
Outbound packet sent to RIO of priority 0	C3-126	Outbound packet sent to RIO of priority 0
Outbound buffer is full	C5-82	Clock cycle occurred in which the outbound buffer is full to any priority. Event asserted for as many clock cycles as this is true.
Outbound buffer is full to priority 0	C1-112	Clock cycle occurred in which the outbound buffer is full to priority 0. Event asserted for as many clock cycles as this is true.
RapidIO Message Unit Events		
Inbound packet received	Ref-60	Packet received of any priority
Inbound priority 0 packet received	C5-77	Packet received of priority 0
Inbound buffer is full	C6-72	Clock cycle occurred in which the inbound buffer is full to any priority (measured from when SOP received to buffer release transferred on OCN). Event asserted for as many clock cycles as this is true.
Inbound buffer is full for priority 0	C7-72	Clock cycle occurred in which the inbound buffer is full to priority 0 (measured from when SOP received to buffer release transferred on OCN). Event asserted for as many clock cycles as this is true.
Packet sent to OCeaN	C8-102	Packet sent to OCN
Packet sent to OCeaN of priority 0	C11-94	Packet sent to OCN of priority 0
Local Bus Events		
Number of correctable FCM ECC errors	C3-121	-
Number of uncorrectable FCM ECC errors	C6-119	-
Atomic reservation time-outs for ECM port	C6-118	-
Cycles for a read to FCM buffer RAM	C7-115	-
Cycles for a write to FCM buffer RAM	C8-115	-
Cycles a read is taking in GPCM	C1-117	-
Cycles a read is taking in UPM	C2-122	-
Cycles a write is taking in GPCM	C4-120	-
Cycles a write is taking in UPM	C5-114	-
L2 Cache/SRAM Events		

Table continues on the next page...

Table 22-43. Performance monitor events (continued)

Event Counted	Number	Description of Event Counted
Core instruction accesses to L2 that hit	C2-123	-
Core instruction accesses to L2 that miss	Ref-23	-
Core data accesses to L2 that hit	C4-121	-
Core data accesses to L2 that miss	C5-115	-
Non-core burst write to L2 (cache external write or SRAM)	C6-120	-
Non-core non-burst write to L2	C7-116	-
Non-core write misses cache external write window and SRAM memory range	Ref-24	-
Non-core read hit in L2	C1-118	-
Non-core read miss in L2	Ref-25	-
L2 allocates, from any source	C2-124	-
L2 retries due to full write queue	C3-122	-
L2 retries due to address collision	C4-122	-
L2 failed lock attempts due to full set	C5-116	-
L2 victimizations of valid lines	C6-121	-
L2 invalidations of lines	C7-117	-
L2 clearing of locks	Ref-22	-
Debug Events		
External event	C3-125	Number of cycles trig_in pin is asserted
Watchpoint monitor hits	C2-125	-
Trace buffer hits	C1-122	-
DUART Events		
UART0 baud rate	C1-127	-
UART1 baud rate	C5-127	-
Chaining Events		
PMC0 carry-out	Ref-1	PMC0[0] 1-to-0 transitions.
PMC1 carry-out	Ref-2	PMC1[0] 1-to-0 transitions. Reserved for PMC1.
PMC2 carry-out	Ref-3	PMC2[0] 1-to-0 transitions. Reserved for PMC2.
PMC3 carry-out	Ref-4	PMC3[0] 1-to-0 transitions. Reserved for PMC3.
PMC4 carry-out	Ref-5	PMC4[0] 1-to-0 transitions. Reserved for PMC4.
PMC5 carry-out	Ref-6	PMC5[0] 1-to-0 transitions. Reserved for PMC5.
PMC6 carry-out	Ref-7	PMC6[0] 1-to-0 transitions. Reserved for PMC6.
PMC7 carry-out	Ref-8	PMC7[0] 1-to-0 transitions. Reserved for PMC7.
PMC8 carry-out	Ref-9	PMC8[0] 1-to-0 transitions. Reserved for PMC8.
PMC9 carry-out	Ref-10	PMC9[0] 1-to-0 transitions. Reserved for PMC9.
PMC10 carry-out	Ref-11	PMC10[0] 1-to-0 transitions. Reserved for PMC10.
PMC11 carry-out	Ref-12	PMC11[0] 1-to-0 transitions. Reserved for PMC11.

22.4.8 Performance monitor examples

Table 22-45 contains sample register settings for the four supported modes.

- Simple event performance monitoring example
- Triggering event performance monitoring example
- Threshold event performance monitoring example
- Burstiness event performance monitoring example

The settings in Table 22-44 are identical for all four examples.

Table 22-44. PMGC0 and PMLCAn Settings

Field	Setting	Reason
PMGC0[FAC]	0	Counters must not be frozen.
PMGC0[PMIE]	1	Performance monitor interrupts are enabled
PMGC0[FCECE]	1	Counters should be frozen when an interrupt is signaled.
PMLCAn[FC]	0	Counters cannot be frozen for counting.
PMLCAn[CE]	1	Overflow condition enable is required to allow interrupt signaling.

For simple event counting, a non-threshold event is selected in PMLCAn[EVENT] and all other features are disabled by clearing all register fields except for CE.

For the triggering example any event can be selected in PMLCAn[EVENT]. All other features are disabled by clearing these register fields except for CE to allow interrupt signaling. If PMLCBn[TRIGONSEL] is 3 and PMLCBn[TRIGOFFSEL] is 5, the counter begins and ends counting based on the conditions in counters three and five. Furthermore, if PMLCBn[TRIGONCNTL] is 1, the counter begins counting when PMC3 changes value. According to the setting in PMLCBn[TRIGOFFCNTL], the counter ends counting when PMC5 overflows. Also, although the register settings for PMC5 is not shown, PMLCAn[CE] for this counter must be cleared so that interrupt signaling is not enabled and the counter does not freeze when it overflows.

For threshold counting, a threshold event must be specified in PMLCAn[EVENT]. For this example, the duration threshold value is scaled by two because PMLCBn[TBMULT] is one. All other features are disabled by clearing the appropriate fields.

Any non-threshold event can use the burstiness feature. For burstiness counting, values for PMLCAn[BFSIZE,BGRAN,BDIST] and PMLCBn[TBMULT] must be specified.

Table 22-45. Register Settings for Counting Examples

Register	Register Field	Simple Event	Triggering	Threshold	Burstiness
PMGC0	FAC	0	0	0	0
	PMIE	1	1	1	1
	FCECE	1	1	1	1
PMLCAn	FC	0	0	0	0
	CE	1	1	1	1
	EVENT	89	68	39	2
	BSIZE	0	0	0	5
	BGRAN	0	0	0	1
	BDIST	0	0	0	8
PMLCBn	TRIGONSEL	0	3	0	0
	TRIGOFFSEL	0	5	0	0
	TRIGONCNTL	0	1	0	0
	TRIGOFFCNTL	0	2	0	0
	TBMULT	0	0	0	0
	THRESHOLD	0	0	3	0

The performance monitor must be reset before event counting sequences. The performance monitor can be reset by first freezing one or more counters and then clearing the freeze condition to allow the counters to count according to the settings in the performance monitor registers. Counters can be frozen individually by setting PMLCAn[FC] bits or simultaneously by setting PMGC0[FAC]. Simply clearing these freeze bits will then allow the performance monitor to begin counting based on the register settings.

Note that using PMLCAn[FC] to reset the performance monitor resets only the specified counter. Performance monitor registers can be configured through reads or writes while the counters are frozen as long as freeze bits are not cleared by the register accesses.



Chapter 23

Debug Features and Watchpoint Facility

This chapter describes all customer-visible debug modes of the SoC integrated device. The debug features on the chip pertain to the enhanced local bus controller (eLBC) and the DDR SDRAM interface.

23.1 Introduction

This chapter describes all customer-visible debug modes of the chip integrated device.

The debug features on the chip pertain to the enhanced local bus controller (eLBC) and the DDR SDRAM interface. In addition to the external interfaces, the chip provides triggering capabilities based on user-programmable events. In addition, the watchpoint monitor and trace buffer provide some visibility into internal buses. This chapter also describes context ID registers, useful for software debug, and describes the JTAG access port signals that comply with the IEEE 1149.1 boundary-scan specification.

The chip provides the following debug features:

- DDR SDRAM transactions ([DDR SDRAM interface debug](#))
- eLBC transactions ([Local bus interface debug](#))
- Watchpoint monitor and trace buffer ([Watchpoint monitor](#), and [Trace buffer](#))

As shown in the figure below, debug information is provided through the eLBC and DDR SDRAM interfaces. Limited visibility, through a 256 entry x 64-bit trace buffer, is also provided for the processor core interface. This visibility into internal device operation is useful for debugging application software through inverse assembly and reconstruction of the fetch stream.

The combination of a source ID (MSRCID[0:4]) and a data-valid signal (MDVAL) indicates that meaningful debug information is visible on either the local bus or DDR SDRAM interfaces. A logic analyzer can be programmed to capture data based on the values of MSRCID[0:4] and MDVAL, which supplement existing DDR and eLBC

signals. Without MSRCID and MDVAL, it is not possible to provide a logic analyzer with the transaction source and validity information. It needs a trigger criteria to filter transactions of interest.

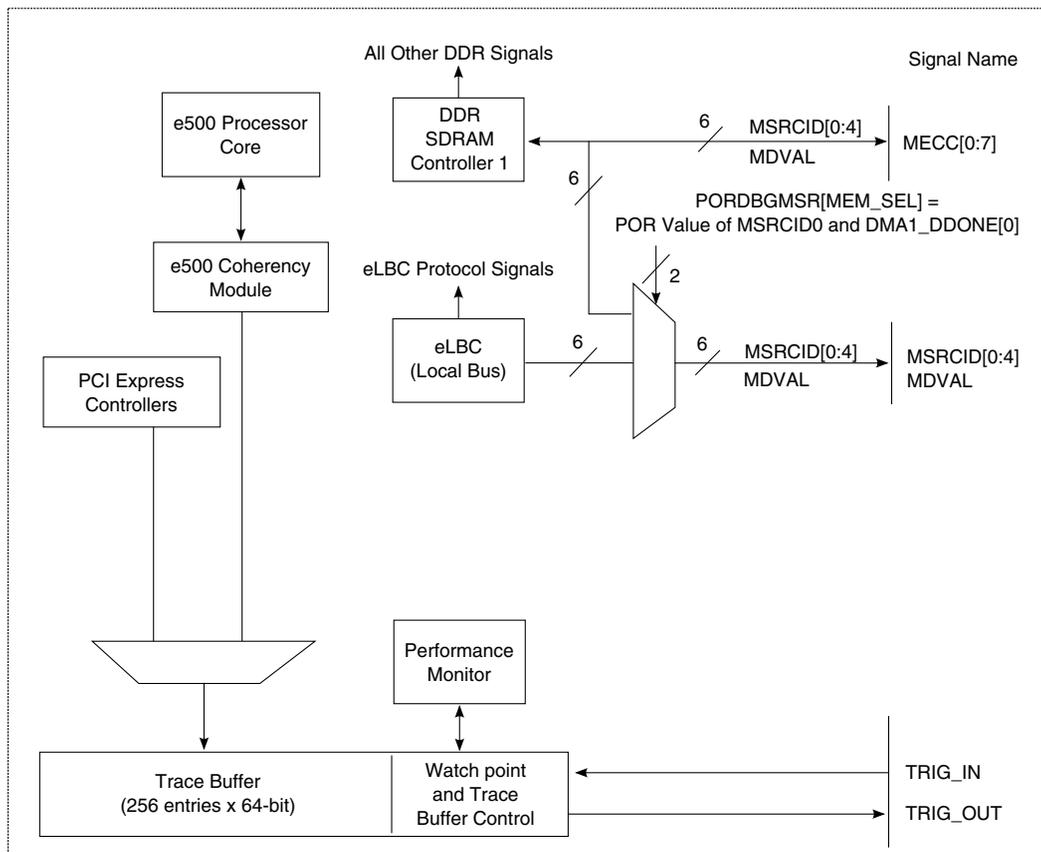


Figure 24-1. Debug and Watchpoint Monitor Block Diagram

Figure 23-1. Debug and watchpoint monitor block diagram

Other system debugging is supported by the programmable triggering of the watchpoint monitor and trace buffer. Both can be triggered from one of the following three sources:

- Each other
- A performance monitor event
- An external source (through TRIG_IN)

The watchpoint monitor can be configured to assert TRIG_OUT when a programmed event occurs. The two context ID registers, described in [Programmed context ID register \(Debug_PCIDR\)](#) and [Current context ID register \(Debug_CCIDR\)](#) are useful for software debug.

23.1.1 Features

The principal features of the debug modes and the watchpoint monitor are as follows:

- eLBC and DDR interface source ID and data-valid indicators that can be selected as follows:
 - Both eLBC or DDR SDRAM source ID can be selected to be driven onto the dedicated MSRCID[0:4] pins
 - DDR SDRAM source ID and data-valid indicators can also be selected to be driven onto the error correcting code (ECC) pins of the DDR interface
- Watchpoint monitor that supports the following:
 - Two-level triggering
 - Programmable external trigger (TRIG_OUT)
 - Interlock with performance monitor to use its large number of counters
- Trace buffer features that support the following:
 - Two-level triggering
 - Programmable external trigger (TRIG_OUT)
 - Interlock with performance monitor to use its large number of counters
 - 256-entry trace buffer, 64 bits each
 - Programmable trace start and stop
 - Ability to function as a second watchpoint monitor
- Context ID registers that can be programmed to trigger events

23.1.2 Modes of operation

The eLBC and DDR SDRAM interfaces have debug modes, which are controlled by values on configuration inputs during the power-on reset (POR) sequence, as shown in this table.

The DDR controller can also drive debug information on either MSRCID[0:4] or MECC[0:7]. See [Source and target ID](#) for additional information about the source ID information driven on the debug signals in these modes.

Note that both the watchpoint monitor and trace buffer also operate in a variety of modes.

Table 23-1. POR configuration settings and debug modes

Configuration signal	POR Value	Effect	Reference
DMA2_DACK0_B	0	Debug information from the eLBC is driven on the MSRCID and MDVAL signals	Memory debug mode (eLBC and DDR)
Default (1)	1	Debug information from the DDR SDRAM controller is driven on the MSRCID and MDVAL signals (default).	
DMA2_DDONE0_B	0	Debug information is driven on the ECC pins instead of normal ECC I/O. ECC signals from memory devices must be disconnected.	DDR SDRAM interface debug mode
Default (1)	1	Debug information is not driven on ECC pins. ECC pins function in their normal mode (default).	

23.1.2.1 Memory debug mode (eLBC and DDR)

Both the eLBC and the DDR SDRAM controller can drive debug information (source ID and data-valid indicator) onto MSRCID[0:4] and MDVAL.

As shown in [Table 23-1](#), the value of DMA2_DACK0_B controls the multiplexing during POR. If DMA2_DACK0_B is low when sampled during POR, the eLBC SDRAM information appears on MSRCID[0:4] and MDVAL. If DMA2_DACK0_B is high when sampled during POR, the debug information from the DDR SDRAM controller is presented on MSRCID[0:4] and MDVAL.

23.1.2.2 DDR SDRAM interface debug mode

DMA2_DDONE0_B is sampled during POR to multiplex either ECC or debug information on the ECC pins of the DDR SDRAM interfaces.

As shown in [Table 23-1](#), if DMA2_DDONE0_B is low during POR, the ECC pins operate in debug mode and provide MSRCID and MDVAL information. DMA2_DDONE0_B must be pulled low during POR to use the ECC pins in debug mode.

If DMA2_DDONE0_B is unconnected, an internal pull-up resistor ensures the ECC pins always source DDR SDRAM error correcting code information as their default power-on reset configuration.

NOTE

If the DDR ECC pins are in debug mode (configured for debug during POR), ECC checking is disabled in the memory controller. In this case, MECC[0:4] do not provide ECC information and must not be connected to SDRAM devices.

23.1.2.3 Watchpoint monitor modes

This table lists the operating modes supported by the watchpoint monitor.

Table 23-2. Watchpoint monitor modes

Operating mode	Trigger logic
Immediate trigger arming (one-level triggering)	The watchpoint monitor triggers as soon as the first trigger event occurs.
Wait for trigger arming (two-level triggering)	The watchpoint monitor waits for a specific event before enabling (arming) the trigger logic. The monitor does not respond to trigger events until after the arming event occurs. This function is similar to two-level triggering on a logic analyzer.
Assert TRIG_OUT on hit	The debug logic can be programmed to assert the TRIG_OUT signal when a programmed watchpoint monitor event occurs. This signal can be used to trigger a logic analyzer.

23.1.2.4 Trace buffer modes

This table lists the operating modes supported by trace buffer.

Table 23-3. Trace buffer modes

Operating mode	Trigger logic
Immediate trigger arming (one-level triggering)	The trace buffer triggers as soon as the first trigger event occurs.
Wait for trigger arming (two-level triggering)	The trace buffer waits for a specific event before enabling (arming) the trigger logic. The trace buffer does not respond to trigger events until after the arming event occurs. This function is similar to two-level triggering on a logic analyzer.
Specific interface selection	The trace buffer can be programmed to trace one of several internal interfaces.
Specific event selection	The trace buffer can be programmed to trace on the occurrence of one or several concurrent events.
Specific trace selection	To facilitate trace data filtering, the trace buffer can be configured to capture data under the following conditions: <ul style="list-style-type: none"> • On every cycle in which a valid transaction is present on the selected interface • Only when the programmed trace event is detected
Programmable trace stop	The trace buffer may be programmed to stop tracing when a programmed stop-tracing event occurs or when the 256-entry buffer is full.

23.2 External signal description

This section provides information about all the external signals associated with the various chip debug functions.

As shown in [Table 23-1](#), the chip has several signals that are sampled during POR to determine the configuration of the phase-locked loop clock mode and the ROM, flash, and dynamic memory. See [Power-on reset configuration](#).

To facilitate system testing, the chip provides a JTAG test access port (TAP) that complies with the IEEE 1149.1 boundary-scan specification. This section also describes JTAG TAP signals.

23.2.1 Overview

[Table 23-4](#) summarizes of all signals associated with device debug features alongside a reference to the page containing further information. [Table 23-5](#) provides the detailed descriptions about the signals. Some signals (the MECC bus for example) are described in other chapters , but are described here also for completeness.

The descriptions in this chapter focus on the signals' debugging utility.

Table 23-4. Debug, Watchpoint, and test signal summary

Name	Description	Functional block	Function	Reset value	I/O	Page
MDVAL	Memory data-valid	Debug	Selectable data-valid signal from either DDR SDRAM controller or eLBC.	0	O	Table 23-5
MECC[0:7]	DDR error correcting code	DDR SDRAM	In debug mode, the high-order six bits carry debug information (transaction source ID and data-valid indication).	0x08	O ¹	Table 23-5
MSRCID[0:1]	Memory source ID	Debug	Selectable transaction source ID from either DDR SDRAM controller or eLBC.	Reset_cfg	O	Table 23-5
MSRCID[2:4]				111	O	Table 23-5
TRIG_IN	Trigger in	Debug	Inbound trigger for various functions in the watchpoint monitor and trace buffer.	1	I	Table 23-6
TRIG_OUT	Trigger out	Debug	Can be used externally for triggering a logic analyzer. Additionally, it can be used for observing system ready indication. Functions are multiplexed onto this signal depending on TOSR[SEL] (see Trigger output source register (Debug_TOSR)).	1	O	Table 23-6
TCK	Test clock	Debug	Clock for JTAG testing.	0	I	Table 23-7
TDI	Test data input	Debug	Serial input for instructions and data to the JTAG test subsystem. Internally pulled up.	1	I	Table 23-7

Table continues on the next page...

Table 23-4. Debug, Watchpoint, and test signal summary (continued)

Name	Description	Functional block	Function	Reset value	I/O	Page
TDO	Test data output	Debug	Serial data output for the JTAG test subsystem. High impedance except when scanning out data.	Hi Z	O	Table 23-7
TMS	Test mode select	Debug	Carries commands to the TAP controller for boundary scan operations. Internally pulled up.	1	I	Table 23-7
TRST_B	Test reset	Debug	Resets the TAP controller asynchronously.	-	I	Table 23-7
TEST_SEL_B	Test select 1	Test	Factory test. Must be negated (pulled high) for normal operation. NOTE: For 2010, TEST_SEL is an active high signal and needs to be pulled low to disable factory test.	-	I	Table 23-7
SCAN_MODE_B	Test	Test	Factory test. Refer to the <i>P2020 Integrated Host Processor Hardware Specifications</i> for proper treatment.	-	I	Table 23-7

1. While these signals are normally bidirectional, when sourcing debug information they are output only.

23.2.2 Detailed signal descriptions

This section describes the details of the debug, watchpoint monitor, and JTAG test signals.

23.2.2.1 Debug signals-details

This table describes all signals associated with device debug modes.

Table 23-5. Debug signals-detailed signal descriptions

Signal	I/O	Description
MDVAL	O	Memory data-valid. Indicates when valid data is available. May be used by a logic analyzer to capture the data on the data bus.
		State meaning Asserted-Indicates that data is valid on the data bus during the current clock cycle. When the DDR SDRAM interface is selected to source information on MDVAL, this signal is valid for every cycle that data is driven or received on the DDR SDRAM interface. When the eLBC is selected, this signal is valid for every cycle that data is driven or received on the local bus interface. The assertion of this signal may be used by a logic analyzer to capture data.
		Timing Asserted/negated-Referenced to the selected interface, (DDR or local bus). Asserts when data is valid. Assertions are held for the duration of the transfer. Read data timing is similar to MA. Write data timing is similar to the output MDQ.

Table continues on the next page...

Table 23-5. Debug signals-detailed signal descriptions (continued)

Signal	I/O	Description	
MECC[0:7]	O	Memory ECC. DDR error checking and correcting. The normally bidirectional operation of the memory ECC (MECC) bus is described in Error checking and correcting (ECC) . This bus is used for debug functions when MSRCID1 is sampled low during POR. In debug mode, the high-order 5 bits (MECC[0:4]) may be used to provide the transaction source ID and MECC5 can be used as the data-valid indicator. In debug mode, MECC[0:5] is constantly driven with debug information and must be disconnected from the DDR memory's ECC pins.	
		State meaning	Asserted/negated-In debug mode, MECC[0:5] is always driven. The source ID values appear during RAS and CAS cycles. A value of 0x1F (all ones) is driven during cycles other than RAS and CAS. The data-valid indicator appears when data is being received or driven on the pins.
		Timing	Driven every cycle in debug mode.
MSRCID[0:4]	O	Memory source ID. Attribute signals associated with the memory interface that indicate the source ID for a transaction on an SDRAM interface. The SDRAM interface, DDR or local bus, to which the debug information applies is specified during POR with MSRCID0 as shown in Table 23-1 . Two of these signals serve as reset configuration input signals.	
		State meaning	Asserted/Negated-In debug mode, always driven with the value of the source ID. The source ID has a value of 0x1F for cycles other than RAS and CAS. The encodings shown in Table 23-27 provide detailed information about a memory transaction.
		Timing	Driven every cycle in debug mode. Similar timing to MA.

23.2.2.2 Watchpoint monitor trigger signals-details

This table shows detailed descriptions of the watchpoint monitor and trace buffer signals.

Table 23-6. Watchpoint and trigger signals-detailed signal descriptions

Signal	I/O	Description	
TRIG_IN	I	Trigger in. Can be used to trigger the watchpoint and trace buffers. Note this is an active-high (rising-edge triggered) signal.	
		State meaning	Asserted-Indicates that a programmed/armed external event has been detected. Assertion may be used internally to trigger trace buffers and watchpoint mechanisms.
		Timing	Assertion/Negation-The chip interprets TRIG_IN as asserted on detection of the rising edge. It may occur at any time. Must remain asserted for at least 3 system clocks to be recognized internally.

Table continues on the next page...

Table 23-6. Watchpoint and trigger signals-detailed signal descriptions (continued)

Signal	I/O	Description				
TRIG_OUT	O	Trigger out. Function determined by TOSR[SEL]. When TOSR[SEL] is non-zero, it can be used for triggering external devices, like a logic analyzer, with either the watchpoint monitor, the trace buffer, or the performance monitor as trigger sources. When TOSR[SEL] is cleared, TRIG_OUT is multiplexed with READY, which indicates the operational readiness of the device (running or in low-power or debug modes). See Reset, Clocking, and Initialization and Global Utilities for more details about reset, low-power, and debug states.				
		<table border="0"> <tr> <td style="vertical-align: top;">State meaning</td> <td>Asserted-When TOSR[SEL] is all zeros, serves as the READY signal, indicating that the device is not in a low-power or debug mode and that it has emerged from reset. SEL ≠ 0 indicates that a programmed trigger event has occurred.</td> </tr> <tr> <td></td> <td>Negation-No final watchpoint match condition</td> </tr> </table>	State meaning	Asserted-When TOSR[SEL] is all zeros, serves as the READY signal, indicating that the device is not in a low-power or debug mode and that it has emerged from reset. SEL ≠ 0 indicates that a programmed trigger event has occurred.		Negation-No final watchpoint match condition
		State meaning	Asserted-When TOSR[SEL] is all zeros, serves as the READY signal, indicating that the device is not in a low-power or debug mode and that it has emerged from reset. SEL ≠ 0 indicates that a programmed trigger event has occurred.			
	Negation-No final watchpoint match condition					
Timing	Assertion may occur at any time. Remains asserted for at least 3 system clocks					

23.2.2.3 JTAG test signals-details

This table shows detailed descriptions of the JTAG test signals.

Table 23-7. JTAG test and other signals-detailed signal descriptions

Signal	I/O	Description		
TCK	I	JTAG test clock.		
		<table border="0"> <tr> <td style="vertical-align: top;">State meaning</td> <td>Asserted/Negated-Should be driven by a free-running clock signal with a 30-70% duty cycle. Input signals to the TAP are clocked in on the rising edge. Changes to the TAP output signals occur on the falling edge. The test logic allows TCK to be stopped. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.</td> </tr> </table>	State meaning	Asserted/Negated-Should be driven by a free-running clock signal with a 30-70% duty cycle. Input signals to the TAP are clocked in on the rising edge. Changes to the TAP output signals occur on the falling edge. The test logic allows TCK to be stopped. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.
		State meaning	Asserted/Negated-Should be driven by a free-running clock signal with a 30-70% duty cycle. Input signals to the TAP are clocked in on the rising edge. Changes to the TAP output signals occur on the falling edge. The test logic allows TCK to be stopped. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.	
Timing	See IEEE 1149.1 standard for more details.			
TDI	I	JTAG test data input.		
		<table border="0"> <tr> <td style="vertical-align: top;">State meaning</td> <td>Asserted/negated-The value present on the rising edge of TCK is clocked into the selected JTAG test instruction or data register. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.</td> </tr> </table>	State meaning	Asserted/negated-The value present on the rising edge of TCK is clocked into the selected JTAG test instruction or data register. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.
		State meaning	Asserted/negated-The value present on the rising edge of TCK is clocked into the selected JTAG test instruction or data register. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.	
Timing	See IEEE 1149.1 standard for more details.			
TDO	O	JTAG test data output.		
		<table border="0"> <tr> <td style="vertical-align: top;">State meaning</td> <td>Asserted/negated-The contents of the selected internal instruction or data register are shifted out on this signal on the falling edge of TCK. Remains in a high-impedance state except when scanning data.</td> </tr> </table>	State meaning	Asserted/negated-The contents of the selected internal instruction or data register are shifted out on this signal on the falling edge of TCK. Remains in a high-impedance state except when scanning data.
		State meaning	Asserted/negated-The contents of the selected internal instruction or data register are shifted out on this signal on the falling edge of TCK. Remains in a high-impedance state except when scanning data.	
Timing	See IEEE 1149.1 standard for more details.			
TMS	I	JTAG test mode select.		
		<table border="0"> <tr> <td style="vertical-align: top;">State meaning</td> <td>Asserted/negated-Decoded by the internal JTAG TAP controller to distinguish the primary operation of the test support circuitry. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.</td> </tr> </table>	State meaning	Asserted/negated-Decoded by the internal JTAG TAP controller to distinguish the primary operation of the test support circuitry. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.
		State meaning	Asserted/negated-Decoded by the internal JTAG TAP controller to distinguish the primary operation of the test support circuitry. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor.	
Timing	See IEEE 1149.1 standard for more details.			

Table continues on the next page...

Table 23-7. JTAG test and other signals-detailed signal descriptions (continued)

Signal	I/O	Description
TRST_B	I	JTAG test reset.
		State meaning Asserted-Causes asynchronous initialization of the internal JTAG TAP controller. Must be asserted during power-on reset in order to properly initialize the JTAG TAP and for normal operation of the chip. An unterminated input appears as a high signal level to the test logic due to an internal pull-up resistor. Negated- Normal operation.
		Timing See IEEE 1149.1 standard for more details.
SCAN_MODE_B	I	Used for factory test. Refer to the device hardware specifications for proper treatment.
TEST_SEL_B	I	Used for factory test. Should be negated (pulled high) for normal operation.

23.3 Debug Memory Map/Register Definition

This table shows the memory-mapped debug and watchpoint registers of the chip. Any undefined 4-byte address spaces within offset 0x000-0xFFF are reserved.

Debug memory map

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
E_2000	Watchpoint monitor control register 0 (Debug_WMCR0)	32	R/W	0000_0000h	23.3.1/1883
E_2004	Watchpoint monitor control register 1 (Debug_WMCR1)	32	R/W	0000_0000h	23.3.2/1885
E_200C	Watchpoint monitor address register (Debug_WMAR)	32	R/W	0000_0000h	23.3.3/1886
E_2014	Watchpoint monitor address mask register (Debug_WMAMR)	32	R/W	0000_0000h	23.3.4/1886
E_2018	Watchpoint monitor transaction mask register (Debug_WMTMR)	32	R/W	0000_0000h	23.3.5/1887
E_201C	Watchpoint monitor status register (Debug_WMSR)	32	R/W	0000_0000h	23.3.6/1888
E_2040	Trace buffer control register 0 (Debug_TBCR0)	32	R/W	0000_0000h	23.3.7/1889
E_2044	Trace buffer control register 1 (Debug_TBCR1)	32	R/W	0000_0000h	23.3.8/1891
E_204C	Trace buffer address register (Debug_TBAR)	32	R/W	0000_0000h	23.3.9/1892
E_2054	Trace buffer address mask register (Debug_TBAMR)	32	R/W	0000_0000h	23.3.10/1892
E_2058	Trace buffer transaction mask register (Debug_TBTMR)	32	R/W	0000_0000h	23.3.11/1893
E_205C	Trace buffer status register (Debug_TBSR)	32	R/W	0000_0000h	23.3.12/1894

Table continues on the next page...

Debug memory map (continued)

Offset address (hex)	Register name	Width (in bits)	Access	Reset value	Section/page
E_2060	Trace buffer access control register (Debug_TBACR)	32	R/W	0000_0000h	23.3.13/1895
E_2064	Trace buffer access data high register (Debug_TBADHR)	32	R/W	0000_0000h	23.3.14/1895
E_2068	Trace buffer access data register (Debug_TBADR)	32	R/W	0000_0000h	23.3.15/1896
E_20A0	Programmed context ID register (Debug_PCIDR)	32	R/W	0000_0000h	23.3.16/1896
E_20A4	Current context ID register (Debug_CCIDR)	32	R/W	0000_0000h	23.3.17/1897
E_20B0	Trigger output source register (Debug_TOSR)	32	R/W	0000_0000h	23.3.18/1897

23.3.1 Watchpoint monitor control register 0 (Debug_WMCR0)

The watchpoint monitor control register 0 (WMCR0) control the specification of watchpoint monitor events.

Address: E_2000h base + 0h offset = E_2000h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R									Reserved							
W	EN	AMD	TMD	ECEN	NECEN	SIDEN	TIDEN									
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved				STRT				Reserved							
W	Reserved				STRT				Reserved							
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Debug_WMCR0 field descriptions

Field	Description
0 EN	Enable 0 Watchpoint monitor events are not flagged. 1 Watchpoint monitor events are flagged.

Table continues on the next page...

Debug_WMCR0 field descriptions (continued)

Field	Description
1 AMD	Address match disable. Qualifies address match as a watchpoint event criterion. 0 Address matching is used to recognize a watchpoint event. 1 Address matching does not affect watchpoint event detection.
2 TMD	Transaction match disable. Qualifies transaction type match (as defined in WMCR1[IFSEL] and WMTMR) as a watchpoint event criterion. 0 A transaction type match is used to recognize watchpoint events. 1 A transaction type match does not affect watchpoint event detection.
3 ECEN	Equal context enable. Qualifies the matching of current context with programmed context as a watchpoint event criterion, as written in the context registers described in Programmed context ID register (Debug_PCIDR) and Current context ID register (Debug_CCIDR) . NOTE: ECEN and NECEN must not be enabled in the same run. If both are set, watchpoint events are inhibited (never occur). 0 Current context match does not affect watchpoint event detection. 1 Watchpoint events are qualified by comparing current context with the programmed context event value.
4 NECEN	Not equal context enable. Qualifies the matching of current context with programmed context as a watchpoint event criterion, as written in the context registers described in Programmed context ID register (Debug_PCIDR) and Current context ID register (Debug_CCIDR) . NOTE: ECEN and NECEN must not be enabled in the same run. If both are set, watchpoint events are inhibited (never occur). 0 The failure of a current context match does not affect watchpoint event detection 1 Watchpoint events are qualified with NOT getting a current context compare with the programmed context event value.
5 SIDEN	Source ID enable 0 Source ID does not affect watchpoint event detection. 1 Watchpoint events are qualified by comparison with the programmed WMCR1(SID) value.
6 TIDEN	Target ID enable 0 Target ID does not affect watchpoint event detection. 1 Watchpoint events are qualified by comparison with the programmed WMCR1(TID) value.
7–20 -	This field is reserved. Reserved
21–23 STRT	Start condition. Specifies the event that arms the watchpoint monitor to start looking for the programmed event. 000 No event. Armed immediately 001 Trace buffer event is detected 010 Performance monitor signals overflow 011 TRIG_IN transitions from 0 to 1 100 TRIG_IN transitions from 1 to 0 101 Current context ID equals programmed context ID 110 Current context ID is not equal to programmed context ID 111 Reserved

Table continues on the next page...

Debug_WMCR0 field descriptions (continued)

Field	Description
24–31 -	This field is reserved. Reserved

23.3.2 Watchpoint monitor control register 1 (Debug_WMCR1)

The watchpoint monitor control register 1 (WMCR1) control the specification of watchpoint monitor events.

Address: E_2000h base + 4h offset = E_2004h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		

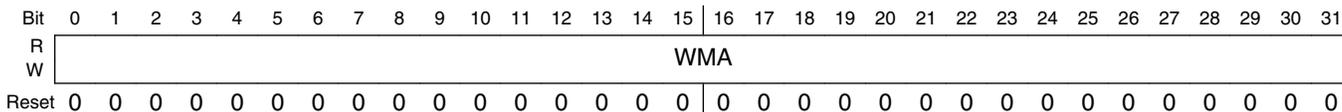
Debug_WMCR1 field descriptions

Field	Description
0–4 -	This field is reserved. Reserved
5–7 IFSEL	Interface selection. Selects the address, transaction type (as defined in WMTMR), and other attributes to be used for comparison 000 Selects e500 coherency module (ECM) dispatch interface 001 Reserved 010 Reserved 011 Reserved 100 Selects internal PCI Express 1 outbound interface 101 Selects internal PCI Express 2 outbound interface 110 Selects internal PCI Express 3 outbound interface 111 Reserved
8–10 -	This field is reserved. Reserved
11–15 SID	Source ID. Specifies the source ID associated with WMCR0[SIDEN]. For a definition of the source ID, see Table 23-27 .
16–26 -	This field is reserved. Reserved
27–31 TID	Target ID. Specifies the target ID associated with WMCR0[TIDEN]. For a definition of the target ID see Table 23-27 .

23.3.3 Watchpoint monitor address register (Debug_WMAR)

The watchpoint monitor address register (WMAR) contains the address to match against if WMCR[AMD] is clear. Note that this address may be further qualified with the bits described in [Watchpoint monitor address mask register \(Debug_WMAMR\)](#).

Address: E_2000h base + Ch offset = E_200Ch



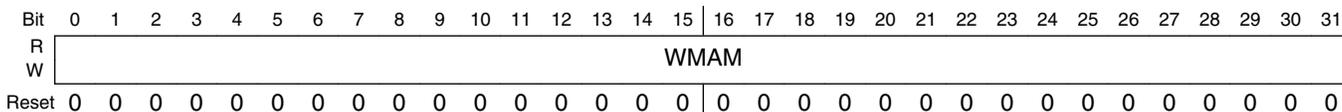
Debug_WMAR field descriptions

Field	Description
0–31 WMA	Watchpoint monitor address.

23.3.4 Watchpoint monitor address mask register (Debug_WMAMR)

The watchpoint monitor address mask register (WMAMR) contains the mask for the address in the WMAR.

Address: E_2000h base + 14h offset = E_2014h



Debug_WMAMR field descriptions

Field	Description
0–31 WMAM	Watchpoint monitor address mask. A value of zero masks the address comparison for the corresponding address bit. These bits only mask the address bits generated by the hardware, but do not affect the bits specified in WMAR. A bit that is masked from the comparison should be set to 0 in WMAR.

23.3.5 Watchpoint monitor transaction mask register (Debug_WMTMR)

The watchpoint monitor transaction mask register (WMTMR) specifies which transaction types to monitor. WMTMR allows users to qualify watchpoint events specifically with any combination of transaction types. As shown in [Table 23-13](#), each bit represents as many as four separate transaction types; one for each interface. Setting a bit enables watchpoint monitoring for the corresponding transaction types.

Because the supported transaction types vary by interface, the type designated by a WMTMR field also depends on the interface specified by WMCR1[IFSEL]. [Table 23-13](#) lists transaction types associated with each WMTMR bit by interface.

This table defines the transactions associated with each transaction mask bit for the different interfaces supported by the watchpoint monitor.

Table 23-13. Transaction Types by Interface

Bit	Description		
	e500 Coherency Module Dispatch	DDR Controller	PCI Express Outbound Transaction
0	Write with local processor snoop	Write	Posted Write
1	Write with no local processor snoop	-	Non-posted Write
2	Write with allocate(L2 stashing)	Write with allocate	-
3	Write with allocate and lock (L2 stashing with locking)	Write with allocate and lock	-
4-7	Reserved	-	-
8	Read with local processor snoop	Read	Read
9	Read with no local processor snoop	-	-
10	Read with unlock	Read with unlock	-
11	Reserved	-	Read Response
12	Reserved	-	-
13-15	Reserved	-	-
16	ATOMIC clear	ATOMIC clear	-
17	ATOMIC set	ATOMIC set	-
18	ATOMIC decrement	ATOMIC decrement	-
19	ATOMIC increment	ATOMIC increment	-
20-24	Reserved	-	-
25	Address only transaction	-	-

Table continues on the next page...

Table 23-13. Transaction Types by Interface (continued)

Bit	Description		
	e500 Coherency Module Dispatch	DDR Controller	PCI Express Outbound Transaction
26-31	Reserved	-	-

Address: E_2000h base + 18h offset = E_2018h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R																																
W																																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Debug_WMTMR field descriptions

Field	Description
0–31 WMTM	Watchpoint monitor transaction mask. Each bit corresponds to a transaction type as defined in Table 23-13 . The transaction associated with any particular bit may be different depending on the interface being monitored. A value of 1 for a given mask bit enables the matching of the transaction associated with that bit. These bits are meaningful only when WMCRO[TMD] = 0.

23.3.6 Watchpoint monitor status register (Debug_WMSR)

The watchpoint monitor status register (WMSR) indicates the state of the watchpoint monitor.

Address: E_2000h base + 1Ch offset = E_201Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R																
W	ACT	TRIG	Reserved													
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved															
W	Reserved															
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Debug_WMSR field descriptions

Field	Description
0 ACT	Active 0 The start triggering event has not occurred; watchpoint monitor is not armed. 1 The start triggering event has occurred; watchpoint monitor is armed.
1 TRIG	Triggered

Table continues on the next page...

Debug_WMSR field descriptions (continued)

Field	Description
	0 The programmed event in WMCR0 has not yet been triggered. 1 The programmed event in WMCR0 has been triggered at least once.
2–31 -	This field is reserved. Reserved

23.3.7 Trace buffer control register 0 (Debug_TBCR0)

The trace buffer control register 0 (TBCR0) specify trace buffer events.

Address: E_2000h base + 40h offset = E_2040h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R									Reserved							
W	EN	AMD	TMD	ECEN	NECEN	SIDEN	TIDEN	HALT	Reserved						MODE	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								Reserved							
W	Reserved				STRT				Reserved						STOP	
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Debug_TBCR0 field descriptions

Field	Description
0 EN	Enable 0 The trace buffer facility is disabled. 1 The trace buffer facility is enabled.
1 AMD	Address match disable 0 The address match is used to qualify a trace buffer event. 1 The address match is ignored when detecting a trace buffer event.
2 TMD	Transaction match disable 0 The transaction type match is used to qualify a trace buffer event. 1 The transaction type match is ignored when detecting a trace buffer event.
3 ECEN	Equal context enable. Qualifies the matching of current context with programmed context as a trace buffer event criterion, as written in the context registers described in Programmed context ID register (Debug_PCIDR) and Current context ID register (Debug_CCIDR) .

Table continues on the next page...

Debug_TBCR0 field descriptions (continued)

Field	Description
	<p>NOTE: ECEN and NECEN must not be enabled in the same run. If both are set, watchpoint events are inhibited (never occur).</p> <p>0 Current context match does not affect trace buffer event detection</p> <p>1 Trace buffer events are qualified by comparing current context with the programmed context event value.</p>
4 NECEN	<p>Not equal context enable. Qualifies the matching of current context with programmed context as a trace buffer event criterion, as written in the context registers described in Programmed context ID register (Debug_PCIDR) and Current context ID register (Debug_CCIDR).</p> <p>NOTE: ECEN and NECEN must not be enabled in the same run. If both are set, watchpoint events are inhibited (never occur).</p> <p>0 The failure of a current context match does not affect trace buffer event detection</p> <p>1 trace buffer events are qualified with NOT getting a current context compare with the programmed context event value.</p>
5 SIDEN	<p>Source ID enable</p> <p>0 Trace buffer events ignore the programmed source ID value.</p> <p>1 Trace buffer events are qualified by comparison with the programmed SID event value.</p>
6 TIDEN	<p>Target ID enable</p> <p>0 Trace buffer events ignore the programmed TID event value.</p> <p>1 Trace buffer events are qualified by comparison with the programmed TID event value. This comparison only applies when the ECM is selected for tracing (TBCR1[IFSEL] is all zeros).</p>
7 HALT	<p>Halt causes the trace buffer to stop tracing immediately. TBSR[ACT] remains set when this bit is set.</p>
8–13 -	<p>This field is reserved. Reserved</p>
14–15 MODE	<p>Trace mode. Specifies one of two trace modes.</p> <p>00 Trace every valid transaction</p> <p>01 Reserved</p> <p>10 Trace only cycles in which a trace event is detected. Note that if EN and other TBCR0 fields are not properly programmed to specify a traceable event, tracing occurs for every valid address.</p> <p>11 Reserved</p>
16–20 -	<p>This field is reserved. Reserved</p>
21–23 STRT	<p>Start condition. Specifies the event that arms the trace buffer to start looking for the programmed event</p> <p>000 No event. Armed immediately</p> <p>001 Watchpoint monitor event is detected</p> <p>010 Trace buffer event is detected</p> <p>011 Performance monitor signals overflow</p> <p>100 TRIG_IN transitions from 0 to 1</p> <p>101 TRIG_IN transitions from 1 to 0</p> <p>110 Current context ID equals programmed context ID</p> <p>111 Current context ID does not equal programmed context ID</p>
24–28 -	<p>This field is reserved. Reserved</p>

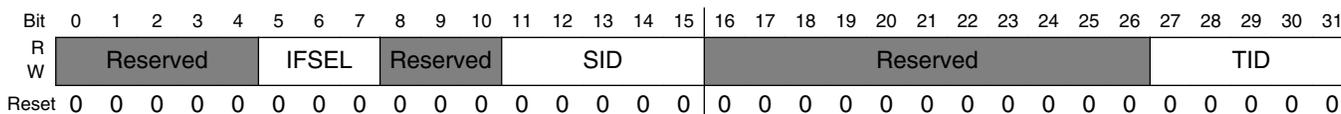
Table continues on the next page...

Debug_TBCR0 field descriptions (continued)

Field	Description
29–31 STOP	Trace stop mode. Specifies the event that stops the updating of the trace buffer after it has been started. Trace buffer only stops after it has been triggered at least once.
000	Buffer is full
001	Watchpoint monitor event is detected
010	Trace buffer event is detected
011	Performance monitor signals overflow
100	TRIG_IN transitions from 0 to 1
101	TRIG_IN transitions from 1 to 0
110	Current context ID equals programmed context ID
111	Current context ID does not equal programmed context ID

23.3.8 Trace buffer control register 1 (Debug_TBCR1)

Address: E_2000h base + 44h offset = E_2044h



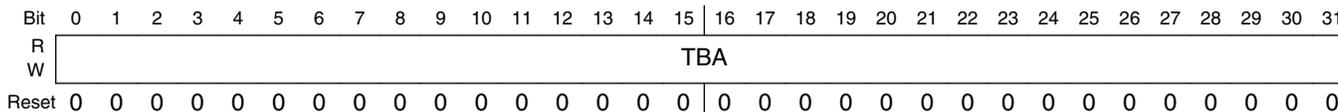
Debug_TBCR1 field descriptions

Field	Description
0–4 -	This field is reserved. Reserved
5–7 IFSEL	Interface selection. Specifies the interface that sources information for both comparison/buffer control and buffer data capture.
000	Selects e500 coherency module (ECM) dispatch interface
001	Reserved
010	Reserved
011	Reserved
100	Selects internal PCI Express 1 outbound interface
101	Selects internal PCI Express 2 outbound interface
110	Selects internal PCI Express 3 outbound interface
111	Reserved
8–10 -	This field is reserved. Reserved
11–15 SID	Source ID. Specifies the source ID associated with TBCR0[SIDEN]. The source ID is defined in Table 23-27 .
16–26 -	This field is reserved. Reserved
27–31 TID	Target ID. Specifies the target ID associated with TBCR0[TIDEN]. The target ID is defined in Table 23-27 .

23.3.9 Trace buffer address register (Debug_TBAR)

The trace buffer address register (TBAR) contains the address to match against (if TBCR0[AMD] is zero). This address may be further qualified by the mask bits defined in [Trace buffer address mask register \(Debug_TBAMR\)](#) ."

Address: E_2000h base + 4Ch offset = E_204Ch



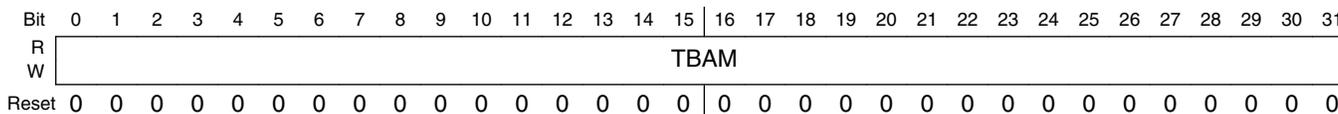
Debug_TBAR field descriptions

Field	Description
0–31 TBAR	Trace buffer address.

23.3.10 Trace buffer address mask register (Debug_TBAMR)

The trace buffer address mask register (TBAMR) contains a mask for the TBAR, which allows excluding address bits from the comparison.

Address: E_2000h base + 54h offset = E_2054h



Debug_TBAMR field descriptions

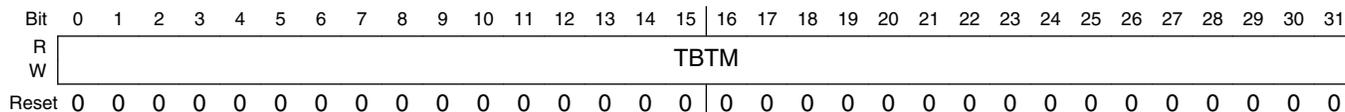
Field	Description
0–31 TBAM	Trace buffer address mask. A value of zero masks the address comparison for the corresponding address bit. These bits only mask the address bits generated by the hardware, but do not affect the bits specified in TBAR. A bit that is masked from the comparison should be set to 0 in TBAR.

23.3.11 Trace buffer transaction mask register (Debug_TBTMR)

The trace buffer transaction mask register (TBTMR) specifies which transaction types to monitor. Each bit in the TBTMR represents a transaction type on the selected interface. The transaction associated with any particular bit depends on the interface being monitored as specified by TBCR1[IFSEL]. Note that the transactions used for defining trace buffer events are the same as those defined for watchpoint monitor events. Thus, [Table 23-13](#) defines the transaction types associated with each interface. Setting a bit enables a hit when this transaction is matched (provided all other match criteria are met and TBCR0[TMD] is clear).

Different interfaces support different transaction types, and the same bit may represent different transaction types depending on the interface.

Address: E_2000h base + 58h offset = E_2058h



Debug_TBTMR field descriptions

Field	Description
0–31 TBTM	Trace buffer transaction mask. Each bit corresponds to a transaction type as defined in Table 23-13 . The transaction associated with a bit depends on the interface being monitored. A value of 1 for a given mask bit enables the matching of the transaction associated with that bit. These bits are meaningful only when TBCR0[TMD] = 0.

23.3.12 Trace buffer status register (Debug_TBSR)

The trace buffer status register (TBSR) indicates the operational state of the trace buffer.

Address: E_2000h base + 5Ch offset = E_205Ch

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R					Reserved											
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								C_INDx							
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Debug_TBSR field descriptions

Field	Description
0 ACT	Active. Indicates trace buffer activity. 0 The start triggering event has not yet occurred. Trace buffer is not armed. 1 The start triggering event has occurred. Trace buffer is armed.
1 TRIG	Triggered. Indicates whether or not a programmed event has been triggered. 0 The programmed event in TBCR0 has not yet been triggered. 1 The programmed event in TBCR0 has been triggered at least once.
2 STP	Stopped. Indicates whether or not a trace buffer stop condition has been detected. 0 No stop condition yet detected. 1 The trace buffer has detected a stop condition and is no longer capturing events.
3 WRAP	Wrapped. Indicates that the trace buffer write pointer has wrapped to the beginning of the buffer at least once. Set when the last entry of the trace buffer is written. 0 Pointer has not yet wrapped. 1 Pointer has wrapped to the beginning at least once.
4–23 -	This field is reserved. Reserved
24–31 C_INDx	Current index. Represents the current value of the write pointer at the time TBSR was read. This value may be written by software to initialize the write pointer; however, software is not allowed to write to the write pointer while the trace buffer is active. Writes are ignored while the trace buffer is active. It is recommended to write the status register before enabling the trace buffer in order to zero out any bits that might have been set during a prior run and to initialize the write pointer to zero.

23.3.13 Trace buffer access control register (Debug_TBACR)

The trace buffer access control register (TBACR) enables software to read or write the trace buffer. Each entry is 64 bits; therefore, it takes one write of TBACR and two reads of the access data register (TBADR and TBADHR) to read one 256-entry array entry. Similarly, it takes one write of TBACR and two writes of TBADR and TBADHR to write one array entry. Software can access any entry by writing the appropriate index into TBACR[INDX]. To read or write the buffer sequentially, starting with entry 0, the index must start with a value of 0 and increment every time a new entry is accessed.

Address: E_2000h base + 60h offset = E_2060h

Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R			Reserved													
W	RD	WR														
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Bit	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Reserved								INDX							
W																
Reset	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Debug_TBACR field descriptions

Field	Description
0 RD	Read command. When set, a trace buffer read is performed using the value of TBACR[INDX]. This bit is automatically cleared when the read is performed.
1 WR	Write command. When set, a trace buffer write is performed using the value of TBACR[INDX]. This bit is automatically cleared when the write is performed. A write occurs only if the trace buffer is not active: write requests are ignored while the buffer is active.
2–23 -	This field is reserved. Reserved
24–31 INDX	Buffer index to read from or write into (0-255). Used in conjunction with TBACR[RD] and TBACR[WR].

23.3.14 Trace buffer access data high register (Debug_TBADHR)

The trace buffer access data high register (TBADHR) contains one of the following:

- High-order 32 bits of the data read from the trace buffer during a software-initiated read command (TBACR[RD])
- Or, data to be written into the trace buffer during a software-initiated write command (TBACR[WR]).

Debug Memory Map/Register Definition

TBACR must be configured to perform a read before TBADHR contains valid data. This register must be initialized by software before configuring the TBACR to perform a write command.

Address: E_2000h base + 64h offset = E_2064h



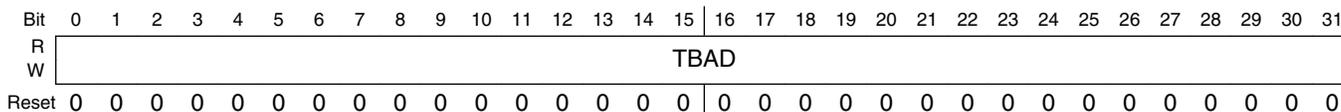
Debug_TBADHR field descriptions

Field	Description
0–31 TBADH	Trace buffer access data high. The higher 32 bits of the data read from or to be written into the trace buffer, depending on whether the array is accessed with a read or a write.

23.3.15 Trace buffer access data register (Debug_TBADR)

The trace buffer access data register (TBADR) contains the low-order 32 bits of the data read from the trace buffer during a software-initiated read command (TBACR[RD]) or the write data to be written into the trace buffer during a software-initiated write command (TBACR[WR]). TBACR must be configured to perform a read before this register contains valid data. This register must be initialized by software before configuring the TBACR to perform a write command.

Address: E_2000h base + 68h offset = E_2068h



Debug_TBADR field descriptions

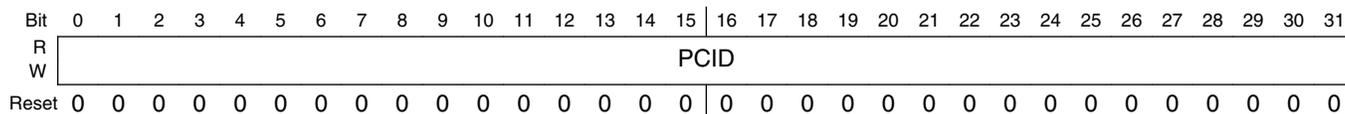
Field	Description
0–31 TBAD	Trace buffer access data. Corresponds to the lower 32 bits of the data read from the trace buffer or to be written into the trace buffer, depending on whether software is accessing the array with a read or a write.

23.3.16 Programmed context ID register (Debug_PCIDR)

This section describes the context ID registers. The current context ID register (CCIDR) and programmed context ID registers (PCIDR) are set by software and facilitate debugging complex software.

The programmed context ID register (PCIDR) contains the user-programmed context ID. This register can be configured to trigger watchpoint events when its value matches the current context ID register (CCIDR), as controlled by WMCR0[ECEN] and WMCR0[NECEN]. See [Watchpoint monitor control register 0 \(Debug_WMCR0\)](#)," for more information.

Address: E_2000h base + A0h offset = E_20A0h



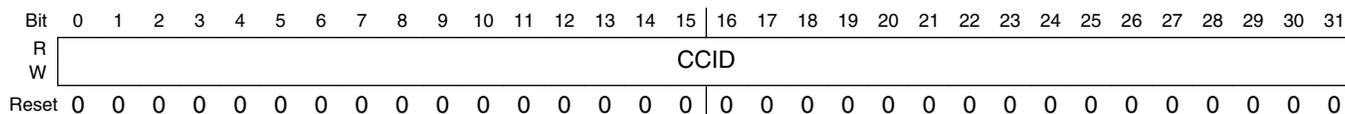
Debug_PCIDR field descriptions

Field	Description
0–31 PCID	Programmed context ID. Contains the user-programmed context ID. Compared with current context ID for context-sensitive event triggering

23.3.17 Current context ID register (Debug_CCIDR)

The current context ID register (CCIDR) contains the current context ID. This register is written by software after a context switch and can be used to trigger events when compared with the programmed context ID register (PCIDR).

Address: E_2000h base + A4h offset = E_20A4h



Debug_CCIDR field descriptions

Field	Description
0–31 CCID	Current context ID. Set by user software. Typically loaded immediately following a context switch. Compared with user-programmed context ID for context-sensitive event triggering

23.3.18 Trigger output source register (Debug_TOSR)

TRIG_OUT provides a convenient mechanism for triggering external system monitors and diagnostic equipment such as logic analyzers. Note that READY is multiplexed with TRIG_OUT. See the last paragraph of [Power-on reset sequence](#) for more information about READY functionality.

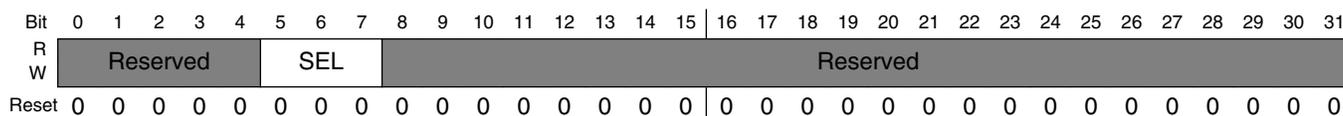
functional description

When the trace buffer hit is selected by TOSR[SEL], TRIG_OUT is only meaningful if the trace buffer control register 0 (TBCR0) is properly configured to hit on a traceable event. The same holds true for the watchpoint monitor when the watchpoint monitor is selected by TOSR[SEL].

The trigger out source register (TOSR) specifies the source for TRIG_OUT. The three event-trigger sources are the following:

- The watchpoint monitor
- The trace buffer
- The performance monitor

Address: E_2000h base + B0h offset = E_20B0h



Debug_TOSR field descriptions

Field	Description
0–4 -	This field is reserved. Reserved
5–7 SEL	Select. Selects the source for TRIG_OUT 000 READY signal. Multiplexed with TRIG_OUT. Basic device state indicator. READY asserts whenever the device is not in reset or not asleep. See Reset, Clocking, and Initialization for more details about the reset sequence, and Global Utilities for more information about power management states. 001 Selects the watchpoint monitor hit indication 010 Selects the trace buffer hit indication 011 Selects the performance monitor overflow indication
8–31 -	This field is reserved. Reserved

23.4 Functional description

The debug features on the chip use the eLBC interfaces and the DDR SDRAM interfaces.

23.4.1 Source and target ID

Debug information that is common to all the interfaces is the source ID (SID).

The transaction source ID provides enough information to determine which block or port originated a transaction including the distinction between instruction and data fetches from the processor core. The table below shows the values and interpretation for the 5-bit

SID field. Note that the table also includes ports that are only slaves, such as local memory. These ports are always targets. As such, the value shown represents a target ID (TID) and not a source ID. For ports that can function in both capacities, the value indicates source ID when mastering transactions, and target ID when responding as slave. The TID field is only meaningful when one of the following participates in the transaction:

- The e500 coherency module (ECM) dispatch bus
- The watchpoint monitor (WMCR1[IFSEL] = 000)
- The trace buffer (TBCR1[IFSEL] = 000)

Table 23-27. Source and target ID values

ID value (Hex)	Transaction source	Transaction target
00	PCI Express 3	PCI Express 3
01	PCI Express 2	PCI Express 2
02	PCI Express 1	PCI Express 1
03	Reserved	Reserved
04	Reserved	Local bus controller
05	Universal serial bus (dual role)	Reserved
06	Reserved	Reserved
07	Security	Reserved
08	Reserved	Configuration space
09	Reserved	Reserved
0A	Boot sequencer	Reserved
0B	eSDHC	Reserved
0C	Serial RapidIO 1	Serial RapidIO 1
0D	Serial RapidIO 2	Serial RapidIO 2
0E	Reserved	Reserved
0F	Reserved	DDR
10	Local processor 0 (instruction fetch)	Reserved
11	Local processor 0 (data fetch)	Reserved
12	Local processor 1 (instruction fetch)	Reserved
13	Local processor 1 (data fetch)	Reserved
14	Reserved	Reserved
15	Direct memory access controller 1 (DMA 1)	Reserved
16	Direct memory access controller 2 (DMA 2)	Reserved
17	System access port (SAP)	Reserved
18	Enhanced three-speed Ethernet controller 1 (eTSEC 1)	Reserved
19	Enhanced three-speed Ethernet controller 2 (eTSEC 2)	Reserved
1A	Enhanced three-speed Ethernet controller 3 (eTSEC 3)	Reserved
1B	Reserved	Reserved

Table continues on the next page...

Table 23-27. Source and target ID values (continued)

ID value (Hex)	Transaction source	Transaction target
1C	Serial RapidIO message unit, Serial RapidIO doorbell unit, Serial RapidIO port-write unit	Reserved
1D	Reserved	Reserved
1E	Reserved	Reserved
1F	Reserved	Reserved

23.4.2 DDR SDRAM interface debug

The DDR interface has two debug modes distinguished by which pins drive the debug information.

In one mode, debug information (source ID, data valid) is multiplexed onto the ECC pins; the other mode uses the debug pins.

23.4.2.1 Debug information on debug pins

If DMA2_DACK0_B is high when sampled during POR, the debug information from the DDR SDRAM controller is driven on MSRCID[0:4] and MDVAL.

This POR value is captured in PORDBGMSR[MEM_SEL] as described in [POR debug mode status register \(GUTS_PORDBGMSR\)](#). In this mode, the source ID appears on MSRCID[0:4] during a RAS or CAS cycle. During any other cycle, the value of MSRCID[0:4] is all ones, which indicates idle cycles on the address/command interface. Similarly, MDVAL is asserted during valid data cycles on the DDR interface.

23.4.2.2 Debug information on ECC pins

If DMA2_DDONE0_B is low when sampled during POR, debug information from the DDR SDRAM interface is selected to appear on MECC[0:5] as shown in [Figure 23-1](#). In this mode, the ID value of the source port (the source ID) appears on MECC[0:4] during a RAS or CAS cycle. During any other cycle, the value of MECC[0:4] is all ones. A data-valid signal (DVAL) is driven on MECC5 during valid DDR SDRAM data cycles.

NOTE

In this mode, MECC[0:5] must be disconnected from all SDRAM devices to prevent contention on those lines.

23.4.3 Local bus interface debug

If DMA2_DACK0_B is low when sampled during POR, the eLBC is selected as the source for the debug information appearing on MSRCID[0:4] and MDVAL.

For more information on this mode, see [Source ID debug mode](#).

23.4.4 Watchpoint monitor

The watchpoint monitor (WM) can be programmed to arm and trigger on many different events including any of the following:

- External event (through TRIG_IN)
- A trace buffer event
- A performance monitor overflow event
- A comparison of the current and programmed context ID registers

A watchpoint event can be used in the following ways:

- Trigger a logic analyzer (using TRIG_OUT)
- Arm or trigger the trace buffer
- Trigger a performance monitor event

The large counters available in the performance monitor block and the interlock between it and the watchpoint monitor support sophisticated debug scenarios.

A WM trigger event may be composed of several events programmed in the watchpoint monitor control registers (WMCR0-WMCR1). Because the watchpoint monitor is disabled by default during POR, these registers must be initialized to make use of this debug feature. Note that the WM address mask register (WMAMR) and the type mask register (WMTMR) are cleared during POR. This means that the watchpoint monitor's default behavior following a power-on reset is to trigger on any address and no transaction type. The reset value of WMCR0[TMD] is 0 which means transaction matching is enabled but since no transaction is selected (WMTMR = 0), a match will never occur. Either the transaction matching must be disabled by setting WMCR0[TMD] to a value of 1, or valid transactions must be selected by setting one or more of the WMTMR bits to a value of 1.

23.4.4.1 Watchpoint monitor performance monitor events

The WM can produce a performance monitor (PM) event with every trigger.

This is accomplished by configuring the performance monitor to count WM events. For more information on this configuration, see the events named 'Number of watchpoint monitor hits' and 'Number of trace buffer hits' in [Table 22-43](#).

Multi-level triggers can be created using the watchpoint monitor, the performance monitor, and the trace buffer combined. For example, the WM can be programmed to trigger on events that also increment a PM counter (the performance monitor must also be programmed to respond to this event), the output of which (`perfmon_overflow`) could trigger the start of tracing in the trace buffer.

23.4.5 Trace buffer

The trace buffer is a 256 entry x 64-bit array that can capture information about the internal processing of transactions to selected interfaces.

The trace buffer controls are a superset of those for the watchpoint monitor. Close inspection of the trace buffer control registers (`TBCR n`) and the watchpoint monitor control registers (`WMCR n`) shows that trace buffer controls not needed for the watchpoint monitor are marked reserved in `WMCR n` . This permits using the trace buffer as a second watchpoint monitor by simply ignoring the trace options.

The trace buffer provides great flexibility about when to start tracing, when to stop tracing, and what to trace. The trace mode field, `TBCR0[MODE]`, indicates when to trace: on every valid cycle, on a watchpoint monitor event, or when all the programmed events in the TBCR are met. This permits a user to program the trace condition in the watchpoint monitor and to program a start or stop condition in the trace buffer control register. The user can also program the TBCR with the conditions in which to stop tracing: on an event, or when the buffer is full. `TBCR0[IFSEL]` specifies which interface transactions are being captured.

The trace buffer can be programmed from e500 coherency module (ECM) to trace the dispatch bus.

Transactions come into the ECM, arbitrate for common resources, and get dispatched to the target port. Information such as transaction types, source ID, and other attributes can be captured in any of the selected interfaces.

Table 23-30. PCI Express trace buffer entry (continued)

R es et	All zeros
---------------	-----------

This table describes the fields of PCI Express trace buffer entries when TBCR1[IFSEL] = 100 or 101 or 110.

Table 23-31. PCI Express trace buffer entry field descriptions

Bits	Name	Function
0-4	PEXTT	Transaction type. Specifies the transaction type as shown in Watchpoint monitor transaction mask register (Debug_WMTMR) . For example, a value of all zeros maps to write.
5-9	PEXSID	Source ID. Identifies the source of the transaction as shown in Table 23-27 . For example, a value of 010101 identifies DMA as the transaction source. For responses, this corresponds to Requestor's ID's bus number bits 3-7.
10-14	PCIBC	Byte count. The size of the transaction. 00000 4 bytes 00001 8 bytes 00010 12 bytes ... 11111 256 bytes
15-33	-	Reserved
34-63	PEXADDR	Address bits 31-2

23.5 Initialization

Configuring the appropriate control register must be the last step in the initialization sequence for either the watchpoint or trace buffer.

That is, all required registers except the corresponding control register must be configured before any control register bits that enable watchpoint or trace events are set.

Appendix A

P2010-A Single Core Version

The P2020 is also available in a single core version known as the P2010.

A.1 Overview

The P2020 is also available in a single core version known as the P2010 and shown in the figure below.

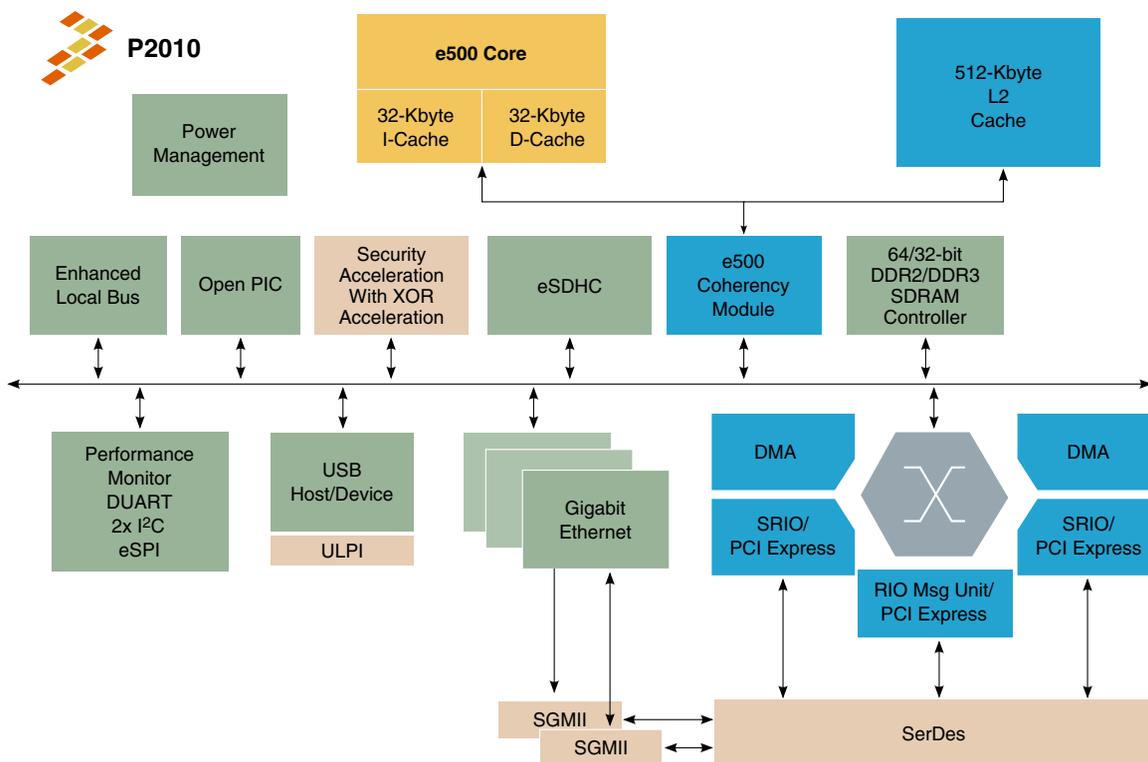


Figure A-1. P2010 Block Diagram

Appendix B

Terminology, Conventions, and Resources

The primary objective of this document is to define the functionality of the *P2020*.

B.1 About this content

The primary objective of this document is to define the functionality of the *P2020*. The *P2020* provides integration of processing power for networking and communications peripherals, resulting in higher device performance. The *P2020* contains dual e500v2 processor core built on *Power Architecture*[®] technology with system logic required for networking. The *e500v2* processor core is a low-power implementation of the family of reduced instruction set computing (RISC) embedded processors that implement the embedded category features of the Power Architecture technology. This document is intended as a companion to the *PowerPC[™] e500 Core Family Reference Manual*.

NOTE

The diagrams in this content are provided to aid in understanding the overall functionality and features of this product. They do not depict the implementation details of the product, which are subject to change.

B.2 Audience

It is assumed that the reader understands operating systems, microprocessor system design, and the basic principles of RISC processing.

B.3 Acronyms and abbreviations

This table describes commonly-used acronyms and abbreviations used in this document.

Table B-1. Acronyms and abbreviations

Acronym/ Abbreviation	Meaning
ADB	Allowable disconnect boundary
ATM	Asynchronous transfer mode
ATMU	Address translation and mapping unit
BD	Buffer descriptor
BIST	Built-in self test
BRI	Basic rate interface
BTB	Branch target buffer
BUID	Bus unit ID
CAM	Content-addressable memory
CCB	Core complex bus
CCSR	Configuration control and status register
CEPT	Conference des administrations Europeanes des Postes et Telecommunications (European Conference of Postal and Telecommunications Administrations)
COL	Collision
CRC	Cyclic redundancy check
CRS	Carrier sense
DDR	Double data-rate
DMA	Direct memory access
DPLL	Digital phase-locked loop
DRAM	Dynamic random access memory
DUART	Dual universal asynchronous receiver/transmitter
EA	Effective address
ECC	Error checking and correction
ECM	e500 coherency module
EEST	Enhanced Ethernet serial transceiver
EHPI	Enhanced host port interface
EPROM	Erasable programmable read-only memory
FCS	Frame-check sequence
FEC	10/100 fast Ethernet controller
GCI	General circuit interface
GMII	Gigabit media-independent interface
GPCM	General-purpose chip-select machine
GPIO	General-purpose I/O
GPR	General-purpose register
GUI	Graphical user interface
HDLC	High-level data link control
I2C	Inter-integrated circuit
IDL	Inter-chip digital link
IEEE	Institute of Electrical and Electronics Engineers

Table continues on the next page...

Table B-1. Acronyms and abbreviations (continued)

Acronym/ Abbreviation	Meaning
IPG	Interpacket gap
IrDA	Infrared Data Association
ISDN	Integrated services digital network
ITLB	Instruction translation lookaside buffer
IU	Integer unit
JTAG	Joint Test Action Group
LAE	Local access error
LAW	Local access window
LBC	Local bus controller
LIFO	Last-in-first-out
LRU	Least recently used
LSB	Least significant byte
lsb	Least significant bit
LSU	Load/store unit
MAC	Multiply accumulate, media access control
MDI	Medium-dependent interface
MESI	Modified/exclusive/shared/invalid-cache coherency protocol
MII	Media independent interface
MMU	Memory management unit
MSB	Most significant byte
msb	Most significant bit
NMSI	Nonmultiplexed serial interface
No-op	No operation
OCeaN	On-chip network
OSI	Open systems interconnection
PCI	Peripheral component interconnect bus
PCMCIA	Personal Computer Memory Card International Association
PCS	Physical coding sublayer
PIC	Programmable interrupt controller
PMA	Physical medium attachment
PMD	Physical medium dependent
POR	Power-on reset
PRI	Primary rate interface
RGMII	Reduced gigabit media-independent interface
RISC	Reduced instruction set computing
RTOS	Real-time operating system
RWITM	Read with intent to modify
RMW	Read-modify-write
Rx	Receiver

Table continues on the next page...

Table B-1. Acronyms and abbreviations (continued)

Acronym/ Abbreviation	Meaning
RxBD	Receive buffer descriptor
SCC	Serial communication controller
SCP	Serial control port
SDLC	Synchronous data link control
SDMA	Serial DMA
SFD	Start frame delimiter
SI	Serial interface
SIU	System interface unit
SMC	Serial management controller
SNA	Systems network architecture
SPI	Serial peripheral interface
SPR	Special-purpose register
SRAM	Static random access memory
TAP	Test access port
TBI	Ten-bit interface
TDM	Time-division multiplexed
TLB	Translation lookaside buffer
TSA	Time-slot assigner
TSEC	Three-speed Ethernet controller
Tx	Transmit
TxBD	Transmit buffer descriptor
UART	Universal asynchronous receiver/transmitter
UPM	User-programmable machine
UTP	Unshielded twisted pair
VA	Virtual address
ZBT	Zero bus turnaround

B.4 Notational conventions

This table shows notational conventions used in this content.

Table B-2. Notational conventions

Convention	Definition
General	
Cleared	When a bit takes the value zero, it is said to be cleared.
Set	When a bit takes the value one, it is said to be set.

Table continues on the next page...

Table B-2. Notational conventions (continued)

Convention	Definition
mnemonics	Instruction mnemonics are shown in lowercase bold.
<i>italics</i>	Italics can indicate the following: <ul style="list-style-type: none"> • Variable command parameters, for example, bcctrx • Titles of publications • Internal signals, for example, <i>core_int</i>
0x	Prefix to denote hexadecimal number
h	Suffix to denote hexadecimal number
0b	Prefix to denote binary number
rA, rB	Instruction syntax used to identify a source GPR
rD	Instruction syntax used to identify a destination GPR
REGISTER[FIELD]	Abbreviations for registers are shown in uppercase text. Specific bits, fields, or ranges appear in brackets. For example, MSR[LE] refers to the little-endian mode enable bit in the machine state register.
x	In some contexts, such as signal encodings, an unitalicized x indicates a don't care.
<i>x</i>	An italicized x indicates an alphanumeric variable.
<i>n</i>	An italicized n indicates either: <ul style="list-style-type: none"> • An integer variable • A general-purpose bitfield unknown
¬	NOT logical operator
&	AND logical operator
	OR logical operator
	Concatenation, for example, TCR[WPEXT] TCR[WP]
Signals	
OVERBAR	An overbar indicates that a signal is active-low.
_B	Suffix to denote a signal is active-low.
<i>lowercase_italics</i>	Lowercase italics is used to indicate internal signals
lowercase_plaintext	Lowercase plain text is used to indicate signals that are used for configuration.
Register access	
Reserved	Ignored for the purposes of determining access type
R/W	Indicates that all non-reserved fields in a register are read/write
R	Indicates that all non-reserved fields in a register are read only
W	Indicates that all non-reserved fields in a register are write only
w1c	Indicates that all non-reserved fields in a register are cleared by writing ones to them

B.5 Related resources

This table shows related resources that may be helpful. Additional literature is published as new processors become available. For current literature, visit www.freescale.com or contact your local FAE.

Table B-3. Related resources

Resource	Purpose
P2020 Security (SEC 3.1) Reference Manual, Rev 0	Provides specific details of security engine (SEC 3.1).
P2020 QorIQ Integrated Processor Hardware Specifications or P2010 QorIQ Integrated Processor Hardware Specifications (for single core chip)	Provides specific data regarding bus timing, signal behavior, and AC, DC, and thermal characteristics, as well as other design considerations.

Appendix C

Substantive Changes from Revision 1 to Revision 2

This appendix contains the revision history for P2020.

C.1 Substantive changes from revision 1 to revision 2

Substantive changes from revision 1 to revision 2 are as follows:

C.1.1 Overview Revision History

Reference	Description
Critical performance parameters	Removed "Operating junction temperature (T_j) range: 0-125 and -40 to 125°C (industrial specification)."

C.1.2 Memory Map Revision History

Reference	Description
Precedence of local access windows	Added the section "Precedence of Local Access Windows."

C.1.3 Signal Descriptions Revision History

Reference	Description
Output signal states during reset	In Section, "Output Signal States During Reset," replaced the word "inactive state" to "tri-state" in the second paragraph.

C.1.4 Reset, Clocking, and Initialization Revision History

Reference	Description
eSPI controller configuration	Removed Figure, “eSPI CS0 Mode Register (SPMODE0) Configuration,” and replace the paragraph preceding the Figure with the following: “For default eSPI CS0 mode register (SPMODE0) configuration, see Section 19.3.1.7, “eSPI CS0 Mode Register (SPMODE0). During ROM boot, the ROM code needs to change the reset value of eSPI CS0 mode register to 0x3117_2210.”
Default e500 addressing during system boot	Updated the second sentence of the second paragraph to say “with the exception that the TLBI Entry 0 is replaced with the following:”
Accessing configuration, control, and status registers	Modified the CCSRBAR value in the fourth sentence of the first paragraph to to 0xFF70_0000 (CCSRBAR = 0x000F_F700).
DDR PLL Ratio	From the second sentence of the first paragraph, deleted the last few words saying “...when clock select is zero”.
Power-on reset sequence	Added a new fourth step to the POR sequence saying “ System negates HRESET after its required hold time and after POR configuration inputs have been valid for at least 4 SYSCLK cycles.”
Figure 4-5	Updated the figure
eTSECn physical interface mode configuration	Updated the table “eTSEC2 POR configuration summary” for MII, GMII, and TBI and removed the row on none. Also, removed the row on none from “eTSEC3 POR configuration summary.”

C.1.5 e500 Core Integration Details Revision History

Reference	Description
Summary of core integration details	Updated ERR_INT_EN[MBEE] must be zero’ as ‘ERR_INT_EN[MBEE] must be set’ in second row of ‘HID1Implementation’ description.
Table 5-1	Deleted the last two sentences of QorIQ implementation of the TCR (timer control register) feature, and added the following sentences instead: 10 Second timeout generates HRESET_REQ output externally 11 Second timeout automatically resets the given processor core
Summary of core integration details	Updated 01 for TCR (timer control register).

C.1.6 L2 Look-Aside Cache/SRAM Revision History

Reference	Description
Allocation of lines	Replaced “(1024 sets by eight ways) to “(2048 sets by eight ways)”.
Accessing the on-chip array as an L2 cache	Updated the image “Physical address usage for L2 cache accesses”.

C.1.7 DDR Memory Controllers Revision History

Reference	Description
Memory data path read capture ECC (DDR_CAPTURE_ECC)	Updated the description for DDR_CAPTURE_ECC[ECE] bit.
- DDR SDRAM control configuration (DDR_DDR_SDRAM_CFG) - Memory error disable (DDR_ERR_DISABLE)	- In DDR_SDRAM_CFG[ECC_EN] description, updated 'ERR_INT_EN[MBEE] must be zero' to 'ERR_INT_EN[MBEE] must be set.' - In ERR_DISABLE[MBED]description, updated 'ERR_INT_EN[MBEE] must be zero' to 'ERR_INT_EN[MBEE] must be set'.
Single-Bit ECC memory error management (DDR_ERR_SBE)	Modified the ERR_SBE[SBEC] description to: "Single-bit error counter. Indicates the number of single-bit errors detected and corrected since the last error report. If single-bit error reporting is enabled, an error is reported and a regular or critical interrupt is generated when this value equals SBET. SBEC is automatically cleared when the threshold value is reached."
Memory error disable (DDR_ERR_DISABLE)	Changed the sentence in the ERR_DISABLE[MBED] description from: "If RFXE is zero and this error occurs, MBED and ERR_INT_EN[MBEE] must be zero and ECC_EN must be one to ensure that an interrupt is generated" to "If RFXE is zero and this error occurs, ERR_DISABLE[MBED] must be zero and ERR_INT_EN[MBEE] and ECC_EN must be one to ensure that an interrupt is generated".
Memory error interrupt enable (DDR_ERR_INT_EN)	Changed the sentence in the ERR_INT_EN[MBEE] description from: "If RFXE is zero and this error occurs, MBEE and ERR_DISABLE[MBED] must be zero and DDR_SDRAM_CFG[ECC_EN] must be set to ensure that an interrupt is generated" to "If RFXE is zero and this error occurs, ERR_DISABLE[MBED] must be zero and ERR_INT_EN[MBEE] and ECC_EN must be one to ensure that an interrupt is generated".
Table 8-80	Added the following action for the Single-bit ECC Threshold error: "Regular or critical interrupt if enabled."
Functional description	Updated the third sentence of the fourth paragraph to read as follows: Using ECC, the DDR memory controller detects and corrects all single/double-bit errors within the data bus, and detects all errors within a nibble.
throughout	Changed the signal ranges representation in the figures from their specific ranges to 'n'.
Detailed signal descriptions	Updated the cross reference from "DDR Control Driver Register 2 (DDRCDR_2)" to "DDR Control Driver Register 1 (DDRCDR_1)" in the description for signal Dn_MDIC[0:1].
Memory data path read capture ECC (DDR_CAPTURE_ECC)	Updated the description for 16–31 (ECE) stating " 16:23 8-bit ECC for the 32 bits in beats 0 & 2 & 4 & 6 in 32-bit bus mode; should be ignored for 64-bit mode 24:31 8-bit ECC for the 32 bits in beats 1 & 3 & 5 & 7 in 32-bit bus mode; should be used for every beat in 64-bit mode."

Table continues on the next page...

Substantive changes from revision 1 to revision 2

Reference	Description
DDR Control Driver Register 1 (DDR_DDRCDR_1)	Changed the sentence from: This can be used to calibrate the DDR drivers to 18 W.. to This can be used to calibrate the DDR drivers to the external calibration resistors on the MDIC pins.
Error management	Changed the bulleted list from: “Generates a critical interrupt if the counter value ERR_SBE[SBEC] equals the programmable threshold ERR_SBE[SBET]” to: “Generates an interrupt if the counter value ERR_SBE[SBEC] equals the programmable threshold ERR_SBE[SBET].”
Error management	Removed the word ‘Critical’ from the second sentence of the third paragraph. The updated sentence reads as follows: “Another error the DDR memory controller detects is a memory select error, which causes the DDR memory controller to log the error and generate a interrupt (if enabled, as described in Section 8.4.1.39, “Memory Error Detect (ERR_DETECT)”).
DDR Self Refresh Counter (DDR_DDR_SR_CNTR)	In Table, “DDR_SR_CNTR Field Descriptions,” added the following note to the description of bit SR_IT (12–15): Note: The term ‘clock’ refers to the data rate rather than the clock frequency of the DDR interface.
DDR SDRAM interface timing	Modified the first sentence of the first paragraph to read as follows:The DDR memory controller supports four- (or eight-) beat bursts to SDRAM.
JEDEC standard DDR SDRAM interface commands	In the section “JEDEC Standard DDR SDRAM Interface Commands,” changed the sentence “Although some SDRAMs provide burst lengths of 4, 8, and page size, this memory controller supports a burst length of 4.” to “This memory controller supports a burst length of 4.”
DDR SDRAM interface operation	Added the following note immediately below Table “Byte Lane to Data Relationship”: NOTE “For a 32-bit bus mode data byte lanes 0-3 in the above table would be used.”
DDR write leveling control (DDR_DDR_WRLVL_CNTL)	In “DDR Write Leveling Control (DDR_WRLVL_CNTL),” modified the bit range for “WRLVL_START” to (28–31) and bits (24–27) are reserved now.
DDR SDRAM refresh	Removed the last sentence of the first paragraph. The sentence reads as follows: “To ensure that the latency caused by a memory transaction does not violate the device refresh period, it is recommended that the programmed value of REFINT be less than that required by the SDRAM.”
Supported DDR SDRAM organizations	Added the last two rows for 8Gbits SDRAM Device configuration.
Table 8-80	Removed “Regular or critical interrupt if enabled” from (Notification category) Action column and updated it same as for ‘Access Error.’
Error management	In the first sentence of the third paragraph, updated “machine check or critical interrupt” as “interrupt.”
DDR SDRAM mode control (DDR_DDR_SDRAM_MD_CNTL)	Added statement to register description: Before issuing a command via the DDR_SDRAM_MD_CNTL register, the DDR interface should be idle. This can be done by setting DDR_SDRAM_CFG[MEM_HALT] and disabling refreshes by clearing DDR_INTERVAL[REFINT]. If there are memory contents that need to be preserved during this time, then software should also force any required refresh commands while DDR_INTERVAL[REFINT] is cleared.
Introduction	Removed “X64” from the first sentence.
Features	Removed the bullet “Memory controller clock frequency of two or four times....”
Memory interface signals	Updated the description of signal MDQS[0:8]/MDQS_B[0:8].
DDR SDRAM interface operation	Updated the last two sentence of the second paragraph”.
DDR SDRAM interface timing	Removed “, except for CASLAT which can be programmed with 1/2 clock granularity.”

Table continues on the next page...

Reference	Description
Error management	In the table "Memory controller errors," updated the action column for Single-bit ECC threshold, Multi-bit ECC error and Memory select error.
Memory interface signals	Updated MRAS_B, MCAS_B, and MWE_B signals in "Timings" to state that this signal is tri-stated when memory controller is idle.
to	Updated last sentence of register description from "This register is used specifically for chip select 1 if...." to "This register is used specifically for chip select 2 if..."
to	Updated last sentence of register description from "This register is used specifically for chip select 1 if...." to "This register is used specifically for chip select 3 if..."
Self-refresh in sleep mode	Added note: Deep sleep mode is not supported for DDR3 registered DIMMS.
DDR training initialization address (DDR_DDR_INIT_ADDR)	Updated the second note to say "...this address should be written to using a 32-byte transaction to avoid possible ECC errors."
Signals overview	<ul style="list-style-type: none"> - Removed bullet for "Debug Signals". - Removed rows for MDVAL and MSRCID signals from Table 8-1 - Removed section "Debug Signals".
DDR training initialization address (DDR_DDR_INIT_ADDR)	In the second note changed from " this address should be written to using an 8- or 32-byte transaction to avoid possible ECC errors" to " this address should be written to using a 32-byte transaction to avoid possible ECC errors."

C.1.8 Programmable Interrupt Controller (PIC) Revision History

Reference	Description
Feature reporting register (PIC_FRR)	Added note for the reset value of NCPU bit field.
Pass-through mode (GCR[M] = 0)	In the following sentence replaced '8259' with 'external'. The PIC does not perform a vector fetch from an 8259 external interrupt controller.
External interrupt n (IRQn) vector/priority register (PIC_EIVPRn)	Added the following footnote to EIVPR[P]: If any of IRQ[0:11] are used to receive INTx signals from one of the PCI Express ports as a root complex, the polarity must be set to be active-high.
Table 9-1	Deleted the line "core_hreset_req can also be caused by a second timer timeout condition as described in "Timer Control Registers (TCRA-TCRB)" in the Sources column of the Reset core interrupt type.
Critical interrupt summary register 0 (PIC_CISR0)	Updated bit 0-19 as Reserved.
Shared message signaled interrupt destination register n (PIC_MSIDRn)	Made EP bit reserved.
IRQ_OUT_B summary register 0 (PIC_IRQSR0)	Updated bit 0-19 as Reserved.
Introduction	Updated the last sentence of the second paragraph.
Overview	Added "...signal IRQ_OUT_B." towards the end of the second last bullet.
PCI Express INTx/IRQn sharing	Added a note below table "PCI Express INTx/IRQn sharing" regarding IRQ and GPIO multiplexing.

Table continues on the next page...

Substantive changes from revision 1 to revision 2

Reference	Description
Interrupt destinations	Added a note below table "PCI Express INTx/IRQn sharing" regarding IRQ and GPIO multiplexing.

C.1.9 Enhanced Local Bus Controller (eLBC) Revision History

Reference	Description
throughout	Updated the chapter to enable the PLL support.
External access termination (LGTA_B)	In the first paragraph, changed "LGTA should be asserted for at least one bus cycle to be effective" to "LGTA must be asserted for two bus cycles to be effective."
Options register 0 layout for FCM Mode (eLBC_ORf0) and Options register 0 layout for UPM Mode (eLBC_ORu0)	Updated the reset values of Options register 0 layout for FCM Mode (ORf0) and Options register 0 layout for UPM Mode (ORu0).
Figure 12-134	Added the figure.
External address latch enable signal (LALE)	Removed the paragraph "If the RCW is loaded by the eLBC, LALE may remain at an unknown value for up to 8 cycles after PORESET negation. Thus, LALE should....."
Table 12-198	For OR0 register field SCY, changed the settings from "From por_cfg_scy[1:3]" to "010."
Data valid and data sample control (UTA)	Added the section.
Figure 12-119	Updated the figure title as "UPM Clock Scheme for LCRR[CLKDIV] = 4 and higher order ratios".
Clock ratio register (eLBC_LCRR)	Added the range from 0-13.
Table 12-201	In "RAM Word Field Descriptions" table, added the following note for LAST bit(bit31): UPM continue to execute RAM words until it finds LAST, irrespective of the assigned region for various patterns like RSS.
External address latch enable signal (LALE)	Removed the second last paragraph starting "If the RCW is loaded by the eLBC, LALE..."
External address latch enable signal (LALE)	Added back a paragraph starting "If the RCW is not loaded by the eLBC..."
eLBC introduction	Removed the following note appearing after the first para at the beginning of the chapter: "This device supports external transceiver control but does not support the LTE pin. References to the LTE pin should be disregarded."
eLBC external signal descriptions	In the Name column, added LGTA_B in the "LGPL4/LFRB_B/LUPWAIT/LPBSE" group.
Table 12-198	Added setting for DECC field.
Figure 12-119	Corrected figure. Removed one of the figures and updated the other one
Table 12-212	Added table for single-beat write access to FPM DRAM.
Interfacing to ZBT SRAM using UPM	Removed the last paragraph.
eLBC initialization/application information	Added the following note: "Note that the data paths for configuration of eLBC and data transfers are different so there should be a delay of 200 clock cycles between register access and data access to avoid synchronization issues."

Table continues on the next page...

Reference	Description
Table 12-207	Added commands FMR and LSOR to the table
External access termination (LGTA_B)	Updated Figure 12-96 <ul style="list-style-type: none"> - Removed the first (the upper one) LCLK clock signal. - Updated "Sample point" as "Sample point for LGTA signal"
Figure 12-73	Updated figure.
eLBC external signal descriptions	Removed all instances of MDVAL and MSRCID debug signals.

C.1.10 DMA Controller Revision History

Reference	Description
DMA <i>n</i> current list descriptor address register (DMA_CLSDAR _{<i>n</i>})	Updated the name of DMA <i>n</i> current list descriptor address register (DMAx_CLSDAR _{<i>n</i>}). It was mislabeled in the last revision.
DMA channel operation	Updated the second sentence of the first paragraph as follows: In both modes, a channel can be activated by clearing and setting MR _{<i>n</i>} [CS], or through the single-write start mode using MR _{<i>n</i>} [CDSM/SWSM] and MR _{<i>n</i>} [SRW], or through an external control mode using MR _{<i>n</i>} [ECS_EN] and the external DMA_DREQ signal.
DMA features summary	Removed "Three priority levels supported for source and destination transactions" from feature list.
Bandwidth control	Added the following text at the end of the paragraph: The maximum performance can be achieved on a single channel by disabling bandwidth control. This is recommended especially if only a single channel is utilized.
Table 13-3	Modified the introductory text to the following: This table describes the external DMA control signals. These signals are only utilized when operating in external master mode. See External control mode transfer .
DMA features summary, DMA overview	Replaced Three high-speed.....to Four high-speed.....
DMA <i>n</i> mode register (DMA_MR _{<i>n</i>})	In the bifield description for DMAx_MR _{<i>n</i>} [XFE], changed "...the new chaining features," to, "striding feature in direct mode or disable list chaining feature in chaining mode"
Basic chaining, single-write start mode	In step 1 of the sequence, changed "Set the mode register current descriptor start mode bit, MR _{<i>n</i>} [CDSM_SWSM], and the extended features enable bit MR _{<i>n</i>} [XFE]." to "Set the mode register current descriptor start mode bit, MR _{<i>n</i>} [CDSM_SWSM], and clear the extended features enable bit MR _{<i>n</i>} [XFE]."

C.1.11 Enhanced Three-Speed Ethernet Controllers Revision History

Reference	Description
eTSEC detailed signal descriptions	Replaced SerDes lane configurations to the following: For SD_TX[n-1]/SD_TX[n-1]_B: <ul style="list-style-type: none"> eTSEC2 utilizes SD_TX[2] and SD_TX[2]_B eTSEC3 utilizes SD_TX[3] and SD_TX[3]_B For SD_RX[n-1]/SD_RX[n-1]_B: <ul style="list-style-type: none"> eTSEC2 utilizes SD_RX[2] and SD_RX[2]_B eTSEC3 utilizes SD_RX[3] and SD_RX[3]_B
Interrupt event register (ETSEC_IEVENT)	Updated the register figure.
Filing rules	Added equation and its example to calculate max rule.
Timer offset high (ETSEC_1588_TMROFF_H) and Timer offset low (ETSEC_1588_TMROFF_L)	Removed the note saying "All TMROFF_H registers in a device should be set to the same value, and all TMROFF_L registers in a device should be set to the same value. Otherwise, the precision time protocol may not work."
Receive control register (ETSEC_RCTRL)	For bit 24–25 (PRSDEP), removed references to fifo interface and updated the bit settings. The updated bit setting are: 00 Parser disabled. Receive frame filter must also be disabled by clearing RCTRL[FILREN] 01 Only L2 (Ethernet) protocols are recognized. 10 L2 and L3 (IP) protocols are recognized over any interface not configured as a FIFO interface. 11 L2, L3, and L4 (TCP/UDP) protocols are recognizedIP15: over any interface not configured as a FIFO interface.
Timer control register (ETSEC_1588_TMR_CTRL)	Updated the bit description of bit 22 (EETP2) and 23 (EETP1). for bit 22, added "The external trigger input signal is TSEC_1588_TRIG_IN2 (see "TMR_ETTS1–2_H/L Register Field Descriptions")" for bit 23, added "The external trigger input signal is TSEC_1588_TRIG_IN1 (see "TMR_ETTS1–2_H/L Register Field Descriptions")"
Timer alarm n high register (ETSEC_1588_TMR_ALARMn_H) and Timer alarm n low register (ETSEC_1588_TMR_ALARMn_L)	Updated the second sentence of the first paragraph to say "This register holds alarm time for comparison with the current counter value plus the offset register." Also, updated the first sentence of the last paragraph under description to say "In FS (Fiper start) mode the alarm trigger is used as an indication to the fiper start down counting."
Table 14-915	Replaced the last requirement from "Second TxBD[Data Length]. FIFO_TX_THR or includes the entire frame" to "Second TxBD[Data Length] includes entire frame or at least FIFO_TX_THR byte size value bytes or greater."
Ethernet control register (ETSEC_ECNTRL)	Modified settings for the "GMIIM" bit as follows: 0 RMII or MII mode interface expected 1 RGMII or GMII mode interface expected
Ethernet control register (ETSEC_ECNTRL)	In the "eTSEC Interface Configurations" table, changed the GMIIM column values for GMII 1Gbps , RGMII 1Gbps, RGMII 100 Mbps, and RGMII 10 Mbps to '1'.
ETSEC_1588 Memory Map/Register Definition	Removed the "TMR_FIPER*3 - Timer fixed period interval" register at offset 0x2_4E88.

Table continues on the next page...

Reference	Description
Receive control register (ETSEC_RCTRL)	<ul style="list-style-type: none"> In the register figure, replaced bit 26 from "PRFSM" to "—" reserved. In the register field description table, updated the second sentence in the description of bit PAL (11–15). In the register field description table, updated the bit settings of bit PRSDEP (24–25).
Time stamp event register (ETSEC_1588_TMR_TEVENT)	In the register figure and field description table, made bit "26" as "Reserved".
Timer fixed period interval (ETSEC_1588_TMR_FIPER _n)	Generalized the heading by removing the fiper numbers and updated the sentence immediately before the register figure to read as follows: The two registers in eTSEC1 are shared for all eTSECs.
Figure 14-883	Updated figure.
Figure 14-883	Updated figure.
Connecting to physical interfaces on Ethernet	Updated the second sentence of the second paragraph to read as follows: (For instance, in the bus TSEC _n _TXD[7:0], bit 7 is the msb and bit 0 is the lsb).
Receive queue filing table property register (ETSEC_RQFPR)	Updated the bit 21 from "FIF" to "Reserved".
Rx timer time stamp register high (ETSEC_TMR_RXTS_H) and Rx timer time stamp register low (ETSEC_TMR_RXTS_L)	Updated TMR_RXTS access from RW to RO.
MAC configuration register 2 (ETSEC_MACCFG2)	Updated the bit description of bit I_F_Mode(22-23).
Ethernet control register (ETSEC_ECNTL)	Updated the reset value of GMIII(25), TBIM(26), RPM(27), RMM(29), and SGMIIIM(30) to n as these are set based on power-on-reset configuration signals.
User initialization	Updated the second and third bullet in the list that shows the order to bring the eTSEC into a functional state (out of reset).
Interrupts	Updated the last requirement in Table "Time-stamp insertion programming requirements".
Ethernet control register (ETSEC_ECNTL)	Added a note in the last bullet regarding MAC loop back is not supported for 10-bit modesl.

C.1.12 Serial RapidIO Interface Revision History

Reference	Description
Error / port-write interrupt status register (SRIO_EPWISR)	Added bit 2-29 as reserved.
Accessing configuration registers using rapidIO packets	Removed the subheading, "Inbound Maintenance Accesses." Also, reorganized and modified the entire section.

C.1.13 PCI Express Interface Controller Revision History

Reference	Description
Error capture registers (inbound error) Error capture registers (outbound error) PCI Express error capture register n (PEX_PEX_ERR_CAP_Rn)	Removed description of PEX_ERR_CAP_Rn from register section and added it to functional description (new sections error_capture_registers_inbound and error_capture_registers_outbound)
PCI Express Advanced Error Capabilities and Control Register (Advanced_Error_Capabilities_and_Control_Register)	Updated reset value from '0x0000_0000' to '0x0000_00B0'.
Memory map/register overview	Revised the addresses for registers in the PCI Express configuration space (all non-memory-mapped registers). These registers had incorrectly specified addresses. The offsets (the last three hexadecimal digits) were correct, but the registers were shown as if they had local memory map addresses instead of PCI Express configuration space addresses.
PCI Express Correctable Error Mask Register (Correctable_Error_Mask_Register) PCI Express Correctable Error Status Register (Correctable_Error_Status_Register)	Added bit 13 ADVNFE Advisory Non-Fatal Error Status to the Correctable_Error_Status_Register and bit 13 ADVNFEM Advisory Non-Fatal Error Mask to the Correctable_Error_Mask_Register
PCI Express Link Status Register (Link_Status_Register)	Changed the access of the Link Status Register[LABS, LBMS] to w1c and added the following note to the bitfield descriptions: 'This bitfield is write-1-to-clear in RC mode; it is read-only in EP mode'.
Type 0 configuration header registers and Type 1 configuration header registers	In the Type 0 header figure, moved the Expansion ROM Base Address to 30h. In the Type 1 header figure, added Expansion ROM Base Address at 38h.
PCI Express Advanced Error Capabilities and Control Register (Advanced_Error_Capabilities_and_Control_Register)	Updated the reset value to 0000_00B0.
PCI Express Command Register (Command_Register)	Revised notes for RC mode in Bus_master and Memory_space bit fields to describe unique treatment of transactions targeting PEXCSRBAR space.
PCI Express I/O Base Upper 16 Bits Register (IO_Base_Upper_16_Bits_Register)	Added the following note to access: Access is read/write in root complex and read only in endpoint.
PCI Express I/O Limit Upper 16 Bits Register (IO_Limit_Upper_16_Bits_Register)	Added the following note to access: Access is read/write in root complex and read only in endpoint.
PCI Express signal descriptions	In second sentence of the first paragraph: TXn and TXn_B updated as SDn_TX and SDn_TX_B respectively. RXn and RXn_B updated as SDn_RX and SDn_RX_B respectively.
PCI Express Link Capabilities Register (Link_Capabilities_Register)	Updated reset value of MAX_LINK_W from 001000 to 00nnnn.
PCI Express Link Capabilities Register (Link_Capabilities_Register)	For bit 9-4, description updated as following: Reset value depends upon the maximum link width of the PCI Express controller is capable of supporting.
Table 16-2	Updated description of SD_RX[3:0] and SD_TX[3:0].

Table continues on the next page...

Reference	Description
PCI Express outbound window attributes register n (PEX_PEXOWAR0), PCI Express outbound window attributes register n (PEX_PEXOWARn), PCI Express outbound window attributes register 3 (PEX_PEXOWAR3), and PCI Express outbound window attributes register 4 (PEX_PEXOWAR4)	Updated description of PEXOWARn[ROE] and PEXOWARn[NS] bits.
Power management	Updated the second sentence of the first paragraph to read as follows: In addition, all link power states are supported with the exception of L2 states.
PCI Express configuration address register (PEX_PEX_CONFIG_ADDR)	Revised the last sentence of register description to state, "the register address is: Extended register number Register number 00b."
PCI Express Interrupt Pin Register (Interrupt_Pin_Register)	Added statement to register description that the register only applies to EP mode. Added footnote to register reset valuedistinguishing RC mode and EP mode reset values.
Table 16-338	Updated Initial credit advertisement to 8 for NPH (Memory Read, IO Read, Cfg Read, Cfg Write).
PCI Express Interrupt Pin Register (Interrupt_Pin_Register)	Added statement to register description that the register only applies to EP mode.
PCI Express Link Capabilities Register (Link_Capabilities_Register)	Updated the description of the following bits: Port Number, L1_EX_LAT, L0s_EX_LAT and ASPM.
Outbound ATMU message generationInbound messages	Moved discussion of Outbound ATMU messages and Inbound messages into the chapter's functional description.

C.1.14 Enhanced Secure Digital Host Controller Revision History

Reference	Description
Present state (eSDHC_PRSTAT)	Changed clock name reference in PRSTAT[SDOFF] from SDHC to SD
Commands for MMC/SD	In the Commands for MMC/SD/SDIO/CE-ATA table, added CMD8 row for SD Cards. In addition, added the following note: "CMD8 differs for MMC and SD cards. For SD cards, CMD8 is referred to as SEND_IF_COND. For MMC cards, CMD8 is referred to as SEND_EXT_CSD"
Reset	Removed reference to SDIO in the figure, figure caption and the pseudocode
Enhanced Secure Digital Host Controller (eSDHC) Memory Map	Reinstated the note regarding register access restriction prior to the memory map table.
Table 17-1	For SDHC_DAT3, updated the "Function" column from "4/8-bit mode: DAT3 line or configured as card detection pin" to "4/8-bit mode: DAT3 line" and updated the "Pull up/Pull down required" column from "10 k–100 kW to OVDD is required during the normal operating conditions. Board should have a 100 KW pull-down resistor if DAT3 is used for an SD card detection. In this case, the pull-up resistor has to be disconnected during the identification phase. The pull-up resistor has to be connected and the pull-down resistor should be switched out after the identification phase" to "Pull up. Do not use DAT3 pin as a CD pin."
Interrupt status (eSDHC_IRQSTAT)	Added a note to IRQSTAT[CCE] and IRQSTAT[CIE]
System control (eSDHC_SYSCTL)	Updated bit setting for SYSCTL(SDCLKFS) and SYSCTL(DVS).

Table continues on the next page...

Substantive changes from revision 1 to revision 2

Reference	Description
System control (eSDHC_SYSCTL)	For SYSCTL(SDCLKFS), in the second sentence of the example paragraph, changed the divisor value from 0xE to 0x2.
Host controller version (eSDHC_HOSTVER)	For HOSTVER(VVN), added bit value definition for 01.
Card insertion and removal detection	Added a note saying "Note that the SDHC_DAT[3] is not used for card detection on P2020."
System control (eSDHC_SYSCTL)	Changed the maximum SD clock frequency to 52 MHz for SYSCTL(16-23) Also, Updated the second example given to yield a clock value of 400 kHz.

C.1.15 Universal Serial Bus Interface Revision History

Reference	Description
Control (USB_CONTROL)	Changed USB_CONTROL[29] bit name to USB_EN.
Port status/control (USB_PORTSC)	For PIC[15–14], added a note saying "The PORTSC[14] is mapped to PCTL0 and PORTSC[15] is mapped to PCTL1."
USB command (USB_USBCMD)	Updated USBCMD[ITC] bit from "Read-only" to "Read/Write."
, Transmit FIFO tuning controls (USB_TXFILLTUNING), and Endpoint flush (USB_ENDPTFLUSH)	Updated the following register names: <ul style="list-style-type: none"> Name of BURSTSIZE from Programmable burst size to Master interface data burst size Name of TXFILLTUNING from Host TT transmit pre-buffer packet tuning to Transmit FIFO tuning controls Name of ENDPTFLUSH from Endpoint de-initialize to Endpoint flush
Capability register length (USB_CAPLENGTH)	Updated the description of CAPLENGTH bit.
Host controller structural parameters (USB_HCSPARAMS)	Added following in description of USB_HCSPARAMS[PPC]: ULPI Mode: 0 – USBDR will write 0 for DrvVbus bit of OTGControl register in PHY. 1 – USBDR will write 1 for DrvVbus bit of OTGControl register in PHY. The OTGControl register is defined in ULPI specification. UTMI Mode: In UTMI mode, USBDR will control the DRVVBUS signal defined in UTMI+ specification. 0 – USBDR will drive utmi DRVVBUS signal LOW. 1 – USBDR will drive utmi DRVVBUS signal HIGH. The operation of Port Power bit is independent of USBDR mode (host or device / USBMODE[CM]).
USB memory map/register definition	Updated reset value of the ID register to 0xE241_AF05.
Device controller initialization	Added the initialization sequence to configure the external ULPI PHY.
Host controller initialization	Added the initialization sequence to configure the external ULPI PHY.

C.1.16 Enhanced Serial Peripheral Interface Revision History

Reference	Description
ESPI detailed signal descriptions	Changed MOSI to I/O (was just O) because it can serve as a second input for Winbond dual output read.
16-bit address example, 24-bit address example	Swapped the 5 and 6th espi intialisation sequence.
eSPI CS0 mode register (ESPI_SPMODE0), eSPI CS1 mode register (ESPI_SPMODE1), eSPI CS2 mode register (ESPI_SPMODE2), and eSPI CS3 mode register (ESPI_SPMODE3)	Modified the description of LENn as follows: "Character length in bits per character."

C.1.17 Global Utilities Revision History

Reference	Description
Figure 21-30	Updated figure to show the handshake for stop in sleep/dpsleep mode.
Reset request status and control register (GUTS_RSTRSCR)	Updated the description of bit 13, WDT1_RR from "...core 0" to "...core1."
Automatic reset status register (GUTS_AUTORSTSR)	Made bit 4 (RST_WRS_P0) and bit 5 (RST_WRS_P1) Reserved.
IO voltage select status register (GUTS_IOVSELSR)	Updated the value of the '11' bit setting for the bits CVDDV (22–23)and BVDDV (26–27) to "3.3V." Also, removed the word "Default" from the '10' bit setting for the "CVDDV" and "BVDDV" bits. Also, replaced all instances of the word "Selects" with "Indicates" in the table.
SRDS Control Register 2 (GUTS_SRDSCR2)	Added "This is a configuration bit and reading this bit do not reflect the current state of lane" to the following bits X3S0,X3S1, X3S2, and X3S3.
SRDS Control Register 4 (GUTS_SRDSCR4)	Updated the bit description of LSTS2(19–21)and LSTS3(27–29).
Device disable control register (GUTS_DEVDISR)	Updated the second sentence of the second paragraph to say "...without a hard reset (except bits TB0 and TB1).

C.1.18 Debug Features and Watchpoint Facility Revision History

Reference	Description
Table 23-4	Removed "Internally pulled up" as there are no internal pull up on the TCK pin. Updated the reset value to 0.

Table continues on the next page...

Reference	Description
Debug information on debug pins	Changed "If MSRCID0 and DMA2_DACK0_B are high when..." to "If DMA2_DACK0_B is high when..."
Debug information on ECC pins	Changed "If MSRCID0 is low when sampled..." to "If DMA2_DACK0_B is low when sampled..."
Local bus interface debug	Changed "If both MSRCID0 and DMA2_DACK0_B are low when..." to "If DMA2_DACK0_B is low when..."
Debug information on ECC pins	Replaced "DMA2_DACK0_B" with DMA2_DDONE0_B in the first sentence of the paragraph
Overview	Updated the reset of MDVAL from "1" to "0".

How to Reach Us:

Home Page:

www.freescale.com

Web Support:

<http://www.freescale.com/support>

USA/Europe or Locations Not Listed:

Freescale Semiconductor
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or +1-480-768-2130
www.freescale.com/support

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064
Japan
0120 191014 or +81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor China Ltd.
Exchange Building 23F
No. 118 Jianguo Road
Chaoyang District
Beijing 100022
China
+86 10 5879 8000
support.asia@freescale.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductors products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation, or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any liability, including without limitation consequential or incidental damages. "Typical" parameters that may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals", must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claims alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

RoHS-compliant and/or Pb-free versions of Freescale products have the functionality and electrical characteristics as their non-RoHS-complaint and/or non-Pb-free counterparts. For further information, see <http://www.freescale.com> or contact your Freescale sales representative.

For information on Freescale's Environmental Products program, go to <http://www.freescale.com/epp>.

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners.

© 2009–2012 Freescale Semiconductor, Inc.