



Intel[®] IXP2800 Network Processor

Hardware Reference Manual

August 2004

Order Number: [278882-010](#)



Revision History

Date	Revision	Description
March 2002	001	First release for IXP2800 Customer Information Book V 0.4
May 2002	002	Update for the IXA SDK 3.0 release.
August 2002	003	Update for the IXA SDK 3.0 Pre-Release 4.
November 2002	004	Update for the IXA SDK 3.0 Pre-Release 5.
May 2003	005	Update for the IXA SDK 3.1 Alpha Release
September 2003	006	Update for the IXA SDK 3.5 Pre-Release 1
October 2003	007	Added information about Receiver and Transmitter Interoperation with Framers and Switch Fabrics.
January 2004	008	Updated for new trademark usage: Intel XScale [®] technology.
May 2004	009	Updated Sections 6.5.2, 8.5.2.2, 9.2.2.1, 9.3.1, 9.3.3.2, 9.5.1.4, 9.5.3.4, and 10.3.1. Updated Figure 123 and Timing Diagrams in Figures 43, 44, 46, 47, 50, 51, 54, and 55. Added Chapter 11, "Performance Monitor Unit".
August 2004	010	Preparation for web posting.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL[®] PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The IXP2800 Network Processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Except as permitted by such license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Intel Corporation.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>.

Intel and XScale are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2004, Intel Corporation.

Contents

1	Introduction	25
1.1	About This Document	25
1.2	Related Documentation	25
1.3	Terminology	26
2	Technical Description	27
2.1	Overview	27
2.2	Intel XScale® Core Microarchitecture	30
2.2.1	ARM* Compatibility	30
2.2.2	Features.....	30
2.2.2.1	Multiply/Accumulate (MAC).....	30
2.2.2.2	Memory Management	30
2.2.2.3	Instruction Cache	30
2.2.2.4	Branch Target Buffer.....	31
2.2.2.5	Data Cache	31
2.2.2.6	Interrupt Controller	31
2.2.2.7	Address Map	32
2.3	Microengines	33
2.3.1	Microengine Bus Arrangement	35
2.3.2	Control Store.....	35
2.3.3	Contexts.....	35
2.3.4	Datapath Registers	37
2.3.4.1	General-Purpose Registers (GPRs)	37
2.3.4.2	Transfer Registers	37
2.3.4.3	Next Neighbor Registers	38
2.3.4.4	Local Memory	39
2.3.5	Addressing Modes	41
2.3.5.1	Context-Relative Addressing Mode	41
2.3.5.2	Absolute Addressing Mode	42
2.3.5.3	Indexed Addressing Mode	42
2.3.6	Local CSRs.....	43
2.3.7	Execution Datapath	43
2.3.7.1	Byte Align	43
2.3.7.2	CAM	45
2.3.8	CRC Unit.....	48
2.3.9	Event Signals	49
2.4	DRAM	50
2.4.1	Size Configuration	50
2.4.2	Read and Write Access	51
2.5	SRAM	51
2.5.1	QDR Clocking Scheme	52
2.5.2	SRAM Controller Configurations.....	52
2.5.3	SRAM Atomic Operations	53
2.5.4	Queue Data Structure Commands	54
2.5.5	Reference Ordering	54
2.5.5.1	Reference Order Tables	54
2.5.5.2	Microengine Software Restrictions to Maintain Ordering	56

2.6	Scratchpad Memory.....	56
2.6.1	Scratchpad Atomic Operations.....	57
2.6.2	Ring Commands.....	57
2.7	Media and Switch Fabric Interface.....	59
2.7.1	SPI-4.....	60
2.7.2	CSIX.....	61
2.7.3	Receive.....	61
2.7.3.1	RBUF.....	62
2.7.3.1.1	SPI-4 and the RBUF.....	62
2.7.3.1.2	CSIX and RBUF.....	63
2.7.3.2	Full Element List.....	63
2.7.3.3	RX_THREAD_FREELIST.....	63
2.7.3.4	Receive Operation Summary.....	64
2.7.4	Transmit.....	65
2.7.4.1	TBUF.....	65
2.7.4.1.1	SPI-4 and TBUF.....	66
2.7.4.1.2	CSIX and TBUF.....	67
2.7.4.2	Transmit Operation Summary.....	67
2.7.5	The Flow Control Interface.....	68
2.7.5.1	SPI-4.....	68
2.7.5.2	CSIX.....	68
2.8	Hash Unit.....	69
2.9	PCI Controller.....	71
2.9.1	Target Access.....	71
2.9.2	Master Access.....	71
2.9.3	DMA Channels.....	71
2.9.3.1	DMA Descriptor.....	72
2.9.3.2	DMA Channel Operation.....	73
2.9.3.3	DMA Channel End Operation.....	74
2.9.3.4	Adding Descriptors to an Unterminated Chain.....	74
2.9.4	Mailbox and Message Registers.....	74
2.9.5	PCI Arbiter.....	75
2.10	Control and Status Register Access Proxy.....	76
2.11	Intel XScale® Core Peripherals.....	76
2.11.1	Interrupt Controller.....	76
2.11.2	Timers.....	77
2.11.3	General Purpose I/O.....	77
2.11.4	Universal Asynchronous Receiver/Transmitter.....	77
2.11.5	Slowport.....	77
2.12	I/O Latency.....	78
2.13	Performance Monitor.....	78
3	Intel XScale® Core.....	79
3.1	Introduction.....	79
3.2	Features.....	80
3.2.1	Multiply/ACcumulate (MAC).....	80
3.2.2	Memory Management.....	80
3.2.3	Instruction Cache.....	81
3.2.4	Branch Target Buffer (BTB).....	81
3.2.5	Data Cache.....	81
3.2.6	Performance Monitoring.....	81

3.2.7	Power Management.....	81
3.2.8	Debugging	81
3.2.9	JTAG.....	81
3.3	Memory Management.....	82
3.3.1	Architecture Model.....	82
3.3.1.1	Version 4 versus Version 5	82
3.3.1.2	Memory Attributes.....	82
3.3.1.2.1	Page (P) Attribute Bit	82
3.3.1.2.2	Instruction Cache	83
3.3.1.2.3	Data Cache and Write Buffer	83
3.3.1.2.4	Details on Data Cache and Write Buffer Behavior.....	83
3.3.1.2.5	Memory Operation Ordering	84
3.3.2	Exceptions	84
3.3.3	Interaction of the MMU, Instruction Cache, and Data Cache.....	85
3.3.4	Control	85
3.3.4.1	Invalidate (Flush) Operation.....	85
3.3.4.2	Enabling/Disabling	85
3.3.4.3	Locking Entries	86
3.3.4.4	Round-Robin Replacement Algorithm	87
3.4	Instruction Cache.....	88
3.4.1	Instruction Cache Operation	89
3.4.1.1	Operation when Instruction Cache is Enabled	89
3.4.1.2	Operation when Instruction Cache is Disabled	90
3.4.1.3	Fetch Policy	90
3.4.1.4	Round-Robin Replacement Algorithm	90
3.4.1.5	Parity Protection.....	91
3.4.1.6	Instruction Cache Coherency.....	91
3.4.2	Instruction Cache Control	92
3.4.2.1	Instruction Cache State at Reset	92
3.4.2.2	Enabling/Disabling	92
3.4.2.3	Invalidating the Instruction Cache.....	92
3.4.2.4	Locking Instructions in the Instruction Cache	92
3.4.2.5	Unlocking Instructions in the Instruction Cache	94
3.5	Branch Target Buffer (BTB).....	94
3.5.1	Branch Target Buffer Operation.....	94
3.5.1.1	Reset.....	95
3.5.2	Update Policy.....	96
3.5.3	BTB Control	96
3.5.3.1	Disabling/Enabling	96
3.5.3.2	Invalidation.....	96
3.6	Data Cache.....	96
3.6.1	Overviews.....	97
3.6.1.1	Data Cache Overview	97
3.6.1.2	Mini-Data Cache Overview	98
3.6.1.3	Write Buffer and Fill Buffer Overview.....	99
3.6.2	Data Cache and Mini-Data Cache Operation	99
3.6.2.1	Operation when Caching is Enabled.....	99
3.6.2.2	Operation when Data Caching is Disabled	99
3.6.2.3	Cache Policies	100
3.6.2.3.1	Cacheability	100
3.6.2.3.2	Read Miss Policy	100
3.6.2.3.3	Write Miss Policy.....	101

	3.6.2.3.4 Write-Back versus Write-Through.....	101
	3.6.2.4 Round-Robin Replacement Algorithm	102
	3.6.2.5 Parity Protection.....	102
	3.6.2.6 Atomic Accesses.....	102
3.6.3	Data Cache and Mini-Data Cache Control	103
	3.6.3.1 Data Memory State After Reset.....	103
	3.6.3.2 Enabling/Disabling	103
	3.6.3.3 Invalidate and Clean Operations.....	103
	3.6.3.3.1 Global Clean and Invalidate Operation	104
3.6.4	Reconfiguring the Data Cache as Data RAM	105
3.6.5	Write Buffer/Fill Buffer Operation and Control	106
3.7	Configuration	106
3.8	Performance Monitoring	107
	3.8.1 Performance Monitoring Events	107
	3.8.1.1 Instruction Cache Efficiency Mode.....	108
	3.8.1.2 Data Cache Efficiency Mode.....	109
	3.8.1.3 Instruction Fetch Latency Mode.....	109
	3.8.1.4 Data/Bus Request Buffer Full Mode	109
	3.8.1.5 Stall/Writeback Statistics.....	110
	3.8.1.6 Instruction TLB Efficiency Mode	111
	3.8.1.7 Data TLB Efficiency Mode	111
	3.8.2 Multiple Performance Monitoring Run Statistics	111
3.9	Performance Considerations	111
	3.9.1 Interrupt Latency.....	112
	3.9.2 Branch Prediction	112
	3.9.3 Addressing Modes.....	113
	3.9.4 Instruction Latencies.....	113
	3.9.4.1 Performance Terms	113
	3.9.4.2 Branch Instruction Timings	115
	3.9.4.3 Data Processing Instruction Timings	115
	3.9.4.4 Multiply Instruction Timings.....	116
	3.9.4.5 Saturated Arithmetic Instructions	117
	3.9.4.6 Status Register Access Instructions	118
	3.9.4.7 Load/Store Instructions	118
	3.9.4.8 Semaphore Instructions	118
	3.9.4.9 Coprocessor Instructions	119
	3.9.4.10 Miscellaneous Instruction Timing.....	119
	3.9.4.11 Thumb Instructions	119
3.10	Test Features.....	119
	3.10.1 IXP2800 Network Processor Endianness.....	120
	3.10.1.1 Read and Write Transactions Initiated by the Intel XScale® Core	121
	3.10.1.1.1 Reads Initiated by the Intel XScale® Core	121
	3.10.1.1.2 The Intel XScale® Core Writing to the IXP2800 Network Processor	123
3.11	Intel XScale® Core Gasket Unit	125
	3.11.1 Overview.....	125
	3.11.2 Intel XScale® Core Gasket Functional Description	127
	3.11.2.1 Command Memory Bus to Command Push/Pull Conversion	127
	3.11.3 CAM Operation	127
	3.11.4 Atomic Operations	128
	3.11.4.1 Summary of Rules for the Atomic Command Regarding I/O	129
	3.11.4.2 Intel XScale® Core Access to SRAM Q-Array.....	129

3.11.5	I/O Transaction	130
3.11.6	Hash Access	130
3.11.7	Gasket Local CSR	131
3.11.8	Interrupt	132
3.12	Intel XScale® Core Peripheral Interface	134
3.12.1	XPI Overview	134
3.12.1.1	Data Transfers	135
3.12.1.2	Data Alignment	135
3.12.1.3	Address Spaces for XPI Internal Devices	136
3.12.2	UART Overview	137
3.12.3	UART Operation	138
3.12.3.1	UART FIFO OPERATION	138
3.12.3.1.1	UART FIFO Interrupt Mode Operation – Receiver Interrupt	138
3.12.3.1.2	FIFO Polled Mode Operation	139
3.12.4	Baud Rate Generator	139
3.12.5	General Purpose I/O (GPIO)	140
3.12.6	Timers	141
3.12.6.1	Timer Operation	141
3.12.7	Slowport Unit	142
3.12.7.1	PROM Device Support	143
3.12.7.2	Microprocessor Interface Support for the Framer	143
3.12.7.3	Slowport Unit Interfaces	144
3.12.7.4	Address Space	145
3.12.7.5	Slowport Interfacing Topology	145
3.12.7.6	Slowport 8-Bit Device Bus Protocols	146
3.12.7.6.1	Mode 0 Single Write Transfer for Fixed-Timed Device	147
3.12.7.6.2	Mode 0 Single Write Transfer for Self-Timing Device	148
3.12.7.6.3	Mode 0 Single Read Transfer for Fixed-Timed Device	149
3.12.7.6.4	Single Read Transfer for a Self-Timing Device	150
3.12.7.7	SONET/SDH Microprocessor Access Support	150
3.12.7.7.1	Mode 1: 16-Bit Microprocessor Interface Support with 16-Bit Address Lines	151
3.12.7.7.2	Mode 2: Interface with 8 Data Bits and 11 Address Bits	155
3.12.7.7.3	Mode 3: Support for the Intel and AMCC* 2488 Mbps SONET/SDH Microprocessor Interface	157
4	Microengines	167
4.1	Overview	167
4.1.1	Control Store	169
4.1.2	Contexts	169
4.1.3	Datapath Registers	171
4.1.3.1	General-Purpose Registers (GPRs)	171
4.1.3.2	Transfer Registers	171
4.1.3.3	Next Neighbor Registers	172
4.1.3.4	Local Memory	172
4.1.4	Addressing Modes	173
4.1.4.1	Context-Relative Addressing Mode	173
4.1.4.2	Absolute Addressing Mode	174
4.1.4.3	Indexed Addressing Mode	174
4.2	Local CSRs	174
4.3	Execution Datapath	174

4.3.1	Byte Align.....	174
4.3.2	CAM.....	176
4.4	CRC Unit.....	179
4.5	Event Signals.....	180
4.5.1	Microengine Endianness	181
4.5.1.1	Read from RBUF (64 Bits).....	181
4.5.1.2	Write to TBUF	182
4.5.1.3	Read/Write from/to SRAM	182
4.5.1.4	Read/Write from/to DRAM.....	182
4.5.1.5	Read/Write from/to SHaC and Other CSRs.....	182
4.5.1.6	Write to Hash Unit.....	183
4.5.2	Media Access	183
4.5.2.1	Read from RBUF	184
4.5.2.2	Write to TBUF	185
4.5.2.3	TBUF to SPI-4 Transfer	186
5	DRAM.....	187
5.1	Overview.....	187
5.2	Size Configuration	188
5.3	DRAM Clocking	189
5.4	Bank Policy.....	190
5.5	Interleaving	191
5.5.1	Three Channels Active (3-Way Interleave).....	191
5.5.2	Two Channels Active (2-Way Interleave)	193
5.5.3	One Channel Active (No Interleave).....	193
5.5.4	Interleaving Across RDRAMs and Banks	194
5.6	Parity and ECC.....	194
5.6.1	Parity and ECC Disabled	194
5.6.2	Parity Enabled	195
5.6.3	ECC Enabled.....	195
5.6.4	ECC Calculation and Syndrome	196
5.7	Timing Configuration.....	196
5.8	Microengine Signals	197
5.9	Serial Port.....	197
5.10	RDRAM Controller Block Diagram.....	198
5.10.1	Commands	199
5.10.2	DRAM Write.....	199
5.10.2.1	Masked Write	199
5.10.3	DRAM Read.....	200
5.10.4	CSR Write.....	200
5.10.5	CSR Read.....	200
5.10.6	Arbitration	201
5.10.7	Reference Ordering	201
5.11	DRAM Push/Pull Arbiter	201
5.11.1	Arbiter Push/Pull Operation	202
5.11.2	DRAM Push Arbiter Description	203
5.12	DRAM Pull Arbiter Description.....	204
6	SRAM Interface.....	207
6.1	Overview.....	207
6.2	SRAM Interface Configurations	208

6.2.1	Internal Interface	209
6.2.2	Number of Channels	209
6.2.3	Coprocessor and/or SRAMs Attached to a Channel.....	209
6.3	SRAM Controller Configurations.....	209
6.4	Command Overview	211
6.4.1	Basic Read/Write Commands.....	211
6.4.2	Atomic Operations	211
6.4.3	Queue Data Structure Commands	213
6.4.3.1	Read_Q_Descriptor Commands.....	216
6.4.3.2	Write_Q_Descriptor Commands.....	216
6.4.3.3	ENQ and DEQ Commands	217
6.4.4	Ring Data Structure Commands.....	217
6.4.5	Journaling Commands.....	217
6.4.6	CSR Accesses	217
6.5	Parity.....	217
6.6	Address Map.....	218
6.7	Reference Ordering	219
6.7.1	Reference Order Tables	219
6.7.2	Microcode Restrictions to Maintain Ordering.....	220
6.8	Coprocessor Mode	221
7	SHaC — Unit Expansion	225
7.1	Overview.....	225
7.1.1	SHaC Unit Block Diagram.....	225
7.1.2	Scratchpad.....	227
7.1.2.1	Scratchpad Description.....	227
7.1.2.2	Scratchpad Interface.....	229
7.1.2.2.1	Command Interface	229
7.1.2.2.2	Push/Pull Interface.....	229
7.1.2.2.3	CSR Bus Interface	229
7.1.2.2.4	Advanced Peripherals Bus Interface (APB).....	229
7.1.2.3	Scratchpad Block Level Diagram.....	229
7.1.2.3.1	Scratchpad Commands	230
7.1.2.3.2	Ring Commands	231
7.1.2.3.3	Clocks and Reset.....	235
7.1.2.3.4	Reset Registers	235
7.1.3	Hash Unit.....	236
7.1.3.1	Hashing Operation	237
7.1.3.2	Hash Algorithm	239
8	Media and Switch Fabric Interface.....	241
8.1	Overview.....	241
8.1.1	SPI-4.....	243
8.1.2	CSIX	246
8.1.3	CSIX/SPI-4 Interleave Mode.....	246
8.2	Receive.....	247
8.2.1	Receive Pins.....	248
8.2.2	RBUF	248
8.2.2.1	SPI-4	250
8.2.2.2	CSIX.....	253
8.2.3	Full Element List	255
8.2.4	Rx_Thread_Freelist_#	255

8.2.5	Rx_Thread_Freelist_Timeout_#	256
8.2.6	Receive Operation Summary	256
8.2.7	Receive Flow Control Status	258
8.2.7.1	SPI-4	258
8.2.7.2	CSIX	259
8.2.7.2.1	Link-Level	259
8.2.7.2.2	Virtual Output Queue	260
8.2.8	Parity	260
8.2.8.1	SPI-4	260
8.2.8.2	CSIX	261
8.2.8.2.1	Horizontal Parity	261
8.2.8.2.2	Vertical Parity	261
8.2.9	Error Cases	261
8.3	Transmit	262
8.3.1	Transmit Pins	262
8.3.2	TBUF	263
8.3.2.1	SPI-4	266
8.3.2.2	CSIX	267
8.3.3	Transmit Operation Summary	268
8.3.3.1	SPI-4	268
8.3.3.2	CSIX	269
8.3.3.3	Transmit Summary	270
8.3.4	Transmit Flow Control Status	270
8.3.4.1	SPI-4	271
8.3.4.2	CSIX	273
8.3.4.2.1	Link-Level	273
8.3.4.2.2	Virtual Output Queue	273
8.3.5	Parity	273
8.3.5.1	SPI-4	273
8.3.5.2	CSIX	274
8.3.5.2.1	Horizontal Parity	274
8.3.5.2.2	Vertical Parity	274
8.4	RBUF and TBUF Summary	274
8.5	CSIX Flow Control Interface	275
8.5.1	TXCSRB and RXCSRB Signals	275
8.5.2	FCIFIFO and FCEFIFO Buffers	276
8.5.2.1	Full Duplex CSIX	277
8.5.2.2	Simplex CSIX	278
8.5.3	TXCDAT/RXCDAT, TXCSOF/RXCSOF, TXCPAR/RXCPAR, and TXCFC/RXCFC Signals	280
8.6	Deskew and Training	280
8.6.1	Data Training Pattern	282
8.6.2	Flow Control Training Pattern	282
8.6.3	Use of Dynamic Training	283
8.7	CSIX Startup Sequence	287
8.7.1	CSIX Full Duplex	287
8.7.1.1	Ingress IXP2800 Network Processor	287
8.7.1.2	Egress IXP2800 Network Processor	287
8.7.1.3	Single IXP2800 Network Processor	288
8.7.2	CSIX Simplex	288
8.7.2.1	Ingress IXP2800 Network Processor	288
8.7.2.2	Egress IXP2800 Network Processor	289

	8.7.2.3	Single IXP2800 Network Processor	289
8.8		Interface to Command and Push and Pull Buses	290
	8.8.1	RBUF or MSF CSR to Microengine S_TRANSFER_IN Register for Instruction:	291
	8.8.2	Microengine S_TRANSFER_OUT Register to TBUF or MSF CSR for Instruction:	291
	8.8.3	Microengine to MSF CSR for Instruction:	291
	8.8.4	From RBUF to DRAM for Instruction:	291
	8.8.5	From DRAM to TBUF for Instruction:	292
8.9		Receiver and Transmitter Interoperation with Framers and Switch Fabrics	292
	8.9.1	Receiver and Transmitter Configurations	293
	8.9.1.1	Simplex Configuration	293
	8.9.1.2	Hybrid Simplex Configuration	294
	8.9.1.3	Dual Network Processor Full Duplex Configuration	295
	8.9.1.4	Single Network Processor Full Duplex Configuration (SPI-4.2)	296
	8.9.1.5	Single Network Processor, Full Duplex Configuration (SPI-4.2 and CSIX-L1)	297
	8.9.2	System Configurations	297
	8.9.2.1	Framer, Single Network Processor Ingress and Egress, and Fabric Interface Chip	298
	8.9.2.2	Framer, Dual Network Processor Ingress, Single Network Processor Egress, and Fabric Interface Chip	298
	8.9.2.3	Framer, Single Network Processor Ingress and Egress, and CSIX-L1 Chips for Translation and Fabric Interface	299
	8.9.2.4	CPU Complex, Network Processor, and Fabric Interface Chip	299
	8.9.2.5	Framer, Single Network Processor, Co-Processor, and Fabric Interface Chip	300
	8.9.3	SPI-4.2 Support	301
	8.9.3.1	SPI-4.2 Receiver	301
	8.9.3.2	SPI-4.2 Transmitter	302
	8.9.4	CSIX-L1 Protocol Support	303
	8.9.4.1	CSIX-L1 Interface Reference Model: Traffic Manager and Fabric Interface Chip	303
	8.9.4.2	Intel® IXP2800 Support of the CSIX-L1 Protocol	304
	8.9.4.2.1	Mapping to 16-Bit Wide DDR LVDS	304
	8.9.4.2.2	Support for Dual Chip, Full-Duplex Operation	305
	8.9.4.2.3	Support for Simplex Operation	306
	8.9.4.2.4	Support for Hybrid Simplex Operation	307
	8.9.4.2.5	Support for Dynamic De-Skew Training	308
	8.9.4.3	CSIX-L1 Protocol Receiver Support	309
	8.9.4.4	CSIX-L1 Protocol Transmitter Support	310
	8.9.4.5	Implementation of a Bridge Chip to CSIX-L1	311
	8.9.5	Dual Protocol (SPI and CSIX-L1) Support	312
	8.9.5.1	Dual Protocol Receiver Support	312
	8.9.5.2	Dual Protocol Transmitter Support	312
	8.9.5.3	Implementation of a Bridge Chip to CSIX-L1 and SPI-4.2	313
	8.9.6	Transmit State Machine	314
	8.9.6.1	SPI-4.2 Transmitter State Machine	314
	8.9.6.2	Training Transmitter State Machine	315
	8.9.6.3	CSIX-L1 Transmitter State Machine	315
	8.9.7	Dynamic De-Skew	316
	8.9.8	Summary of Receiver and Transmitter Signals	317

9	PCI Unit	319
9.1	Overview.....	319
9.2	PCI Pin Protocol Interface Block.....	321
9.2.1	PCI Commands.....	322
9.2.2	IXP2800 Network Processor Initialization.....	323
9.2.2.1	Initialization by the Intel XScale® Core.....	324
9.2.2.2	Initialization by a PCI Host.....	324
9.2.3	PCI Type 0 Configuration Cycles.....	325
9.2.3.1	Configuration Write.....	325
9.2.3.2	Configuration Read.....	325
9.2.4	PCI 64-Bit Bus Extension.....	325
9.2.5	PCI Target Cycles.....	326
9.2.5.1	PCI Accesses to CSR.....	326
9.2.5.2	PCI Accesses to DRAM.....	326
9.2.5.3	PCI Accesses to SRAM.....	326
9.2.5.4	Target Write Accesses from the PCI Bus.....	326
9.2.5.5	Target Read Accesses from the PCI Bus.....	327
9.2.6	PCI Initiator Transactions.....	327
9.2.6.1	PCI Request Operation.....	327
9.2.6.2	PCI Commands.....	328
9.2.6.3	Initiator Write Transactions.....	328
9.2.6.4	Initiator Read Transactions.....	328
9.2.6.5	Initiator Latency Timer.....	328
9.2.6.6	Special Cycle.....	329
9.2.7	PCI Fast Back-to-Back Cycles.....	329
9.2.8	PCI Retry.....	329
9.2.9	PCI Disconnect.....	329
9.2.10	PCI Built-In System Test.....	329
9.2.11	PCI Central Functions.....	330
9.2.11.1	PCI Interrupt Inputs.....	330
9.2.11.2	PCI Reset Output.....	330
9.2.11.3	PCI Internal Arbiter.....	331
9.3	Slave Interface Block.....	332
9.3.1	CSR Interface.....	332
9.3.2	SRAM Interface.....	333
9.3.2.1	SRAM Slave Writes.....	333
9.3.2.2	SRAM Slave Reads.....	334
9.3.3	DRAM Interface.....	334
9.3.3.1	DRAM Slave Writes.....	334
9.3.3.2	DRAM Slave Reads.....	335
9.3.4	Mailbox and Doorbell Registers.....	336
9.3.5	PCI Interrupt Pin.....	339
9.4	Master Interface Block.....	340
9.4.1	DMA Interface.....	340
9.4.1.1	Allocation of the DMA Channels.....	341
9.4.1.2	Special Registers for Microengine Channels.....	341
9.4.1.3	DMA Descriptor.....	342
9.4.1.4	DMA Channel Operation.....	343
9.4.1.5	DMA Channel End Operation.....	344
9.4.1.6	Adding Descriptor to an Unterminated Chain.....	344
9.4.1.7	DRAM to PCI Transfer.....	344
9.4.1.8	PCI to DRAM Transfer.....	345

- 9.4.2 Push/Pull Command Bus Target Interface..... 345
 - 9.4.2.1 Command Bus Master Access to Local Configuration Registers 345
 - 9.4.2.2 Command Bus Master Access to Local Control and Status Registers..... 346
 - 9.4.2.3 Command Bus Master Direct Access to PCI Bus 346
 - 9.4.2.3.1 PCI Address Generation for IO and MEM Cycles..... 346
 - 9.4.2.3.2 PCI Address Generation for Configuration Cycles..... 347
 - 9.4.2.3.3 PCI Address Generation for Special and IACK Cycles..... 347
 - 9.4.2.3.4 PCI Enables 347
 - 9.4.2.3.5 PCI Command 347
- 9.5 PCI Unit Error Behavior 348
 - 9.5.1 PCI Target Error Behavior 348
 - 9.5.1.1 Target Access Has an Address Parity Error 348
 - 9.5.1.2 Initiator Asserts PCI_PERR_L in Response to One of Our Data Phases 348
 - 9.5.1.3 Discard Timer Expires on a Target Read..... 348
 - 9.5.1.4 Target Access to the PCI_CSR_BAR Space Has Illegal Byte Enables 348
 - 9.5.1.5 Target Write Access Receives Bad Parity PCI_PAR with the Data 349
 - 9.5.1.6 SRAM Responds with a Memory Error on One or More Data Phases on a Target Read 349
 - 9.5.1.7 DRAM Responds with a Memory Error on One or More Data Phases on a Target Read 349
 - 9.5.2 As a PCI Initiator During a DMA Transfer 349
 - 9.5.2.1 DMA Read from DRAM (Memory-to-PCI Transaction) Gets a Memory Error 349
 - 9.5.2.2 DMA Read from SRAM (Descriptor Read) Gets a Memory Error 350
 - 9.5.2.3 DMA from DRAM Transfer (Write to PCI) Receives PCI_PERR_L on PCI Bus..... 350
 - 9.5.2.4 DMA To DRAM (Read from PCI) Has Bad Data Parity 350
 - 9.5.2.5 DMA Transfer Experiences a Master Abort (Time-Out) on PCI 351
 - 9.5.2.6 DMA Transfer Receives a Target Abort Response During a Data Phase 351
 - 9.5.2.7 DMA Descriptor Has a 0x0 Word Count (Not an Error) 351
 - 9.5.3 As a PCI Initiator During a Direct Access from the Intel XScale® Core or Microengine 351
 - 9.5.3.1 Master Transfer Experiences a Master Abort (Time-Out) on PCI 351
 - 9.5.3.2 Master Transfer Receives a Target Abort Response During a Data Phase 351
 - 9.5.3.3 Master from the Intel XScale® Core or Microengine Transfer (Write to PCI) Receives PCI_PERR_L on PCI Bus 352
 - 9.5.3.4 Master Read from PCI (Read from PCI) Has Bad Data Parity 352
 - 9.5.3.5 Master Transfer Receives PCI_SERR_L from the PCI Bus 352
 - 9.5.3.6 Intel XScale® Core Microengine Requests Direct Transfer when the PCI Bus is in Reset 352
- 9.6 PCI Data Byte Lane Alignment 352
 - 9.6.1 Endian for Byte Enable 355
- 10 Clocks and Reset..... 359**
 - 10.1 Clocks 359
 - 10.2 Synchronization Between Frequency Domains 363
 - 10.3 Reset 364
 - 10.3.1 Hardware Reset Using nRESET or PCI_RST_L 364

10.3.2	PCI-Initiated Reset.....	366
10.3.3	Watchdog Timer-Initiated Reset	366
10.3.3.1	Slave Network Processor (Non-Central Function).....	367
10.3.3.2	Master Network Processor (PCI Host, Central Function)	367
10.3.3.3	Master Network Processor (Central Function).....	367
10.3.4	Software-Initiated Reset	367
10.3.5	Reset Removal Operation Based on CFG_PROM_BOOT.....	368
10.3.5.1	When CFG_PROM_BOOT is 1 (BOOT_PROM is Present).....	368
10.3.5.2	When CFG_PROM_BOOT is 0 (BOOT_PROM is Not Present)	368
10.3.6	Strap Pins	368
10.3.7	Powerup Reset Sequence	370
10.4	Boot Mode	370
10.4.1	Flash ROM.....	372
10.4.2	PCI Host Download	372
10.5	Initialization	373
11	Performance Monitor Unit	375
11.1	Introduction	375
11.1.1	Motivation for Performance Monitors.....	375
11.1.2	Motivation for Choosing CHAP Counters	376
11.1.3	Functional Overview of CHAP Counters.....	377
11.1.4	Basic Operation of the Performance Monitor Unit	378
11.1.5	Definition of CHAP Terminology	379
11.1.6	Definition of Clock Domains.....	380
11.2	Interface and CSR Description	380
11.2.1	APB Peripheral	381
11.2.2	CAP Description	381
11.2.2.1	Selecting the Access Mode.....	381
11.2.2.2	PMU CSR	381
11.2.2.3	CAP Writes	381
11.2.2.4	CAP Reads	381
11.2.3	Configuration Registers	382
11.3	Performance Measurements	382
11.4	Events Monitored in Hardware	385
11.4.1	Queue Statistics Events.....	385
11.4.1.1	Queue Latency.....	385
11.4.1.2	Queue Utilization.....	385
11.4.2	Count Events	385
11.4.2.1	Hardware Block Execution Count	385
11.4.3	Design Block Select Definitions	386
11.4.4	Null Event	387
11.4.5	Threshold Events.....	388
11.4.6	External Input Events.....	389
11.4.6.1	XPI Events Target ID(000001) / Design Block #(0100)	389
11.4.6.2	SHaC Events Target ID(000010) / Design Block #(0101).....	393
11.4.6.3	IXP2800 Network Processor MSF Events Target ID(000011) / Design Block #(0110).....	396
11.4.6.4	Intel XScale® Core Events Target ID(000100) / Design Block #(0111).....	402
11.4.6.5	PCI Events Target ID(000101) / Design Block #(1000)	405
11.4.6.6	ME00 Events Target ID(100000) / Design Block #(1001).....	409

11.4.6.7 ME01 Events Target ID(100001) / Design Block #(1001)..... 410

11.4.6.8 ME02 Events Target ID(100010) / Design Block #(1001)..... 411

11.4.6.9 ME03 Events Target ID(100011) / Design Block #(1001)..... 411

11.4.6.10 ME04 Events Target ID(100100) / Design Block #(1001)..... 412

11.4.6.11 ME05 Events Target ID(100101) / Design Block #(1001)..... 412

11.4.6.12 ME06 Events Target ID(100110) / Design Block #(1001)..... 413

11.4.6.13 ME07 Events Target ID(100111) / Design Block #(1001)..... 413

11.4.6.14 ME10 Events Target ID(110000) / Design Block #(1010)..... 414

11.4.6.15 ME11 Events Target ID(110001) / Design Block #(1010)..... 414

11.4.6.16 ME12 Events Target ID(110010) / Design Block #(1010)..... 415

11.4.6.17 ME13 Events Target ID(110011) / Design Block #(1010)..... 415

11.4.6.18 ME14 Events Target ID(110100) / Design Block #(1010)..... 416

11.4.6.19 ME15 Events Target ID(110101) / Design Block #(1010)..... 416

11.4.6.20 ME16 Events Target ID(100110) / Design Block #(1010)..... 417

11.4.6.21 ME17 Events Target ID(110111) / Design Block #(1010)..... 417

11.4.6.22 SRAM DP1 Events Target ID(001001) / Design Block #(0010)..... 418

11.4.6.23 SRAM DP0 Events Target ID(001010) / Design Block #(0010)..... 418

11.4.6.24 SRAM CH3 Events Target ID(001011) / Design Block #(0010)..... 420

11.4.6.25 SRAM CH2 Events Target ID(001100) / Design Block #(0010)..... 421

11.4.6.26 SRAM CH1 Events Target ID(001101) / Design Block #(0010)..... 421

11.4.6.27 SRAM CH0 Events Target ID(001110) / Design Block #(0010)..... 422

11.4.6.28 DRAM DPLA Events Target ID(010010) / Design Block #(0011)..... 423

11.4.6.29 DRAM DPSA Events Target ID(010011) / Design Block #(0011)..... 424

11.4.6.30 IXP2800 Network Processor DRAM CH2 Events Target ID(010100) /
Design Block #(0011)..... 425

11.4.6.31 IXP2800 Network Processor DRAM CH1 Events Target ID(010101) /
Design Block #(0011)..... 429

11.4.6.32 IXP2800 Network Processor DRAM CH0 Events Target ID(010110) /
Design Block #(0011)..... 429

Figures

1	IXP2800 Network Processor Functional Block Diagram	28
2	IXP2800 Network Processor Detailed Diagram.....	29
3	Intel XScale® Core 4-GB (32-Bit) Address Space	32
4	Microengine Block Diagram.....	34
5	Context State Transition Diagram	36
6	Byte-Align Block Diagram.....	44
7	CAM Block Diagram	46
8	Echo Clock Configuration	52
9	Logical View of Rings	57
10	Example System Block Diagram	59
11	Full-Duplex Block Diagram	60
12	Simplified MSF Receive Section Block Diagram	61
13	Simplified Transmit Section Block Diagram.....	65
14	Hash Unit Block Diagram	70
15	DMA Descriptor Reads.....	72
16	Intel XScale® Core Architecture Features	80
17	Example of Locked Entries in TLB	88
18	Instruction Cache Organization	89
19	Locked Line Effect on Round Robin Replacement.....	93
20	BTB Entry	95
21	Branch History	95
22	Data Cache Organization	97
23	Mini-Data Cache Organization	98
24	Byte Steering for Read and Byte-Enable Generation by the Intel XScale® Core	122
25	Intel XScale® Core-Initiated Write to the IXP2800 Network Processor.....	124
26	Intel XScale® Core-Initiated Write to the IXP2800 Network Processor (Continued).....	125
27	Global Buses Connection to the Intel XScale® Core Gasket.....	126
28	Flow Through the Intel XScale® Core Interrupt Controller.....	132
29	Interrupt Mask Block Diagram	133
30	XPI Interfaces for IXP2800 Network Processor.....	135
31	UART Data Frame.....	138
32	GPIO Functional Diagram	140
33	Timer Control Unit Interfacing Diagram	141
34	Timer Internal Logic Diagram	142
35	Slowport Unit Interface Diagram.....	144
36	Address Space Hole Diagram	145
37	Slowport Example Application Topology	146
38	Mode 0 Single Write Transfer for a Fixed-Timed Device.....	147
39	Mode 0 Single Write Transfer for a Self-Timing Device	148
40	Mode 0 Single Read Transfer for Fixed-Timed Device	149
41	Mode 0 Single Read Transfer for a Self-Timing Device	150
42	An Interface Topology with Lucent* TDAT042G5 SONET/SDH.....	152
43	Mode 1 Single Write Transfer for Lucent* TDAT042G5 Device (B0)	153
44	Mode 1 Single Read Transfer for Lucent* TDAT042G5 Device (B0)	154
45	An Interface Topology with PMC-Sierra* PM5351 S/UNI-TETRA*	155
46	Mode 2 Single Write Transfer for PMC-Sierra* PM5351 Device (B0)	156
47	Mode 2 Single Read Transfer for PMC-Sierra* PM5351 Device (B0).....	157

48	An Interface Topology with Intel / AMCC* SONET/SDH Device	158
49	Mode 3 Second Interface Topology with Intel / AMCC* SONET/SDH Device.....	159
50	Mode 3 Single Write Transfer Followed by Read (B0)	160
51	Mode 3 Single Read Transfer Followed by Write (B0)	161
52	An Interface Topology with Intel / AMCC* SONET/SDH Device in Motorola* Mode	162
53	Second Interface Topology with Intel / AMCC* SONET/SDH Device.....	163
54	Mode 4 Single Write Transfer (B0)	164
55	Mode 4 Single Read Transfer (B0).....	165
56	Microengine Block Diagram.....	168
57	Context State Transition Diagram.....	170
58	Byte Align Block Diagram	175
59	CAM Block Diagram	177
60	Read from RBUF (64 Bits).....	181
61	Write to TBUF (64 Bits).....	182
62	48-, 64-, and 128-Bit Hash Operand Transfers	183
63	Bit, Byte, and Longword Organization in One RBUF Element.....	184
64	Write to TBUF	185
65	MSF Interface	186
66	Clock Configuration	189
67	IXP2800 Clocking for RDRAM at 400 MHz	190
68	IXP2800 Clocking for RDRAM at 508 MHz	190
69	Address Mapping Flow	191
70	RDRAM Controller Block Diagram.....	198
71	DRAM Push/Pull Arbiter Functional Blocks	202
72	DRAM Push Arbiter Functional Blocks	204
73	DRAM Pull Arbiter Functional Blocks	205
74	SRAM Controller/Chassis Block Diagram.....	208
75	SRAM Clock Connection on a Channel.....	210
76	External Pipeline Registers Block Diagram	211
77	Queue Descriptor with Four Links	213
78	Enqueueing One Buffer at a Time	213
79	Previously Linked String of Buffers.....	214
80	First Step to Enqueue a String of Buffers to a Queue (ENQ_Tail_and_Link).....	214
81	Second Step to Enqueue a String of Buffers to a Queue (ENQ_Tail)	214
82	Connection to a Coprocessor Though Standard QDR Interface	221
83	Coprocessor with Memory Mapped FIFO Ports	222
84	SHaC Top Level Diagram.....	226
85	Scratchpad Block Diagram	228
86	Ring Communication Logic Diagram	231
87	Hash Unit Block Diagram.....	236
88	Example System Block Diagram	242
89	Full-Duplex Block Diagram	243
90	Receive and Transmit Clock Generation	245
91	Simplified Receive Section Block Diagram.....	247
92	RBUF Element State Diagram.....	257
93	Extent of DIP-4 Codes	260
94	Simplified Transmit Section Block Diagram.....	262
95	TBUF State Diagram	270
96	Tx Calendar Block Diagram.....	271
97	CSIX Flow Control Interface — TXCSRB and RXCSRB.....	275

98	CSIX Flow Control Interface — FCIFIFO and FCEFIFO in Full Duplex Mode	277
99	CSIX Flow Control Interface — FCIFIFO and FCEFIFO in Simplex Mode	278
100	MSF to Command and Push and Pull Buses Interface Block Diagram	290
101	Basic I/O Capability of the Intel® IXP2800 Network Processor	292
102	Simplex Configuration	293
103	Hybrid Simplex Configuration	294
104	Dual Network Processor, Full Duplex Configuration	295
105	Single Network Processor, Full Duplex Configuration (SPI-4.2 Protocol)	296
106	Single Network Processor, Full Duplex Configuration (SPI-4.2 and CSIX-L1 Protocols)	297
107	Framer, Single Network Processor Ingress and Egress, and Fabric Interface Chip	298
108	Framer, Dual Processor Ingress, Single Processor Egress, and Fabric Interface Chip	298
109	Framer, Single Network Processor Ingress, Single Network Processor Egress, CSIX-L1 Translation Chip and CSIX-L1 Fabric Interface Chip	299
110	CPU Complex, Network Processor, and Fabric Interface Chips	299
111	Framer, Single Network Processor, Co-Processor, and Fabric Interface Chip	300
112	SPI-4.2 Interface Reference Model with Receiver and Transmitter Labels Corresponding to Link Layer Device Functions	301
113	CSIX-L1 Interface Reference Model with Receiver and Transmitter Labels Corresponding to Fabric Interface Chip Functions	303
114	Reference Model for IXP2800 Support of the Simplex Configuration Using Independent Ingress and Egress Interfaces	306
115	Reference Model for Hybrid Simplex Operation	307
116	Block Diagram of Dual Protocol (SPI-4.2 and CSIX-L1) Bridge Chip	313
117	Summary of Receiver and Transmitter Signaling	317
118	PCI Functional Blocks	320
119	Data Access Paths	321
120	PCI Arbiter Configuration Using CFG_PCI_ARB(GPIO[2])	331
121	Example of Target Write to SRAM of 68 Bytes	333
122	Example of Target Write to DRAM of 68 Bytes	335
123	Example of Target Read from DRAM Using 64-Byte Burst	336
124	Generation of the Doorbell Interrupts to PCI	337
125	Generation of the Doorbell Interrupts to the Intel XScale® Core	338
126	PCI Interrupts	339
127	DMA Descriptor Reads	342
128	PCI Address Generation for Command Bus Master to PCI	346
129	PCI Address Generation for Command Bus Master to PCI Configuration Cycle	347
130	Overall Clock Generation and Distribution	360
131	IXP2800 Network Processor Clock Generation	363
132	Synchronization Between Frequency Domains	364
133	Reset Out Behavior	365
134	Reset Generation	366
135	Boot Process	371
136	Performance Monitor Interface Block Diagram	376
137	Block Diagram of a Single CHAP Counter	378
138	Basic Block Diagram of IXP2800 Network Processor with PMU	379
139	CAP Interface to the APB	380
140	Conceptual Diagram of Counter Array	382

Tables

1	Data Terminology	26
2	Longword Formats	26
3	IXP2800 Network Processor Microengine Bus Arrangement	35
4	Next Neighbor Write as a Function of CTX_ENABLE[NN_MODE]	38
5	Registers Used By Contexts in Context-Relative Addressing Mode	41
6	Align Value and Shift Amount	43
7	Register Contents for Example 10	44
8	Register Contents for Example 11	45
9	RDRAM Sizes	50
10	SRAM Controller Configurations	52
11	Total Memory per Channel	53
12	Address Reference Order	55
13	Q_array Entry Reference Order	55
14	Ring Full Signal Use – Number of Contexts and Length versus Ring Size	58
15	TBUF SPI-4 Control Definition	66
16	TBUF CSIX Control Definition	67
17	DMA Descriptor Format	72
18	Doorbell Interrupt Registers	75
19	I/O Latency	78
20	Data Cache and Buffer Behavior when X = 0	83
21	Data Cache and Buffer Behavior when X = 1	83
22	Memory Operations that Impose a Fence	84
23	Valid MMU and Data/Mini-Data Cache Combinations	85
24	Performance Monitoring Events	107
25	Some Common Uses of the PMU	108
26	Branch Latency Penalty	112
27	Latency Example	114
28	Branch Instruction Timings (Predicted by the BTB)	115
29	Branch Instruction Timings (Not Predicted by the BTB)	115
30	Data Processing Instruction Timings	115
31	Multiply Instruction Timings	116
32	Multiply Implicit Accumulate Instruction Timings	117
33	Implicit Accumulator Access Instruction Timings	117
34	Saturated Data Processing Instruction Timings	117
35	Status Register Access Instruction Timings	118
36	Load and Store Instruction Timings	118
37	Load and Store Multiple Instruction Timings	118
38	Semaphore Instruction Timings	118
39	CP15 Register Access Instruction Timings	119
40	CP14 Register Access Instruction Timings	119
41	SWI Instruction Timings	119
42	Count Leading Zeros Instruction Timings	119
43	Little-Endian Encoding	120
44	Big-Endian Encoding	120
45	Byte-Enable Generation by the Intel XScale® Core for Byte Transfers in Little- and Big-Endian Systems	121
46	Byte-Enable Generation by the Intel XScale® Core for 16-Bit Data Transfers in Little- and Big-Endian Systems	123

47	Byte-Enable Generation by the Intel XScale® Core for Byte Writes in Little- and Big-Endian Systems	123
48	Byte-Enable Generation by the Intel XScale® Core for Word Writes in Little- and Big-Endian Systems	124
49	CMB Write Command to CPP Command Conversion	127
50	IXP2800 Network Processor SRAM Q-Array Access Alias Addresses	129
51	GCSR Address Map (0xd700 0000).....	131
52	Data Transaction Alignment	136
53	Address Spaces for XPI Internal Devices.....	136
54	8-Bit Flash Memory Device Density	143
55	SONET/SDH Devices.....	143
56	Next Neighbor Write as a Function of CTX_Enable[NN_Mode]	172
57	Registers Used by Contexts in Context-Relative Addressing Mode.....	173
58	Align Value and Shift Amount.....	174
59	Register Contents for Example 23.....	175
60	Register Contents for Example 24.....	176
61	RDRAM Loading.....	188
62	RDRAM Sizes.....	188
63	Address Rearrangement for 3-Way Interleave (Sheet 1 of 2)	192
64	Address Rearrangement for 3-Way Interleave (Sheet 2 of 2) (Rev B).....	193
65	Address Bank Interleaving.....	194
66	RDRAM Timing Parameter Settings.....	196
67	Ordering of Reads and Writes to the Same Address for DRAM.....	201
68	DRAM Push Arbiter Operation	203
69	DPLA Description	204
70	SRAM Controller Configurations	209
71	Total Memory per Channel	210
72	Atomic Operations	212
73	Queue Format	215
74	Ring/Journal Format	216
75	Ring Size Encoding	216
76	Address Map	218
77	Address Reference Order.....	219
78	Q_array Entry Reference Order	220
79	Ring Full Signal Use – Number of Contexts and Length versus Ring Size	232
80	Head/Tail, Base, and Full Threshold – by Ring Size	233
81	Intel XScale® Core and Microengine Instructions	235
82	S_Transfer Registers Hash Operands	237
83	SPI-4 Control Word Format.....	244
84	Order of Bytes within the SPI-4 Data Burst	245
85	CFrame Types.....	246
86	Receive Pins Usage by Protocol	248
87	Order in which Received Data Is Stored in RBUF	248
88	Mapping of Received Data to RBUF Partitions	249
89	Number of Elements per RBUF Partition.....	249
90	RBUF SPIF-4 Status Definition	252
91	RBUF CSIX Status Definition	254
92	Rx_Thread_Freelist Use.....	255
93	Summary of SPI-4 and CSIX RBUF Operations	258
94	Transmit Pins Usage by Protocol	262

95	Order in which Data is Transmitted from TBUF	263
96	Mapping of TBUF Partitions to Transmit Protocol	263
97	Number of Elements per TBUF Partition	264
98	TBUF SPI-4 Control Definition	266
99	TBUF CSIX Control Definition	267
100	Transmit SPI-4 Control Word	268
101	Transmit CSIX Header	269
102	Summary of RBUF and TBUF Operations	274
103	SRB Definition by Clock Phase Number	276
104	Data Deskew Functions	281
105	Calendar Deskew Functions	281
106	Flow Control Deskew Functions	281
107	Data Training Sequence	282
108	Flow Control Training Sequence	282
109	Calendar Training Sequence	283
110	IXP2800 Network Processor Requires Data Training	284
111	Switch Fabric or SPI-4 Framer Requires Data Training	285
112	IXP2800 Network Processor Requires Flow Control Training	286
113	Switch Fabric Requires Flow Control Training	286
114	SPI-4.2 Transmitter State Machine Transitions on 16-Bit Bus Transfers	314
115	Training Transmitter State Machine Transitions on 16-Bit Bus Transfers	315
116	CSIX-L1 Transmitter State Machine Transitions on CWord Boundaries	315
117	PCI Block FIFO Sizes	322
118	Maximum Loading	322
119	PCI Commands	322
120	PCI BAR Programmable Sizes	324
121	PCI BAR Sizes with PCI Host Initialization	324
122	Legal Combinations of the Strap Pin Options	330
123	Slave Interface Buffer Sizes	332
124	Doorbell Interrupt Registers	337
125	IRQ Interrupt Options by Stepping	339
126	DMA Descriptor Format	342
127	PCI Maximum Burst Size	345
128	Command Bus Master Configuration Transactions	347
129	Command Bus Master Address Space Map to PCI	347
130	Byte Lane Alignment for 64-Bit PCI Data In (64 Bits PCI Little-Endian to Big-Endian with Swap)	353
131	Byte Lane Alignment for 64-Bit PCI Data In (64 Bits PCI Big-Endian to Big-Endian without Swap)	353
132	Byte Lane Alignment for 32-Bit PCI Data In (32 Bits PCI Little-Endian to Big-Endian with Swap)	353
133	Byte Lane Alignment for 32-Bit PCI Data In (32 Bits PCI Big-Endian to Big-Endian without Swap)	353
134	Byte Lane Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Little Endian with Swap)	354
135	Byte Lane Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Big-Endian without Swap)	354
136	Byte Lane Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Little Endian with Swap)	354
137	Byte Lane Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Big-Endian without Swap)	354

138	Byte Enable Alignment for 64-Bit PCI Data In (64 Bits PCI Little-Endian to Big-Endian with Swap).....	355
139	Byte Enable Alignment for 64-Bit PCI Data In (64 Bits PCI Big-Endian to Big-Endian without Swap).....	355
140	Byte Enable Alignment for 32-Bit PCI Data In (32 bits PCI Little-Endian to Big-Endian with Swap).....	355
141	Byte Enable Alignment for 32-Bit PCI Data In (32 Bits PCI Big-Endian to Big-Endian without Swap).....	356
142	Byte Enable Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Little Endian with Swap).....	356
143	Byte Enable Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Big Endian without Swap).....	356
144	Byte Enable Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Little Endian with Swap).....	356
145	Byte Enable Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Big Endian without Swap).....	357
146	PCI I/O Cycles with Data Swap Enable.....	358
147	Clock Usage Summary.....	360
148	Clock Rates Examples.....	362
149	IXP2800 Network Processor Strap Pins.....	369
150	Supported Strap Combinations.....	370
151	APB Usage.....	381
152	Hardware Blocks and Their Performance Measurement Events.....	383
153	PMU Design Unit Selection.....	386
154	Chap Counter Threshold Events (Design Block # 0001).....	388
155	XPI PMU Event List.....	389
156	SHaC PMU Event List.....	393
157	IXP2800 Network Processor MSF PMU Event List.....	396
158	Intel XScale® Core Gasket PMU Event List.....	402
159	PCI PMU Event List.....	405
160	ME00 PMU Event List.....	409
161	ME01 PMU Event List.....	410
162	ME02 PMU Event List.....	411
163	ME03 PMU Event List.....	411
164	ME04 PMU Event List.....	412
165	ME05 PMU Event List.....	412
166	ME06 PMU Event List.....	413
167	ME07 PMU Event List.....	413
168	ME10 PMU Event List.....	414
169	ME11 PMU Event List.....	414
170	ME12 PMU Event List.....	415
171	ME13 PMU Event List.....	415
172	ME14 PMU Event List.....	416
173	ME15 PMU Event List.....	416
174	ME16 PMU Event List.....	417
175	ME17 PMU Event List.....	417
176	SRAM DP1 PMU Event List.....	418
177	SRAM DP0 PMU Event List.....	418
178	SRAM CH3 PMU Event List.....	420
179	SRAM CH3 PMU Event List.....	421
180	SRAM CH3 PMU Event List.....	421

181 SRAM CH0 PMU Event List	422
182 IXP2800 Network Processor Dram DPLA PMU Event List.....	423
183 IXP2800 Network Processor Dram DPSA PMU Event List	424
184 IXP2800 Network Processor Dram CH2 PMU Event List.....	425
185 IXP2800 Network Processor Dram CH1 PMU Event List.....	429
186 IXP2800 Network Processor Dram CH0 PMU Event List.....	429

1.1 About This Document

This document is the hardware reference manual for the Intel® IXP2800 Network Processor. This information is intended for use by developers and is organized as follows:

Section 2, “[Technical Description](#)” contains a hardware overview.

Section 3, “[Intel XScale® Core](#)” describes the embedded core.

Section 4, “[Microengines](#)” describes Microengine operation.

Section 5, “[DRAM](#)” describes the DRAM Unit.

Section 6, “[SRAM Interface](#)” describes the SRAM Unit.

Section 7, “[SHaC — Unit Expansion](#)” describes the Scratchpad, Hash Unit, and CSRs (SHaC).

Section 8, “[Media and Switch Fabric Interface](#)” describes the Media and Switch Fabric (MSF) Interface used to connect the network processor to a physical layer device.

Section 9, “[PCI Unit](#)” describes the PCI Unit.

Section 10, “[Clocks and Reset](#)” describes the clocks, reset and initialization sequence.

Section 11, “[Performance Monitor Unit](#)” describes the PMU.

1.2 Related Documentation

Further information on the IXP2800 is available in the following documents:

IXP2800 Network Processor Datasheet – Contains summary information on the IXP2800 Network Processor including a functional description, signal descriptions, electrical specifications, and mechanical specifications.

IXP2400 and IXP2800 Network Processor Programmer's Reference Manual – Contains detailed programming information for designers.

IXP2400/IXP2800 Network Processor Development Tools User's Guide – Describes the Developer Workbench and the development tools you can access through the use of the Workbench GUI.



1.3 Terminology

Table 1 and Table 2 list the terminology used in this manual.

Table 1. Data Terminology

Term	Words	Bytes	Bits
Byte	½	1	8
Word	1	2	16
Longword	2	4	32
Quadword	4	8	64

Table 2. Longword Formats

Endian Type	32-Bit	64-Bit
Little-Endian	(0x12345678) arranged as {12 34 56 78}	64-bit data 0x12345678 9ABCDE56 arranged as {12 34 56 78 9A BC DE 56}
Big-Endian	(0x12345678) arranged as {78 56 34 12}	64-bit data 0x12345678 9ABCDE56 arranged as {78 56 34 12, 56 DE BC 9A}



Technical Description

2

2.1 Overview

This section provides a brief overview of the IXP2800 Network Processor internal hardware. This section is intended as an overall hardware introduction to the network processor.

The major blocks are:

- Intel XScale® core — General purpose 32-bit RISC processor (ARM* Version 5 Architecture compliant) used to initialize and manage the network processor, and can be used for higher layer network processing tasks.
- Intel XScale® technology Peripherals (XPI) — Interrupt Controller, Timers, UART, General Purpose I/O (GPIO) and interface to low-speed off chip peripherals (such as maintenance port of network devices) and Flash ROM.
- Microengines (MEs) — Sixteen 32-bit programmable engines specialized for Network Processing. Microengines do the main data plane processing per packet.
- DRAM Controllers — Three independent controllers for Rambus* DRAM. Typically DRAM is used for data buffer storage.
- SRAM Controllers — Four independent controllers for QDR SRAM. Typically SRAM is used for control information storage.
- Scratchpad Memory — 16 Kbytes storage for general purpose use.
- Hash Unit — Polynomial hash accelerator. The Intel XScale® core and Microengines can use it to offload hash calculations.
- Control and Status Register Access Proxy (CAP) — These provide special inter-processor communication features to allow flexible and efficient inter-Microengine and Microengine to Intel XScale® core communication.
- Media and Switch Fabric Interface (MSF) — Interface for network framers and/or Switch Fabric. Contains receive and transmit buffers.
- PCI Controller — *PCI Local Bus Specification, Version 2.2** interface for 64-bit 66-MHz I/O. PCI can be used to either connect to a Host processor, or to attach PCI-compliant peripheral devices.
- Performance Monitor — Counters that can be programmed to count selected internal chip hardware events, which can be used to analyze and tune performance.

[Figure 1](#) is a simple block diagram of the network processor showing the major internal hardware blocks. [Figure 2](#) is a detailed diagram of the network processor units and buses.

Figure 1. IXP2800 Network Processor Functional Block Diagram

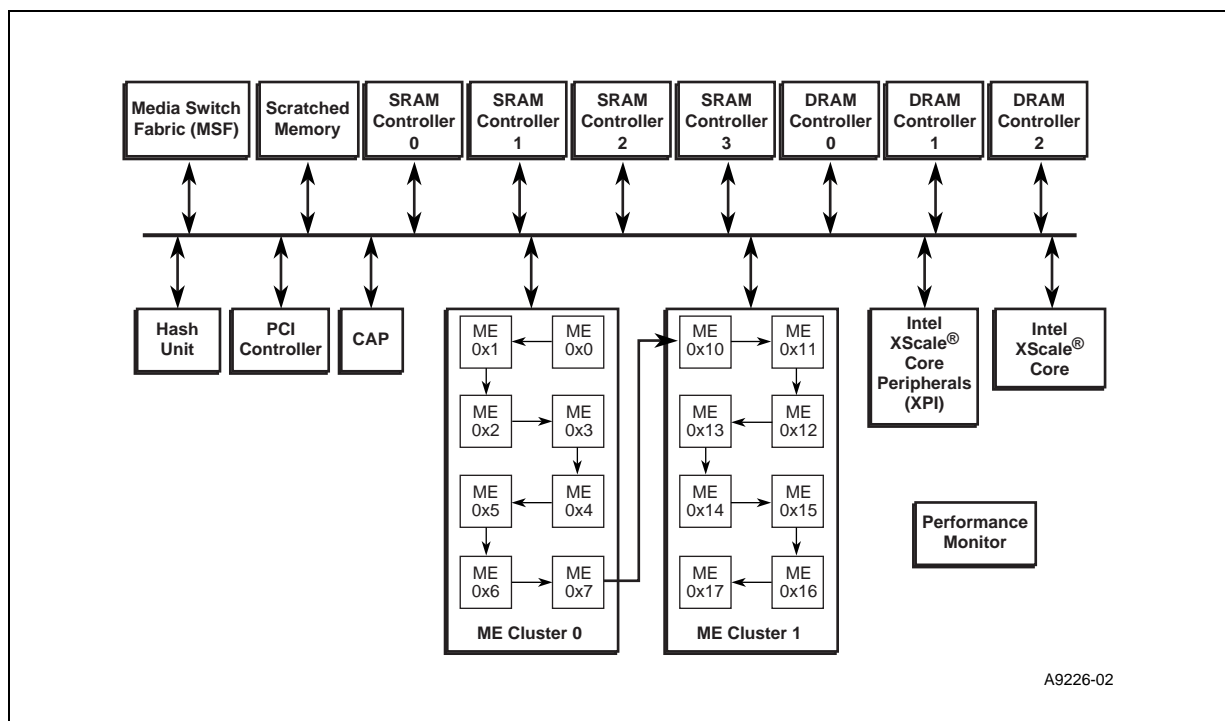
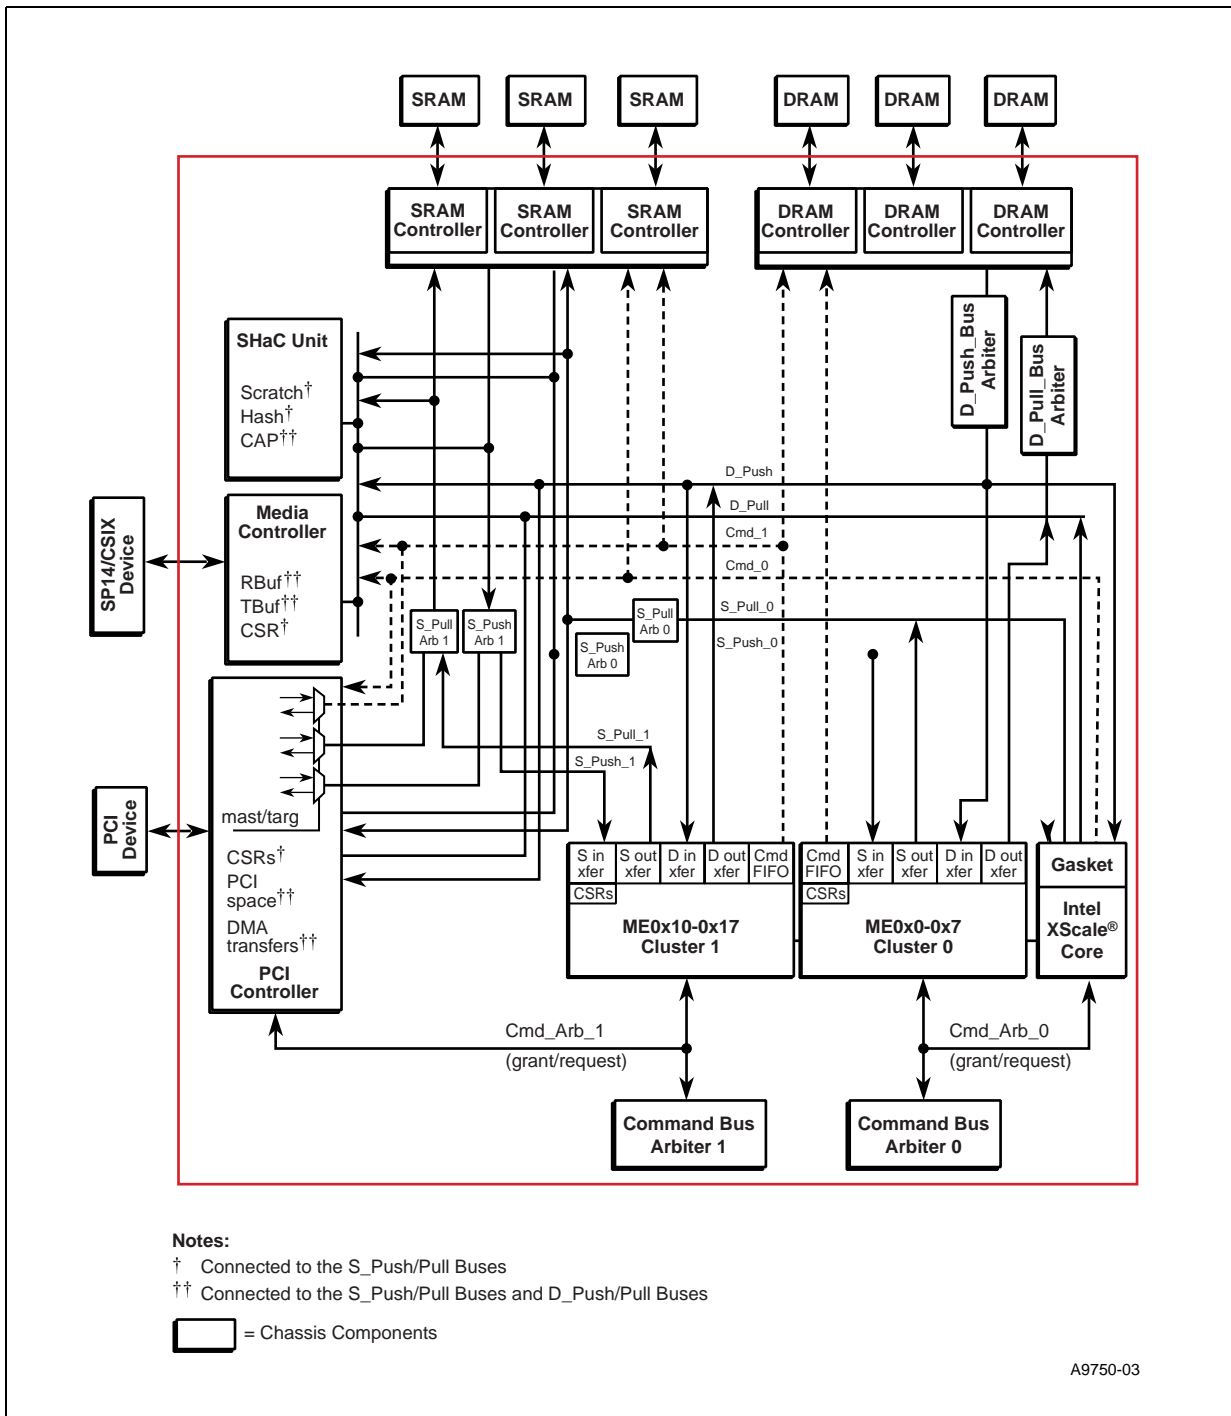


Figure 2. IXP2800 Network Processor Detailed Diagram





2.2 Intel XScale® Core Microarchitecture

The Intel XScale® microarchitecture consists of a 32-bit general purpose RISC processor that incorporates an extensive list of architecture features that allows it to achieve high performance.

2.2.1 ARM* Compatibility

The Intel XScale® microarchitecture is ARM* Version 5 (V5) Architecture compliant. It implements the integer instruction set of ARM* V5, but does not provide hardware support of the floating point instructions.

The Intel XScale® microarchitecture provides the Thumb instruction set (ARM V5T) and the ARM V5E DSP extensions.

Backward compatibility with the first generation of StrongARM* products is maintained for user-mode applications. Operating systems may require modifications to match the specific hardware features of the Intel XScale® microarchitecture and to take advantage of the performance enhancements added to the Intel XScale® core.

2.2.2 Features

2.2.2.1 Multiply/Accumulate (MAC)

The MAC unit supports early termination of multiplies/accumulates in two cycles and can sustain a throughput of a MAC operation every cycle. Several architectural enhancements were made to the MAC to support audio coding algorithms, which include a 40-bit accumulator and support for 16-bit packed values.

2.2.2.2 Memory Management

The Intel XScale® microarchitecture implements the Memory Management Unit (MMU) Architecture specified in the *ARM Architecture Reference Manual*. The MMU provides access protection and virtual to physical address translation.

The MMU Architecture also specifies the caching policies for the instruction cache and data memory. These policies are specified as page attributes and include:

- identifying code as cacheable or non-cacheable
- selecting between the mini-data cache or data cache
- write-back or write-through data caching
- enabling data write allocation policy
- and enabling the write buffer to coalesce stores to external memory

2.2.2.3 Instruction Cache

The Intel XScale® microarchitecture implements a 32-Kbyte, 32-way set associative instruction cache with a line size of 32 bytes. All requests that “miss” the instruction cache generate a 32-byte read request to external memory. A mechanism to lock critical code within the cache is also provided.



2.2.2.4 Branch Target Buffer

The Intel XScale® microarchitecture provides a Branch Target Buffer (BTB) to predict the outcome of branch type instructions. It provides storage for the target address of branch type instructions and predicts the next address to present to the instruction cache when the current instruction address is that of a branch.

The BTB holds 128 entries.

2.2.2.5 Data Cache

The Intel XScale® microarchitecture implements a 32-Kbyte, 32-way set associative data cache and a 2-Kbyte, 2-way set associative mini-data cache. Each cache has a line size of 32 bytes, and supports write-through or write-back caching.

The data/mini-data cache is controlled by page attributes defined in the MMU Architecture and by coprocessor 15.

The Intel XScale® microarchitecture allows applications to reconfigure a portion of the data cache as data RAM. Software may place special tables or frequently used variables in this RAM.

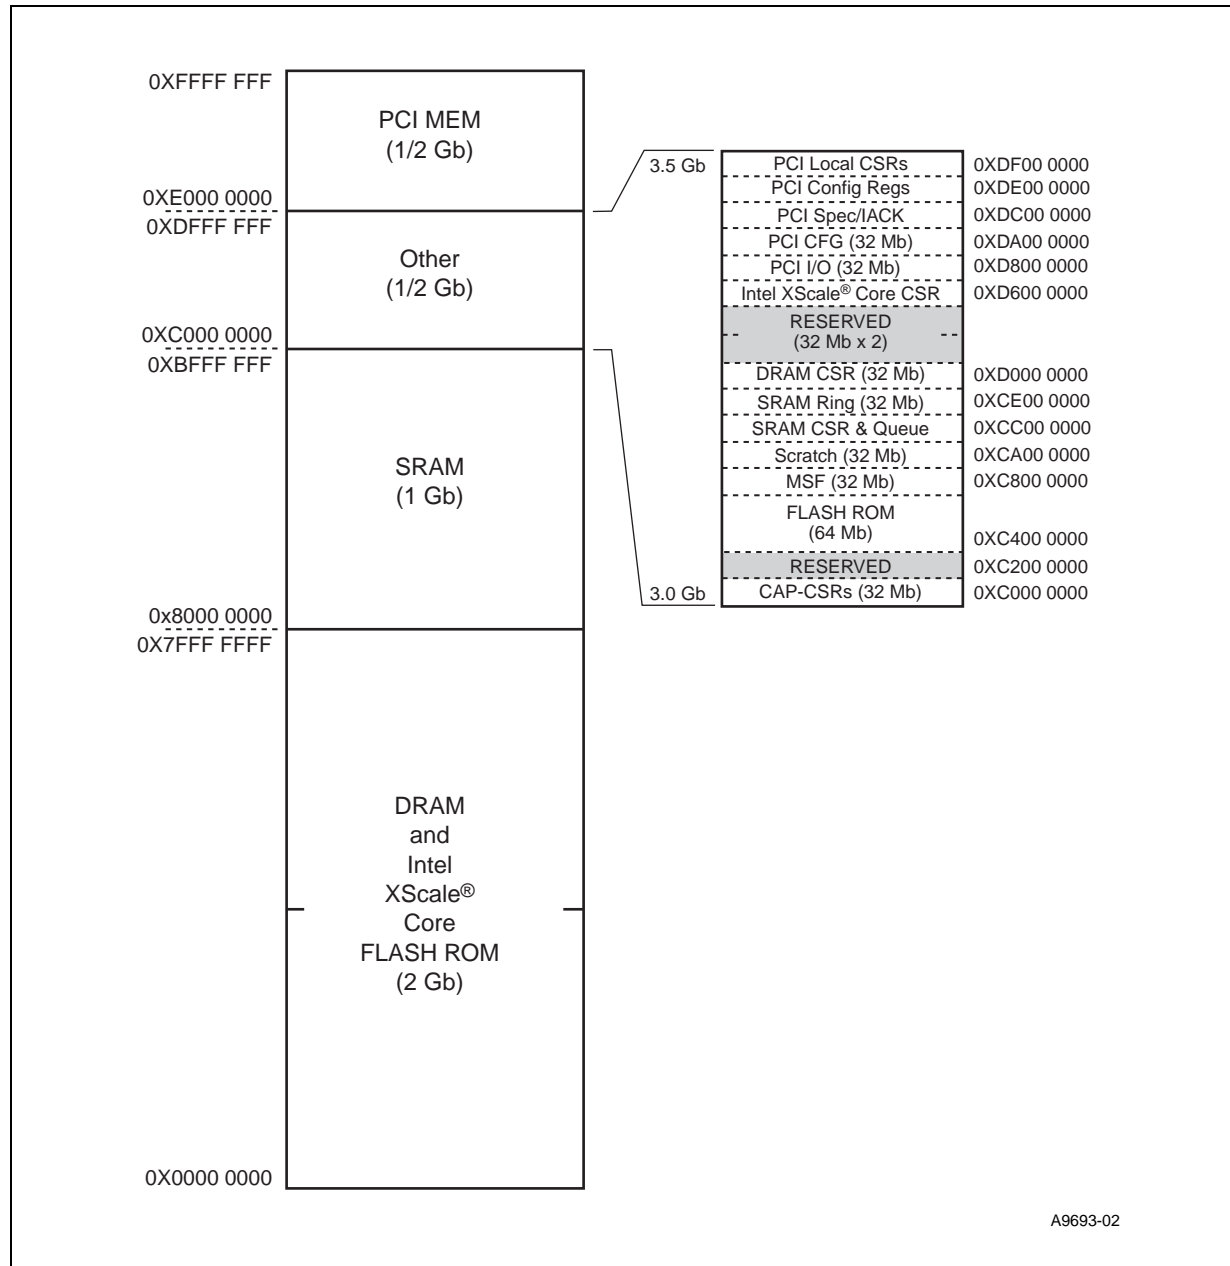
2.2.2.6 Interrupt Controller

The Intel XScale® microarchitecture provides two levels of interrupt, IRQ and FIQ. They can be masked via coprocessor 13. Note that there is also a memory-mapped interrupt controller described with the Intel XScale® technology peripherals (see [Section 3.12](#)), which is used to mask and steer many chip-wide interrupt sources.

2.2.2.7 Address Map

Figure 3 shows the partitioning of the Intel XScale® core microarchitecture 4-Gbyte address space.

Figure 3. Intel XScale® Core 4-GB (32-Bit) Address Space



A9693-02



2.3 Microengines

The Microengines do most of the programmable pre-packet processing in the IXP2800 Network Processor. There are 16 Microengines, connected as shown in [Figure 1](#). The Microengines have access to all shared resources (SRAM, DRAM, MSF, etc.) as well as private connections between adjacent Microengines (referred to as “next neighbors”).

The block diagram in [Figure 4](#) is used in the Microengine description. Note that this block diagram is simplified for clarity; some blocks and connectivity have been omitted to make the diagram more readable. Also, this block diagram does not show any pipeline stages, rather it shows the logical flow of information.

Microengines provide support for software-controlled multi-threaded operation. Given the disparity in processor cycle times versus external memory times, a single thread of execution often blocks, waiting for external memory operations to complete. Multiple threads allow for thread-interleave operation, as there is often at least one thread ready to run while others are blocked.



2.3.1 Microengine Bus Arrangement

The IXP2800 Network Processor supports a single D_Push/D_Pull bus, and both Microengine clusters interface to the same bus. Also, it supports two command buses, and two sets of S_Push/S_Pull buses connected as shown in Table 3, which also shows the next neighbor relationship between the Microengine.

Table 3. IXP2800 Network Processor Microengine Bus Arrangement

Microengine Cluster	Microengine Number	Next Neighbor	Previous Neighbor	Command Bus	S_Push and S_Pull Bus
0	0x00	0x01	NA	0	0
	0x01	0x02	0x00		
	0x02	0x03	0x01		
	0x03	0x04	0x02		
	0x04	0x05	0x03		
	0x05	0x06	0x04		
	0x06	0x07	0x05		
	0x07	0x10	0x06		
1	0x10	0x11	0x07	1	1
	0x11	0x12	0x10		
	0x12	0x13	0x11		
	0x13	0x14	0x12		
	0x14	0x15	0x13		
	0x15	0x16	0x14		
	0x16	0x17	0x15		
	0x17	NA	0x16		

2.3.2 Control Store

The Control Store is a RAM that holds the program that is executed by the Microengine. It holds 8192 instructions, each of which is 40 bits wide. It is initialized by the Intel XScale® core, which writes to USTORE_ADDR and USTORE_DATA Local CSRs.

The Control Store is protected by parity against soft errors. Parity checking is enabled by CTX_ENABLE[CONTROL STORE PARITY ENABLE]. A parity error on an instruction read will halt the Microengine and assert an interrupt to the Intel XScale® core.

2.3.3 Contexts

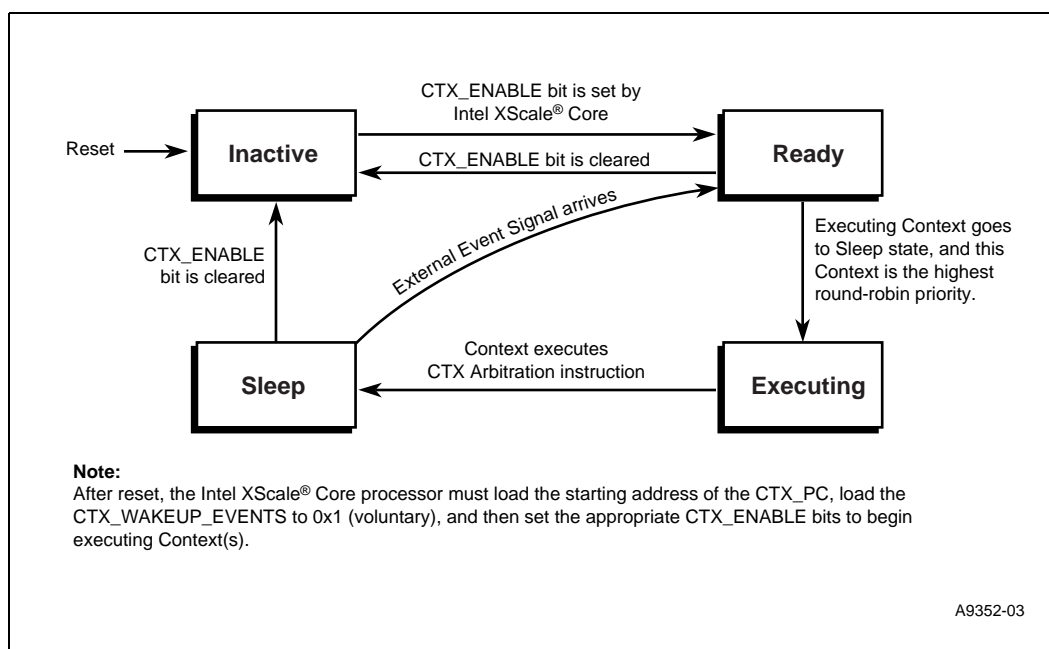
There are eight hardware Contexts available in the Microengine. To allow for efficient context swapping, each Context has its own register set, Program Counter, and Context specific Local registers. Having a copy per Context eliminates the need to move Context specific information to/from shared memory and Microengine registers for each Context swap. Fast context swapping allows a Context to do computation while other Contexts wait for I/O (typically external memory accesses) to complete or for a signal from another Context or hardware unit. (A context swap is similar to a taken branch in timing.)

Each of the eight Contexts is in one of four states.

1. Inactive — Some applications may not require all eight contexts. A Context is in the Inactive state when its CTX_ENABLE CSR enable bit is a 0.
2. Executing — A Context is in Executing state when its context number is in ACTIVE_CTX_STS CSR. The executing Context's PC is used to fetch instructions from the Control Store. A Context will stay in this state until it executes an instruction that causes it to go to Sleep state (there is no hardware interrupt or preemption; Context swapping is completely under software control). At most one Context can be in Executing state at any time.
3. Ready — In this state, a Context is ready to execute, but is not because a different Context is executing. When the Executing Context goes to the Sleep state, the Microengine's context arbiter selects the next Context to go to the Executing state from among all the Contexts in the Ready state. The arbitration is round robin.
4. Sleep — Context is waiting for external event(s) specified in the INDIRECT_WAKEUP_EVENTS CSR to occur (typically, but not limited to, an I/O access). In this state the Context does not arbitrate to enter the Executing state.

The state diagram in Figure 5 illustrates the Context state transitions. Each of the eight Contexts will be in one of these states. At most one Context can be in Executing state at a time; any number of Contexts can be in any of the other states.

Figure 5. Context State Transition Diagram



The Microengine is in Idle state whenever no Context is running (all Contexts are in either Inactive or Sleep states). This state is entered:

1. After reset (CTX_ENABLE Local CSR is clear, putting all Contexts into Inactive states).
2. When a context swap is executed, but no context is ready to wake up.
3. When a `ctx_arb[bpt]` instruction is executed by the Microengine (this is a special case of condition 2 above, since the `ctx_arb[bpt]` clears CTX_ENABLE, putting all Contexts into Inactive states).

The Microengine provides the following functionality during the Idle state:

1. The Microengine continuously checks if a Context is in Ready state. If so, a new Context begins to execute. If no Context is Ready, the Microengine remains in the Idle state.
2. Only the ALU instructions are supported. They are used for debug via special hardware defined in number 3 below.
3. A write to the USTORE_ADDR Local CSR with the USTORE_ADDR[ECS] bit set, causing the Microengine to repeatedly execute the instruction pointed by the address specified in the USTORE_ADDR CSR. Only the ALU instructions are supported in this mode. Also, the result of the execution is written to the ALU_OUT Local CSR rather than a destination register.
4. A write to the USTORE_ADDR Local CSR with the USTORE_ADDR[ECS] bit set, followed by a write to the USTORE_DATA Local CSR loads an instruction into the Control Store. After the Control Store is loaded, execution proceeds as described in number 3 above.

2.3.4 Datapath Registers

As shown in the block diagram in [Figure 4](#), each Microengine contains four types of 32-bit datapath registers:

1. 256 General Purpose registers
2. 512 Transfer registers
3. 128 Next Neighbor registers
4. 640 32-bit words of Local Memory

2.3.4.1 General-Purpose Registers (GPRs)

GPRs are used for general programming purposes. They are read and written exclusively under program control. GPRs, when used as a source in an instruction, supply operands to the execution datapath. When used as a destination in an instruction, they are written with the result of the execution datapath. The specific GPRs selected are encoded in the instruction.

The GPRs are physically and logically contained in two banks, GPR A, and GPR B, defined in [Table 5](#).

2.3.4.2 Transfer Registers

Transfer (abbreviated as Xfer) registers are used for transferring data to and from the Microengine and locations external to the Microengine, (for example DRAMs, SRAMs etc.). There are four types of transfer registers.

- S_TRANSFER_IN
- S_TRANSFER_OUT
- D_TRANSFER_IN
- D_TRANSFER_OUT

TRANSFER_IN registers, when used as a source in an instruction, supply operands to the execution datapath. The specific register selected is either encoded in the instruction, or selected indirectly via T_INDEX. TRANSFER_IN registers are written by external units (A typical case is when the external unit returns data in response to read instructions. However, there are other



methods to write TRANSFER_IN registers, for example a read instruction executed by one Microengine may cause the data to be returned to a different Microengine. Details are covered in the instruction set descriptions).

TRANSFER_OUT registers, when used as a destination in an instruction, are written with the result from the execution datapath. The specific register selected is encoded in the instruction, or selected indirectly via T_INDEX. TRANSFER_OUT registers supply data to external units (for example, write data for an SRAM write).

The S_TRANSFER_IN and S_TRANSFER_OUT registers connect to the S_PUSH and S_PULL buses, respectively.

The D_TRANSFER_IN and D_TRANSFER_OUT Transfer registers connect to the D_PUSH and D_PULL buses, respectively.

Typically, the external units access the Transfer registers in response to instructions executed by the Microengines. However, it is possible for an external unit to access a given Microengine’s Transfer registers either autonomously, or under control of a different Microengine, or the Intel XScale® core, etc. The Microengine interface signals controlling writing/reading of the TRANSFER_IN and TRANSFER_OUT registers are independent of the operation of the rest of the Microengine, therefore the data movement does not stall or impact other instruction processing (it is the responsibility of software to synchronize usage of read data).

2.3.4.3 Next Neighbor Registers

Next Neighbor registers, when used as a source in an instruction, supply operands to the execution datapath. They are written in two different ways:

1. By an adjacent Microengine (the “Previous Neighbor”).
2. By the same Microengine they are in, as controlled by CTX_ENABLE[NN_MODE].

The specific register is selected in one of two ways:

1. Context-relative, the register number is encoded in the instruction.
2. As a Ring, selected via NN_GET and NN_PUT CSR registers.

The usage is configured in CTX_ENABLE[NN_MODE].

- When CTX_ENABLE[NN_MODE] is ‘0’ — when Next Neighbor is a destination in an instruction, the result is sent out of the Microengine, to the Next Neighbor Microengine.
- When CTX_ENABLE[NN_MODE] is ‘1’ — when Next Neighbor is used as a destination in an instruction, the instruction result data is written to the selected Next Neighbor register in the same Microengine. Note that there is a 5-instruction latency until the newly written data may be read. The data is not sent out of the Microengine as it would be when CTX_ENABLE[NN_MODE] is ‘0’.

Table 4. Next Neighbor Write as a Function of CTX_ENABLE[NN_MODE]

NN_MODE	Where the Write Goes	
	External?	NN Register in this Microengine?
0	Yes	No
1	No	Yes



2.3.4.4 Local Memory

Local Memory is addressable storage within the Microengine. Local Memory is read and written exclusively under program control. Local Memory supplies operands to the execution datapath as a source, and receives results as a destination. The specific Local Memory location selected is based on the value in one of the LM_ADDR registers, which are written by local_csr_wr instructions. There are two LM_ADDR registers per Context and a working copy of each. When a Context goes to the Sleep state, the value of the working copies is put into the Context's copy of LM_ADDR. When the Context goes to the Executing state, the value in its copy of LM_ADDR are put into the working copies. The choice of LM_ADDR_0 or LM_ADDR_1 is selected in the instruction.

It is also possible to make use of both or one LM_ADDRs as global by setting CTX_ENABLE[LM_ADDR_0_GLOBAL] and/or CTX_ENABLE[LM_ADDR_1_GLOBAL]. When used globally, all Contexts use the working copy of LM_ADDR in place of their own Context specific one; the Context specific ones are unused. There is a three-instruction latency when writing a new value to the LM_ADDR, as shown in [Example 1](#).

Example 1. Three-Cycle Latency when Writing a New Value to LM_ADDR

```

;some instruction to compute the address into gpr_m
local_csr_wr[INDIRECT_LM_ADDR_0, gpr_m]; put gpr_m into lm_addr
;unrelated instruction 1
;unrelated instruction 2
;unrelated instruction 3
alu[dest_reg, *l$index0, op, src_reg]
;dest_reg can be used as a source in next instruction

```

LM_ADDR can also be incremented or decremented in parallel with use as a source and/or destination (using the notation *l\$index#++ and *l\$index#--), as shown in [Example 2](#), where three consecutive Local Memory locations are used in three consecutive instructions.

Example 2. Using LM_ADDR in Consecutive Instructions

```

alu[dest_reg1, src_reg1, op, *l$index0++]
alu[dest_reg2, src_reg2, op, *l$index0++]
alu[dest_reg3, src_reg3, op, *l$index0++]

```

Local Memory is written by selecting it as a destination. [Example 3](#) shows copying a section of Local Memory to another section. Each instruction accesses the next sequential Local Memory location from the previous instruction.

Example 3. Copying One Section of Local Memory to Another Section

```

alu[*l$index1++, --, B, *l$index0++]
alu[*l$index1++, --, B, *l$index0++]
alu[*l$index1++, --, B, *l$index0++]

```

[Example 4](#) shows loading and using both Local Memory addresses.

Example 4. Loading and Using Both Local Memory Addresses

```

local_csr_wr[INDIRECT_LM_ADDR_0, gpr_m]
local_csr_wr[INDIRECT_LM_ADDR_1, gpr_n]
;unrelated instruction 1
;unrelated instruction 2
alu[dest_reg1, *l$index0, op, src_reg1]
alu[dest_reg2, *l$index1, op, src_reg2]

```



As shown in [Example 1](#), there is a latency in loading LM_ADDR. Until the new value is loaded, the old value is still usable. [Example 5](#) shows the maximum pipelined usage of LM_ADDR.

Example 5. Maximum Pipelined Usage of LM_ADDR

```
local_csr_wr[INDIRECT_LM_ADDR_0, gpr_m]
local_csr_wr[INDIRECT_LM_ADDR_0, gpr_n]
local_csr_wr[INDIRECT_LM_ADDR_0, gpr_o]
local_csr_wr[INDIRECT_LM_ADDR_0, gpr_p]
alu[dest_reg1, *1$index0, op, src_reg1] ; uses address from gpr_m
alu[dest_reg2, *1$index0, op, src_reg2] ; uses address from gpr_n
alu[dest_reg3, *1$index0, op, src_reg3] ; uses address from gpr_o
alu[dest_reg4, *1$index0, op, src_reg4] ; uses address from gpr_p
```

LM_ADDR can also be used as the base of a 16 32-bit word region of memory, with the instruction specifying the offset from that base, as shown in [Example 6](#). The source and destination can use different offsets.

Example 6. LM_ADDR Used as Base of a 16 32-Bit Word Region of Local Memory

```
alu[*1$index0[3], *1$index0[4], +, 1]
```

Note: Local Memory has 640 32-bit words. The local memory pointers (LM_ADDR) have an addressing range of up to 1K longwords. However, only 640 longwords are currently populated with RAM. Therefore:

0 – 639 (0x0 – 0x27F) are addressable as local memory.

640 – 1023 (0x280 – 0x3FF) are addressable, but not populated with RAM.

To the programmer, all instructions using Local Memory act as follows, including read/modify/write instructions like `immed_w0`, `ld_field`, etc.

1. Read LM_ADDR location (if LM_ADDR is specified as source).
2. Execute logic function.
3. Write LM_ADDR location (if LM_ADDR is specified as destination).
4. If specified, increment or decrement LM_ADDR.
5. Proceed to next instruction.

[Example 7](#) is legal because `lm_addr_0[2]` does not post-modify LM_ADDR.

Example 7. LM_ADDR Use as Source and Destination

```
alu[*1$index0[2], --, ~B, *1$index0]
```

In [Example 7](#), the programmer sees:

1. Read Local Memory memory location pointed to by LM_ADDR.
2. Invert the data.
3. Write the data into the address pointed to by LM_ADDR with the value of 2 that is OR'ed into the lower bits.
4. Increment LM_ADDR.
5. Proceed to next instruction.

In [Example 8](#), the second instruction will access the Local Memory location one past the source/destination of the first.

Example 8. LM_ADDR Post-Increment

```
alu[*1$index0++, --, ~B, gpr_n]
alu[gpr_m, --, ~B, *1$index0]
```

2.3.5 Addressing Modes

GPRs can be accessed in either a context-relative or an absolute addressing mode. Some instructions can specify either mode; other instructions can specify only Context-Relative mode.

Transfer and Next Neighbor registers can be accessed in Context-Relative and Indexed modes, and Local Memory is accessed in Indexed mode. The addressing mode in use is encoded directly into each instruction, for each source and destination specifier.

2.3.5.1 Context-Relative Addressing Mode

The GPRs are logically subdivided into equal regions such that each Context has relative access to one of the regions. The number of regions is configured in the CTX_ENABLE CSR, and can be either 4 or 8. Thus a Context-Relative register number is actually associated with multiple different physical registers. The actual register to be accessed is determined by the Context making the access request (the Context number is concatenated with the register number specified in the instruction). Context-Relative addressing is a powerful feature that enables eight (or four) different contexts to share the same code image, yet maintain separate data.

[Table 5](#) shows how the Context number is used in selecting the register number in relative mode. The register number in [Table 5](#) is the Absolute GPR address, or Transfer or Next Neighbor Index number to use to access the specific Context-Relative register. For example, with eight active Contexts, Context-Relative Register 0 for Context 2 is Absolute Register Number 32.

Table 5. Registers Used By Contexts in Context-Relative Addressing Mode

Number of Active Contexts	Active Context Number	GPR Absolute Register Numbers		S_Transfer or Neighbor Index Number	D_Transfer Index Number
		A Port	B Port		
8 (Instruction always specifies registers in range 0 – 15)	0	0 – 15	0 – 15	0 – 15	0 – 15
	1	16 – 31	16 – 31	16 – 31	16 – 31
	2	32 – 47	32 – 47	32 – 47	32 – 47
	3	48 – 63	48 – 63	48 – 63	48 – 63
	4	64 – 79	64 – 79	64 – 79	64 – 79
	5	80 – 95	80 – 95	80 – 95	80 – 95
	6	96 – 111	96 – 111	96 – 111	96 – 111
	7	112 – 127	112 – 127	112 – 127	112 – 127
4 (Instruction always specifies registers in range 0 – 31)	0	0 – 31	0 – 31	0 – 31	0 – 31
	2	32 – 63	32 – 63	32 – 63	32 – 63
	4	64 – 95	64 – 95	64 – 95	64 – 95
	6	96 – 127	96 – 127	96 – 127	96 – 127



2.3.5.2 Absolute Addressing Mode

With Absolute addressing, any GPR can be read or written by any of the eight Contexts in a Microengine. Absolute addressing enables register data to be shared among all of the Contexts, e.g., for global variables or for parameter passing. All 256 GPRs can be read by Absolute address.

2.3.5.3 Indexed Addressing Mode

With Indexed addressing, any Transfer or Next Neighbor register can be read or written by any one of the eight Contexts in a Microengine. Indexed addressing enables register data to be shared among all of the Contexts. For indexed addressing the register number comes from the T_INDEX register for Transfer registers or NN_PUT and NN_GET registers (for Next Neighbor registers). [Example 9](#) shows the Index Mode usage. Assume that the numbered bytes have been moved into the S_TRANSFER_IN registers as shown.

Example 9. Use of Indexed Addressing Mode

Transfer Register Number	Data			
	31:24	23:16	15:8	7:0
0	0x00	0x01	0x02	0x03
1	0x04	0x05	0x06	0x07
2	0x08	0x09	0x0a	0x0b
3	0x0c	0x0d	0x0e	0x0f
4	0x10	0x11	0x12	0x13
5	0x14	0x15	0x16	0x17
6	0x18	0x19	0x1a	0x1b
7	0x1c	0x1d	0x1e	0x1f

If the software wants to access a specific byte that is known at compile-time, it will normally use context-relative addressing. For example to access the word in transfer register 3:

```
alu[dest, --, B, $xfer3] ; move the data from s_transfer 3 to gpr dest
```

If the location of the data is found at run-time, indexed mode can be used, e.g., if the start of an encapsulated header depends on an outer header value (the outer header byte is in a fixed location).

```
; Check byte 2 of transfer 0
; If value==5 header starts on byte 0x9, else byte 0x14
br=byte[$0, 2, 0x5, L1#], defer_[1]
local_csr_wr[t_index_byte_index, 0x09]
local_csr_wr[t_index_byte_index, 0x14]
nop ; wait for index registers to be loaded
L1#:
; Move bytes right justified into destination registers
nop ; wait for index registers to be loaded
nop ;
byte_align_be[dest1, *$index++]
byte_align_be[dest2, *$index++] ;etc.
; The t_index and byte_index registers are loaded by the same instruction.
```



2.3.6 Local CSRs

Local Control and Status registers (CSRs) are external to the Execution Datapath, and hold specific data. They can be read and written by special instructions (`local_csr_rd` and `local_csr_wr`) and are accessed less frequently than datapath registers.

Because Local CSRs are not built in the datapath, there is a write-to-use delay of three instructions, and a read-to-consume penalty of two instructions.

2.3.7 Execution Datapath

The Execution Datapath can take one or two operands, perform an operation, and optionally write back a result. The sources and destinations can be GPRs, Transfer registers, Next Neighbor registers, and Local Memory. The operations are shifts, add/subtract, logicals, multiply, byte align, and find first one bit.

2.3.7.1 Byte Align

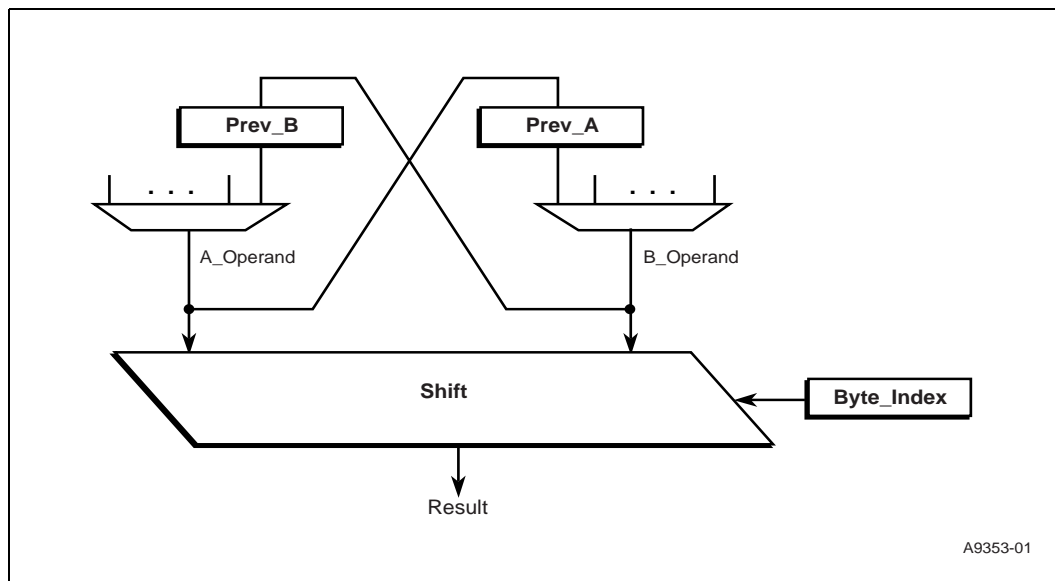
The datapath provides a mechanism to move data from source register(s) to any destination register(s) with byte aligning. Byte aligning takes four consecutive bytes from two concatenated values (8 bytes), starting at any of four byte boundaries (0, 1, 2, 3), and based on the endian-type (which is defined in the instruction opcode), as shown in [Example 5](#). The four bytes are taken from two concatenated values. Four bytes are always supplied from a temporary register that always holds the A or B operand from the previous cycle, and the other four bytes from the B or A operand of the Byte Align instruction.

The operation is described below, using the block diagram in [Figure 6](#). The alignment is controlled by the two LSBs of the `BYTE_INDEX` Local CSR.

Table 6. Align Value and Shift Amount

Align Value (in <code>Byte_Index[1:0]</code>)	Right Shift Amount (Number of Bits) (Decimal)	
	Little-Endian	Big-Endian
0	0	32
1	8	24
2	16	16
3	24	8

Figure 6. Byte-Align Block Diagram



Example 10 shows a big-endian align sequence of instructions and the value of the various operands. Table 7 shows the data in the registers for this example. The value in BYTE_INDEX[1:0] CSR (which controls the shift amount) for this example is 2.

Table 7. Register Contents for Example 10

Register	Byte 3 [31:24]	Byte 2 [23:16]	Byte 1 [15:8]	Byte 0 [7:0]
0	0	1	2	3
1	4	5	6	7
2	8	9	A	B
3	C	D	E	F

Example 10. Big-Endian Align

Instruction	Prev B	A Operand	B Operand	Result
Byte_align_be[--, r0]	--	--	0123	--
Byte_align_be[dest1, r1]	0123	0123	4567	2345
Byte_align_be[dest2, r2]	4567	4567	89AB	6789
Byte_align_be[dest3, r3]	89AB	89AB	CDEF	ABCD

NOTE: A Operand comes from Prev_B register during byte_align_be instructions.

Example 11 shows a little-endian sequence of instructions and the value of the various operands. Table 8 shows the data in the registers for this example. The value in BYTE_INDEX[1:0] CSR (which controls the shift amount) for this example is 2.

Table 8. Register Contents for Example 11

Register	Byte 3 [31:24]	Byte 2 [23:16]	Byte 1 [15:8]	Byte 0 [7:0]
0	3	2	1	0
1	7	6	5	4
2	B	A	9	8
3	F	E	D	C

Example 11. Little-Endian Align

Instruction	A Operand	B Operand	Prev A	Result
Byte_align_le[--, r0]	3210	--	--	--
Byte_align_le[dest1, r1]	7654	3210	3210	5432
Byte_align_le[dest2, r2]	BA98	7654	7654	9876
Byte_align_le[dest3, r3]	FEDC	BA98	BA98	DCBA
NOTE: B Operand comes from Prev_A register during byte_align_le instructions.				

As the examples show, byte aligning “n” words takes “n+1” cycles due to the first instruction needed to start the operation.

Another mode of operation is to use the T_INDEX register with post-increment, to select the source registers. T_INDEX operation is described later in this chapter.

2.3.7.2 CAM

The block diagram in Figure 7 is used to explain the CAM operation.

The CAM has 16 entries. Each entry stores a 32-bit value, which can be compared against a source operand by instruction:


```
CAM_Lookup[dest_reg, source_reg]
```

All entries are compared in parallel, and the result of the lookup is a 9-bit value that is written into the specified destination register in bits 11:3, with all other bits of the register 0 (the choice of bits 11:3 is explained below). The result can also optionally be written into either of the LM_Addr registers (see below in this section for details).

The 9-bit result consists of four State bits (dest_reg[11:8]), concatenated with a 1-bit Hit/Miss indication (dest_reg[7]), concatenated with 4-bit entry number (dest_reg[6:3]). All other bits of dest_reg are written with 0. Possible results of the lookup are:

- miss (0) — lookup value is not in CAM, entry number is Least Recently Used entry (which can be used as a suggested entry to replace), and State bits are 0000.
- hit (1) — lookup value is in CAM, entry number is entry that has matched; State bits are the value from the entry that has matched.

Figure 7. CAM Block Diagram



Note: The State bits are data associated with the entry. The use is only by software. There is no implication of ownership of the entry by any Context. The State bits hardware function is:

- the value is set by software (at the time the entry is loaded, or changed in an already loaded entry).
- its value is read out on a lookup that hits, and used as part of the status written into the destination register.
- its value can be read out separately (normally only used for diagnostic or debug).

The LRU (Least Recently Used) Logic maintains a time-ordered list of CAM entry usage. When an entry is loaded, or matches on a lookup, it is marked as MRU (Most Recently Used). Note that a lookup that misses does *not* modify the LRU list.

The CAM is loaded by instruction:

```
CAM_Write[entry_reg, source_reg, state_value]
```

The value in the register specified by `source_reg` is put into the Tag field of the entry specified by `entry_reg`. The value for the State bits of the entry is specified in the instruction as `state_value`.



The value in the State bits for an entry can be written, without modifying the Tag, by instruction:

```
CAM_Write_State[entry_reg, state_value]
```

Note: CAM_Write_State does *not* modify the LRU list.

One possible way to use the result of a lookup is to dispatch to the proper code using instruction:

```
jump[register, label#],defer [3]
```

where the register holds the result of the lookup. The State bits can be used to differentiate cases where the data associated with the CAM entry is in flight, or is pending a change, etc. Because the lookup result was loaded into bits[11:3] of the destination register, the jump destinations are spaced eight instructions apart. This is a balance between giving enough space for many applications to complete their task without having to jump to another region, versus consuming too much Control Store. Another way to use the lookup result is to branch on just the hit miss bit, and use the entry number as a base pointer into a block of Local Memory.

When enabled, the CAM lookup result is loaded into Local_Addr as follows:

```
LM_Addr[5:0] = 0 ([1:0] are read-only bits)  
LM_Addr[9:6] = lookup result [6:3] (entry number)  
LM_Addr[11:10] = constant specified in instruction
```

This function is useful when the CAM is used as a cache, and each entry is associated with a block of data in Local Memory. Note that the latency from when CAM_Lookup executes until the LM_Addr is loaded is the same as when LM_Addr is written by a Local_CSR_Wr instruction.

The Tag and State bits for a given entry can be read by instructions:

```
CAM_Read_Tag[dest_reg, entry_reg]  
CAM_Read_State[dest_reg, entry_reg]
```

The Tag value and State bits value for the specified entry is written into the destination register, respectively for the two instructions (the State bits are placed into bits [11:8] of dest_reg, with all other bits 0). Reading the tag is useful in the case where an entry needs to be evicted to make room for a new value—the lookup of the new value results in a miss, with the LRU entry number returned as a result of the miss. The CAM_Read_Tag instruction can then be used to find the value that was stored in that entry. An alternative would be to keep the tag value in a GPR. These two instructions can also be used by debug and diagnostic software. Neither of these modify the state of the LRU pointer.

Note: The following rules must be adhered to when using the CAM.

- CAM is not reset by Microengine reset. Software must either do a CAM_clear prior to using the CAM to initialize the LRU and clear the tags to 0, or explicitly write all entries with CAM_write.
- No two tags can be written to have same value. If this rule is violated, the result of a lookup that matches that value will be unpredictable, and LRU state is unpredictable.

The value 0x00000000 can be used as a valid lookup value. However, note that CAM_clear instruction puts 0x00000000 into all tags. To avoid violating rule 2 after doing CAM_clear, it is necessary to write all entries to unique values prior to doing a lookup of 0x00000000.

An algorithm for debug software to find out the contents of the CAM is shown in [Example 12](#).

Example 12. Algorithm for Debug Software to Find out the Contents of the CAM

```
; First read each of the tag entries. Note that these reads
; don't modify the LRU list or any other CAM state.
tag[0] = CAM_Read_Tag(entry_0);
.....
tag[15] = CAM_Read_Tag(entry_15);
; Now read each of the state bits
state[0] = CAM_Read_State(entry_0);
...
state[15] = CAM_Read_State(entry_15);
; Knowing what tags are in the CAM makes it possible to
; create a value that is not in any tag, and will therefore
; miss on a lookup.

; Next loop through a sequence of 16 lookups, each of which will
; miss, to obtain the LRU values of the CAM.
for (i = 0; i < 16; i++)
    BEGIN_LOOP
        ; Do a lookup with a tag not present in the CAM. On a
        ; miss, the LRU entry will be returned. Since this lookup
        ; missed the LRU state is not modified.
        LRU[i] = CAM_Lookup(some_tag_not_in_cam);
        ; Now do a lookup using the tag of the LRU entry. This
        ; lookup will hit, which makes that entry MRU.
        ; This is necessary to allow the next lookup miss to
        ; see the next LRU entry.
        junk = CAM_Lookup(tag[LRU[i]]);
    END_LOOP
; Because all entries were hit in the same order as they were
; LRU, the LRU list is now back to where it started before the
; loop executed.
; LRU[0] through LRU[15] holds the LRU list.
```

The CAM can be cleared with CAM_Clear instruction. This instruction writes 0x00000000 simultaneously to all entries tag, clears all the state bits, and puts the LRU into an initial state (where entry 0 is LRU, ..., entry 15 is MRU).

2.3.8 CRC Unit

The CRC Unit operates in parallel with the Execution Datapath. It takes two operands, performs a CRC operation, and writes back a result. CRC-CCITT, CRC-32, CRC-10, CRC-5, and iSCSI polynomials are supported. One of the operands is the CRC_Remainder Local CSR, and the other is a GPR, Transfer_In register, Next Neighbor, or Local Memory, specified in the instruction and passed through the Execution Datapath to the CRC Unit.

The instruction specifies the CRC operation type, whether to swap bytes and or bits, and which bytes of the operand to include in the operation. The result of the CRC operation is written back into CRC_Remainder. The source operand can also be written into a destination register (however the byte/bit swapping and masking do not affect the destination register; they only affect the CRC computation). This allows moving data, for example, from S_TRANSFER_IN registers to S_TRANSFER_OUT registers at the same time as computing the CRC.



2.3.9 Event Signals

Event Signals are used to coordinate a program with completion of external events. For example, when a Microengine executes an instruction to an external unit to read data (which will be written into a Transfer_In register), the program must insure that it does not try to use the data until the external unit has written it. This time is not deterministic due to queuing delays and other uncertainty in the external units (for example, DRAM refresh). There is no hardware mechanism to flag that a register write is pending, and then prevent the program from using it. Instead the coordination is under software control, with hardware support.

In the instructions that use external units (i.e., SRAM, DRAM, etc.) there are fields that direct the external unit to supply an indication (called an Event Signal) that the command has been completed. There are 15 Event Signals per Context that can be used, and Local CSRs per Context to track which Event Signals are pending and which have been returned. The Event Signals can be used to move a Context from Sleep state to Ready state, or alternatively, the program can test and branch on the status of Event Signals.

Event Signals can be set in nine different ways.

1. When data is written into S_TRANSFER_IN registers
2. When data is written into D_TRANSFER_IN registers
3. When data is taken from S_TRANSFER_OUT registers
4. When data is taken from D_TRANSFER_OUT registers
5. By a write to INTERTHREAD_SIGNAL register
6. By a write from Previous Neighbor Microengine to NEXT_NEIGHBOR_SIGNAL
7. By a write from Next Neighbor Microengine to PREVIOUS_NEIGHBOR_SIGNAL
8. By a write to SAME_ME_SIGNAL Local CSR
9. By Internal Timer

Any or all Event Signals can be set by any of the above sources.

When a Context goes to the Sleep state (executes a `ctx_arb` instruction, or an instruction with `ctx_swap` token), it specifies which Event Signal(s) it requires to be put in Ready state. The `ctx_arb` instruction also specifies if the logical AND or logical OR of the Event Signal(s) is needed to put the Context into Ready state.

When all of the Context's Event Signals arrive, the Context goes to Ready state, and then eventually to Executing state. In the case where the Event Signal is linked to moving data into or out of Transfer registers (numbers 1 through 4 in the list above), the code can safely use the Transfer register as the first instruction (for example, using a Transfer_In register as a source operand will get the new read data). The same is true when the Event Signal is tested for branches (`br=signal` or `br!signal` instructions).

The `ctx_arb` instruction, `CTX_SIG_EVENTS`, and `ACTIVE_CTX_WAKEUP_#_EVENTS` Local CSR descriptions provide details.



2.4 DRAM

The IXP2800 Network Processor has controllers for three Rambus* DRAM (RDRAM) channels. Each of the controllers independently accesses its own RDRAMs, and can operate concurrently with the other controllers (i.e., they are not operating as a single, wider memory). DRAM provides high-density, high-bandwidth storage and is typically used for data buffers.

- RDRAM sizes of 64, 128, 256, or 512 Mbytes, and 1 Gbyte are supported; however, each of the channels must have the same number, size, and speed of RDRAMs populated. Refer to [Section 5.2](#) for supported size and loading configurations.
- Up to two Gbytes of DRAM is supported. If less than two Gbytes of memory is present, the upper part of the address space is not used. It is also possible, for system cost and area savings, to have Channels 0 and 1 populated with Channel 2 empty, or Channel 0 populated with Channels 1 and 2 empty.
- Reads and writes to RDRAM are generated by Microengines, The Intel XScale® core, and PCI (external Bus Masters and DMA Channels). The controllers also do refresh and calibration cycles to the RDRAMs, transparently to software.
- RDRAM Powerdown and Nap modes are not supported.
- Hardware interleaving (also known as striping) of addresses is done to provide balanced access to all populated channels. The interleave size is 128 bytes. Interleaving helps to maintain utilization of available bandwidth by spreading consecutive accesses to multiple channels. The interleaving is done in the hardware in such a way that the three channels appear to software as a single contiguous memory space.
- ECC (Error Correcting Code) is supported, but can be disabled. Enabling ECC requires that x18 RDRAMs be used. If ECC is disabled x16 RDRAMs can be used. ECC can detect and correct all single-bit errors, and detect all double-bit errors. When ECC is enabled, partial writes (writes of less than 8 bytes) must be done as read-modify-writes.

2.4.1 Size Configuration

Each channel can be populated with anywhere from one-to-four RDRAMs (Short Channel Mode). Refer to [Section 5.2](#) for supported size and loading configurations. The RAM technology used will determine the increment size and maximum memory per channel as shown in [Table 9](#).

Table 9. RDRAM Sizes

RDRAM Technology ¹	Increment Size	Maximum per Channel
64/72 MB	8 MB	256 MB
128/144 MB	16 MB	512 MB
256/288 MB	32 MB	1 GB ²
512/576 MB	64 MB	2 GB ²
NOTES: 1. The two numbers shown for each technology indicate x16 parts and x18 parts. 2. The maximum memory that can be addressed across all channels is 2 GB. This limitation is based on the partitioning of the 4-GB address space (32-bit addresses). Therefore, if all three channels are used, each can be populated up to a maximum of 768 MB. Two channels can be populated to a maximum of 1 GB each. A single channel can be populated to a maximum of 2 GB.		

RDRAMs with 1 x 16 or 2 x 16 dependent banks, and 4 independent banks are supported.



2.4.2 Read and Write Access

The minimum DRAM physical access length is 16 bytes. Software (and PCI) can read or write as little as a single byte, however the time (and bandwidth) taken at the DRAMs is the same as for an access of 16 bytes. Therefore, the best utilization of DRAM bandwidth will be for accesses that are multiples of 16 bytes.

If ECC is enabled, writes of less than 8 bytes must do read-modify-writes, which take two 16-byte time accesses (one for the read and one for the write).

2.5 SRAM

The IXP2800 Network Processor has four independent SRAM controllers, which each support pipelined QDR synchronous static RAM (SRAM) and/or a coprocessor that adheres to QDR signaling. Any or all controllers can be left unpopulated if the application does not need to use them. SRAM are accessible by the Microengines, the Intel XScale® core, and the PCI Unit (external bus masters and DMA).

The memory is logically four bytes (32-bits) wide; physically the data pins are two bytes wide and are double clocked. Byte parity is supported. Each of the four bytes has a parity bit, which is written when the byte is written and checked when the data is read. There are byte-enables that select which bytes to write for writes of less than 32 bits.

Each of the 4 QDR ports are QDR and QDRII compatible. Each port implements the “_K” and “_C” output clocks and “_CQ” as an input and their inversions. (Note: the “_C” and “_CQ” clocks are optional). Extensive work has been performed providing impedance controls within the IXP2800 Network Processor for processor-initiated signals driving to QDR parts. Providing a clean signaling environment is critical to achieving 200 – 250 MHz QDRII data transfers.

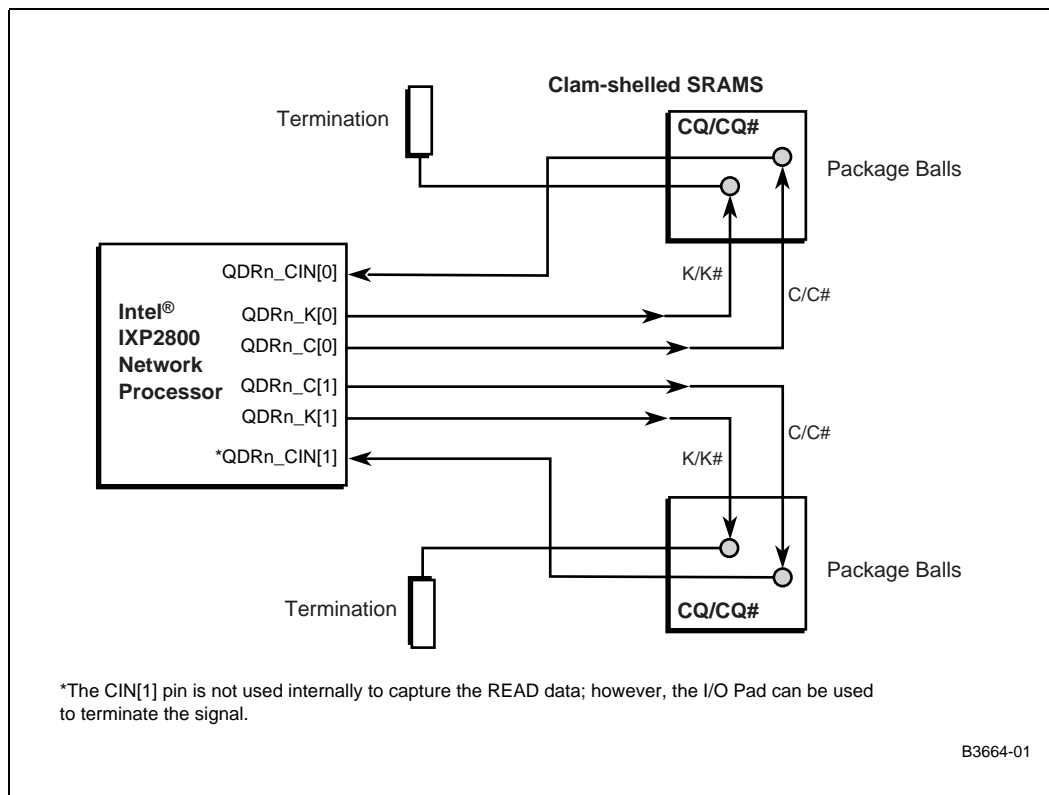
The configuration assumptions for the IXP2800 Network Processor I/O driver/receiver development includes four QDR loads and the IXP2800 Network Processor. The IXP2800 Network Processor supports bursts of two SRAMs, but does not support bursts of four SRAMs.

The SRAM controller can also be configured to interface to an external coprocessor that adheres to the QDR electricals and protocol. Each SRAM controller may also interface to an external coprocessor through its standard QDR interface. This interface enables the cohabitation of both SRAM devices and coprocessors to operate on the same bus. The coprocessor behaves as a memory-mapped device on the SRAM bus.

2.5.1 QDR Clocking Scheme

The controller drives out two pairs of K clock (K and K#). It also drives out two pairs of C clock (C and C#). Both C/C# clocks externally return to the controller for reading data. Figure 8 shows the clock diagram if the clocking scheme for QDR interface driving four SRAM chips.

Figure 8. Echo Clock Configuration



2.5.2 SRAM Controller Configurations

Each channel has enough address pins (24) to support up to 64 Mbytes of SRAM. The SRAM controllers can directly generate multiple port enables (up to four pairs) to allow for depth expansion. Two pairs of pins are dedicated for port enables. Smaller RAMs use fewer address signals than the number provided to accommodate the largest RAMs, so some address pins (23:20) are configurable as either address or port enable based on CSR setting as shown in Table 10. Note that all of the SRAMs on a given channel must be the same size.

Table 10. SRAM Controller Configurations (Sheet 1 of 2)

SRAM Configuration	SRAM Size	Addresses Needed to Index SRAM	Addresses Used as Port Enables	Total Number of Port Select Pairs Available
512K x 18	1 MB	17:0	23:22, 21:20	4
1M x 18	2 MB	18:0	23:22, 21:20	4
2M x 18	4 MB	19:0	23:22, 21:20	4
4M x 18	8 MB	20:0	23:22	3

Table 10. SRAM Controller Configurations (Sheet 2 of 2)

SRAM Configuration	SRAM Size	Addresses Needed to Index SRAM	Addresses Used as Port Enables	Total Number of Port Select Pairs Available
8M x 18	16 MB	21:0	23:22	3
16M x 18	32 MB	22:0	None	2
32M x 18	64 MB	23:0	None	2

Each channel can be expanded by depth according to the number of port enables available. If external decoding is used, then the number of SRAMs used is not limited by the number of port enables generated by the SRAM controller.

Note: Doing external decoding may require external pipeline registers to account for the decode time, depending on the desired frequency.

Maximum SRAM system sizes are shown in Table 11. Shaded entries require external decoding, because they use more port enables than the SRAM controller can supply directly.

Table 11. Total Memory per Channel

SRAM Size	Number of SRAMs on Channel							
	1	2	3	4	5	6	7	8
512K x 18	1 MB	2 MB	3 MB	4 MB	5 MB	6 MB	7 MB	8 MB
1M x 18	2 MB	4 MB	6 MB	8 MB	10 MB	12 MB	14 MB	16 MB
2M x 18	4 MB	8 MB	12 MB	16 MB	20 MB	24 MB	28 MB	32 MB
4M x 18	8 MB	16 MB	24 MB	32 MB	64 MB	NA	NA	NA
8M x 18	16 MB	32 MB	48 MB	64 MB	NA	NA	NA	NA
16M x 18	32 MB	64 MB	NA	NA	NA	NA	NA	NA
32M x 18	64 MB	NA	NA	NA	NA	NA	NA	NA

2.5.3 SRAM Atomic Operations

In addition to normal reads and writes, SRAM supports the following atomic operations. Microengines have specific instructions to do each atomic operation; Intel XScale® microarchitecture uses aliased address regions to do atomic operations.

- bit set
- bit clear
- increment
- decrement
- add
- swap

The SRAM does read-modify-writes for the atomic operations, the pre-modified data can also be returned if desired. The atomic operations operate on a single 32-bit word.

2.5.4 Queue Data Structure Commands

The ability to enqueue and dequeue data buffers at a fast rate is key to meeting line-rate performance. This is a difficult problem as it involves dependent memory references that must be turned around very quickly. The SRAM controller includes a data structure (called the Q_array) and associated control logic to perform efficient enqueue and dequeue operations. The Q_array has 64 entries, each of which can be used in one of four ways.

- Linked-list queue descriptor (resident queues)
- Cache of recently used linked-list queue descriptors (backing store for the cache is in SRAM)
- Ring descriptor
- Journal

The commands provided are:

For Linked-list queues or Cache of recently used linked-list queue descriptors

- Read_Q_Descriptor_Head(address, length, entry, xfer_addr)
- Read_Q_Descriptor_Tail(address, length, entry)
- Read_Q_Descriptor_Other(address, entry)
- Write_Q_Descriptor(address, entry)
- Write_Q_Descriptor_Count(address, entry)
- ENQ(buff_desc_adr, cell_count, EOP, entry)
- ENQ_tail(buff_desc_adr, entry)
- DEQ(entry, xfer_addr)

For Rings

- Get(entry, length, xfer_addr)
- Put(entry, length, xfer_addr)

For Journals

- Journal(entry, length, xfer_addr)
- Fast_journal(entry)

Note: The Read_Q_Descriptor_Head, Read_Q_Descriptor_Tail, etc.) are used to initialize the rings and journals but not used to perform the ring and journal function.

2.5.5 Reference Ordering

This section covers the ordering between accesses to any one SRAM controller.

2.5.5.1 Reference Order Tables

Table 12 shows the architectural guarantees of order to access to the SAME SRAM address between a reference of any given type (shown in the column labels) and a subsequent reference of any given type (shown in the row labels). The definition of first and second is defined by the order they are received by the SRAM controller.

Note: A given Network Processor version may implement a superset of these order guarantees. However, that superset may not be supported in future implementations.

Verification is required to test only the order rules shown in [Table 12](#) and [Table 13](#)).

Note: A blank entry in [Table 12](#) means that no order is enforced.

Table 12. Address Reference Order

1 st ref → 2 nd ref ↓	Memory Read	CSR Read	Memory Write	CSR Write	Memory RMW	Queue / Ring / Q_Descr Commands
Memory Read			Order			
CSR Read				Order		
Memory Write			Order			
CSR Write				Order		
Memory RMW					Order	
Queue / Ring / Q_Descr Commands						See Table 13 .

[Table 13](#) shows the architectural guarantees of order to access to the SAME SRAM Q_array entry between a reference of any given type (shown in the column labels) and a subsequent reference of any given type (shown in the row labels). The definition of first and second is defined by the order they are received by the SRAM controller. The same caveats apply as for [Table 12](#).

Table 13. Q_array Entry Reference Order

1 st ref → 2 nd ref ↓	Read_Q_Descr head, tail	Read_Q_Descr other	Write_Q_Descr	Enqueue	Dequeue	Put	Get	Journal
Read_Q_Descr head,tail			Order					
Read_Q_Descr other	Order							
Write_Q_Descr								
Enqueue	Order	Order		Order				
Dequeue	Order	Order			Order			
Put						Order		
Get							Order	
Journal								Order

2.5.5.2 Microengine Software Restrictions to Maintain Ordering

It is the Microengine programmer's job to ensure order where the program flow finds order to be necessary and where the architecture does not guarantee that order. The signaling mechanism can be used to do this. For example, say that microcode needs to update several locations in a table. A location in SRAM is used to "lock" access to the table. [Example 13](#) is the code for the table update.

Example 13. Table Update Code

```
IMMED [$xfer0, 1]
SRAM [write, $xfer0, flag_address, 0, 1], ctx_swap [SIG_DONE_2]
; At this point, the write to flag_address has passed the point of coherency. Do
the table updates.
SRAM [write, $xfer1, table_base, offset1, 2] , sig_done [SIG_DONE_3]
SRAM [write, $xfer3, table_base, offset2, 2] , sig_done [SIG_DONE_4]
CTX_ARB [SIG_DONE_3, SIG_DONE_4]
; At this point, the table writes have passed the point of coherency. Clear the
flag to allow access by other threads.
IMMED [$xfer0, 0]
SRAM [write, $xfer0, flag_address, 0, 1, ctx_swap [SIG_DONE_2]
```

Other rules:

- All accesses to atomic variables should be via read-modify-write instructions.
- If the flow must know that a write is completed (actually in the SRAM itself), follow the write with a read to the same address. The write is guaranteed to be complete when the read data has been returned to the Microengine.
- With the exception of initialization, never do WRITE commands to the first three longwords of a queue_descriptor data structure (these are the longwords that hold head, tail, and count, etc.). All accesses to this data must be via the Q commands.
- To initialize the Q_array registers, perform a memory write of at least three longwords, followed by a memory read to the same address (to guarantee that the write completed). Then, for each entry in the Q_array, perform a read_q_descriptor_head followed by a read_q_descriptor_other using the address of the same three longwords.

2.6 Scratchpad Memory

The IXP2800 Network Processor contains a 16 Kbytes of Scratchpad Memory, organized as 4K 32-bit words, that is accessible by Microengines and the Intel XScale® core. The Scratchpad Memory provides the following operations:

- Normal reads and writes. 1–16 32-bit words can be read/written with a single Microengine instruction. Note that Scratchpad is not byte-writable (each write must write all four bytes).
- Atomic read-modify-write operations, bit-set, bit-clear, increment, decrement, add, subtract, and swap. The RMW operations can also optionally return the pre-modified data.
- Sixteen Hardware Assisted Rings for interprocess communication. (A ring is a FIFO that uses a head and tail pointer to store/read information in Scratchpad memory.)

Scratchpad Memory is provided as a third memory resource (in addition to SRAM and DRAM) that is shared by the Microengines and the Intel XScale® core. The Microengines and the Intel XScale® core can distribute memory accesses between these three types of memory resources to provide a greater number of memory accesses occurring in parallel.



2.6.1 Scratchpad Atomic Operations

In addition to normal reads and writes, the Scratchpad Memory supports the following atomic operations. Microengines have specific instructions to do each atomic operation; the Intel XScale® microarchitecture uses aliased address regions to do atomic operations.

- bit set
- bit clear
- increment
- decrement
- add
- subtract
- swap

The Scratchpad Memory does read-modify-writes for the atomic operations, the pre-modified data can also be returned if desired. The atomic operations operate on a single 32-bit word.

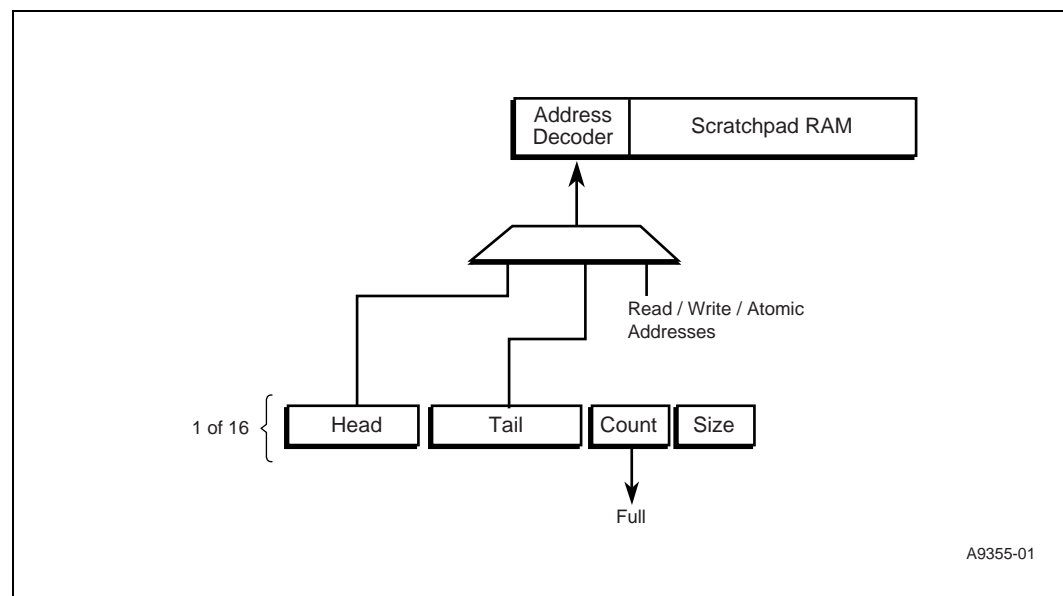
2.6.2 Ring Commands

The Scratchpad Memory provides sixteen Rings used for interprocess communication. The rings provide two operations.

- `Get(ring, length)`
- `Put(ring, length)`

Ring is the number of the ring (0 through 15) to get or put from, and length specifies the number of 32-bit words to transfer. A logical view of one of the rings is shown in Table 9.

Figure 9. Logical View of Rings





Head, Tail, and Size are registers in the Scratchpad Unit. Head and Tail point to the actual ring data, which is stored in the Scratchpad RAM. The count of how many entries are on the Ring is determined by hardware using the Head and Tail. For each Ring in use, a region of Scratchpad RAM must be reserved for the ring data.

Note: The reservation is by software convention. The hardware does not prevent other accesses to the region of Scratchpad Memory used by the Ring. Also the regions of Scratchpad Memory allocated to different Rings must not overlap.

Head points to the next address to be read on a get, and Tail points to the next address to be written on a put. The size of each Ring is selectable from the following choices: 128, 256, 512, or 1024 32-bit words.

Note: The region of Scratchpad used for a Ring is naturally aligned to its size.

When the Ring is near full, it asserts an output signal, which is used as a state input to the Microengines. They must use that signal to test (by doing Branch on Input State) for room on the Ring before putting data onto it. There is a lag in time from a put instruction executing to the Full signal being updated to reflect that put. To guarantee that a put will not overflow the ring there is a bound on the number of Contexts and the number of 32-bit words per write based on the size of the ring, as shown in Table 14. Each Context should test the Full signal, then do the put if not Full, and then wait until the Context has been signaled that the data has been pulled before testing the Full signal again.

An alternate usage method is to have Contexts allocate and deallocate entries from a shared count variable, using the atomic `subtract` to allocate and atomic `add` to deallocate. In this case the Full signal is not used.

Table 14. Ring Full Signal Use – Number of Contexts and Length versus Ring Size

Number of Contexts	Ring Size			
	128	256	512	1024
1	16	16	16	16
2	16	16	16	16
4	8	16	16	16
8	4	12	16	16
16	2	6	14	16
24	1	4	9	16
32	1	3	7	15
40	Illegal	2	5	12
48	Illegal	2	4	10
64	Illegal	1	3	7
128	Illegal	Illegal	1	3

NOTES:
 1. Number in each table entry is the largest length that should be put. 16 is the largest length that a single put instruction can generate.
 2. Illegal -- With that number of Contexts, even a length of one could cause the Ring to overflow.



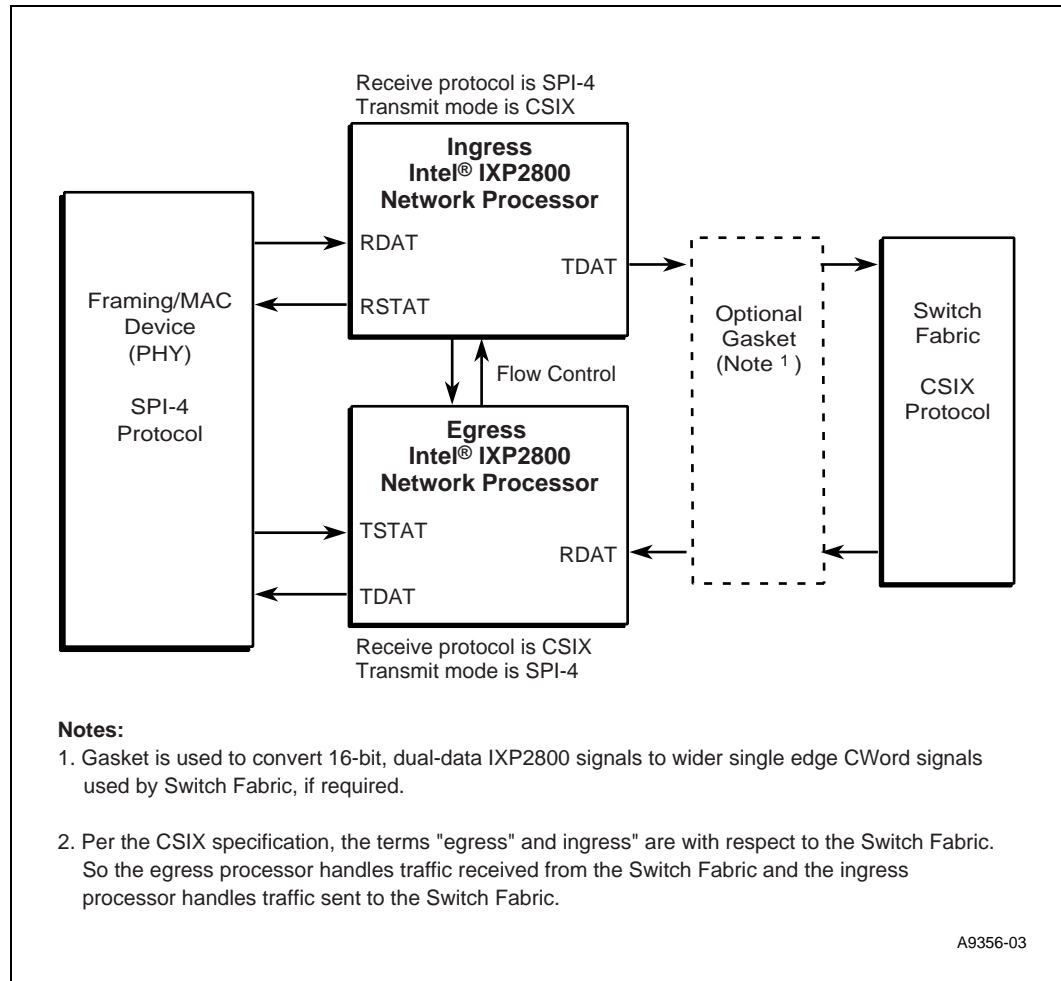
2.7 Media and Switch Fabric Interface

The Media and Switch Fabric (MSF) Interface is used to connect the IXP2800 Network Processor to a physical layer device (PHY) and/or to a Switch Fabric. The MSF consists of separate receive and transmit interfaces. Each of the receive and transmit interfaces can be separately configured for either SPI-4 Phase 2 (System Packet Interface) for PHY devices or CSIX-L1 protocol for Switch Fabric Interfaces.

The receive and transmit ports are unidirectional and independent of each other. Each port has 16 data signals, a clock, a control signal, and a parity signal, all of which use LVDS (differential) signaling, and are sampled on both edges of the clock. There is also a flow control port consisting of a clock, data, and ready status bits, and used to communicate between two IXP2800 Network Processors, or the IXP2800 Network Processor chip and a Switch Fabric Interface. These are also LVDS, dual-edge data transfer. All of the high speed LVDS interfaces support dynamic deskew training.

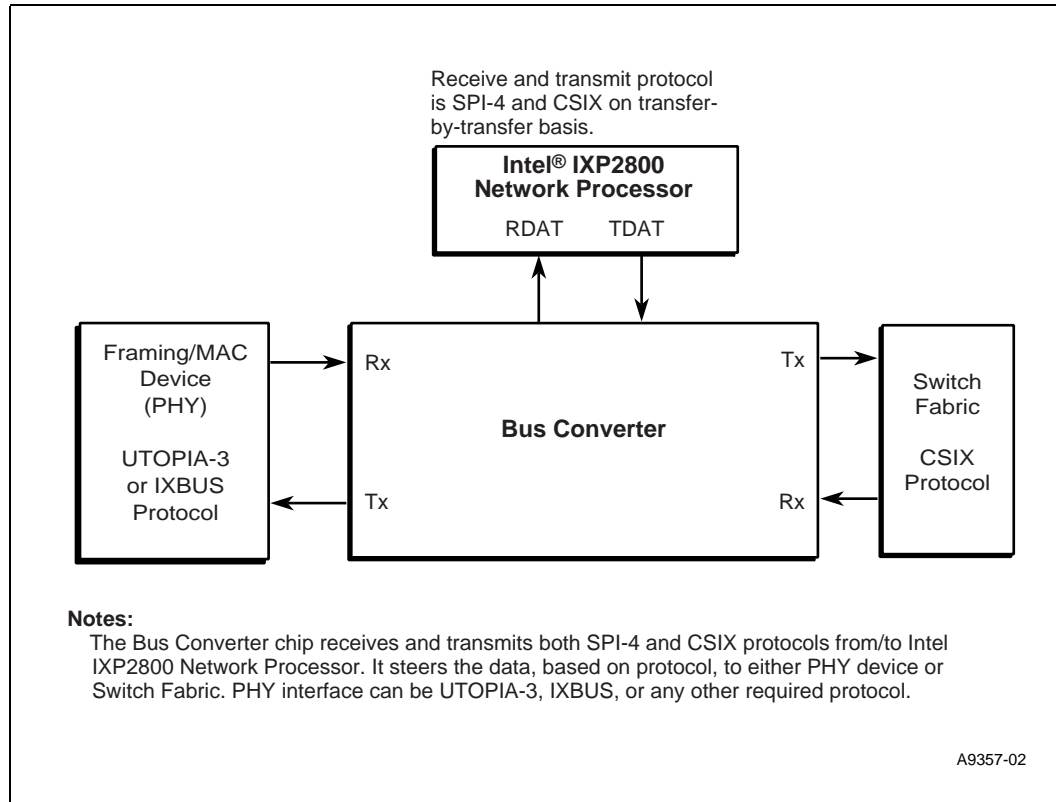
The block diagram in Figure 10 shows a typical configuration.

Figure 10. Example System Block Diagram



An alternate system configuration is shown in the block diagram in Figure 11. In this case, a single IXP2800 Network Processor is used for both Ingress and Egress. The bit rate supported would be less than in Figure 10. A hypothetical Bus Converter chip, external to the IXP2800 Network Processor is used. The block diagram in Figure 11 is only an illustrative example.

Figure 11. Full-Duplex Block Diagram



2.7.1 SPI-4

SPI-4 is an interface for packet and cell transfer between a physical layer (PHY) device and a link layer device (the IXP2800 Network Processor), for aggregate bandwidths of OC-192 ATM and Packet over SONET/SDH (POS), as well as 10 Gb/s Ethernet applications.

The Optical Internetworking Forum (OIF), www.oiforum.com, controls the SPI-4 Implementation Agreement document.

SPI-4 protocol transfers data in variable length bursts. Associated with each burst is information such as Port number (for a multi-port device such as a 10 x 1 GbE), SOP, and EOP. This information is collected by the MSF and passed to the Microengines.



2.7.2 CSIX

CSIX-L1 (Common Switch Interface) defines an interface between a Traffic Manager (TM) and a Switch Fabric (SF) for ATM, IP, MPLS, Ethernet, and similar data communications applications.

The Network Processor Forum (NPF) www.npforum.org, controls the CSIX-L1 specification.

The basic unit of information transferred between Traffic Managers and Switch Fabrics is called a CFrame. There are three categories of CFrames:

- Data
- Control
- Flow Control


Associated with each CFrame is information such as length, type, address. This information is collected by MSF and passed to Microengines.

MSF also contains a number of hardware features related to flow control.

2.7.3 Receive

Figure 12 is a simplified block diagram of the MSF receive section.

Figure 12. Simplified MSF Receive Section Block Diagram





2.7.3.1 RBUF

RBUF is a RAM that holds received data. It stores received data in sub-blocks (referred to as elements), and is accessed by Microengines or the Intel XScale® core reading the received information. Details of how RBUF elements are allocated and filled is based on the receive data protocol. When data is received, the associated status is put into the FULL_ELEMENT_LIST FIFO and subsequently sent to Microengines to process. FULL_ELEMENT_LIST insures that received elements are sent to Microengines in the order that the data was received.

RBUF contains a total of 8 KB of data. The element size is programmable as either 64 bytes, 128 bytes, or 256 bytes per element. In addition, RBUF can be programmed to be split into one, two, or three partitions depending on application. For receiving SPI-4, one partition would be used. For receiving CSIX, two partitions are used (Control CFrames and Data CFrames). When both protocols are being used, the RBUF can be split into three partitions. For both SPI-4 and CSIX, three partitions are used.

Microengines can read data from the RBUF to Microengine S_TRANSFER_IN registers using the `msf[read]` instruction where they specify the starting byte number (which must be aligned to 4 bytes), and number of 32-bit words to read. The number in the instruction can be either the number of 32-bit words, or number of 32-bit word pairs, using the single and double instruction modifiers, respectively.

Microengines can move data from RBUF to DRAM using the `dram` instruction where they specify the starting byte number (which must be aligned to 4 bytes), the number of 32-bit words to read, and the address in DRAM to write the data.

For both types of RBUF read, reading an element does not modify any RBUF data, and does not free the element, so buffered data can be read as many times as desired. This allows, for example, a processing pipeline to have different Microengines handle different protocol layers, with each Microengine reading only the specific header information it requires.

2.7.3.1.1 SPI-4 and the RBUF

SPI-4 data is placed into RBUF with each SPI-4 burst allocating an element. If a SPI-4 burst is larger than the element size, another element is allocated. The status information for the element contains the following information:

3	3	2	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	9	8	7	6	5	4	3	2	1	0
RPROT	Element						Byte Count						SOP	EOP	Err	Len Err	Par Err	Abort Err	Null	Type	ADR										

6	6	6	6	5	5	5	5	5	5	5	5	5	5	4	4	4	4	4	4	4	4	4	4	3	3	3	3	3	3	3	3
3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2
Reserved														Checksum																	

The definitions of the fields are shown in [Table 90, “RBUF SPIF-4 Status Definition”](#) on page 252.



2.7.3.1.2 CSIX and RBUF

CSIX CFrames are placed into either RBUF with each CFrame allocating an element. Unlike SPI-4, a single CFrame must not spill over into another element. Since CSIX spec specifies a maximum CFrame size of 256 bytes, this can be done by programming the element size to 256 bytes. However, if the Switch Fabric uses a smaller CFrame size, then a smaller RBUF element size can be used.

Flow Control CFrames are put into the FCEFIFO, to be sent to the Ingress IXP2800 Network Processor where a Microengine will read them to manage flow control information to the Switch Fabric.

The status information for the element contains the following information:

3	3	2	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
RPR0T	Element						Payload Length						CR	P	Er	Len Er	HP Er	VP Er	Null	Reserved					Type						
6	6	6	6	5	5	5	5	5	5	5	5	5	4	4	4	4	4	4	4	4	4	4	3	3	3	3	3	3	3	3	3
3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2
Extension Header																															

The definitions of the fields are shown in [Table 91, “RBUF CSIX Status Definition” on page 254](#).

2.7.3.2 Full Element List

Receive control hardware maintains the FULL_ELEMENT_LIST to hold the status of valid RBUF elements, in the order in which they were received. When an RBUF element is filled, its status is added to the tail of the FULL_ELEMENT_LIST. When a Microengine is notified of element arrival (by having the status written to its S_Transfer register), it is removed from the head of the FULL_ELEMENT_LIST.

2.7.3.3 RX_THREAD_FREELIST

RX_THREAD_FREELIST is a FIFO that indicates Microengine Contexts that are awaiting an RBUF element to process. This allows the Contexts to indicate their ready status prior to the reception of the data, as a way to eliminate latency. Each entry added to a Freelist also has an associated S_TRANSFER register and signal number. There are three RX_THREAD_FREELISTS that correspond to the RBUF partitions.

To be added as ready to receive an element, a Microengine does an `msf[write]` or an `msf[fast_write]` to the RX_THREAD_FREELIST address; the write data is the Microengine/CONTEXT/S_TRANSFER register number to add to the Freelist.

When there is valid status at the head of the Full Element List, it will be pushed to a Microengine. The receive control logic pushes the status information (which includes the element number) to the Microengine in the head entry of RX_THREAD_FREELIST, and sends an Event Signal to the Microengine. It then removes that entry from the RX_THREAD_FREELIST, and removes the status from Full Element List.



Each RX_THREAD_FREELIST has an associated countdown timer. If the timer expires and no new receive data is available yet, the receive logic will autopush a Null Receive Status Word to the next thread on the RX_THREAD_FREELIST. A Null Receive Status Word has the “Null” bit set, and does not have any data or RBUF entry associated with it.

The RX_THREAD_FREELIST timer is useful for certain applications. Its primary purpose is to keep the receive processing pipeline (implemented as code running on the Microengines) moving even when the line has gone idle.

It is especially useful if the pipeline is structured to handle mpackets in groups, i.e., eight mpackets at a time. If seven mpackets are received, then the line goes idle, then the timeout will trigger the autopush of a null Receive Status Word, filling the eighth slot and allowing the pipeline to advance. Another example is if one valid mpacket is received before the line goes idle for a long period; seven null Receive Status Words will be autopushed, allowing the pipeline to proceed. Typically the timeout interval is programmed to be slightly larger than the minimum arrival time of the incoming cells or packets.

The timer is controlled using the RX_THREAD_FREELIST_TIMEOUT_# CSR. The timer may be enabled or disabled, and the timeout value specified using this CSR.

2.7.3.4 Receive Operation Summary

During receive processing, received CFrames, and SPI-4 cells and packets (which in this context are all called mpackets) are placed into the RBUF, and then handed off to a Microengine to process. Normally, by application design, some number of Microengine Contexts will be assigned to receive processing. Those Contexts will have their number added to the proper RX_THREAD_FREELIST (via `msf[write]` or `msf[fast_write]`), and then will go to sleep to wait for arrival of an mpacket (or alternatively poll waiting for arrival of an mpacket).

When an mpacket arrives, MSF receive control logic will autopush eight bytes of information for the element to the Microengine/CONTEXT/S_TRANSFER registers at the head of RX_THREAD_FREELIST. The information pushed is:

- Status Word (SPI-4) or Header Status (CSIX) — see [Table 90, “RBUF SPIF-4 Status Definition” on page 252](#) for more information.
- Checksum (SPI-4) or Extension Header (CSIX) — see [Table 91, “RBUF CSIX Status Definition” on page 254](#) for more information.

To handle the case where the receive Contexts temporarily fall behind and RX_THREAD_FREELIST is empty, all received element numbers are held in the FULL_ELEMENT_LIST. In that case, as soon as an RX_THREAD_FREELIST entry is entered, the status of the head element of FULL_ELEMENT_LIST will be pushed to it.

The Microengines may read part of (or the entire) RBUF element to their S_TRANSFER registers (via an `msf[read]` instruction) for header processing, etc., and may also move the element data to DRAM (via a `dram[rbuf_rd]` instruction).

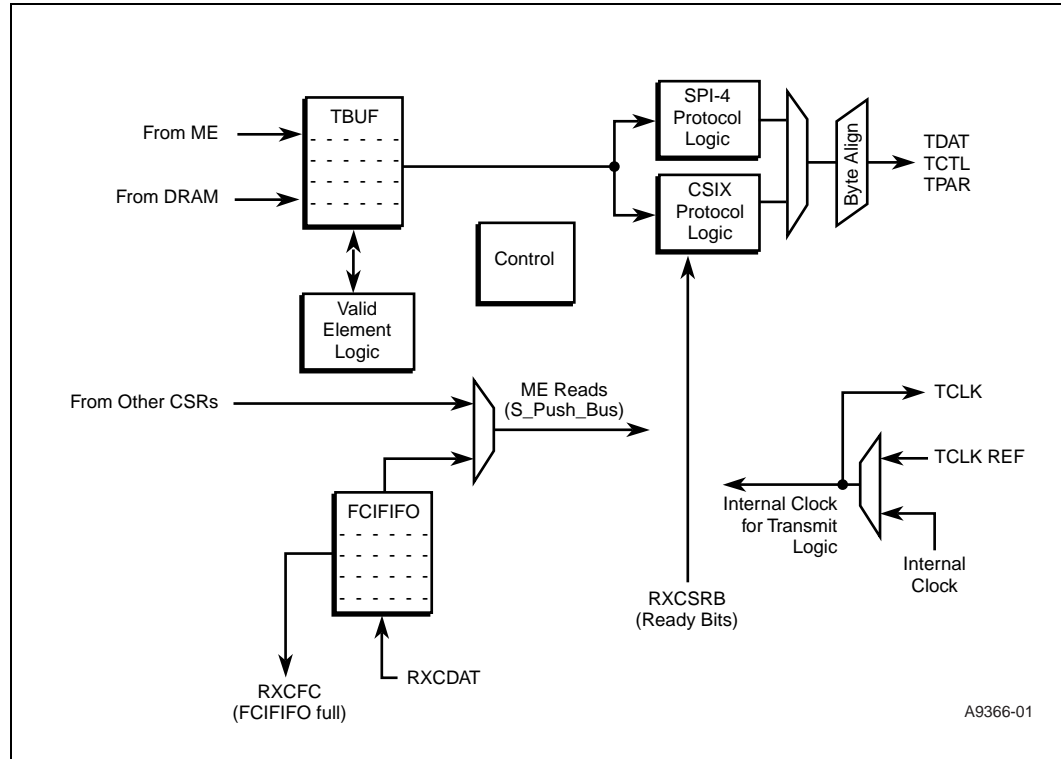
When a Context is done with an element, it does an `msf[write]` or `msf[fast_write]` to RBUF_ELEMENT_DONE address; the write data is the element number. This marks the element as free and available to be re-used. There is no restriction on the order in which elements are freed; Contexts can do different amounts of processing per element based on the contents of the element — therefore elements can be returned in a different order than they were handed to Contexts.



2.7.4 Transmit

Figure 13 is a simplified Block Diagram of the MSF transmit section.

Figure 13. Simplified Transmit Section Block Diagram



2.7.4.1 TBUF

TBUF is a RAM that holds data and status to be transmitted. The data is written into sub-blocks referred to as elements, by Microengines or the Intel XScale® core.

TBUF contains a total of 8 Kbytes of data. The element size is programmable as either 64 bytes, 128 bytes, or 256 bytes per element. In addition, TBUF can be programmed to be split into one, two, or three partitions depending on application. For transmitting SPI-4, one partition would be used. For transmitting CSIX, two partitions are used (Control CFrames and Data CFrames). For both SPI-4 and CSIX, three partitions are used.

Microengines can write data from Microengine S_TRANSFER_OUT registers to the TBUF using the `msf[write]` instruction where they specify the starting byte number (which must be aligned to 4 bytes), and number of 32-bit words to write. The number in the instruction can be either the number of 32-bit words, or number of 32-bit word pairs, using the single and double instruction modifiers, respectively.

Microengines can move data from DRAM to TBUF using the `dram` instruction where they specify the starting byte number (which must be aligned to 4 bytes), the number of 32-bit words to write, and the address in DRAM of the data.



All elements within a TBUF partition are transmitted in the order. Control information associated with the element defines which bytes are valid. The data from the TBUF will be shifted and byte aligned as required to be transmitted.

2.7.4.1.1 SPI-4 and TBUF

For SPI-4, data is put into the data portion of the element, and information for the SPI-4 Control Word that will precede the data is put into the Element Control Word.

When the Element Control Word is written, the information is:

3	3	2	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	9	8	7	6	5	4	3	2	1	0
Payload Length										Prepend Offset		Prepend Length				Payload Offset		Res	Skip	Res	SOP	EOP	ADR								
6	6	6	6	5	5	5	5	5	5	5	5	5	5	4	4	4	4	4	4	4	4	4	4	3	3	3	3	3	3	3	3
3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2
Res																															

The definitions of the fields are shown in [Table 15](#).

Table 15. TBUF SPI-4 Control Definition

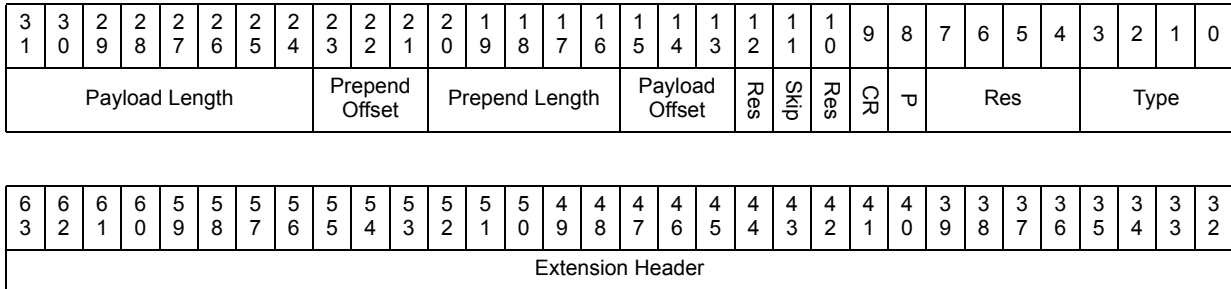
Field	Definition
Payload Length	Indicates the number of Payload bytes, from 1 to 256, in the element. The value of 0x00 means 256 bytes. The sum of Prepend Length and Payload Length will be sent. That value will also control the EOPS field (1 or 2 bytes valid indicated) of the Control Word that will succeed the data transfer. Note 1.
Prepend Offset	Indicates the first valid byte of Prepend, from 0 to 7
Prepend Length	Indicates the number of bytes in Prepend, from 0 to 31.
Payload Offset	Indicates the first valid byte of Payload, from 0 to 7.
Skip	Allows software to allocate a TBUF element and then not transmit any data from it. 0—transmit data according to other fields of Control Word. 1—free the element without transmitting any data.
SOP	Indicates if the element is the start of a packet. This field will be sent in the SOPC field of the Control Word that will precede the data transfer.
EOP	Indicates if the element is the end of a packet. This field will be sent in the EOPS field of the Control Word that will succeed the data transfer. Note 1.
ADR	The port number to which the data is directed. This field will be sent in the ADR field of the Control Word that will precede the data transfer.
NOTE: 1. Normally EOPS is sent on the next Control Word (along with ADR and SOP) to start the next element. If there is no valid element pending at the end of sending the data, the transmit logic will insert an Idle Control Word with the EOPS information.	



2.7.4.1.2 CSIX and TBUF

For CSIX, payload information is put into the data area of the element, and Base and Extension Header information is put into the Element Control Word.

When the Element Control Word is written, the information is:



The definitions of the fields are shown in Table 16.

Table 16. TBUF CSIX Control Definition

Field	Definition
Payload Length	Indicates the number of Payload bytes, from 1 to 256, in the element. The value of 0x00 means 256 bytes. The sum of Prepend Length and Payload Length will be sent, and also put into the CSIX Base Header Payload Length field. Note that this length does not include any padding that may be required. Padding is inserted by transmit hardware as needed.
Prepend Offset	Indicates the first valid byte of Prepend, from 0 to 7.
Prepend Length	Indicates the number of bytes in Prepend, from 0 to 31.
Payload Offset	Indicates the first valid byte of Payload, from 0 to 7.
Skip	Allows software to allocate a TBUF element and then not transmit any data from it. 0—transmit data according to other fields of Control Word. 1—free the element without transmitting any data.
CR	CR (CSIX Reserved) bit to put into the CSIX Base Header.
P	P (Private) bit to put into the CSIX Base Header.
Type	Type Field to put into the CSIX Base Header. Idle type is <i>not</i> legal here.
Extension Header	The Extension Header to be sent with the CFrame. The bytes are sent in big-endian order; byte 0 is in bits 63:56, byte 1 is in bits 55:48, byte 2 is in bits 47:40, and byte 3 is in bits 39:32.

2.7.4.2 Transmit Operation Summary

During transmit processing data to be transmitted is placed into the TBUF under Microengine control. The Microengine allocates an element in software; the transmit hardware processes TBUF elements within a partition in strict sequential order so the software can track which element to allocate next.

Microengines may write directly into an element by an `msf[write]` instruction, or have data from DRAM written into the element by a `dram[tbuf_wr]` instruction. Data can be merged into the element by doing both.



There is a Transmit Valid bit per element, that marks the element as ready to be transmitted. Microengines move all data into the element, by either or both of `msf[write]` and `dram[tbuf_wr]` instructions to the TBUF. Microengines also write the element Transmit Control Word with information about the element. When all of the data movement is complete, the Microengine sets the element valid bit.

1. Move data into TBUF by either or both of `msf[write]` and `dram[tbuf_wr]` instructions to the TBUF.
2. Wait for 1 to complete.
3. Write Transmit Control Word at `TBUF_ELEMENT_CONTROL_#` address. Using this address sets the Transmit Valid bit.

2.7.5 The Flow Control Interface

The MSF provides flow control support for SPI-4 and CSIX.

2.7.5.1 SPI-4

SPI-4 uses a FIFO Status Channel to provide flow control information. MSF receives the information from the PHY device and stores it so that Microengines can read the information on a per-port basis. It can then use that information to determine when to transmit data to a given port.

The MSF also sends status to the PHY based on the amount of available space in the RBUF — i.e., done by hardware without Microengines.

2.7.5.2 CSIX

CSIX provides two types of flow control — link level and per queue.

- The link level control is handled by hardware. MSF will stop transmission in response to link level flow control received from the Switch Fabric. MSF will assert link level flow control based on the amount of available space in the RBUF.
- Per queue flow control information is put into the FCIFIFO and handled by Microengine software. Also, if required, Microengines can send Flow Control CFrames to the Switch Fabric under software control.

In both cases, for a full-duplex configuration, information is passed from the Switch Fabric to the Egress IXP2800 Network Processor, which then passes it to the Ingress IXP2800 Network Processor over a proprietary flow control interface.



2.8 Hash Unit

The IXP2800 Network Processor contains a Hash Unit that can take 48-, 64-, or 128-bit data and produce a 48-, 64-, or a 128-bit hash index, respectively. The Hash Unit is accessible by the Microengines and the Intel XScale® core, and is useful in doing table searches with large keys, for example L2 addresses. Figure 14 is a block diagram of the Hash Unit.

Up to three hash indexes can be created using a single Microengine instruction. This helps to minimize command overhead. The Intel XScale® core can only do a single hash at a time.

A Microengine initiates a hash operation by writing the hash operands into a contiguous set of S_TRANSFER_OUT registers and then executing the hash instruction. The Intel XScale® core initiates a hash operation by writing a set of memory-mapped HASH_OP registers, which are built in the Intel XScale® core gasket, with the data to be used to generate the hash index. There are separate registers for 48-, 64-, and 128-bit hashes. The data is written from MSB to LSB, with the write to LSB triggering the Hash Operation. In both cases, the Hash Unit reads the operand into an input buffer, performs the hash operation, and returns the result.

The Hash Unit uses a hard-wired polynomial algorithm and a programmable hash multiplier to create hash indexes. Three separate multipliers are supported, one for 48-bit hash operations, one for 64-bit hash operations and one for 128-bit hash operations. The multiplier is programmed through Control registers in the Hash Unit.

The multiplicand is shifted into the hash array, 16 bits at a time. The hash array performs a 1's-complement multiply and polynomial divide, using the multiplier and 16 bits of the multiplicand. The result is placed into an output buffer register and also feeds back into the array. This process is repeated three times for a 48-bit hash (16 bits x 3 = 48), four times for a 64-bit hash (16 bits x 4 = 64), and eight times for a 128-bit hash (16 x 8 = 128). After the multiplicand has been passed through the hash array, the resulting hash index is placed into a two-stage output buffer.


After each hash index is completed, the Hash Unit returns the hash index to the Microengines' S_TRANSFER_IN registers, or the Intel XScale® core HASH_OP registers. For Microengine initiated hash operations, the Microengine is signaled after all the hashes specified in the instruction have been completed.

For the Intel XScale® core initiated hash operations, the Intel XScale® core reads the results from the memory-mapped HASH_OP registers. The addresses of Hash Results are the same as the HASH_OP registers. Because of queuing delays at the Hash Unit, the time to complete an operation is not fixed. The Intel XScale® core can do one of two operations to get the hash results.

- Poll the HASH_DONE register. This register is cleared when the HASH_OP registers are written. Bit [0] of HASH_DONE register is set when the HASH_OP registers get the return result from the Hash Unit (when the last word of the result is returned). The Intel XScale® core software can poll on HASH_DONE, and read HASH_OP when HASH_DONE is equal to 0x00000001.
- Read HASH_OP directly. The interface hardware will acknowledge the read only when the result is valid. This method will result in the Intel XScale® core stalling if the result is not valid when the read happens.

The number of clock cycles required to perform a single hash operation equals: two or four cycles through the input buffers, three, four or eight cycles through the hash array, and two or four cycles through the output buffers. Because of the pipeline characteristics of the Hash Unit, performance is improved if multiple hash operations are initiated with a single instruction rather than separate hash instructions for each hash operation.

Figure 14. Hash Unit Block Diagram





2.9 PCI Controller

The PCI Controller provides a 64-bit, 66 MHz capable *PCI Local Bus Revision 2.2* interface, and is compatible to 32-bit or 33 MHz PCI devices. The PCI controller provides the following functions:

- Target Access (external Bus Master access to SRAM, DRAM, and CSRs)
- Master Access (the Intel XScale® core access to PCI Target devices)
- Two DMA Channels
- Mailbox and Doorbell registers for the Intel XScale® core to Host communication
- PCI arbiter

The IXP2800 Network Processor can be configured to act as PCI central function (for use in a stand-alone system), where it provides the PCI reset signal, or as an add-in device, where it uses the PCI reset signal as the chip reset input. The choice is made by connecting the `cfg_rst_dir` input pin low or high.

2.9.1 Target Access

There are three Base Address Registers (BARs) to allow PCI Bus Masters to access SRAM, DRAM, and CSRs, respectively. Examples of PCI Bus Masters include a Host Processor (for example a Pentium® processor), or an I/O device such as an Ethernet controller, SCSI controller, or encryption coprocessor.

The SRAM BAR can be programmed to sizes of 16, 32, 64, 128, or 256 Mbytes, or no access.

The DRAM BAR can be programmed to sizes of 128, 256, or 512 Mbytes or 1 Gbyte, or no access.

The CSR BAR is 8 KB.

PCI Boot Mode is supported, in which the Host downloads the Intel XScale® core boot image into DRAM, while holding the Intel XScale® core in reset. Once the boot image has been loaded, the Intel XScale® core reset is deasserted. The alternative is to provide the boot image in a Flash ROM attached to the Slowport.

2.9.2 Master Access

The Intel XScale® core and Microengines can directly access the PCI bus. The Intel XScale® core can do loads and stores to specific address regions to generate all PCI command types. Microengines use PCI instruction, and also use address regions to generate different PCI commands.

2.9.3 DMA Channels

There are two DMA Channels, each of which can move blocks of data from DRAM to the PCI or from the PCI to DRAM. The DMA channels read parameters from a list of descriptors in SRAM, perform the data movement to or from DRAM, and stop when the list is exhausted. The descriptors are loaded from predefined SRAM entries or may be set directly by CSR writes to DMA Channel registers. There is no restriction on byte alignment of the source address or the destination address.

For PCI to DRAM transfers, the PCI command is Memory Read, Memory Read line, or Memory Read Multiple. For DRAM to PCI transfers, the PCI command is Memory Write. Memory Write Invalidate is not supported.

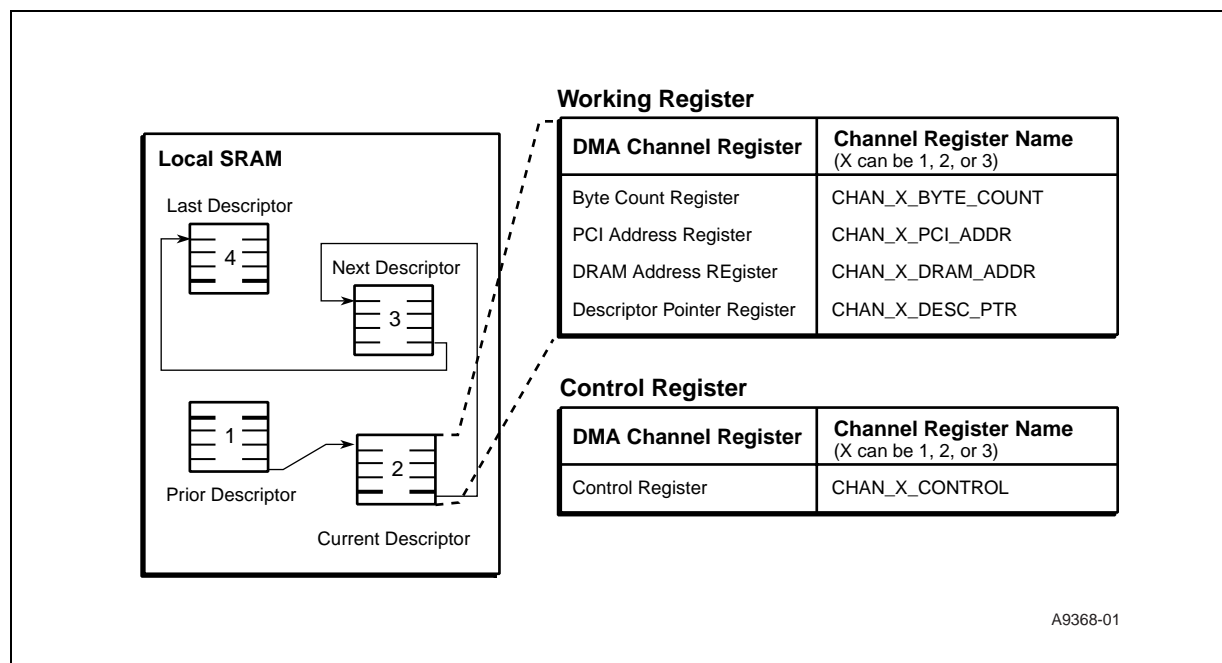
Up to two DMA channels are running at a time with three descriptors outstanding. Effectively, the active channels interleave bursts to or from the PCI Bus.

Interrupts are generated at the end of DMA operation for the Intel XScale® core. However, Microengines do not provide an interrupt mechanism. The DMA Channel will instead use an Event Signal to notify the particular Microengine on completion of DMA.

2.9.3.1 DMA Descriptor

Each descriptor uses four 32-bit words in SRAM, aligned on a 16-byte boundary. The DMA channels read the descriptors from SRAM into working registers once the control register has been set to initiate the transaction. This control must be set *explicitly*; this starts the DMA transfer. Register names for DMA channels are listed in Figure 15 and Table 17 lists the descriptor contents.

Figure 15. DMA Descriptor Reads



After a descriptor is processed, the next descriptor is loaded in the working registers. This process repeats until the chain of descriptors is terminated (i.e., the End of Chain bit is set).

Table 17. DMA Descriptor Format

Offset from Descriptor Pointer	Description
0x0	Byte Count
0x4	PCI Address
0x8	DRAM Address
0xC	Next Descriptor Address



2.9.3.2 DMA Channel Operation

The DMA channel can be set up to read the first descriptor in SRAM, or with the first descriptor written directly to the DMA channel registers. When descriptors and the descriptor list are in SRAM, the procedure is as follows:

1. The DMA channel owner writes the address of the first descriptor into the DMA Channel Descriptor Pointer register (DESC_PTR).
2. The DMA channel owner writes the DMA Channel Control register (CONTROL) with miscellaneous control information and also sets the channel enable bit (bit 0). The channel initial descriptor bit (bit 4) in the CONTROL register must also be cleared to indicate that the first descriptor is in SRAM.
3. Depending on the DMA channel number, the DMA channel reads the descriptor block into the corresponding DMA registers, BYTE_COUNT, PCI_ADDR, DRAM_ADDR, and DESC_PTR.
4. The DMA channel transfers the data until the byte count is exhausted, and then sets the channel transfer done bit in the CONTROL register.
5. If the end of chain bit (bit 31) in the BYTE_COUNT register is clear, the channel checks the Chain Pointer value. If the Chain Pointer value is not equal to 0, it reads the next descriptor and transfers the data (step 3 and 4 above). If the Chain Pointer value is equal to 0, it waits for the Descriptor Added bit of the Channel Control register to be set before reading the next descriptor and transfers the data (step 3 and 4 above). If bit 31 is set, the channel sets the channel chain done bit in the CONTROL register and then stops.
6. Proceed to the Channel End Operation.

When single descriptors are written into the DMA channel registers, the procedure is as follows:

1. The DMA channel owner writes the descriptor values directly into the DMA channel registers. The end of chain bit (bit 31) in the BYTE_COUNT register must be set, and the value in the DESC_PTR register is not used.
2. The DMA channel owner writes the base address of the DMA transfer into the PCI_ADDR to specify the PCI starting address.
3. When the first descriptor is in the BYTE_COUNT register, the DRAM_ADDR register must be written with the address of the data to be moved.
4. The DMA channel owner writes the CONTROL register with miscellaneous control information, along with setting the channel enable bit (bit 0). The channel initial descriptor in register bit (bit 4) in the CONTROL register must also be set to indicate that the first descriptor is already in the channel descriptor registers.
5. The DMA channel transfers the data until the byte count is exhausted, and then sets the channel transfer done bit (bit 2) in the CONTROL register.
6. Since the end of the chain bit (bit 31) in the BYTE_COUNT register is set, the channel sets the channel chain done bit (bit 7) in the CONTROL register and then stops.
7. Proceed to the Channel End Operation.



2.9.3.3 DMA Channel End Operation

1. Channel owned by PCI:
If not masked via the PCI Outbound Interrupt Mask register, the DMA channel interrupts the PCI host after the setting of the DMA done bit in the CHAN_X_CONTROL register, which is readable in the PCI Outbound Interrupt Status register.
2. Channel owned by the Intel XScale® core:
If enabled via the Intel XScale® core Interrupt Enable registers, the DMA channel interrupts the Intel XScale® core by setting the DMA channel done bit in the CHAN_X_CONTROL register, which is readable in the Intel XScale® core Interrupt Status register.
3. Channel owned by Microengine:
If enabled via the Microengine Auto-Push Enable registers, the DMA channel signals the Microengine after setting the DMA channel done bit in the CHAN_X_CONTROL register, which is readable in the Microengine Auto-Push Status register.

2.9.3.4 Adding Descriptors to an Unterminated Chain

It is possible to add a descriptor to a chain while a channel is running. To do so, the chain should be left unterminated, i.e., the last descriptor should have End of Chain clear, and the Chain Pointer value equal to 0. A new descriptor (or linked list of descriptors) can be added to the chain by overwriting the Chain Pointer value of the unterminated descriptor (in SRAM) with the Local Memory address of the (first) added descriptor (the added descriptor must actually be valid in Local Memory prior to that). After updating the Chain Pointer field, the software must write a 1 to the Descriptor Added bit of the Channel Control register. This is necessary for the case where the channel was paused to reactivate the channel. However, software need not check the state of the channel before writing that bit; there is no side-effect of writing that bit in the case where the channel had not yet read the unlinked descriptor.

If the channel was paused or had read an unlinked Pointer, it will re-read the last descriptor processed (i.e., the one that originally had the 0 value for Chain Pointer) to get the address of the newly added descriptor.

A descriptor cannot be added to a descriptor that has End of Chain set.

2.9.4 Mailbox and Message Registers

Mailbox and Doorbell registers provide hardware support for communication between the Intel XScale® core and a device on the PCI Bus.

Four 32-bit mailbox registers are provided so that messages can be passed between the Intel XScale® core and a PCI device. All four registers can be read and written with byte resolution from both the Intel XScale® core and PCI. How the registers are used is application dependent and the messages are not used internally by the PCI Unit in any way. The mailbox registers are often used with the Doorbell interrupts.

Doorbell interrupts provide an efficient method of generating an interrupt as well as encoding the purpose of the interrupt. The PCI Unit supports a 32-bit the Intel XScale® core DOORBELL register that is used by a PCI device to generate an the Intel XScale® core interrupt, and a separate 32-bit PCI DOORBELL register that is used by the Intel XScale® core to generate a PCI interrupt. A source generating the Doorbell interrupt can write a software defined bitmap to the register to indicate a specific purpose. This bitmap is translated into a single interrupt signal to the destination



(either a PCI interrupt or an Intel XScale® core interrupt). When an interrupt is received, the DOORBELL registers can be read and the bit mask can be interpreted. If a larger bit mask is required than that is provided by the DOORBELL register, the MAILBOX registers can be used to pass up to 16 bytes of data.

The doorbell interrupts are controlled through the registers shown in [Table 18](#).

Table 18. Doorbell Interrupt Registers

Register Name	Description
XSCALE DOORBELL	Used to generate the Intel XScale® core Doorbell interrupts.
XSCALE DOORBELL SETUP	Used to initialize the Intel XScale® core Doorbell register and for diagnostics.
PCI DOORBELL	Used to generate the PCI Doorbell interrupts.
PCI DOORBELL SETUP	Used to initialize the PCI Doorbell register and for diagnostics.

2.9.5 PCI Arbiter

The PCI unit contains a PCI bus arbiter that supports two external masters in addition to the PCI Unit's initiator interface. If more than two external masters are used in the system, the arbiter can be disabled and an external (to the IXP2800 Network Processor) used. In that case, the IXP2800 Network Processor will provide its PCI request signal to the external arbiter, and use that arbiter's grant signal.

The arbiter uses a simple round-robin priority algorithm; it asserts the grant signal corresponding to the next request in the round-robin during the current executing transaction on the PCI bus (this is also called hidden arbitration). If the arbiter detects that an initiator has failed to assert frame 1 after 16 cycles of both grant assertion and PCI bus idle condition, the arbiter deasserts the grant. That master does not receive any more grants until it deasserts its request for at least one PCI clock cycle. Bus parking is implemented in that the last bus grant will stay asserted if no request is pending.

To prevent bus contention, if the PCI bus is idle, the arbiter never asserts one grant signal in the same PCI cycle in which it deasserts another. It deasserts one grant, and then asserts the next grant after one full PCI clock cycle has elapsed to provide for bus driver turnaround.



2.10 Control and Status Register Access Proxy

The Control and Status Register Access Proxy (CAP) contains a number of chip-wide control and status registers. Some provide miscellaneous control and status, while others are used for inter-Microengine or Microengine to the Intel XScale® core communication (note that rings in Scratchpad Memory and SRAM can also be used for inter-process communication). These include:

- **INTERTHREAD SIGNAL** — Each thread (or context) on a Microengine can send a signal to any other thread by writing to InterThread_Signal register. This allows a thread to go to sleep waiting completion of a task by a different thread.
- **THREAD MESSAGE** — Each thread has a message register where it can post a software-specific message. Other Microengine threads, or the Intel XScale® core, can poll for availability of messages by reading the THREAD_MESSAGE_SUMMARY register. Both the THREAD_MESSAGE and corresponding THREAD_MESSAGE_SUMMARY clear upon a read of the message; this eliminates a race condition when there are multiple message readers. Only one reader will get the message.
- **SELF DESTRUCT** — This register provides another type of communication. Microengine software can atomically set individual bits in the SELF_DESTRUCT registers; the registers clear upon read. The meaning of each bit is software-specific. Clearing the register upon read eliminates a race condition when there are multiple readers.
- **THREAD INTERRUPT** — Each thread can interrupt the Intel XScale® core on two different interrupts; the usage is software-specific. Having two interrupts allows for flexibility, for example, one can be assigned to normal service requests and one can be assigned to error conditions. If more information needs to be associated with the interrupt, mailboxes or Rings in Scratchpad Memory or SRAM could be used.
- **REFLECTOR** — CAP provides a function (called “reflector”) where any Microengine thread can move data between its registers and those of any other thread. In response to a single write or read instruction (with the address in the specific reflector range) CAP will get data from the source Microengine and put it into the destination Microengine. Both the sending and receiving threads can optionally be signaled upon completion of the data movement.

2.11 Intel XScale® Core Peripherals

2.11.1 Interrupt Controller

The Interrupt Controller provides the ability to enable or mask interrupts from a number of chip wide sources, for example:

- Timers (normally used by Real-Time Operating System).
- Interrupts generated by Microengine software to request services from the Intel XScale® core.
- External agents such as PCI devices.
- Error conditions, such as DRAM ECC error, or SPI-4 parity error.

Interrupt status is read as memory mapped registers; the state of an interrupt signal can be read even if it is masked from interrupting. Enabling and masking of interrupts is done as writes to memory mapped registers.



2.11.2 Timers

The IXP2800 Network Processor contains four programmable 32-bit timers, which can be used for software support. Each timer can be clocked by the internal clock, by a divided version of the clock, or by a signal on an external GPIO pin. Each timer can be programmed to generate a periodic interrupt after a programmed number of clocks. The range is from several ns to several minutes depending on the clock frequency.

In addition, timer 4 can be used as a watchdog timer. In this use, software must periodically reload the timer value; if it fails to do so and the timer counts to 0, it will reset the chip. This can be used to detect if software “hangs” or for some other reason fails to reload the timer.

2.11.3 General Purpose I/O

The IXP2800 Network Processor contains eight General Purpose I/O (GPIO) pins. These can be programmed as either input or output and can be used for slow speed I/O such as LEDs or input switches. They can also be used as interrupts to the Intel XScale® core, or to clock the programmable timers.

2.11.4 Universal Asynchronous Receiver/Transmitter

The IXP2800 Network Processor contains a standard RS-232 compatible Universal Asynchronous Receiver/Transmitter (UART), which can be used for communication with a debugger or maintenance console. Modem controls are not supported; if they are needed, GPIO pins can be used for that purpose.

The UART performs serial-to-parallel conversion on data characters received from a peripheral device and parallel-to-serial conversion on data characters received from the processor. The processor can read the complete status of the UART at any time during operation. Available status information includes the type and condition of the transfer operations being performed by the UART and any error conditions (parity, overrun, framing or break interrupt).

The serial ports can operate in either FIFO or non-FIFO mode. In FIFO mode, a 64-byte transmit FIFO holds data from the processor to be transmitted on the serial link and a 64-byte receive FIFO buffers data from the serial link until read by the processor.

The UART includes a programmable baud rate generator that is capable of dividing the internal clock input by divisors of 1 to $2^{16} - 1$ and produces a 16X clock to drive the internal transmitter logic. It also drives the receive logic. The UART can be operated in polled or in interrupt driven mode as selected by software.

2.11.5 Slowport

The Slowport is an external interface to the IXP2800 Network Processor, used for Flash ROM access and 8, 16, or 32-bit asynchronous device access. It allows the Intel XScale® core do read/write data transfers to these slave devices.

The address bus and data bus are multiplexed to reduce the pin count. In addition, 24 bits of address are shifted out on three clock cycles. Therefore, an external set of buffers is needed to latch the address. Two chip selects are provided.



The access is asynchronous. Insertion of delay cycles for both data setup and hold time is programmable via internal Control registers. The transfer can also wait for a handshake acknowledge signal from the external device.

2.12 I/O Latency

Table 19 shows the latencies for transferring data between the Microengine and the other sub-system components. The latency is measured in 1.4 GHz cycles.

Table 19. I/O Latency

	Sub-system			
	DRAM (RDR)	SRAM (QDR)	Scratch	MSF
Transfer Size	8 bytes – 16 bytes (note 2)	4 bytes	4 bytes	8 bytes
Average Read Latency	~ 295 cycles (note 3)	100 (light load) – 160 (heavy load)	~ 100 cycles (range 53 – 152)	range 53 – 120 (RBUF)
Average Write Latency	~ 53 cycles	~ 53 cycles	~ 40 cycles	~ 48 cycles (TBUF)
Note1: RDR, QDR, MSF, and Scratch values are extracted from a simulation model. Note 2: Minimum DRAM burst size on pins is 16 bytes. Transfers less than 16 bytes incur the same as a 16-byte transfer. Note 3: At 1016 MHz, read latency should be ~ 240 cycles.				

2.13 Performance Monitor

The Intel XScale® core hardware provides two 32-bit performance counters that allow two unique events to be monitored simultaneously. In addition, the Intel XScale® core implements a 32-bit clock counter that can be used in conjunction with the performance counters; its sole purpose is to count the number of core clock cycles, which is useful in measuring total execution time.

This section contains information describing the Intel XScale® core, Intel XScale® core gasket, and Intel XScale® core Peripherals (XPI).

For additional information about the Intel XScale® architecture refer to the *Intel XScale® Core Developers Manual* available on Intel's Developers web site (<http://www.developer.intel.com>).

3.1 Introduction

The Intel XScale® core is an ARM* V5TE compliant microprocessor. It has been designed for high performance and low-power; leading the industry in mW/MIPs. The Intel XScale® core incorporates an extensive list of architecture features that allows it to achieve high performance. Many of the architectural features added to the Intel XScale® core help hide memory latency that often is a serious impediment to high performance processors.

This includes:

- The ability to continue instruction execution even while the data cache is retrieving data from external memory.
- A write buffer.
- Write-back caching.
- Various data cache allocation policies that can be configured different for each application.
- Cache locking.

All these features improve the efficiency of the memory bus external to the core.

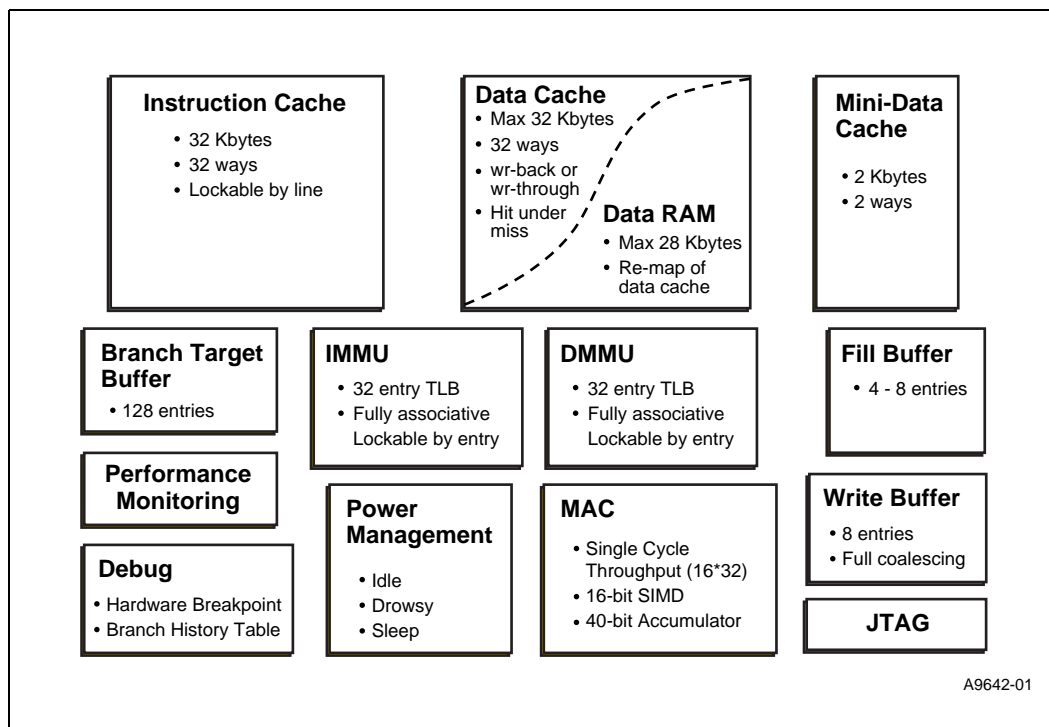
ARM* Version 5 (V5) Architecture added floating point instructions to ARM* Version 4. The Intel XScale® core implements the integer instruction set architecture of ARM* V5, but does not provide hardware support of the floating point instructions.

The Intel XScale® core provides the Thumb instruction set (ARM* V5T) and the ARM* V5E DSP extensions.

3.2 Features

Figure 16 shows the major functional blocks of the Intel XScale® core.

Figure 16. Intel XScale® Core Architecture Features



3.2.1 Multiply/ACcumulate (MAC)

The MAC unit supports early termination of multiplies/accumulates in two cycles and can sustain a throughput of a MAC operation every cycle. Architectural enhancements to the MAC support audio coding algorithms, including a 40-bit accumulator and support for 16-bit packed data.

3.2.2 Memory Management

The Intel XScale® core implements the Memory Management Unit (MMU) Architecture specified in the *ARM* Architecture Reference Manual* (see the ARM* website at <http://www.arm.com>). The MMU provides access protection and virtual to physical address translation. The MMU Architecture also specifies the caching policies for the instruction cache and data memory. These policies are specified as page attributes and include:

- identifying code as cacheable or non-cacheable
- selecting between the mini-data cache or data cache
- write-back or write-through data caching
- enabling data write allocation policy
- and enabling the write buffer to coalesce stores to external memory



3.2.3 Instruction Cache

The Intel XScale® core implements a 32-Kbyte, 32-way set associative instruction cache with a line size of 32 bytes. All requests that “miss” the instruction cache generate a 32-byte read request to external memory. A mechanism to lock critical code within the cache is also provided.

3.2.4 Branch Target Buffer (BTB)

The Intel XScale® core provides a Branch Target Buffer to predict the outcome of branch type instructions. It provides storage for the target address of branch type instructions and predicts the next address to present to the instruction cache when the current instruction address is that of a branch.

The BTB holds 128 entries.

3.2.5 Data Cache

The Intel XScale® core implements a 32-Kbyte, a 32-way set associative data cache and a 2-Kbyte, 2-way set associative mini-data cache. Each cache has a line size of 32 bytes, and supports write-through or write-back caching.

The data/mini-data cache is controlled by page attributes defined in the MMU Architecture and by coprocessor 15. The Intel XScale® core allows applications to reconfigure a portion of the data cache as data RAM. Software may place special tables or frequently used variables in this RAM.

3.2.6 Performance Monitoring

Two performance monitoring counters have been added to the Intel XScale® core that can be configured to monitor various events. These events allow a software developer to measure cache efficiency, detect system bottlenecks, and reduce the overall latency of programs.

3.2.7 Power Management

The Intel XScale® core incorporates a power and clock management unit that can assist in controlling clocking and managing power.

3.2.8 Debugging

The Intel XScale® core supports software debugging through two instruction address breakpoint registers, one data-address breakpoint register, one data-address/mask breakpoint register, and a trace buffer.

3.2.9 JTAG

Testability is supported on the Intel XScale® core through the Test Access Port (TAP) Controller implementation, which is based on IEEE 1149.1 (JTAG) Standard Test Access Port and Boundary-Scan Architecture. The purpose of the TAP controller is to support test logic internal and external to the Intel XScale® core such as built-in self-test, boundary-scan, and scan.



3.3 Memory Management

The Intel XScale® core implements the Memory Management Unit (MMU) Architecture specified in the *ARM Architecture Reference Manual*. To accelerate virtual to physical address translation, the Intel XScale® core uses both an instruction Translation Look-aside Buffer (TLB) and a data TLB to cache the latest translations. Each TLB holds 32 entries and is fully-associative. Not only do the TLBs contain the translated addresses, but also the access rights for memory references.

If an instruction or data TLB miss occurs, a hardware translation-table-walking mechanism is invoked to translate the virtual address to a physical address. Once translated, the physical address is placed in the TLB along with the access rights and attributes of the page or section. These translations can also be locked down in either TLB to guarantee the performance of critical routines.

The Intel XScale® core allows system software to associate various attributes with regions of memory:

- cacheable
- bufferable
- line allocate policy
- write policy
- I/O
- mini Data Cache
- Coalescing
- P bit

Note: The virtual address with which the TLBs are accessed may be remapped by the PID register.

3.3.1 Architecture Model

3.3.1.1 Version 4 versus Version 5

ARM* MMU Version 5 Architecture introduces the support of tiny pages, which are 1 Kbyte in size. The reserved field in the first-level descriptor (encoding 0b11) is used as the fine page table base address.

3.3.1.2 Memory Attributes

The attributes associated with a particular region of memory are configured in the memory management page table and control the behavior of accesses to the instruction cache, data cache, mini-data cache and the write buffer. These attributes are ignored when the MMU is disabled.

To allow compatibility with older system software, the new Intel XScale® core attributes take advantage of encoding space in the descriptors that was formerly reserved.

3.3.1.2.1 Page (P) Attribute Bit

The P bit assigns a page attribute to a memory region. Refer to the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual* for details about the P bit.



3.3.1.2.2 Instruction Cache

When examining these bits in a descriptor, the Instruction Cache only utilizes the C bit. If the C bit is clear, the Instruction Cache considers a code fetch from that memory to be non-cacheable, and will not fill a cache entry. If the C bit is set, then fetches from the associated memory region will be cached.

3.3.1.2.3 Data Cache and Write Buffer

All of these descriptor bits affect the behavior of the Data Cache and the Write Buffer.

If the X bit for a descriptor is 0 (see Table 20), the C and B bits operate as mandated by the ARM* architecture. If the X bit for a descriptor is one, the C and B bits' meaning is extended, as detailed in Table 21.

Table 20. Data Cache and Buffer Behavior when X = 0

C B	Cacheable?	Bufferable?	Write Policy	Line Allocation Policy	Notes
0 0	N	N	—	—	Stall until complete ¹
0 1	N	Y	—	—	
1 0	Y	Y	Write Through	Read Allocate	
1 1	Y	Y	Write Back	Read Allocate	

1. Normally, the processor will continue executing after a data access if no dependency on that access is encountered. With this setting, the processor will stall execution until the data access completes. This guarantees to software that the data access has taken effect by the time execution of the data access instruction completes. External data aborts from such accesses will be imprecise.

Table 21. Data Cache and Buffer Behavior when X = 1

C B	Cacheable?	Bufferable?	Write Policy	Line Allocation Policy	Notes
0 0	—	—	—	—	Unpredictable; do not use
0 1	N	Y	—	—	Writes will not coalesce into buffers ¹
1 0	(Mini Data Cache)	—	—	—	Cache policy is determined by MD field of Auxiliary Control register
1 1	Y	Y	Write Back	Read/Write Allocate	

1. Normally, bufferable writes can coalesce with previously buffered data in the same address range

3.3.1.2.4 Details on Data Cache and Write Buffer Behavior

If the MMU is disabled all data accesses will be non-cacheable and non-bufferable. This is the same behavior as when the MMU is enabled, and a data access uses a descriptor with X, C, and B all set to 0.

The X, C, and B bits determine when the processor should place new data into the Data Cache. The cache places data into the cache in lines (also called blocks). Thus, the basis for making a decision about placing new data into the cache is called a “Line Allocation Policy.”



If the Line Allocation Policy is read-allocate, all load operations that miss the cache request a 32-byte cache line from external memory and allocate it into either the data cache or mini-data cache (this is assuming the cache is enabled). Store operations that miss the cache will not cause a line to be allocated.

If read/write-allocate is in effect, load *or* store operations that miss the cache will request a 32-byte cache line from external memory if the cache is enabled.

The other policy determined by the X, C, and B bits is the Write Policy. A write-through policy instructs the Data Cache to keep external memory coherent by performing stores to both external memory and the cache. A write-back policy only updates external memory when a line in the cache is cleaned or needs to be replaced with a new line. Generally, write-back provides higher performance because it generates less data traffic to external memory.

3.3.1.2.5 Memory Operation Ordering

A *fence* memory operation (memop) is one that guarantees all memops issued prior to the fence will execute before any memop issued after the fence. Thus software may issue a fence to impose a partial ordering on memory accesses.

Table 22 shows the circumstances in which memops act as fences.

Any swap (SWP or SWPB) to a page that would create a fence on a load or store is a fence.

Table 22. Memory Operations that Impose a Fence

operation	X	C	B
load	—	0	—
store	1	0	1
load or store	0	0	0

3.3.2 Exceptions

The MMU may generate prefetch aborts for instruction accesses and data aborts for data memory accesses.

Data address alignment checking is enabled by setting bit 1 of the Control register (CP15, register 1). Alignment faults are still reported even if the MMU is disabled. All other MMU exceptions are disabled when the MMU is disabled.



3.3.3 Interaction of the MMU, Instruction Cache, and Data Cache

The MMU, instruction cache, and data/mini-data cache may be enabled/disabled independently. The instruction cache can be enabled with the MMU enabled or disabled. However, the data cache can only be enabled when the MMU is enabled. Therefore only three of the four combinations of the MMU and data/mini-data cache enables are valid (see [Table 23](#)). The invalid combination will cause undefined results.

Table 23. Valid MMU and Data/Mini-Data Cache Combinations

MMU	Data/Mini-data Cache
Off	Off
On	Off
On	On

3.3.4 Control

3.3.4.1 Invalidate (Flush) Operation

The entire instruction and data TLB can be invalidated at the same time with one command or they can be invalidated separately. An individual entry in the data or instruction TLB can also be invalidated.

Globally invalidating a TLB will not affect locked TLB entries. However, the invalidate-entry operations can invalidate individual locked entries. In this case, the locked remains in the TLB, but will never “hit” on an address translation. Effectively, a hole is in the TLB. This situation may be rectified by unlocking the TLB.

3.3.4.2 Enabling/Disabling

The MMU is enabled by setting bit 0 in coprocessor 15, register 1 (Control register). When the MMU is disabled, accesses to the instruction cache default to cacheable and all accesses to data memory are made non-cacheable. A recommended code sequence for enabling the MMU is shown in [Example 14](#).

Example 14. Enabling the MMU

```

; This routine provides software with a predictable way of enabling the MMU.
; After the CPWAIT, the MMU is guaranteed to be enabled. Be aware
; that the MMU will be enabled sometime after MCR and before the instruction
; that executes after the CPWAIT.
; Programming Note: This code sequence requires a one-to-one virtual to
; physical address mapping on this code since
; the MMU may be enabled part way through. This would allow the instructions
; after MCR to execute properly regardless the state of the MMU.

MRC P15,0,R0,C1,C0,0; Read CP15, register 1
ORR R0, R0, #0x1; Turn on the MMU
MCR P15,0,R0,C1,C0,0; Write to CP15, register 1

; The MMU is guaranteed to be enabled at this point; the next instruction or
; data address will be translated.

```



3.3.4.3 Locking Entries

Individual entries can be locked into the instruction and data TLBs. If a lock operation finds the virtual address translation already resident in the TLB, the results are unpredictable. An invalidate by entry command before the lock command will ensure proper operation. Software can also accomplish this by invalidating all entries, as shown in [Example 15](#).

Locking entries into either the instruction TLB or data TLB reduces the available number of entries (by the number that was locked down) for hardware to cache other virtual to physical address translations.

A procedure for locking entries into the instruction TLB is shown in [Example 15](#).

If a MMU abort is generated during an instruction or data TLB lock operation, the Fault Status register is updated to indicate a Lock Abort, and the exception is reported as a data abort.

Example 15. Locking Entries into the Instruction TLB

```
; R1, R2 and R3 contain the virtual addresses to translate and lock into
; the instruction TLB.

; The value in R0 is ignored in the following instruction.
; Hardware guarantees that accesses to CP15 occur in program order

MCR P15,0,R0,C8,C5,0 ; Invalidate the entire instruction TLB

MCR P15,0,R1,C10,C4,0 ; Translate virtual address (R1) and lock into
; instruction TLB
MCR P15,0,R2,C10,C4,0 ; Translate
; virtual address (R2) and lock into instruction TLB
MCR P15,0,R3,C10,C4,0 ; Translate virtual address (R3) and lock into
; instruction TLB

CPWAIT

; The MMU is guaranteed to be updated at this point; the next instruction will
; see the locked instruction TLB entries.
```

Note: If exceptions are allowed to occur in the middle of this routine, the TLB may end up caching a translation that is about to be locked. For example, if R1 is the virtual address of an interrupt service routine and that interrupt occurs immediately after the TLB has been invalidated, the lock operation will be ignored when the interrupt service routine returns back to this code sequence. Software should disable interrupts (FIQ or IRQ) in this case.

As a general rule, software should avoid locking in all other exception types.



The proper procedure for locking entries into the data TLB is shown in [Example 16](#).

Example 16. Locking Entries into the Data TLB

```
; R1, and R2 contain the virtual addresses to translate and lock into the data TLB

MCR P15,0,R1,C8,C6,1      ; Invalidate the data TLB entry specified by the
                          ; virtual address in R1
MCR P15,0,R1,C10,C8,0    ; Translate virtual address (R1) and lock into
                          ; data TLB

; Repeat sequence for virtual address in R2
MCR P15,0,R2,C8,C6,1      ; Invalidate the data TLB entry specified by the
                          ; virtual address in R2
MCR P15,0,R2,C10,C8,0    ; Translate virtual address (R2) and lock into
                          ; data TLB

CPWAIT                    ; wait for locks to complete

; The MMU is guaranteed to be updated at this point; the next instruction will
; see the locked data TLB entries.
```

Note: Care must be exercised here when allowing exceptions to occur during this routine whose handlers may have data that lies in a page that is trying to be locked into the TLB.

3.3.4.4 Round-Robin Replacement Algorithm

The line replacement algorithm for the TLBs is round-robin; there is a round-robin pointer that keeps track of the next entry to replace. The next entry to replace is the one sequentially after the last entry that was written. For example, if the last virtual to physical address translation was written into entry 5, the next entry to replace is entry 6.

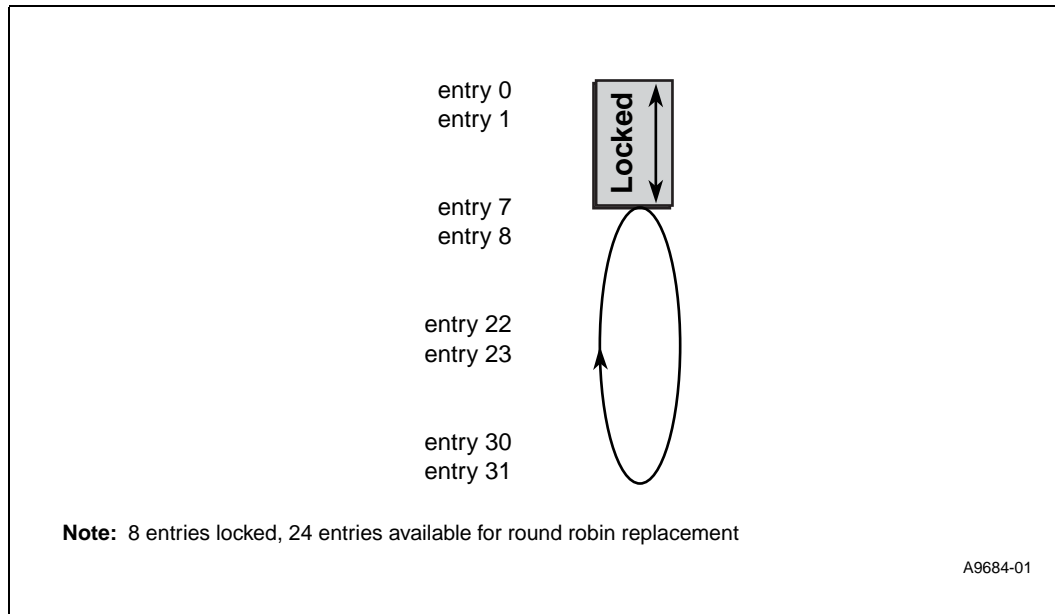
At reset, the round-robin pointer is set to entry 31. Once a translation is written into entry 31, the round-robin pointer gets set to the next available entry, beginning with entry 0 if no entries have been locked down. Subsequent translations move the round-robin pointer to the next sequential entry until entry 31 is reached, where it will wrap back to entry 0 upon the next translation.

A lock pointer is used for locking entries into the TLB and is set to entry 0 at reset. A TLB lock operation places the specified translation at the entry designated by the lock pointer, moves the lock pointer to the next sequential entry, and resets the round-robin pointer to entry 31. Locking entries into either TLB effectively reduces the available entries for updating. For example, if the first three entries were locked down, the round-robin pointer would be entry 3 after it rolled over from entry 31.

Only entries 0 through 30 can be locked in either TLB; entry 31 can never be locked. If the lock pointer is at entry 31, a lock operation will update the TLB entry with the translation and ignore the lock. In this case, the round-robin pointer will stay at entry 31.

Figure 17 illustrates locked entries in TLB.

Figure 17. Example of Locked Entries in TLB



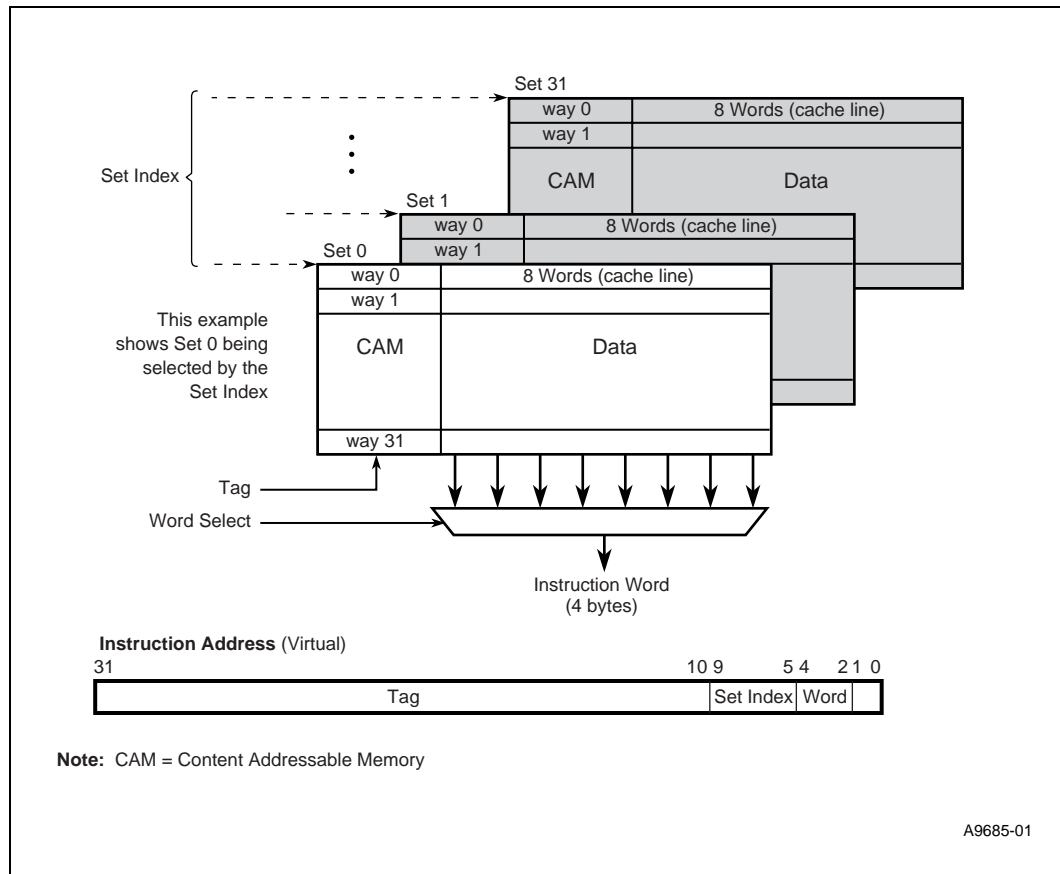
3.4 Instruction Cache

The Intel XScale® core instruction cache enhances performance by reducing the number of instruction fetches from external memory. The cache provides fast execution of cached code. Code can also be locked down when guaranteed or fast access time is required.

Figure 18 shows the cache organization and how the instruction address is used to access the cache.

The instruction cache is a 32-Kbyte, 32-way set associative cache; this means there are 32 sets with each set containing 32 ways. Each way of a set contains eight 32-bit words and one valid bit, which is referred to as a *line*. The replacement policy is a round-robin algorithm and the cache also supports the ability to lock code in at a line granularity.

Figure 18. Instruction Cache Organization



The instruction cache is virtually addressed and virtually tagged. The virtual address presented to the instruction cache may be remapped by the PID register.

3.4.1 Instruction Cache Operation

3.4.1.1 Operation when Instruction Cache is Enabled

When the cache is enabled, it compares every instruction request address to the addresses of instructions that it is holding in cache. If the requested instruction is found, the access “hits” the cache, which returns the requested instruction. If the instruction is not found, the access “misses” the cache, which requests a fetch from external memory of the 8-word line (32 bytes) that contains the instruction (using the fetch policy). As the fetch returns instructions to the cache, they are put in one of two fetch buffers and the requested instruction is delivered to the instruction decoder. A fetched line is written into the cache if it is cacheable (code is cacheable if the MMU is disabled or if the MMU is enabled and the cacheable (C) bit is set to 1 in its corresponding page).

Note: An instruction fetch may “miss” the cache but “hit” one of the fetch buffers. If this happens, the requested instruction is delivered to the instruction decoder in the same manner as a cache “hit.”



3.4.1.2 Operation when Instruction Cache is Disabled

Disabling the cache prevents any lines from being written into the instruction cache. Although the cache is disabled, it is still accessed and may generate a “hit” if the data is already in the cache.

Disabling the instruction cache *does not* disable instruction buffering that may occur within the instruction fetch buffers. Two 8-word instruction fetch buffers will always be enabled in the cache disabled mode. As instruction fetches continue to “hit” within either buffer (even in the presence of forward and backward branches), no external fetches for instructions are generated. A miss causes one or the other buffer to be filled from external memory using the fill policy.

3.4.1.3 Fetch Policy

An instruction-cache “miss” occurs when the requested instruction is not found in the instruction fetch buffers or instruction cache; a fetch request is then made to external memory. The instruction cache can handle up to two “misses.” Each external fetch request uses a fetch buffer that holds 32-bytes and eight valid bits, one for each word. A miss causes the following:

1. A fetch buffer is allocated.
2. The instruction cache sends a fetch request to the external bus. This request is for a 32-byte line.
3. Instructions words are returned back from the external bus, at a maximum rate of 1 word per core cycle. As each word returns, the corresponding valid bit is set for the word in the fetch buffer.
4. As soon as the fetch buffer receives the requested instruction, it forwards the instruction to the instruction decoder for execution.
5. When all words have returned, the fetched line will be written into the instruction cache if it is cacheable and if the instruction cache is enabled. The line chosen for update in the cache is controlled by the round-robin replacement algorithm. This update may evict a valid line at that location.
6. Once the cache is updated, the eight valid bits of the fetch buffer are invalidated.

3.4.1.4 Round-Robin Replacement Algorithm

The line replacement algorithm for the instruction cache is round-robin. Each set in the instruction cache has a round-robin pointer that keeps track of the next line (in that set) to replace. The next line to replace in a set is the one after the last line that was written. For example, if the line for the last external instruction fetch was written into way 5-set 2, the next line to replace for that set would be way 6. None of the other round-robin pointers for the other sets are affected in this case.

After reset, way 31 is pointed to by the round-robin pointer for all the sets. Once a line is written into way 31, the round-robin pointer points to the first available way of a set, beginning with way 0 if no lines have been locked into that particular set. Locking lines into the instruction cache effectively reduces the available lines for cache updating. For example, if the first three lines of a set were locked down, the round-robin pointer would point to the line at way 3 after it rolled over from way 31.



3.4.1.5 Parity Protection

The instruction cache is protected by parity to ensure data integrity. Each instruction cache word has 1 parity bit. (The instruction cache tag is *not* parity protected.) When a parity error is detected on an instruction cache access, a prefetch abort exception occurs if the Intel XScale® core attempts to execute the instruction. Before servicing the exception, hardware place a notification of the error in the Fault Status register (Coprocessor 15, register 5).

A software exception handler can recover from an instruction cache parity error. This can be accomplished by invalidating the instruction cache and the branch target buffer and then returning to the instruction that caused the prefetch abort exception. A simplified code example is shown in [Example 17](#). A more complex handler might choose to invalidate the specific line that caused the exception and then invalidate the BTB.

Example 17. Recovering from an Instruction Cache Parity Error

```

; Prefetch abort handler
MCR P15,0,R0,C7,C5,0 ; Invalidate the instruction cache and branch target
                        ; buffer

CPWAIT                ; wait for effect
                        ;

SUBS PC,R14,#4        ; Returns to the instruction that generated the
                        ; parity error

; The Instruction Cache is guaranteed to be invalidated at this point

```

If a parity error occurs on an instruction that is locked in the cache, the software exception handler needs to unlock the instruction cache, invalidate the cache and then re-lock the code in before it returns to the faulting instruction.

3.4.1.6 Instruction Cache Coherency

The instruction cache does not detect modification to program memory by loads, stores or actions of other bus masters. Several situations may require program memory modification, such as uploading code from disk.

The application program is responsible for synchronizing code modification and invalidating the cache. In general, software must ensure that modified code space is not accessed until modification and invalidating are completed.

To achieve cache coherence, instruction cache contents can be invalidated after code modification in external memory is complete.

If the instruction cache is not enabled, or code is being written to a non-cacheable region, software must still invalidate the instruction cache before using the newly-written code. This precaution ensures that state associated with the new code is not buffered elsewhere in the processor, such as the fetch buffers or the BTB.

Naturally, when writing code as data, care must be taken to force it completely out of the processor into external memory before attempting to execute it. If writing into a non-cacheable region, flushing the write buffers is sufficient precaution. If writing to a cacheable region, then the data cache should be submitted to a Clean/Invalidate operation to ensure coherency.



3.4.2 Instruction Cache Control

3.4.2.1 Instruction Cache State at Reset

After reset, the instruction cache is always disabled, unlocked, and invalidated (flushed).

3.4.2.2 Enabling/Disabling

The instruction cache is enabled by setting bit 12 in coprocessor 15, register 1 (Control register). This process is illustrated in [Example 18](#).

Example 18. Enabling the Instruction Cache

```
; Enable the ICache
MRC P15, 0, R0, C1, C0, 0      ; Get the control register
ORR R0, R0, #0x1000           ; set bit 12 -- the I bit
MCR P15, 0, R0, C1, C0, 0      ; Set the control register

CPWAIT
```

3.4.2.3 Invalidating the Instruction Cache

The entire instruction cache along with the fetch buffers are invalidated by writing to coprocessor 15, register 7. This command does not unlock any lines that were locked in the instruction cache nor does it invalidate those locked lines. To invalidate the entire cache including locked lines, the unlock instruction cache command needs to be executed before the invalidate command.

There is an inherent delay from the execution of the instruction cache invalidate command to where the next instruction will see the result of the invalidate. The routine in [Example 19](#) can be used to guarantee proper synchronization.

Example 19. Invalidating the Instruction Cache

```
MCR P15,0,R1,C7,C5,0      ; Invalidate the instruction cache and branch
                           ; target buffer

CPWAIT

; The instruction cache is guaranteed to be invalidated at this point; the next
; instruction sees the result of the invalidate command.
```

The Intel XScale® core also supports invalidating an individual line from the instruction cache.

3.4.2.4 Locking Instructions in the Instruction Cache

Software has the ability to lock performance critical routines into the instruction cache. Up to 28 lines in each set can be locked; hardware will ignore the lock command if software is trying to lock all the lines in a particular set (i.e., ways 28 – 31 can never be locked). When this happens, the line is still allocated into the cache, but the lock will be ignored. The round-robin pointer will stay at way 31 for that set.

Lines can be locked into the instruction cache by initiating a write to coprocessor 15. Register *Rd* contains the virtual address of the line to be locked into the cache.

There are several requirements for locking down code:


1. The routine used to lock lines down in the cache must be placed in non-cacheable memory, which means the MMU is enabled. As a corollary: no fetches of cacheable code should occur while locking instructions into the cache.
2. The code being locked into the cache must be cacheable.
3. The instruction cache must be enabled and invalidated prior to locking down lines.

Failure to follow these requirements will produce unpredictable results when accessing the instruction cache.

System programmers should ensure that the code to lock instructions into the cache does not reside closer than 128 bytes to a non-cacheable/cacheable page boundary. If the processor fetches ahead into a cacheable page, then the first requirement noted above could be violated.

Lines are locked into a set starting at way 0 and may progress up to way 27; which set a line gets locked into depends on the set index of the virtual address. [Figure 19](#) is an example of where lines of code may be locked into the cache along with how the round-robin pointer is affected.

Figure 19. Locked Line Effect on Round Robin Replacement



Software can lock down several different routines located at different memory locations. This may cause some sets to have more locked lines than others as shown in [Figure 19](#).



Example 20 shows how a routine, called “lockMe” in this example, might be locked into the instruction cache. Note that it is possible to receive an exception while locking code.

Example 20. Locking Code into the Cache

```
lockMe:                ; This is the code that will be locked into the cache
    mov r0, #5
    add r5, r1, r2
    . . .
lockMeEnd:
    . . .

codeLock:              ; here is the code to lock the "lockMe" routine
    ldr r0, =(lockMe AND NOT 31); r0 gets a pointer to the first line we
    should lock
    ldr r1, =(lockMeEnd AND NOT 31); r1 contains a pointer to the last line we
    should lock

lockLoop:
    mcr p15, 0, r0, c9, c1, 0; lock next line of code into ICache
    cmp r0, r1          ; are we done yet?
    add r0, r0, #32     ; advance pointer to next line
    bne lockLoop       ; if not done, do the next line
```

3.4.2.5 Unlocking Instructions in the Instruction Cache

The Intel XScale® core provides a global unlock command for the instruction cache. Writing to coprocessor 15, register 9 unlocks all the locked lines in the instruction cache and leaves them valid. These lines then become available for the round-robin replacement algorithm.

3.5 Branch Target Buffer (BTB)

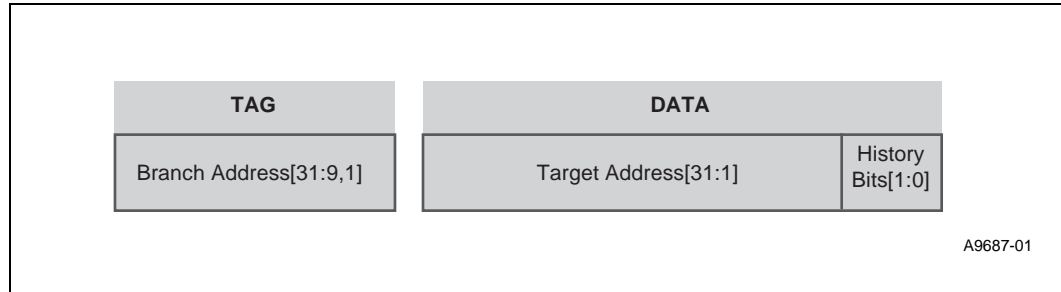
The Intel XScale® core uses dynamic branch prediction to reduce the penalties associated with changing the flow of program execution. The Intel XScale® core features a branch target buffer that provides the instruction cache with the target address of branch type instructions. The branch target buffer is implemented as a 128-entry, direct mapped cache.

3.5.1 Branch Target Buffer Operation

The BTB stores the history of branches that have executed along with their targets. Figure 20 shows an entry in the BTB, where the tag is the instruction address of a previously executed branch and the data contains the target address of the previously executed branch along with two bits of history information.



Figure 20. BTB Entry



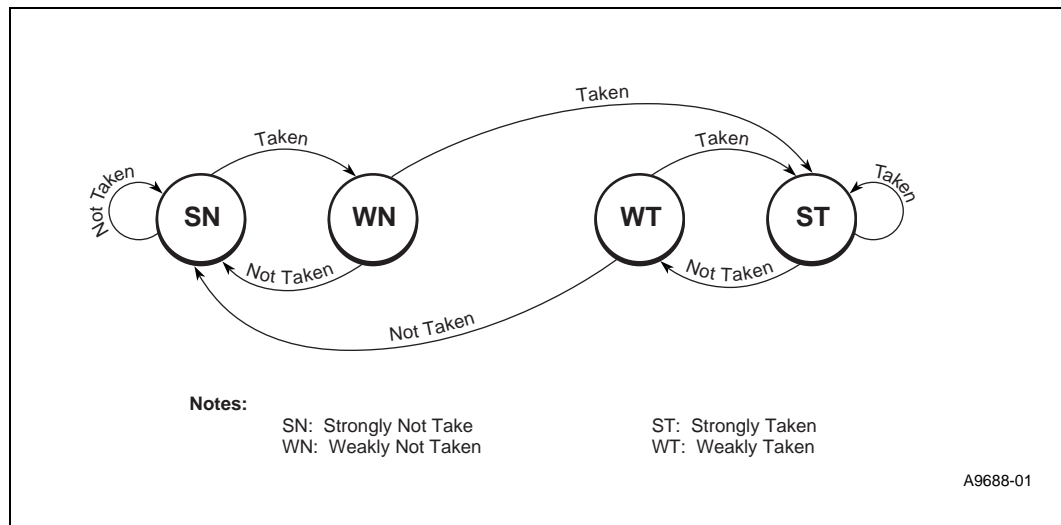
The BTB takes the current instruction address and checks to see if this address is a branch that was previously seen. It uses bits [8:2] of the current address to read out the tag and then compares this tag to bits [31:9,1] of the current instruction address. If the current instruction address matches the tag in the cache and the history bits indicate that this branch is usually taken in the past, the BTB uses the data (target address) as the next instruction address to send to the instruction cache.

Bit[1] of the instruction address is included in the tag comparison to support Thumb execution. This organization means that two consecutive Thumb branch (B) instructions, with instruction address bits[8:2] the same, will contend for the same BTB entry. Thumb also requires 31 bits for the branch target address. In ARM* mode, bit[1] is 0.

The history bits represent four possible prediction states for a branch entry in the BTB. Figure 21 shows these states along with the possible transitions. The initial state for branches stored in the BTB is Weakly-Taken (WT). Every time a branch that exists in the BTB is executed, the history bits are updated to reflect the latest outcome of the branch, either taken or not-taken.

The BTB does not have to be managed explicitly by software; it is disabled by default after reset and is invalidated when the instruction cache is invalidated.

Figure 21. Branch History



3.5.1.1 Reset

After Processor Reset, the BTB is disabled and all entries are invalidated.



3.5.2 Update Policy

A new entry is stored into the BTB when the following conditions are met:

- The branch instruction has executed
- The branch was taken
- The branch is not currently in the BTB

The entry is then marked valid and the history bits are set to WT. If another valid branch exists at the same entry in the BTB, it will be evicted by the new branch.

Once a branch is stored in the BTB, the history bits are updated upon every execution of the branch as shown in [Figure 21](#).

3.5.3 BTB Control

3.5.3.1 Disabling/Enabling

The BTB is always disabled with Reset. Software can enable the BTB through a bit in a coprocessor register.

Before enabling or disabling the BTB, software must invalidate it (described in the following section). This action will ensure correct operation in case stale data is in the BTB. Software should not place any branch instruction between the code that invalidates the BTB and the code that enables/disables it.

3.5.3.2 Invalidation

There are four ways the contents of the BTB can be invalidated.

1. Reset.
2. Software can directly invalidate the BTB via a CP15, register 7 function.
3. The BTB is invalidated when the Process ID register is written.
4. The BTB is invalidated when the instruction cache is invalidated via CP15, register 7 functions.

3.6 Data Cache

The Intel XScale® core data cache enhances performance by reducing the number of data accesses to and from external memory. There are two data cache structures in the Intel XScale® core, a 32-Kbyte data cache and a 2-Kbyte mini-data cache. An eight entry write buffer and a four entry fill buffer are also implemented to decouple the Intel XScale® core instruction execution from external memory accesses, which increases overall system performance.



3.6.1 Overviews

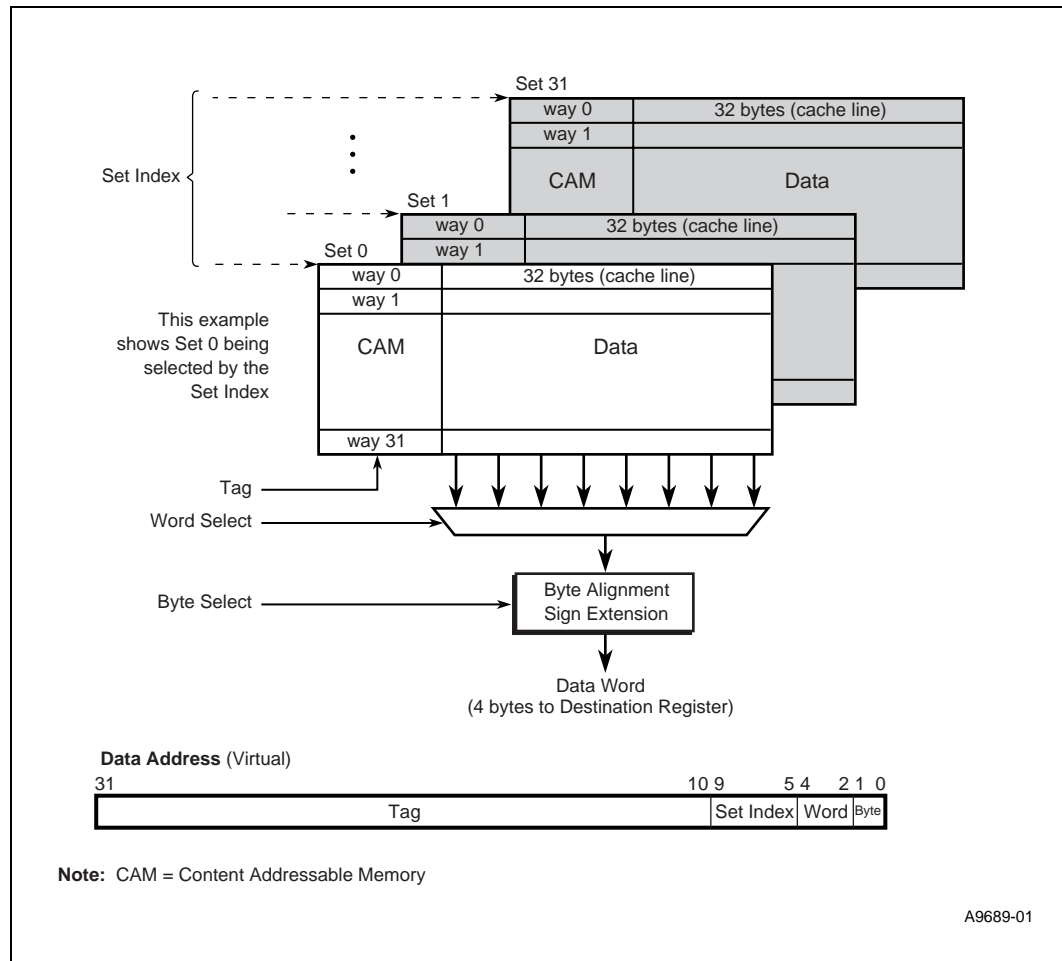
3.6.1.1 Data Cache Overview

The data cache is a 32-Kbyte, 32-way set associative cache, i.e., there are 32 sets and each set has 32 ways. Each way of a set contains 32 bytes (one cache line) and one valid bit. There also exist two dirty bits for every line, one for the lower 16 bytes and the other one for the upper 16 bytes. When a store hits the cache, the dirty bit associated with it is set. The replacement policy is a round-robin algorithm and the cache also supports the ability to reconfigure each line as data RAM.

Figure 22 shows the cache organization and how the data address is used to access the cache. Cache policies may be adjusted for particular regions of memory by altering page attribute bits in the MMU descriptor that controls that memory.

The data cache is virtually addressed and virtually tagged. It supports write-back and write-through caching policies. The data cache always allocates a line in the cache when a cacheable read miss occurs and will allocate a line into the cache on a cacheable write miss when write allocate is specified by its page attribute. Page attribute bits determine whether a line gets allocated into the data cache or mini-data cache.

Figure 22. Data Cache Organization



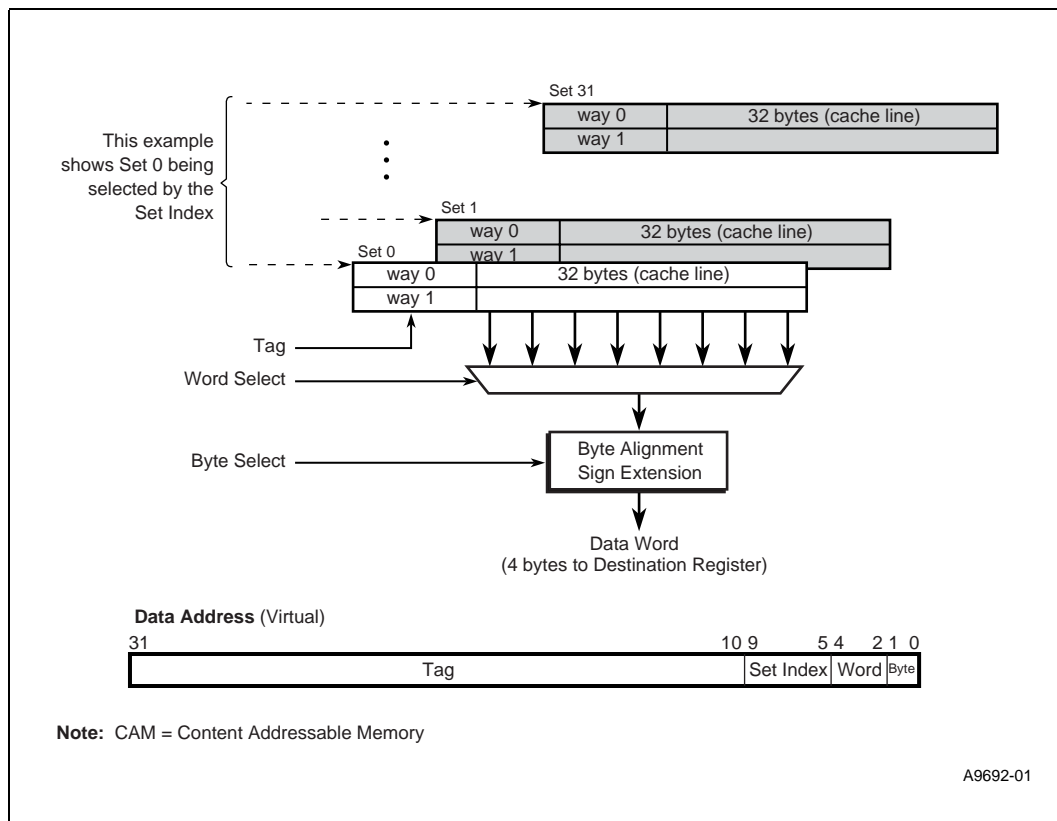
3.6.1.2 Mini-Data Cache Overview

The mini-data cache is a 2-Kbyte, 2-way set associative cache; this means there are 32 sets with each set containing 2 ways. Each way of a set contains 32 bytes (one cache line) and one valid bit. There also exist 2 dirty bits for every line, one for the lower 16 bytes and the other one for the upper 16 bytes. When a store hits the cache, the dirty bit associated with it is set. The replacement policy is a round-robin algorithm.

Figure 23 shows the cache organization and how the data address is used to access the cache.

The mini-data cache is virtually addressed and virtually tagged and supports the same caching policies as the data cache. However, lines cannot be locked into the mini-data cache.

Figure 23. Mini-Data Cache Organization





3.6.1.3 Write Buffer and Fill Buffer Overview

The Intel XScale® core employs an eight entry write buffer, each entry containing 16 bytes. Stores to external memory are first placed in the write buffer and subsequently taken out when the bus is available. The write buffer supports the coalescing of multiple store requests to external memory. An incoming store may coalesce with any of the eight entries.

The fill buffer holds the external memory request information for a data cache or mini-data cache fill or non-cacheable read request. Up to four 32-byte read request operations can be outstanding in the fill buffer before the Intel XScale® core needs to stall.

The fill buffer has been augmented with a four-entry pend buffer that captures data memory requests to outstanding fill operations. Each entry in the pend buffer contains enough data storage to hold one 32-bit word, specifically for store operations. Cacheable load or store operations that hit an entry in the fill buffer get placed in the pend buffer and are completed when the associated fill completes. Any entry in the pend buffer can be pended against any of the entries in the fill buffer; multiple entries in the pend buffer can be pended against a single entry in the fill buffer. Pended operations complete in program order.

3.6.2 Data Cache and Mini-Data Cache Operation

The following discussions refer to the data cache and mini-data cache as one cache (data/mini-data) since their behavior is the same when accessed.

3.6.2.1 Operation when Caching is Enabled

When the data/mini-data cache is enabled for an access, the data/mini-data cache compares the address of the request against the addresses of data that it is currently holding. If the line containing the address of the request is resident in the cache, the access “hits” the cache. For a load operation the cache returns the requested data to the destination register and for a store operation the data is stored into the cache. The data associated with the store may also be written to external memory if write-through caching is specified for that area of memory. If the cache does not contain the requested data, the access ‘misses’ the cache, and the sequence of events that follows depends on the configuration of the cache, the configuration of the MMU and the page attributes.

3.6.2.2 Operation when Data Caching is Disabled

The data/mini-data cache is still accessed even though it is disabled. If a load hits the cache it will return the requested data to the destination register. If a store hits the cache, the data is written into the cache. Any access that misses the cache will not allocate a line in the cache when it’s disabled, even if the MMU is enabled and the memory region’s cacheability attribute is set.



3.6.2.3 Cache Policies

3.6.2.3.1 Cacheability

Data at a specified address is cacheable given the following:

- The MMU is enabled
- The cacheable attribute is set in the descriptor for the accessed address
- The data/mini-data cache is enabled

3.6.2.3.2 Read Miss Policy

The following sequence of events occurs when a cacheable load operation misses the cache:

1. The fill buffer is checked to see if an outstanding fill request already exists for that line.
 - If so, the current request is placed in the pending buffer and waits until the previously requested fill completes, after which it accesses the cache again, to obtain the request data and returns it to the destination register.
 - If there is no outstanding fill request for that line, the current load request is placed in the fill buffer and a 32-byte external memory read request is made. If the pending buffer or fill buffer is full, the Intel XScale® core will stall until an entry is available.
2. A line is allocated in the cache to receive the 32 bytes of fill data. The line selected is determined by the round-robin pointer (see [Section 3.6.2.4](#)). The line chosen may contain a valid line previously allocated in the cache. In this case both dirty bits are examined and if set, the four words associated with a dirty bit that's asserted will be written back to external memory as a 4-word burst operation.
3. When the data requested by the load is returned from external memory, it is immediately sent to the destination register specified by the load. A system that returns the requested data back first, with respect to the other bytes of the line, will obtain the best performance.
4. As data returns from external memory, it is written into the cache in the previously allocated line.

A load operation that misses the cache and is *not* cacheable makes a request from external memory for the exact data size of the original load request. For example, **LDRH** requests exactly two bytes from external memory, **LDR** requests four bytes from external memory, etc. This request is placed in the fill buffer until, the data is returned from external memory, which is then forwarded back to the destination register(s).



3.6.2.3.3 Write Miss Policy

A write operation that misses the cache, requests a 32-byte cache line from external memory if the access is cacheable and write allocation is specified in the page; then, the following events occur:

1. The fill buffer is checked to see if an outstanding fill request already exists for that line.
 - If so, the current request is placed in the pending buffer and waits until the previously requested fill completes, after which it writes its data into the recently allocated cache line.
 - If there is no outstanding fill request for that line, the current store request is placed in the fill buffer and a 32-byte external memory read request is made. If the pending buffer or fill buffer is full, the Intel XScale® core will stall until an entry is available.
2. The 32 bytes of data can be returned back to the Intel XScale® core in any word order, i.e., the eight words in the line can be returned in any order. Note that it does not matter, for performance reasons, which order the data is returned to the Intel XScale® core since the store operation has to wait until the entire line is written into the cache before it can complete.
3. When the entire 32-byte line has returned from external memory, a line is allocated in the cache, selected by the round-robin pointer (see [Section 3.6.2.4](#)). The line to be written into the cache may replace a valid line previously allocated in the cache. In this case both dirty bits are examined and if any are set, the four words associated with a dirty bit that's asserted will be written back to external memory as a 4-word burst operation. This write operation will be placed in the write buffer.
4. The line is written into the cache along with the data associated with the store operation.

If the above condition for requesting a 32-byte cache line is not met, a write miss will cause a write request to external memory for the exact data size specified by the store operation, assuming the write request does not coalesce with another write operation in the write buffer.

3.6.2.3.4 Write-Back versus Write-Through

The Intel XScale® core supports write-back caching or write-through caching, controlled through the MMU page attributes. When write-through caching is specified, all store operations are written to external memory even if the access hits the cache. This feature keeps the external memory coherent with the cache, i.e., no dirty bits are set for this region of memory in the data/mini-data cache. This however does not guarantee that the data/mini-data cache is coherent with external memory, which is dependent on the system level configuration, specifically if the external memory is shared by another master.

When write-back caching is specified, a store operation that hits the cache will not generate a write to external memory, thus reducing external memory traffic.



3.6.2.4 Round-Robin Replacement Algorithm

The line replacement algorithm for the data cache is round-robin. Each set in the data cache has a round-robin pointer that keeps track of the next line (in that set) to replace. The next line to replace in a set is the next sequential line after the last one that was just filled. For example, if the line for the last fill was written into way 5-set 2, the next line to replace for that set would be way 6. None of the other round-robin pointers for the other sets are affected in this case.

After reset, way 31 is pointed to by the round-robin pointer for all the sets. Once a line is written into way 31, the round-robin pointer points to the first available way of a set, beginning with way 0 if no lines have been reconfigured as data RAM in that particular set. Reconfiguring lines as data RAM effectively reduces the available lines for cache updating. For example, if the first three lines of a set were reconfigured, the round-robin pointer would point to the line at way 3 after it rolled over from way 31. Refer to [Section 3.6.4](#) for more details on data RAM.

The mini-data cache follows the same round-robin replacement algorithm as the data cache except that there are only two lines the round-robin pointer can point to such that the round-robin pointer always points to the least recently filled line. A least recently used replacement algorithm is not supported because the purpose of the mini-data cache is to cache data that exhibits low temporal locality, i.e., data that is placed into the mini-data cache is typically modified once and then written back out to external memory.

3.6.2.5 Parity Protection

The data cache and mini-data cache are protected by parity to ensure data integrity; there is one parity bit per byte of data. (The tags are *not* parity protected.) When a parity error is detected on a data/mini-data cache access, a data abort exception occurs. Before servicing the exception, hardware will set bit 10 of the Fault Status register.

A data/mini-data cache parity error is an imprecise data abort, meaning R14_ABORT (+8) may not point to the instruction that caused the parity error. If the parity error occurred during a load, the targeted register may be updated with incorrect data.

A data abort due to a data/mini-data cache parity error may not be recoverable if the data address that caused the abort occurred on a line in the cache that has a write-back caching policy. Prior updates to this line may be lost; in this case the software exception handler should perform a “clean and clear” operation on the data cache, ignoring subsequent parity errors, and restart the offending process. This operation is shown in [Section 3.6.3.3.1](#).

3.6.2.6 Atomic Accesses

The **SWP** and **SWPB** instructions generate an atomic load and store operation allowing a memory semaphore to be loaded and altered without interruption. These accesses may hit or miss the data/mini-data cache depending on configuration of the cache, configuration of the MMU, and the page attributes. Refer to [Section 3.11.4](#) for more information.



3.6.3 Data Cache and Mini-Data Cache Control

3.6.3.1 Data Memory State After Reset

After processor reset, both the data cache and mini-data cache are disabled, all valid bits are set to 0 (invalid), and the round-robin bit points to way 31. Any lines in the data cache that were configured as data RAM before reset are changed back to cacheable lines after reset, i.e., there are 32 KBytes of data cache and 0 bytes of data RAM.

3.6.3.2 Enabling/Disabling

The data cache and mini-data cache are enabled by setting bit 2 in coprocessor 15, register 1 (Control register).

[Example 21](#) shows code that enables the data and mini-data caches. Note that the MMU must be enabled to use the data cache.

Example 21. Enabling the Data Cache

```
enableDCache:

    MCR p15, 0, r0, c7, c10, 4; Drain pending data operations...
    ;
    MRC p15, 0, r0, c1, c0, 0; Get current control register
    ORR r0, r0, #4 ; Enable DCache by setting 'C' (bit 2)
    MCR p15, 0, r0, c1, c0, 0; And update the Control register
```

3.6.3.3 Invalidate and Clean Operations

Individual entries can be invalidated and cleaned in the data cache and mini-data cache via coprocessor 15, register 7. Note that a line locked into the data cache remains locked even after it has been subjected to an invalidate-entry operation. This will leave an unusable line in the cache until a global unlock has occurred. For this reason, do not use these commands on locked lines.

This same register also provides the command to invalidate the entire data cache and mini-data cache. These global invalidate commands have no effect on lines locked in the data cache. Locked lines must be unlocked before they can be invalidated. This is accomplished by the Unlock Data Cache command.



3.6.3.3.1 Global Clean and Invalidate Operation

A simple software routine is used to globally clean the data cache. It takes advantage of the line-allocate data cache operation, which allocates a line into the data cache. This allocation evicts any cache dirty data back to external memory. [Example 22](#) shows how data cache can be cleaned.

Example 22. Global Clean Operation

```

; Global Clean/Invalidate THE DATA CACHE
; R1 contains the virtual address of a region of cacheable memory reserved for
; this clean operation
; R0 is the loop count; Iterate 1024 times which is the number of lines in the
; data cache

;; Macro ALLOCATE performs the line-allocation cache operation on the
;; address specified in register Rx.
;;
MACRO ALLOCATE Rx
    MCR P15, 0, Rx, C7, C2, 5
ENDM

MOV R0, #1024
LOOP1:
ALLOCATE R1          ; Allocate a line at the virtual address
                    ; specified by R1.
ADD R1, R1, #32     ; Increment the address in R1 to the next cache line
SUBS R0, R0, #1     ; Decrement loop count
BNE LOOP1
;
; Clean the Mini-data Cache
; Can't use line-allocate command, so cycle 2KB of unused data through.
; R2 contains the virtual address of a region of cacheable memory reserved for
; cleaning the Mini-data Cache
; R0 is the loop count; Iterate 64 times which is the number of lines in the
; Mini-data Cache.

MOV R0, #64
LOOP2:
LDR R3,[R2],#32 ; Load and increment to next cache line
SUBS R0, R0, #1 ; Decrement loop count
BNE LOOP2
;
; Invalidate the data cache and mini-data cache
MCR P15, 0, R0, C7, C6, 0
;

```

The line-allocate operation does not require physical memory to exist at the virtual address specified by the instruction, since it does not generate a load/fill request to external memory. Also, the line-allocate operation does not set the 32 bytes of data associated with the line to any known value. Reading this data will produce unpredictable results.

The line-allocate command will not operate on the mini Data Cache, so system software must clean this cache by reading two Kbytes of contiguous unused data into it. This data must be unused and reserved for this purpose so that it will not already be in the cache. It must reside in a page that is marked as mini Data Cache cacheable.

The time it takes to execute a global clean operation depends on the number of dirty lines in cache.



3.6.4 Reconfiguring the Data Cache as Data RAM

Software has the ability to lock tags associated with 32-byte lines in the data cache, thus creating the appearance of data RAM. Any subsequent access to this line will always hit the cache unless it is invalidated. Once a line is locked into the data cache it is no longer available for cache allocation on a line fill. Up to 28 lines in each set can be reconfigured as data RAM, such that the maximum data RAM size is 28 Kbytes.

Hardware does not support locking lines into the mini-data cache; any attempt to do this will produce unpredictable results.

There are two methods for locking tags into the data cache; the method of choice depends on the application. One method is used to lock data that resides in external memory into the data cache and the other method is used to reconfigure lines in the data cache as data RAM. Locking data from external memory into the data cache is useful for lookup tables, constants, and any other data that is frequently accessed. Reconfiguring a portion of the data cache as data RAM is useful when an application needs scratch memory (bigger than the register file can provide) for frequently used variables. These variables may be strewn across memory, making it advantageous for software to pack them into data RAM memory.

Refer to the *Intel XScale® Core Developers Manual* for code examples.

Tags can be locked into the data cache by enabling the data cache lock mode bit located in coprocessor 15, register 9. Once enabled, any new lines allocated into the data cache will be locked down.

Note that the **PLD** instruction will not affect the cache contents if it encounters an error while executing. For this reason, system software should ensure the memory address used in the **PLD** is correct. If this cannot be ascertained, replace the **PLD** with a **LDR** instruction that targets a scratch register.

Lines are locked into a set starting at way 0 and may progress up to way 27; which set a line gets locked into depends on the set index of the virtual address of the request. [Figure 19](#) is an example of where lines of code may be locked into the cache along with how the round-robin pointer is affected.

Software can lock down data located at different memory locations. This may cause some sets to have more locked lines than others as shown in [Figure 19](#).

Lines are unlocked in the data cache by performing an unlock operation.

Before locking, the programmer must ensure that no part of the target data range is already resident in the cache. The Intel XScale® core will not refetch such data, which will result in it not being locked into the cache. If there is any doubt as to the location of the targeted memory data, the cache should be cleaned and invalidated to prevent this scenario. If the cache contains a locked region that the programmer wishes to lock again, then the cache must be unlocked before being cleaned and invalidated.



3.6.5 Write Buffer/Fill Buffer Operation and Control

The write buffer is always enabled, which means stores to external memory will be buffered. The K bit in the Auxiliary Control register (CP15, register 1) is a global enable/disable for allowing coalescing in the write buffer. When this bit disables coalescing, no coalescing will occur regardless the value of the page attributes. If this bit enables coalescing, the page attributes X, C, and B are examined to see if coalescing is enabled for each region of memory.

All reads and writes to external memory occur in program order when coalescing is disabled in the write buffer. If coalescing is enabled in the write buffer, writes may occur out of program order to external memory. Program correctness is maintained in this case by comparing all store requests with all the valid entries in the fill buffer.

The write buffer and fill buffer support a drain operation, such that before the next instruction executes, all the Intel XScale® core data requests to external memory have completed.

Writes to a region marked non-cacheable/non-bufferable (page attributes C, B, and X all 0) will cause execution to stall until the write completes.

If software is running in a privileged mode, it can explicitly drain all buffered writes.

3.7 Configuration

The System Control Coprocessor (CP15) configures the MMU, caches, buffers and other system attributes. Where possible, the definition of CP15 follows the definition of the StrongARM* products. Coprocessor 14 (CP14) contains the performance monitor registers and the trace buffer registers.

CP15 is accessed through MRC and MCR coprocessor instructions and allowed only in privileged mode. Any access to CP15 in user mode or with LDC or STC coprocessor instructions will cause an undefined instruction exception.

CP14 registers can be accessed through MRC, MCR, LDC, and STC coprocessor instructions and allowed only in privileged mode. Any access to CP14 in user mode will cause an undefined instruction exception.

The Intel XScale® core Coprocessors, CP15 and CP14, do not support access via CDP, MRRC, or MCRR instructions. An attempt to access these coprocessors with these instructions will result in an Undefined Instruction exception.

Many of the MCR commands available in CP15 modify hardware state sometime after execution. A software sequence is available for those wishing to determine when this update occurs.

Like certain other ARM* architecture products, the Intel XScale® core includes an extra level of virtual address translation in the form of a PID (Process ID) register and associated logic. Privileged code needs to be aware of this facility because, when interacting with CP15, some addresses are modified by the PID and others are not.

An address that has yet to be modified by the PID (“PIDified”) is known as a *virtual address* (VA). An address that has been through the PID logic, but not translated into a physical address, is a *modified virtual address* (MVA). Non-privileged code always deals with VAs, while privileged code that programs CP15 occasionally needs to use MVAs. For details refer to the *Intel XScale® Core Developers Manual*.



3.8 Performance Monitoring

The Intel XScale® core hardware provides two 32-bit performance counters that allow two unique events to be monitored simultaneously. In addition, the Intel XScale® core implements a 32-bit clock counter that can be used in conjunction with the performance counters; its sole purpose is to count the number of core clock cycles, which is useful in measuring total execution time.

The Intel XScale® core can monitor either occurrence events or duration events. When counting occurrence events, a counter is incremented each time a specified event takes place and when measuring duration, a counter counts the number of processor clocks that occur while a specified condition is true. If any of the three counters overflow, an IRQ or FIQ will be generated if it's enabled. Each counter has its own interrupt enable. The counters continue to monitor events even after an overflow occurs, until disabled by software. Refer to the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual* for more detail.

Each of these counters can be programmed to monitor any one of various events.

To further augment performance monitoring, the Intel XScale® core clock counter can be used to measure the executing time of an application. This information combined with a duration event can feedback a percentage of time the event occurred with respect to overall execution time.

Each of the three counters and the performance monitoring control register are accessible through Coprocessor 14 (CP14), registers 0-3. Access is allowed in privileged mode only.

The following are a few notes about controlling the performance monitoring mechanism:

- An interrupt will be reported when a counter's overflow flag is set and its associated interrupt enable bit is set in the PMNC register. The interrupt will remain asserted until software clears the overflow flag by writing a one to the flag that is set. *Note:* the product specific interrupt unit and the CPSR must have enabled the interrupt in order for software to receive it.
- The counters continue to record events even after they overflow.

3.8.1 Performance Monitoring Events

Table 24 lists events that may be monitored by the PMU. Each of the Performance Monitor Count registers (PMN0 and PMN1) can count any listed event. Software selects which event is counted by each PMNx register by programming the evtCountx fields of the PMNC register.

Table 24. Performance Monitoring Events (Sheet 1 of 2)

Event Number (evtCount0 or evtCount1)	Event Definition
0x0	Instruction cache miss requires fetch from external memory.
0x1	Instruction cache cannot deliver an instruction. This could indicate an ICache miss or an ITLB miss. This event will occur every cycle in which the condition is present.
0x2	Stall due to a data dependency. This event will occur every cycle in which the condition is present.
0x3	Instruction TLB miss.
0x4	Data TLB miss.
0x5	Branch instruction executed, branch may or may not have changed program flow.
0x6	Branch mispredicted. (B and BL instructions only.)



Table 24. Performance Monitoring Events (Sheet 2 of 2)

Event Number (evtCount0 or evtCount1)	Event Definition
0x7	Instruction executed.
0x8	Stall because the data cache buffers are full. This event will occur every cycle in which the condition is present.
0x9	Stall because the data cache buffers are full. This event will occur once for each contiguous sequence of this type of stall.
0xA	Data cache access, not including Cache Operations
0xB	Data cache miss, not including Cache Operations
0xC	Data cache write-back. This event occurs once for each ½ line (four words) that are written back from the cache.
0xD	Software changed the PC. This event occurs any time the PC is changed by software and there is not a mode change. For example, a mov instruction with PC as the destination will trigger this event. Executing a swi from User mode will not trigger this event, because it will incur a mode change.
0x10 — 0x17	Refer to the <i>Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual</i> for more details.
all others	Reserved, unpredictable results

Some typical combination of counted events are listed in this section and summarized in [Table 25](#). In this section, we call such an event combination a *mode*.

Table 25. Some Common Uses of the PMU

Mode	PMNC.evtCount0	PMNC.evtCount1
Instruction Cache Efficiency	0x7 (instruction count)	0x0 (ICache miss)
Data Cache Efficiency	0xA (Dcache access)	0xB (DCache miss)
Instruction Fetch Latency	0x1 (ICache cannot deliver)	0x0 (ICache miss)
Data/Bus Request Buffer Full	0x8 (DBuffer stall duration)	0x9 (DBuffer stall)
Stall/Writeback Statistics	0x2 (data stall)	0xC (DCache writeback)
Instruction TLB Efficiency	0x7 (instruction count)	0x3 (ITLB miss)
Data TLB Efficiency	0xA (Dcache access)	0x4 (DTLB miss)

3.8.1.1 Instruction Cache Efficiency Mode

PMN0 totals the number of instructions that were executed, which does not include instructions fetched from the instruction cache that were never executed. This can happen if a branch instruction changes the program flow; the instruction cache may retrieve the next sequential instructions after the branch, before it receives the target address of the branch.

PMN1 counts the number of instruction fetch requests to external memory. Each of these requests loads 32 bytes at a time.

Statistics derived from these two events:

- Instruction cache miss-rate. This is derived by dividing PMN1 by PMN0.
- *The average number of cycles it took to execute an instruction or commonly referred to as cycles-per-instruction (CPI)*. CPI can be derived by dividing CCNT by PMN0, where CCNT was used to measure total execution time.



3.8.1.2 Data Cache Efficiency Mode

PMN0 totals the number of data cache accesses, which includes cacheable and non-cacheable accesses, mini-data cache access and accesses made to locations configured as data RAM.

Note that **STM** and **LDM** will each count as several accesses to the data cache depending on the number of registers specified in the register list. **LDRD** will register two accesses.

PMN1 counts the number of data cache and mini-data cache misses. Cache operations do not contribute to this count.

The statistic derived from these two events is:

- Data cache miss-rate. This is derived by dividing PMN1 by PMN0.

3.8.1.3 Instruction Fetch Latency Mode

PMN0 accumulates the number of cycles when the instruction-cache is not able to deliver an instruction to the Intel XScale® core due to an instruction-cache miss or instruction-TLB miss. This event means that the processor core is stalled.

PMN1 counts the number of instruction fetch requests to external memory. Each of these requests loads 32 bytes at a time. This is the same event as measured in instruction cache efficiency mode and is included in this mode for convenience so that only one performance monitoring run is need.

Statistics derived from these two events:

- *The average number of cycles the processor stalled waiting for an instruction fetch from external memory to return.* This is calculated by dividing PMN0 by PMN1. If the average is high then the Intel XScale® core may be starved of the bus external to the Intel XScale® core.
- *The percentage of total execution cycles the processor stalled waiting on an instruction fetch from external memory to return.* This is calculated by dividing PMN0 by CCNT, which was used to measure total execution time.

3.8.1.4 Data/Bus Request Buffer Full Mode

The Data Cache has buffers available to service cache misses or uncacheable accesses. For every memory request that the Data Cache receives from the processor core, a buffer is speculatively allocated in case an external memory request is required or temporary storage is needed for an unaligned access. If no buffers are available, the Data Cache will stall the processor core.

The frequency of Data Cache stalls depends on the performance of the bus external to the Intel XScale® core and what the memory access latency is for Data Cache miss requests to external memory. If the Intel XScale® core memory access latency is high (possibly due to starvation) these Data Cache buffers will become full. This performance monitoring mode is provided to determine whether the Intel XScale® core is being starved of the bus external to the Intel XScale® core — which affects the performance of the application running on the Intel XScale® core.

PMN0 accumulates the number of clock cycles by which the processor is stalled due to this condition and PMN1 monitors the number of times this condition occurs.



Statistics derived from these two events:

- *The average number of cycles the processor stalled on a data-cache access that may overflow the data-cache buffers.*

This is calculated by dividing PMN0 by PMN1. This statistic lets you know if the duration event cycles are due to many requests or are attributed to just a few requests. If the average is high, the Intel XScale® core may be starved of the bus external to the Intel XScale® core.

- *The percentage of total execution cycles the processor stalled because a Data Cache request buffer was not available.*

This is calculated by dividing PMN0 by CCNT, which was used to measure total execution time.

3.8.1.5 Stall/Writeback Statistics

When an instruction requires the result of a previous instruction and that result is not yet available, the Intel XScale® core stalls, to preserve the correct data dependencies. PMN0 counts the number of stall cycles due to data dependencies. Not all data dependencies cause a stall; only the following dependencies cause such a stall penalty:

- **Load-use penalty:** attempting to use the result of a load before the load completes. To avoid the penalty, software should delay using the result of a load until it's available. This penalty shows the latency effect of data-cache access.
- **Multiply/Accumulate-use penalty:** attempting to use the result of a multiply or multiply-accumulate operation before the operation completes. Again, to avoid the penalty, software should delay using the result until it's available.
- **ALU use penalty:** there are a few isolated cases where back-to-back ALU operations may result in one cycle delay in the execution.

PMN1 counts the number of writeback operations emitted by the data cache. These writebacks occur when the data cache evicts a dirty line of data to make room for a newly requested line or as the result of clean operation (CP15, register 7).

Statistics derived from these two events:

- *The percentage of total execution cycles the processor stalled because of a data dependency.*

This is calculated by dividing PMN0 by CCNT, which was used to measure total execution time. Often, a compiler can reschedule code to avoid these penalties when given the right optimization switches.

- Total number of data writeback requests to external memory can be derived solely with PMN1.



3.8.1.6 Instruction TLB Efficiency Mode

PMN0 totals the number of instructions that were executed, which does not include instructions that were translated by the instruction TLB and never executed. This can happen if a branch instruction changes the program flow; the instruction TLB may translate the next sequential instructions after the branch, before it receives the target address of the branch.

PMN1 counts the number of instruction TLB table-walks that occurs when there is a TLB miss. If the instruction TLB is disabled, PMN1 will not increment.

Statistics derived from these two events:

- Instruction TLB miss-rate. This is derived by dividing PMN1 by PMN0.
- *The average number of cycles it took to execute an instruction or commonly referred to as cycles-per-instruction (CPI).*
CPI can be derived by dividing CCNT by PMN0, where CCNT was used to measure total execution time.

3.8.1.7 Data TLB Efficiency Mode

PMN0 totals the number of data cache accesses, which includes cacheable and non-cacheable accesses, mini-data cache access and accesses made to locations configured as data RAM.

Note that **STM** and **LDM** will each count as several accesses to the data TLB depending on the number of registers specified in the register list. **LDRD** will register two accesses.

PMN1 counts the number of data TLB table-walks, which occurs when there is a TLB miss. If the data TLB is disabled PMN1 will not increment.

The statistic derived from these two events is:

- Data TLB miss-rate. This is derived by dividing PMN1 by PMN0.

3.8.2 Multiple Performance Monitoring Run Statistics

Even though only two events can be monitored at any given time, multiple performance monitoring runs can be done, capturing different events from different modes. For example, the first run could monitor the number of writeback operations (PMN1 of mode, Stall/Writeback) and the second run could monitor the total number of data cache accesses (PMN0 of mode, Data Cache Efficiency). From the results, a percentage of writeback operations to the total number of data accesses can be derived.

3.9 Performance Considerations

This section describes relevant performance considerations that compiler writers, application programmers, and system designers need to be aware of to efficiently use the Intel XScale® core. Performance numbers discussed here include interrupt latency, branch prediction, and instruction latencies.



3.9.1 Interrupt Latency

Minimum Interrupt Latency is defined as the minimum number of cycles from the assertion of any interrupt signal (IRQ or FIQ) to the execution of the instruction at the vector for that interrupt. The point at which the assertion begins is TBD. This number assumes best case conditions exist when the interrupt is asserted, e.g., the system isn't waiting on the completion of some other operation.

A useful number to work with is the *Maximum Interrupt Latency*. This is typically a complex calculation that depends on what else is going on in the system at the time the interrupt is asserted. Some examples that can adversely affect interrupt latency are:

- The instruction currently executing could be a 16-register LDM.
- The processor could fault just when the interrupt arrives.
- The processor could be waiting for data from a load, doing a page table walk, etc.
- There are high core-to-system (bus) clock ratios.

Maximum Interrupt Latency can be reduced by:

- Ensuring that the interrupt vector and interrupt service routine are resident in the instruction cache. This can be accomplished by locking them down into the cache.
- Removing or reducing the occurrences of hardware page table walks. This also can be accomplished by locking down the application's page table entries into the TLBs, along with the page table entry for the interrupt service routine.

3.9.2 Branch Prediction

The Intel XScale® core implements dynamic branch prediction for the ARM* instructions **B** and **BL** and for the Thumb instruction **B**. Any instruction that specifies the PC as the destination is predicted as not taken. For example, an **LDR** or a **MOV** that loads or moves directly to the PC will be predicted not taken and incur a branch latency penalty.

These instructions -- ARM **B**, ARM **BL** and Thumb **B** -- enter into the branch target buffer when they are "taken" for the first time. (A "taken" branch refers to when they are evaluated to be true.) Once in the branch target buffer, the Intel XScale® core dynamically predicts the outcome of these instructions based on previous outcomes. [Table 26](#) shows the branch latency penalty when these instructions are correctly predicted and when they are not. A penalty of 0 for correct prediction means that the Intel XScale® core can execute the next instruction in the program flow in the cycle following the branch.

Table 26. Branch Latency Penalty

Core Clock Cycles		Description
ARM*	Thumb	
+0	+0	Predicted Correctly. The instruction is in the branch target cache and is correctly predicted.
+4	+5	Mispredicted. There are three occurrences of branch misprediction, all of which incur a 4-cycle branch delay penalty. <ol style="list-style-type: none"> 1. The instruction is in the branch target buffer and is predicted not-taken, but is actually taken. 2. The instruction is not in the branch target buffer and is a taken branch. 3. The instruction is in the branch target buffer and is predicted taken, but is actually not-taken



3.9.3 Addressing Modes

All load and store addressing modes implemented in the Intel XScale® core do not add to the instruction latencies numbers.

3.9.4 Instruction Latencies

The latencies for all the instructions are shown in the following sections with respect to their functional groups: branch, data processing, multiply, status register access, load/store, semaphore, and coprocessor. The following section explains how to read these tables.

3.9.4.1 Performance Terms

- Issue Clock (cycle 0)
The first cycle when an instruction is decoded *and* allowed to proceed to further stages in the execution pipeline (i.e., when the instruction is actually issued).
- Cycle Distance from A to B
The cycle distance from cycle *A* to cycle *B* is $(B-A)$ – that is, the number of cycles from the start of cycle *A* to the start of cycle *B*. Example: the cycle distance from cycle 3 to cycle 4 is one cycle.
- Issue Latency
The cycle distance *from* the first issue clock of the current instruction *to* the issue clock of the next instruction. The actual number of cycles can be influenced by cache-misses, resource-dependency stalls, and resource availability conflicts.
- Result Latency
The cycle distance *from* the first issue clock of the current instruction *to* the issue clock of the first instruction that can use the result without incurring a resource dependency stall. The actual number of cycles can be influenced by cache-misses, resource-dependency stalls, and resource availability conflicts.
- Minimum Issue Latency (without Branch Misprediction)
The minimum cycle distance *from* the issue clock of the current instruction *to* the first possible issue clock of the next instruction assuming best case conditions (i.e., that the issuing of the next instruction is not stalled due to a resource dependency stall; the next instruction is immediately available from the cache or memory interface; the current instruction does not incur resource dependency stalls during execution that cannot be detected at issue time; and if the instruction uses dynamic branch prediction, correct prediction is assumed).
- Minimum Result Latency
The required minimum cycle distance *from* the issue clock of the current instruction *to* the issue clock of the first instruction that can use the result without incurring a resource dependency stall assuming best case conditions (i.e., that the issuing of the next instruction is not stalled due to a resource dependency stall; the next instruction is immediately available from the cache or memory interface; and the current instruction does not incur resource dependency stalls during execution that cannot be detected at issue time).
- Minimum Issue Latency (with Branch Misprediction)
The minimum cycle distance *from* the issue clock of the current branching instruction *to* the first possible issue clock of the next instruction. This definition is identical to *Minimum Issue Latency* except that the branching instruction has been mispredicted. It is calculated by adding



Minimum Issue Latency (without Branch Misprediction) to the minimum branch latency penalty number from [Table 26](#), which is four cycles.

- **Minimum Resource Latency**
The minimum cycle distance from the issue clock of the current multiply instruction to the issue clock of the next multiply instruction assuming the second multiply does not incur a data dependency and is immediately available from the instruction cache or memory interface.

[Example 23](#) contains a code fragment and an example of computing latencies.

Example 23. Computing Latencies

```

UMLALr6,r8,r0,r1
ADD r9,r10,r11
SUB r2,r8,r9
MOV r0,r1
    
```

[Table 27](#) shows how to calculate Issue Latency and Result Latency for each instruction. Looking at the issue column, the **UMLAL** instruction starts to issue on cycle 0 and the next instruction, **ADD**, issues on cycle 2, so the Issue Latency for **UMLAL** is two. From the code fragment, there is a result dependency between the **UMLAL** instruction and the **SUB** instruction. In [Table 27](#), **UMLAL** starts to issue at cycle 0 and the **SUB** issues at cycle 5; so the Result Latency is 5.

Table 27. Latency Example

Cycle	Issue	Executing
0	umlal (1st cycle)	--
1	umlal (2nd cycle)	umlal
2	add	umlal
3	sub (stalled)	umlal & add
4	sub (stalled)	umlal
5	sub	umlal
6	mov	sub
7	--	mov



3.9.4.2 Branch Instruction Timings

Table 28. Branch Instruction Timings (Predicted by the BTB)

Mnemonic	Minimum Issue Latency when Correctly Predicted by the BTB	Minimum Issue Latency with Branch Misprediction
B	1	5
BL	1	5

Table 29. Branch Instruction Timings (Not Predicted by the BTB)

Mnemonic	Minimum Issue Latency when the branch is not taken	Minimum Issue Latency when the branch is taken
BLX(1)	N/A	5
BLX(2)	1	5
BX	1	5
Data Processing Instruction with PC as the destination	Same as Table 30	4 + numbers in Table 30
LDR PC,<>	2	8
LDM with PC in register list	3 + numreg ¹	10 + max (0, numreg-3)

1. numreg is the number of registers in the register list including the PC.

3.9.4.3 Data Processing Instruction Timings

Table 30. Data Processing Instruction Timings

Mnemonic	<shifter operand> is <i>not</i> a Shift/Rotate by Register		<shifter operand> is a Shift/Rotate by Register or <shifter operand> is RRX	
	Minimum Issue Latency	Minimum Result Latency ¹	Minimum Issue Latency	Minimum Result Latency ¹
ADC	1	1	2	2
ADD	1	1	2	2
AND	1	1	2	2
BIC	1	1	2	2
CMN	1	1	2	2
CMP	1	1	2	2
EOR	1	1	2	2
MOV	1	1	2	2
MVN	1	1	2	2
ORR	1	1	2	2
RSB	1	1	2	2
RSC	1	1	2	2
SBC	1	1	2	2
SUB	1	1	2	2
TEQ	1	1	2	2
TST	1	1	2	2

1. If the next instruction needs to use the result of the data processing for a shift by immediate or as Rn in a QDADD or QDSUB, one extra cycle of result latency is added to the number listed.



3.9.4.4 Multiply Instruction Timings

Table 31. Multiply Instruction Timings (Sheet 1 of 2)

Mnemonic	Rs Value (Early Termination)	S-Bit Value	Minimum Issue Latency	Minimum Result Latency ¹	Minimum Resource Latency (Throughput)
MLA	Rs[31:15] = 0x00000 or Rs[31:15] = 0x1FFFF	0	1	2	1
		1	2	2	2
	Rs[31:27] = 0x00 or Rs[31:27] = 0x1F	0	1	3	2
		1	3	3	3
	all others	0	1	4	3
		1	4	4	4
MUL	Rs[31:15] = 0x00000 or Rs[31:15] = 0x1FFFF	0	1	2	1
		1	2	2	2
	Rs[31:27] = 0x00 or Rs[31:27] = 0x1F	0	1	3	2
		1	3	3	3
	all others	0	1	4	3
		1	4	4	4
SMLAL	Rs[31:15] = 0x00000 or Rs[31:15] = 0x1FFFF	0	2	RdLo = 2; RdHi = 3	2
		1	3	3	3
	Rs[31:27] = 0x00 or Rs[31:27] = 0x1F	0	2	RdLo = 3; RdHi = 4	3
		1	4	4	4
	all others	0	2	RdLo = 4; RdHi = 5	4
		1	5	5	5
SMLALxy	N/A	N/A	2	RdLo = 2; RdHi = 3	2
SMLAWy	N/A	N/A	1	3	2
SMLAxy	N/A	N/A	1	2	1
SMULL	Rs[31:15] = 0x00000 or Rs[31:15] = 0x1FFFF	0	1	RdLo = 2; RdHi = 3	2
		1	3	3	3
	Rs[31:27] = 0x00 or Rs[31:27] = 0x1F	0	1	RdLo = 3; RdHi = 4	3
		1	4	4	4
	all others	0	1	RdLo = 4; RdHi = 5	4
		1	5	5	5
SMULWy	N/A	N/A	1	3	2
SMULxy	N/A	N/A	1	2	1
UMLAL	Rs[31:15] = 0x00000	0	2	RdLo = 2; RdHi = 3	2
		1	3	3	3
	Rs[31:27] = 0x00	0	2	RdLo = 3; RdHi = 4	3
		1	4	4	4
	all others	0	2	RdLo = 4; RdHi = 5	4
		1	5	5	5

Table 31. Multiply Instruction Timings (Sheet 2 of 2)

Mnemonic	Rs Value (Early Termination)	S-Bit Value	Minimum Issue Latency	Minimum Result Latency ¹	Minimum Resource Latency (Throughput)
UMULL	Rs[31:15] = 0x00000	0	1	RdLo = 2; RdHi = 3	2
		1	3	3	3
	Rs[31:27] = 0x00	0	1	RdLo = 3; RdHi = 4	3
		1	4	4	4
	all others	0	1	RdLo = 4; RdHi = 5	4
		1	5	5	5

1. If the next instruction needs to use the result of the multiply for a shift by immediate or as Rn in a QDADD or QDSUB, one extra cycle of result latency is added to the number listed.

Table 32. Multiply Implicit Accumulate Instruction Timings

Mnemonic	Rs Value (Early Termination)	Minimum Issue Latency	Minimum Result Latency	Minimum Resource Latency (Throughput)
MIA	Rs[31:16] = 0x0000 or Rs[31:16] = 0xFFFF	1	1	1
	Rs[31:28] = 0x0 or Rs[31:28] = 0xF	1	2	2
	all others	1	3	3
MIAXy	N/A	1	1	1
MIAPH	N/A	1	2	2

Table 33. Implicit Accumulator Access Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency	Minimum Resource Latency (Throughput)
MAR	2	2	2
MRA	1	(RdLo = 2; RdHi = 3) ¹	2

1. If the next instruction needs to use the result of the MRA for a shift by immediate or as Rn in a QDADD or QDSUB, one extra cycle of result latency is added to the number listed.

3.9.4.5 Saturated Arithmetic Instructions

Table 34. Saturated Data Processing Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
QADD	1	2
QSUB	1	2
QDADD	1	2
QDSUB	1	2



3.9.4.6 Status Register Access Instructions

Table 35. Status Register Access Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
MRS	1	2
MSR	2 (6 if updating mode bits)	1

3.9.4.7 Load/Store Instructions

Table 36. Load and Store Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
LDR	1	3 for load data; 1 for writeback of base
LDRB	1	3 for load data; 1 for writeback of base
LDRBT	1	3 for load data; 1 for writeback of base
LDRD	1 (+1 if Rd is R12)	3 for Rd; 4 for Rd+1; 2 for writeback of base
LDRH	1	3 for load data; 1 for writeback of base
LDRSB	1	3 for load data; 1 for writeback of base
LDRSH	1	3 for load data; 1 for writeback of base
LDRT	1	3 for load data; 1 for writeback of base
PLD	1	N/A
STR	1	1 for writeback of base
STRB	1	1 for writeback of base
STRBT	1	1 for writeback of base
STRD	2	1 for writeback of base
STRH	1	1 for writeback of base
STRT	1	1 for writeback of base

Table 37. Load and Store Multiple Instruction Timings

Mnemonic	Minimum Issue Latency ¹	Minimum Result Latency
LDM	3 – 23	1 – 3 for load data; 1 for writeback of base
STM	3 – 18	1 for writeback of base

1. LDM issue latency is 7 + N if R15 is in the register list and 2 + N if it is not. STM issue latency is calculated as 2 + N. N is the number of registers to load or store.

3.9.4.8 Semaphore Instructions

Table 38. Semaphore Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
SWP	5	5
SWPB	5	5



3.9.4.9 Coprocessor Instructions

Table 39. CP15 Register Access Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
MRC	4	4
MCR	2	N/A

Table 40. CP14 Register Access Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
MRC	7	7
MCR	7	N/A
LDC	10	N/A
STC	7	N/A

3.9.4.10 Miscellaneous Instruction Timing

Table 41. SWI Instruction Timings

Mnemonic	Minimum latency to first instruction of SWI exception handler
SWI	6

Table 42. Count Leading Zeros Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
CLZ	1	1

3.9.4.11 Thumb Instructions

The timing of Thumb instructions are the same as their equivalent ARM* instructions. This mapping can be found in the *ARM* Architecture Reference Manual*. The only exception is the Thumb BL instruction when H = 0; the timing in this case would be the same as an ARM* data processing instruction.

3.10 Test Features

This section gives a brief overview of the Intel XScale® core JTAG features. The Intel XScale® core provides test features compatible with the IEEE Standard Test Access Port and Boundary Scan Architecture (IEEE Std. 1149.1). These features include a TAP controller, a 5-bit instruction register, and test data registers to support software debug. The Intel XScale® core also provides support for a boundary-scan register, device ID register, and other data test registers.

A full description of these features can be found in the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.



3.10.1 IXP2800 Network Processor Endianness

Endianness defines the way bytes are addressed within a word. A little-endian system is one in which byte 0 is the least significant byte (LSB) in the word and byte 3 is the most significant byte (MSB). A big-endian system is one in which byte 0 is the MSB and byte 3 is the LSB. For example, the value of 0x12345678 at address 0x0 in a 32-bit little-endian system looks like this:

Table 43. Little-Endian Encoding

Address/Byte Lane	0x0/ByteLane 3	0x0/ByteLane 2	0x0/ByteLane 1	0x0/ByteLane 0
Byte Value	0x12	0x34	0x56	0x78

The same value stored in a big-endian system is shown in Table 44:

Table 44. Big-Endian Encoding

Address/Byte Lane	0x0/ByteLane 3	0x0/ByteLane 2	0x0/ByteLane 1	0x0/ByteLane 0
Byte Value	0x78	0x56	0x34	0x12

Bits within a byte are always in little-endian order. The least significant bit resides at bit location 0 and the most significant bit resides at bit location 7 (7:0).

The following conventions are used in this document:

- 1 Byte: 8-bit data
- 1 Word: 16-bit data
- 1 Longword: 32-bit data
- Longword Little-Endian Format (LWLE): 32-bit data (0x12345678) arranged as {12 34 56 78}
- Longword Big-Endian format (LWBE): 32-bit data (0x12345678) arranged as {78 56 34 12}
- 64-bit data 0x12345678 9ABCDE56 arranged as {12 34 56 78 9A BC DE 56}
- 64-bit data 0x12345678 9ABCDE56 arranged as {78 56 34 12, 56 DE BC 9A}

Endianness for the IXP2800 network processor can be divided into three major categories:

- Read and write transactions initiated by the Intel XScale® core:
 - Reads initiated by the Intel XScale® core
 - Writes initiated by the Intel XScale® core
- SRAM and DRAM access:
 - 64-bit Data transfer between DRAM and the Intel XScale® core
 - Byte, word, or longword transfer between SRAM/DRAM and the Intel XScale® core
 - Data transfer between SRAM/DRAM and PCI
 - Microengine-initiated access to SRAM and DRAM
- PCI Accesses
 - Intel XScale® core generated reads/writes to PCI in memory space
 - Intel XScale® core generated read/write of external/internal PCI configuration registers



3.10.1.1 Read and Write Transactions Initiated by the Intel XScale® Core

The Intel XScale® core may be used in either a little-endian or big-endian configuration. The configuration affects the entire system in which the Intel XScale® core microarchitecture exists. Software and hardware must agree on the byte ordering to be used. In software, a system's byte order is configured with CP15 register 1, the control register. Bit 7 of this register, the B bit, informs the processor of the byte order in use by the system. Note that this bit takes effect even if the MMU is not otherwise in use or enabled. The state of this bit is reflected in the `cbiBigEndian` signal.

Although it is the responsibility of system hardware to assign correct byte lanes to each byte field in the data bus, in the IXP2800 network processor, it is left to the software to interpret byte lanes in accordance with the endianness of the system. As shown in [Figure 24](#), system byte lanes 0 – 3 are connected directly to the Intel XScale® core byte lanes 0 – 3. This means that byte lane 0 (M[7:0]) of the system is connected to byte lane 0 (X[7:0]) of the Intel XScale® core, byte lane 1 (M[15:8]) of the system is connected to byte lane 1 (X[15:8]) of the Intel XScale® core, etc.

Interface operation of the Intel XScale® core and the rest of the IXP2800 network processor can be divided into two parts:

- Intel XScale® core reading from the IXP2800 network processor
- Intel XScale® core writing to the IXP2800 network processor

3.10.1.1.1 Reads Initiated by the Intel XScale® Core

Intel XScale® core reads can be one of the following three types:

- Byte read
- 16-bits (word) read
- 32-bits (longword) read

Byte Read

When reading a byte, the Intel XScale® core generates the `byte_enable` that corresponds to the proper byte lane as defined by the endianness setting. [Table 45](#) summarizes byte-enable generation for this mode.

Table 45. Byte-Enable Generation by the Intel XScale® Core for Byte Transfers in Little- and Big-Endian Systems

Byte Number to be Read	Byte-Enables for Little-Endian System				Byte-Enables for Big-Endian System			
	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]
Byte 0	1	0	0	0	0	0	0	1
Byte 1	0	1	0	0	0	0	1	0
Byte 2	0	0	1	0	0	1	0	0
Byte 3	0	0	0	1	1	0	0	0

The 4-to-1 multiplexer steers the byte read into the byte lane 0 location of the read register inside the Intel XScale® core. Select signals for the multiplexer are generated based on endian setting and `ByteEnable` generated by the Intel XScale® core as defined in [Figure 24](#).



16-Bit (Word) Read

When reading a word, the Intel XScale® core generates the byte_enable that corresponds to the proper byte lane as defined by the endianness setting. Figure 25 summarizes byte enable generation for this mode.

The 4-to-1 multiplexer steers byte lane 0 or byte lane 2 into the byte 0 location of the read register inside the Intel XScale® core. The 2-to-1 multiplexer steers byte lane 1 or byte lane 3 into the byte 1 location of the read register inside the Intel XScale® core. The Intel XScale® core does not allow word access to an odd-byte address. Select signals for the multiplexer are generated based on endian setting and ByteEnable generated by the Intel XScale® core, as defined in Figure 24. Table 46 summarizes byte-enable generation for this mode.

Figure 24. Byte Steering for Read and Byte-Enable Generation by the Intel XScale® Core

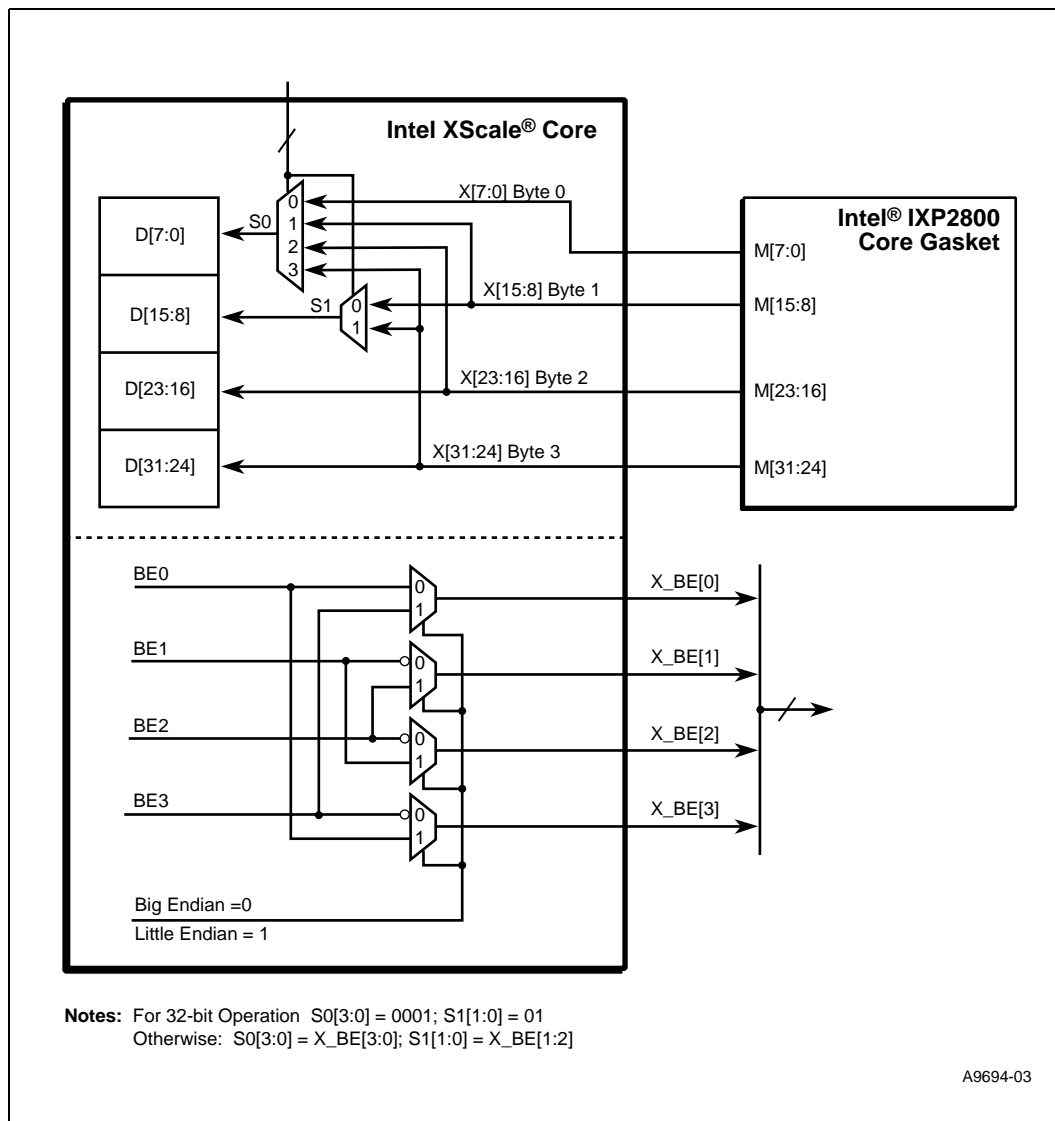


Table 46. Byte-Enable Generation by the Intel XScale® Core for 16-Bit Data Transfers in Little- and Big-Endian Systems

Word to be Read	Byte-Enables for Little-Endian System				Byte-Enables for Big-Endian System			
	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]
Byte 0, Byte 1	1	1	0	0	0	0	1	1
Byte 2, Byte 3	0	0	1	1	1	1	0	0

32-Bit (Longword) Read

32-bit (longword) reads are independent of endianness. Byte lane 0 from the Intel XScale® core’s data bus gets into the byte 0 location of the read register inside the Intel XScale® core, byte lane 1 from the Intel XScale® core’s data bus gets into the byte 1 location of the read register inside the Intel XScale® core, etc. The software determines byte location, based on the endian setting.

3.10.1.1.2 The Intel XScale® Core Writing to the IXP2800 Network Processor

Writes by the Intel XScale® core can also be divided into the following three categories:

- Byte Write
- Word Write (16 bits)
- Longword write (32 bits)

Byte Write

When the Intel XScale® core writes a single byte to external memory, it puts the byte in the byte lane where it intends to write it, along with the byte enable for that byte turned ON, based on the endian setting of the system. Intel XScale® core register bits [7:0] always contain the byte to be written, regardless of the B-bit setting.

For example, if the Intel XScale® core wants to write to byte 0 in the little-endian system, it puts the byte in byte lane 0 and turns X_BE[0] to ON. If the system is big-endian, the Intel XScale® core puts byte 0 in byte lane 3 and turns X_BE[3] to ON. Other possible combinations of byte lanes and byte enables are shown in the Table 47. Byte lanes other than the one currently being driven by the Intel XScale® core, contain undefined data.

Table 47. Byte-Enable Generation by the Intel XScale® Core for Byte Writes in Little- and Big-Endian Systems

Byte Number to be Written	Byte-Enables for Little-Endian Systems				Byte-Enables for Big-Endian Systems			
	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]
Byte 0	1	0	0	0	0	0	0	1
Byte 1	0	1	0	0	0	0	1	0
Byte 2	0	0	1	0	0	1	0	0
Byte 3	0	0	0	1	1	0	0	0



Word Write (16-Bits Write)

When the Intel XScale® core writes a 16-bit word to external memory, it puts the bytes in the byte lanes where it intends to write them along with the byte enables for those bytes turned ON based on the endian setting of the system. The Intel XScale® core does not allow a word write on an odd-byte address. The Intel XScale® core register bits [15:0] always contain the word to be written regardless of the B-bit setting.

For example, if the Intel XScale® core wants to write one word to a little-endian system at address 0x0002, it will copy byte 0 to byte lane 2 and byte 1 to byte lane 3 along with X_BE[2] and X_BE[3] turned ON. If the Intel XScale® core wants to write one word to a big-endian system at address 0x0002, it will copy byte 0 to byte lane 0 and byte 1 to byte lane 1 along with X_BE[0] and X_BE[1] turned ON. Table 48 shows other possible combinations of byte lanes and byte enables. Byte lanes other than those currently driven by the Intel XScale® core contain undefined data.

Table 48. Byte-Enable Generation by the Intel XScale® Core for Word Writes in Little- and Big-Endian Systems

Word to be Written	Byte-Enables for Little-Endian Systems				Byte-Enables for Big-Endian Systems			
	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]	X_BE[0]	X_BE[1]	X_BE[2]	X_BE[3]
Byte 0, Byte 1	1	1	0	0	0	0	1	1
Byte 2, Byte 3	0	0	1	1	1	1	0	0

Longword (32-Bits) Write

The longword to be written is put on the Intel XScale® core’s data bus with byte 0 on X[7:0], byte 1 on X[15:8], byte 2 on X[23:16], and byte 4 on X[31:24] (see Figure 25). All of the byte enables are turned ON. A 32-bit longword write (0x12345678) by the Intel XScale® core to address 0x0000 regardless of endianness, causes byte 0 (0x78) to be written to address 0x0000, byte 1 (0x56) to address 0x0001, byte 2 (0x34) to address 0x0002, and byte 3 (0x12) to address 0x0003.

Figure 25. Intel XScale® Core-Initiated Write to the IXP2800 Network Processor

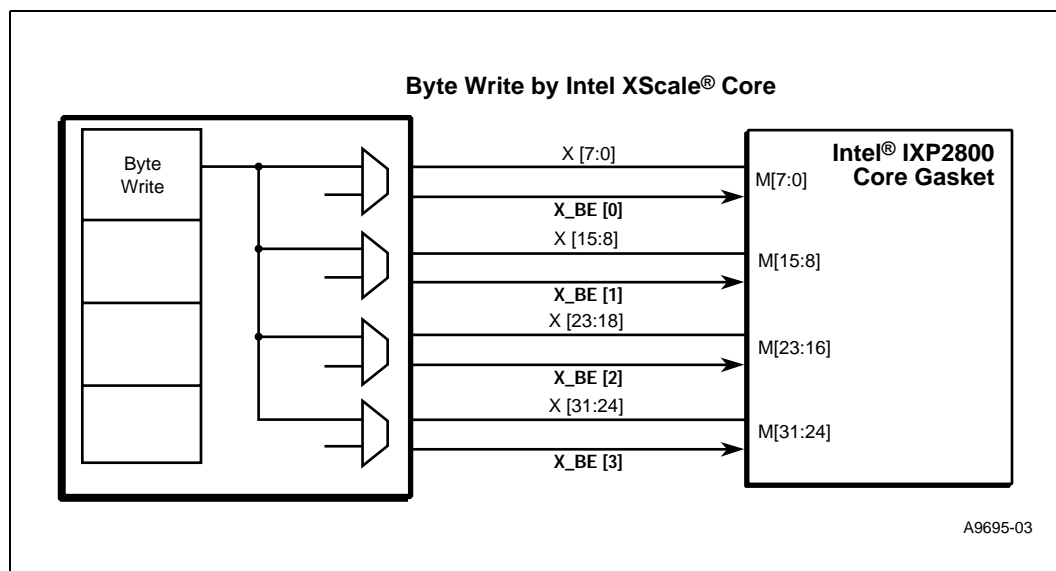
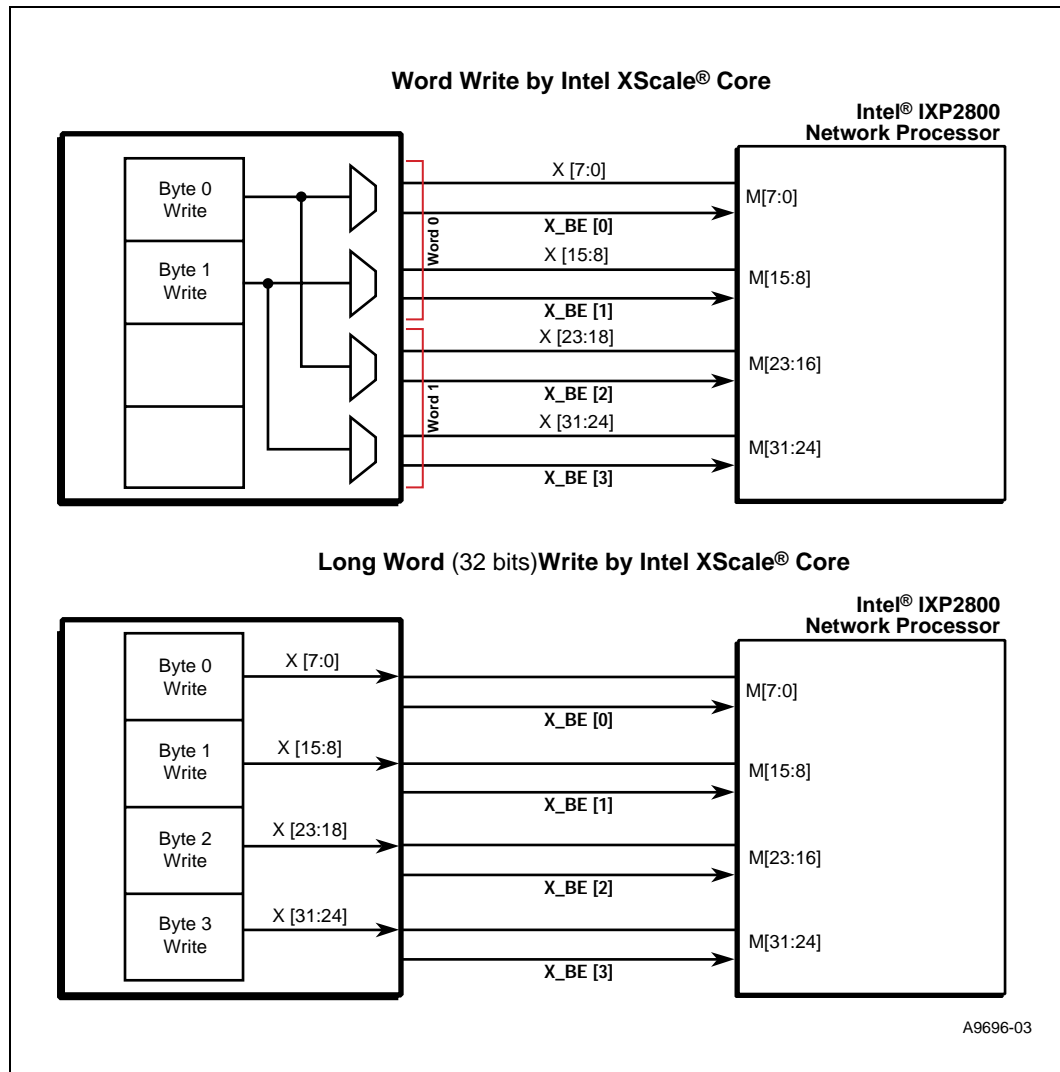


Figure 26. Intel XScale® Core-Initiated Write to the IXP2800 Network Processor (Continued)



3.11 Intel XScale® Core Gasket Unit

3.11.1 Overview

The Intel XScale® core uses the Core Memory Bus (CMB) to communicate with the functional blocks. The rest of the IXP2800 Network Processor functional blocks use the Command Push Pull (CPP) as the global bus to pass data. Therefore, the gasket is needed to translate Core Memory Bus commands to Command Push Pull commands.

This gasket has a set of local CSRs, including interrupt registers. These registers can be accessed by the Intel XScale® core via the gasket internal bus. The CSR Access Proxy (CAP) is only allowed to do a set on these interrupt registers.

The Intel XScale® core coprocessor bus is not used in the IXP2800 Network Processors, therefore all accesses are only through the Command Memory Bus.

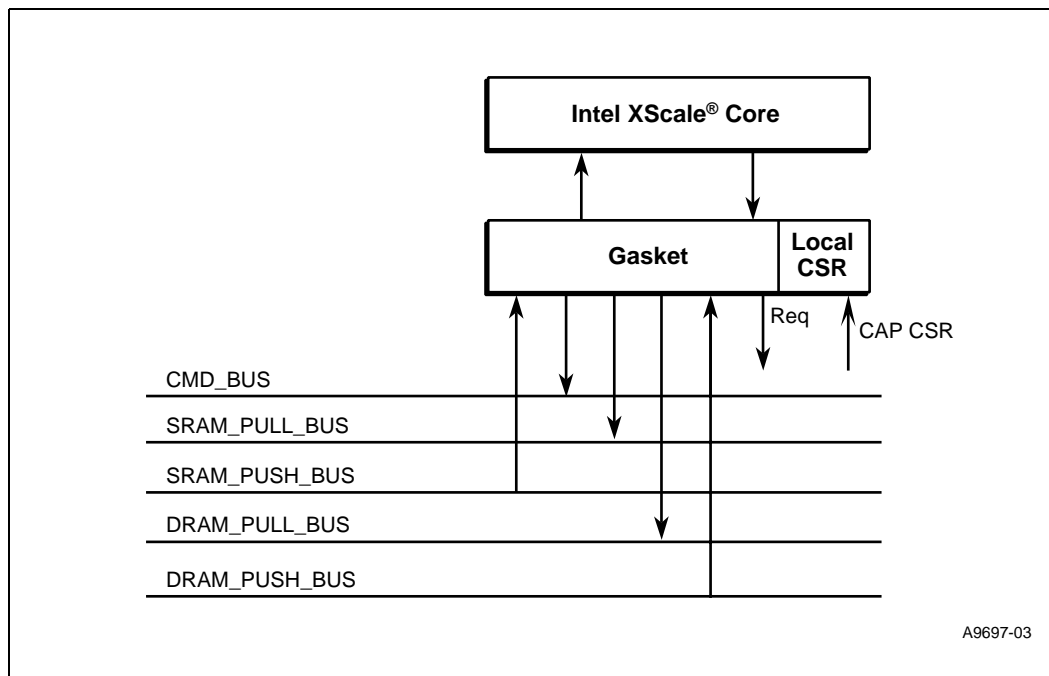
Figure 27 shows the block diagram of the global bus connections to the gasket.

The gasket unit has the following features:

- Interrupts are sent to the Intel XScale® core via the gasket, with the interrupt controller registers used for masking the interrupts.
- The gasket converts CMB reads and writes to CPP format.
- All the atomic operations are applied on SRAM and SCRATCH only, not DRAM.
- There is a stepping-stone sitting between the Intel XScale® core and the gasket. The Intel XScale® core runs at 600 – 700 MHz. The gasket currently supports a 1:1 (IXP2800 Network Processor) clock ratio. For a 2:1 ratio, the Command Push Pull bus will be running at ½ of the frequency of the Intel XScale® core.
- In IXP2800 memory controllers, read after write ordering is enforced. There is no write after read enforcement for the Intel XScale® core. The gasket will perform enforcement by employing Content Addressable Memory (CAM) to detect a write to an address with read pending. This only applies for writes to SRAM.
- The gasket CPP interface contains one command bus, one D_Push bus, one D_Pull bus, one S_Push bus, and one S_Pull bus, each with a 32-bit data width.

A maximum four outstanding reads and four outstanding writes from the Intel XScale® core are allowed.

Figure 27. Global Buses Connection to the Intel XScale® Core Gasket



A9697-03



3.11.2 Intel XScale® Core Gasket Functional Description

3.11.2.1 Command Memory Bus to Command Push/Pull Conversion

The primary function of the Intel XScale® core gasket unit is to translate commands initiated from the Intel XScale® core in the Intel XScale® core command bus format, into the IXP2800 internal command format (Command Push/Pull format).

Table 49 shows how many CPP commands are generated by the gasket from each CMB command. Write data is guaranteed to be 32-bit (longword) aligned. Table 49 shows only the Store command. In the Load case, the gasket simply converts it to the CPP format. No command splitting is required. A Load can only be a byte (8 bits), a word (16 bits), longword (32 bits), or eight longwords (8x32).

Table 49. CMB Write Command to CPP Command Conversion

Store Length	CPP SRAM Cmd Count	CPP DRAM Cmd Count	Remark
Byte, word, longword	1	1	SRAM uses 4-bit mask, and DRAM uses an 8-bit mask.
2 longwords	1 or 2	1 or 2	SRAM: If there is any mask bit detected as '0', two commands will be generated. DRAM: If it starts with odd word address, two commands will be generated.
3 longwords	1 or 3	2	SRAM: If there is a mask bit of '0' detected, Three SRAM commands will be generated. DRAM: always two DRAM commands.
4 longwords	1 or 4	1 or 2	SRAM: If there is a mask bit of '0' detected, four commands will be generated. DRAM: If there is a mask bit of '0' detected, two commands will be generated.
8 longwords			Not allowed in a write.

3.11.3 CAM Operation

In the SRAM controller, access ordering is guaranteed only for a read coming after a write. The gasket enforces order rules in the following two cases.

1. Write coming after a read.
2. Read-Modify-Write coming after read.

The address CAMing is on 8-word boundaries. The SRAM effective address is 28 bits. Deduct five bits (two bits for the word address and three bits for eight words), and the tag width for the CAM is 23 bits wide. The CAM only operates on SRAM accesses.



3.11.4 Atomic Operations

The Intel XScale® core has Swap (SWP) and Swap Byte (SWPB) instructions that generate an atomic read-write pair to a single address. These instructions are supported for the SRAM and Scratch space, and also to any other address space if it is done by a Read command followed by Write command.

cbiIO is asserted when a data cache request is initiated to a memory region with cacheable and bufferable bits in the translation table first-level descriptor set to 0. Also, if cbiIO is asserted during the CMB read portion of the SWP, then it also does a Read Command followed by Write Command, regardless of address. In those cases, the SWP/SWPB is atomic with respect to processes running on the Intel XScale® core, but not with respect to the Microengines.

The following summarizes the Atomic operation.

Address Space	cbiIO	Operation
SRAM/Scratch	0	RMW Command
Not SRAM/Scratch	x	Read Command followed by Write Command
Any	1	Read Command followed by Write Command

When the Intel XScale® core presents the read command portion of the SWP, it will assert the cbiLock signal. The gasket will “ack” the read and save the BufID in the push_ff. It will not arbitrate for the command bus at that time; rather it will wait for the corresponding write of the SWP (which will also have cbiLock asserted). At that time the gasket will arbitrate for the command bus to send a command with the atomic operation in the command field (the command is based on the address space chosen for the SRAM/Scratch, which has multiple aliased address ranges).

The SRAM or Scratch controller will pull the data, do the atomic read-modify-write, and then push the read data back. The gasket will use the saved BufID when returning the data to CMB.

Note: Unrelated reads, such as instruction and Page Table fetches, can come in the interval between the read-lock and write-unlock, and will be handled by the gasket. No other data reads or writes will come in that interval. Also, the Intel XScale® core will not wait for the SWP read data before presenting the write data.

The gasket uses address aliases to generate the following atomic operations.

- Bit Set
- Bit Clear
- Add
- Subtract
- Swap

For the alias address type of atomic operation, the Intel XScale® core will issue a SWP command with an alias address if it needs data return. The Intel XScale® core will use the write command with an alias address if it does not need data return.

Xscale_IF will not check the second address, as long as it detects one write after one read, both with cbiLock enabled. It will take the write address and put it in the command.



3.11.4.1 Summary of Rules for the Atomic Command Regarding I/O

The following rules summarize the Atomic command, regarding I/O.

- SWP to SRAM/Scratch and Not cbiIO, Xscale_IF generates an Atomic operation command.
- SWP to all other Addresses that are not SRAM/Scratch, will be treated as separate read and write commands. No Atomic command is generated.
- SWP to SRAM/Scratch and cbiIO, will be treated as separate read and write commands. No Atomic command is generated.

3.11.4.2 Intel XScale® Core Access to SRAM Q-Array

The Intel XScale® core can access the SRAM controllers queue function to do buffer allocation and freeing. Allocation does a SRAM dequeue (deq) operation, and freeing does a SRAM enqueue (enq) operation. Alias addresses are used as shown in Table 50 to access the freelist. Each SRAM channel supports up to 64 lists, so there are 64 addresses per channel.

Table 50. IXP2800 Network Processor SRAM Q-Array Access Alias Addresses

Channel	Address Range
0	0xCC00 0100 – 0xCC00 01FC
1	0xCC40 0100 – 0xCC40 01FC
2	0xCC80 0100 – 0xCC80 01FC
3	0xCCC0 0100 – 0xCCC0 01FC

Address 7:2 selects which Queue_Array entry within the SRAM channel is used.

Doing a *load* to an address in the table will do a *deq*, and the SRAM controller returns the dequeued information (i.e., the buffer pointer) as the load data; a *store* to an address in the table will do an *enq*, and the data to be enqueued is taken from the Intel XScale® core store data.

The gasket generates command fields as follows, based on address and cbiLd:

```

Target_ID = SRAM (00 0010)
Command = deq (1011) if cbiLd, enq (1100) if ~cbiLd
Token[1:0] = 0x0
Byte_Mask = 0xFF
Length = 0x1
Address = {XScale_Address[23:22], XScale_Address[7:2], XScale_Write_Data[25:2]}
    
```

Note: On the command bus, address[31:30] selects the SRAM channel, address[29:24] is the Q_Array number, and address[23:0] is the SRAM longword address. For Dequeue, the SRAM controller ignores address[23:0].



3.11.5 I/O Transaction

The Intel XScale® core can request an I/O transaction by asserting `xsoCBI_IO` concurrently with `xsoCBI_Req`. The value of `xsoCBI_IO` is undefined when `xsoCBI_Req` is not asserted. When the gasket sees an I/O request with `xsoCBI_IO` asserted, it will raise `xsiCBR_Ack` but will not acknowledge future requests until the IO transaction is complete. The gasket will check if all of the command FIFOs and write data FIFOs are empty or not. It will also check if the command counters (SRAM and DRAM) are equal to 0. All of these checks are to guarantee that:

- Writes are issued to the target, and targets have pulled the data.
- Pending reads have their data all back to the gasket.

When the gasket sees that all of the conditions are satisfied, it will assert `xsiCBR_SynchDone` to the Intel XScale® core. `XsiCBR_SynchDone` is one cycle long and does not need to coincide with `xsiCBR_DataValid`.

3.11.6 Hash Access

Hash accesses are accomplished by the gasket `Local_CSR` accesses from the Intel XScale® core. There are two sets of registers in the gasket that are involved in Hash accesses.

- Four 32-bit `XG_GCSR_Hash[3:0]` registers for holding the data to be hashed and index returned as well.
- A `XG_GCSR_CTR0(valid)` register to hold the status of the Hash Access.

The procedure for the Intel XScale® core to setup a Hash access is as follows.

1. The Intel XScale® core writes data to `XG_GCSR_Hash` by `Local_CSR` access, using address `[X:yy:zz]`. `X` selects Hash register set, `yy` selects `hash_48`, `hash_64`, or `hash_128` mode, and `zz` selects one of four `Hash_Data` registers.
2. The data write order is 3-2-1-0 (for `hash_128`) and 1-0 (for `hash_48` or `hash_64`). When the data write to `Hash_Data[0]` is performed, it triggers the Hash request to go out on the CPP bus. At the same time, `XG_GCSR_Hash(valid)` is cleared by hardware.
3. The Intel XScale® core starts to poll `Hash_Result_Valid` periodically by `Local_CSR` read.
4. After a period of time, the `Hash_Result` is returned to `XG_GCSR_Hash`, and `XG_GCSR_CTR0(valid)` is set to indicate that `Hash_Result` is ready to be retrieved.
5. The Intel XScale® core issues a `Local_CSR` read to read back the `Hash_Result`.

Note: Each Hash command requests only one index returned.

The Hash CSR is in the gasket local CSR space. See [Section 3.11.7](#).



3.11.7 Gasket Local CSR

There are two sets of Control and Status registers residing in the gasket Local CSR space. ICSR refers to the Interrupt CSR. The ICSR address range is 0xd600_0000 – 0xd6ff_ffff. The Gasket CSR (GCSR) refers to the Hash CSRs and debug CSR. It has a range of 0xd700_0000 – 0xd7ff_ffff. GCSR is shown in [Table 51](#).

Note: The Gasket registers are defined in the *IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

Table 51. GCSR Address Map (0xd700 0000)

Bits	Name	R/W	Description	Address Offset
[31:0]	XG_GCSR_HASH0	R/W	Hash word 0 Write from Intel XScale® core. Rd/Wr from CPP.	0x00—for 48-bit Hash 0x10—for 64-bit Hash 0x20—for 128-bit Hash
[31:0]	XG_GCSR_HASH1	R/W	Hash word 1 Write from Intel XScale® core. Rd/Wr from CPP.	0x04—for 48-bit Hash 0x14—for 64-bit Hash 0x24—for 128-bit Hash
[31:0]	XG_GCSR_HASH2	R/W	Hash word 2 Write from Intel XScale® core. Rd/Wr from CPP.	0x28—for 128-bit Hash
[31:0]	XG_GCSR_HASH3	R/W	Hash word 3 Write from Intel XScale® core. Rd/Wr from CPP.	0x2c—for 128-bit Hash
[31:0]	XG_GCSR_CTRL0	R	[31:1] reserved. [0] hash valid flag. Read from Intel XScale® core. Set by LCSR control.	0x30
[31:0]	XG_GCSR_CTRL1	R/W	[31:1] reserved. [0] Break_Function When set to 1, the debug break signal is used to stop the clocks. When set to 0, the debug break signal is used to cause an Intel XScale® core debug breakpoint	0x3c

3.11.8 Interrupt

The Intel XScale® core CSR controller contains local CSR(s) and interrupts inputs from multiple sources. The diagram in Figure 28 shows the flow through the controller.

Within the Interrupt/CSR register block there are raw status registers, enable registers, and local CSR(s). The raw status registers are the un-masked interrupt status. These interrupt status are masked or steered to the Intel XScale® core's IRQ or FIQ inputs by multiple levels of enable registers.

Refer to Figure 29.

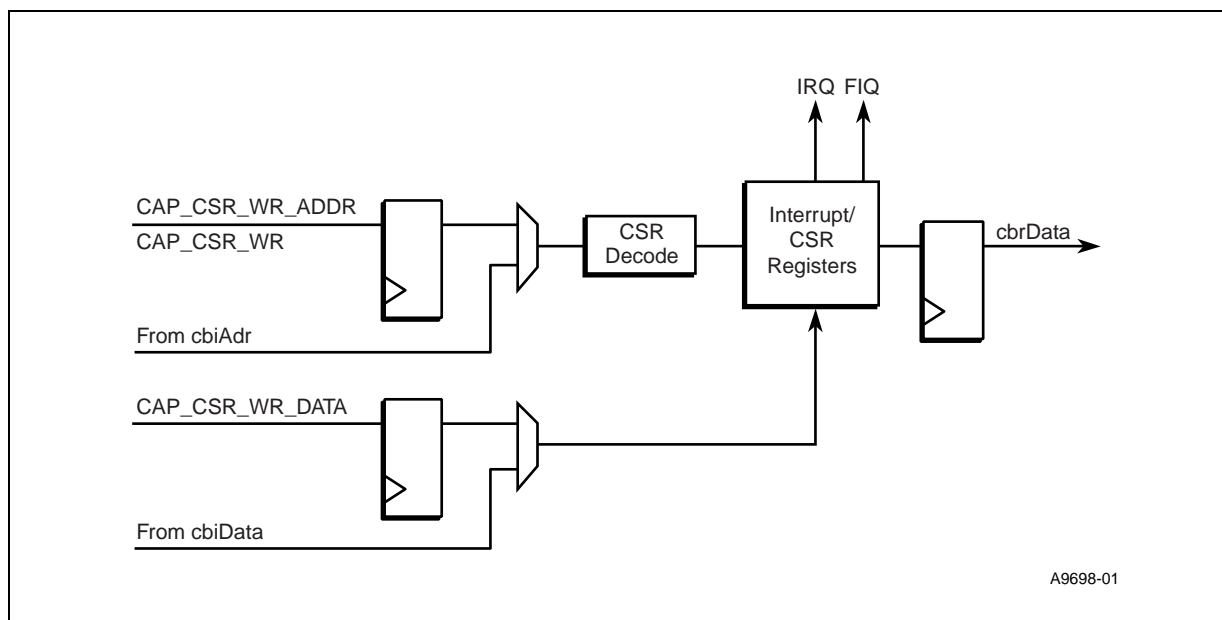
- $\{IRQ,FIQ\}Status = (RawStatus \& \{IRQ,FIQ\}Enable)$
- $\{IRQ,FIQ\}ErrorStatus = (ErrorRawStatus \& \{IRQ,FIQ\}ErrorEnable)$
- $\{IRQ,FIQ\}ThreadStatus_ \$ \# = (\{IRQ,FIQ\}ThreadRawStatus_ \$ \# \& \{IRQ,FIQ\}ThreadEnable_ \$ \#)$

Each interrupt input is visible in the RawStatus register and is masked or steered by two level of interrupt enable registers. The error and thread status are masked by one level of enable registers. Their combination along with other interrupt sources contributes to the RawStatusReg. The RawStatus is masked via IRQEnable/FIQEnable to trigger the IRQ and FIQ interrupt to the Intel XScale® core.

The enable register's bits are set and cleared through EnableSet and EnableClear registers. The Status, RawStatus, and Enable registers are read-only, and EnableSet and EnableClear are write-only. Also, Enable and EnableSet share the same address for reads and writes respectively.

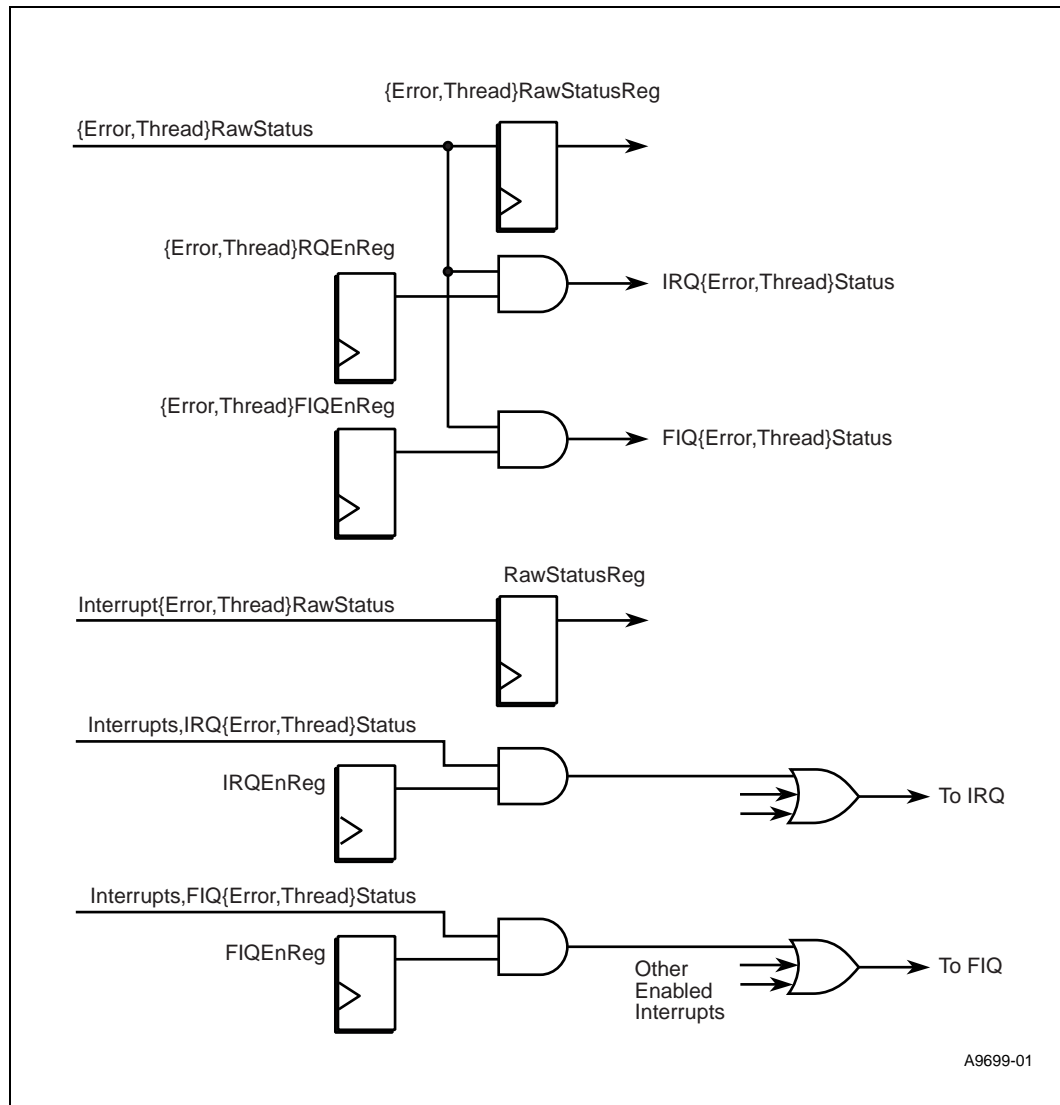
Note that software needs to take into account the delay between the clearing of an interrupt condition and having its status updated in the RawStatus registers. Also in the case of simultaneous writes to the same registers, the value of the last write is recorded.

Figure 28. Flow Through the Intel XScale® Core Interrupt Controller



A9698-01

Figure 29. Interrupt Mask Block Diagram





3.12 Intel XScale® Core Peripheral Interface

This section describes the Intel XScale® core Peripheral Interface unit (XPI). The XPI is the block that connects to all the slow and serial interfaces that communicate with the Intel XScale® core through the APB. These can also be accessed by the Microengines and PCI unit.

This section does not describe the Intel XScale® core interface protocol, only how to interface with the peripheral devices connected to the core. The I/O units described are:

- UART
- Watchdog timers
- GPIO
- Slowport

All the peripheral units are memory mapped from the Intel XScale® core point of view.

3.12.1 XPI Overview

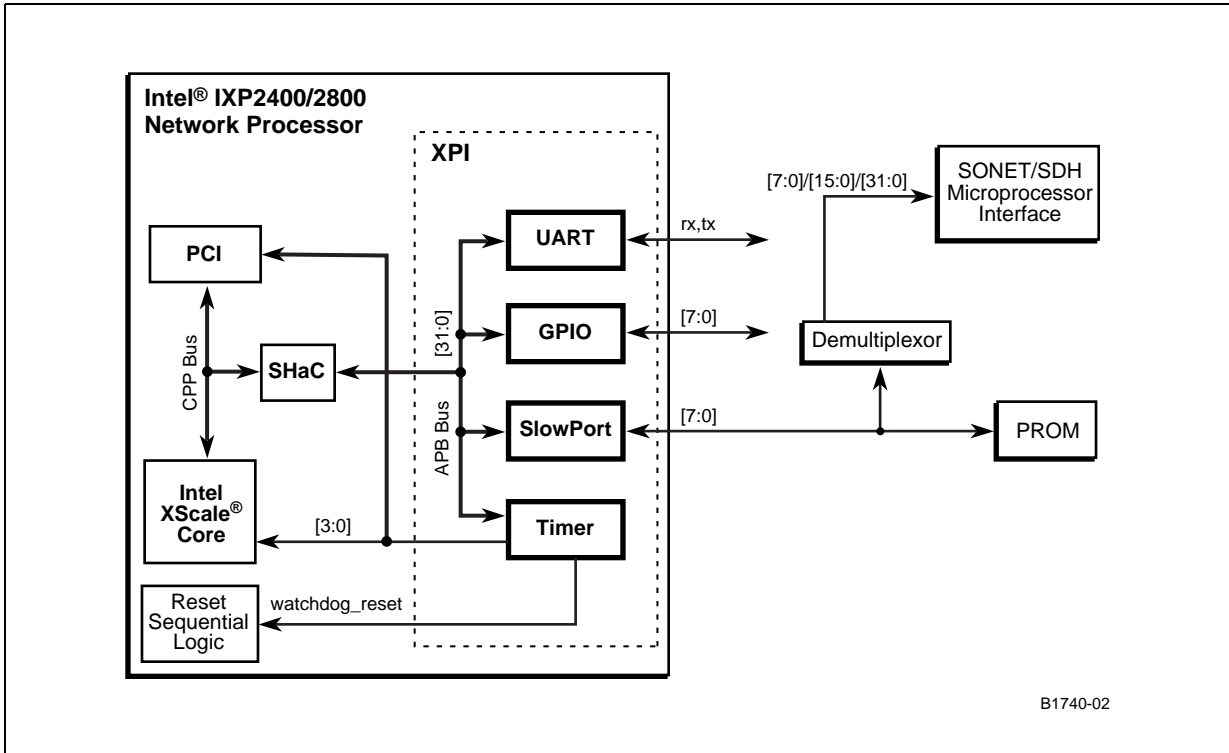
Figure 30 shows the XPI location in the IXP2800 Network Processor. The XPI receives read and write commands from the Command Push Pull bus to addresses the memory has mapped to I/O devices.

The SHaC (Scratchpad, Hash Unit, and CSRs) acts like a bridge to control the access from the Intel XScale® core or other host (like the PCI Unit). The extended APB is used to communicate between the XPI and the SHaC. The extended APB has only one signal, APB_RDY, added. This signal is used to tell the SHaC when the transaction should be terminated.

The XPI is responsible for passing the data between the extended APB and the internal blocks, like the UART, GPIO, Timer, and Slowport, which will in turn pass these data to an external peripheral device with a corresponding bus format.

The XPI is always a master on the Slowport bus and all the Slowport devices act like slaves. On the other side, the SHaC is always the master and the XPI is the slave with respect to the APB.

Figure 30. XPI Interfaces for IXP2800 Network Processor



3.12.1.1 Data Transfers

The current rate for data transfers is four bytes, except for the Slowport. The 8-bit and 16-bit accesses are only available in the Slowport bus. The devices connected to the Slowport dictate this data width. The user has to configure the data width register resident in the Slowport to run a different type of data transaction. There is no burst to Slowport.

3.12.1.2 Data Alignment

For all the CSR accesses, a 32-bit data bus is assumed. Therefore, the lower two bits of the address bus are ignored.

For the Slowport accesses, 8-, 16-, or 32-bit data access is dictated by the external device connected to the Slowport. The APB Bus should be able to match the data width according to which devices it is talking to.

See [Table 52](#) for additional details on data alignment.



Table 52. Data Transaction Alignment

Interface Units	APB Bus	Read	Write
GRegs	32 bits	32 bits	32 bits
UART	32 bits	32 bits	32 bits
GPIO	32 bits	32 bits	32 bits
Timer	32 bits	32 bits	32 bits
Slowport Microprocessor Access	8 bits	8 bits	8 bits
	16 bits	16 bits	16 bits
	32 bits	32 bits	32 bits
Slowport Flash Memory Access ¹	32 bits for 32-bit read mode, 8 bits for register read mode; 8 bits for write;	Assemble 8 bits into 32-bit data for 32-bit read mode; 8 bits for register read mode (8-bit read mode).	8 bits
CSR Access	32 bits	32 bits	32 bits

1. The flash memory interface only supports 8-bit wide flash devices. APB write transactions are assumed to be 8-bits wide, and correspond to one write cycle at the flash interface. APB read transactions are assumed to be 32-bits wide, and correspond to four flash read cycles for the 32-bit read mode set in the SP_FRM register. However, for the flash register read mode (8-bit read mode), it only needs one flash read cycle of 8-bit data and passes it back to APB directly. By default, the 32-bit read mode is set. It is advisable to stay in this mode most of the time and not change them dynamically during accesses.

3.12.1.3 Address Spaces for XPI Internal Devices

Table 53 shows the address space assignment for XPI devices.

Table 53. Address Spaces for XPI Internal Devices

Units	Starting Address	Ending Address
GPIO	0xC0010000	0xC0010040
TIMER	0xC0020000	0xC0020040
UART	0xC0030000	0xC003001C
PMU	0xC0050000	0xC0050E00
Slowport CSR	0xC0080000	0xC0080028
Slowport Device	0xC4000000	0xC7FFFFFF



3.12.2 UART Overview

The UART performs serial-to-parallel conversion on data characters received from a peripheral device and parallel-to-serial conversion on data characters received from the network processor. The processor can read the complete status of the UART at any time during the functional operation. Available status information includes the type and condition of the transfer operations being performed by the UART and any error conditions (parity, overrun, framing, or break interrupt).

The serial ports can operate in either FIFO or non-FIFO mode. In FIFO mode, a 64-byte transmit FIFO holds data from the processor to be transmitted on the serial link and a 64-byte receive FIFO buffers data from the serial link until read by the processor.

The UART includes a programmable baud rate generator that is capable of dividing the clock input by divisors of 1 to $2^{16} - 1$ and produces a 16X clock to drive the internal transmitter logic. It also drives the receive logic. The UART has a processor interrupt system. The UART can be operated in polled or in interrupt driven mode as selected by software.

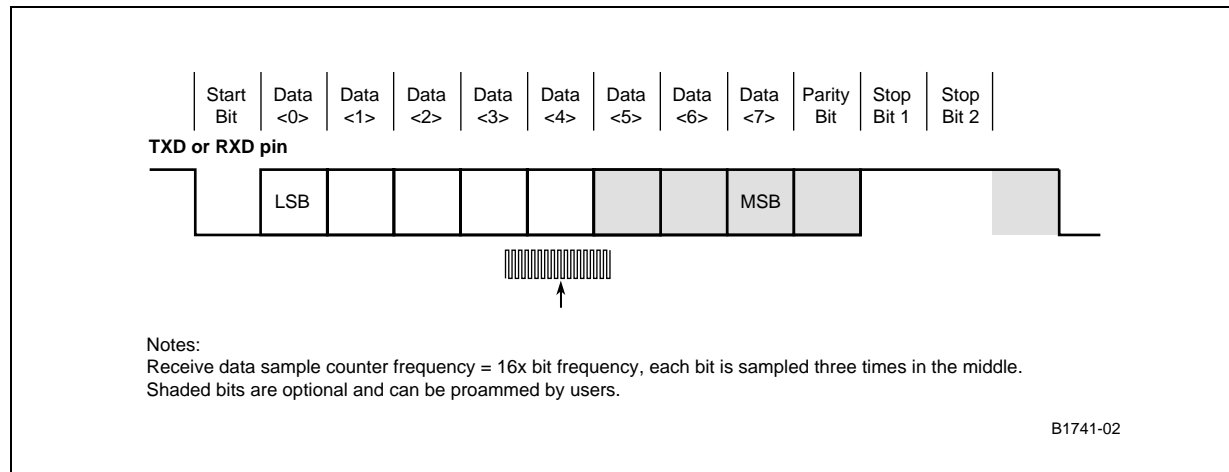
The UART has the following features

- Functionally compatible with National Semiconductor*'s PC16550D for basic receive and transmit.
- Adds or deletes standard asynchronous communications bits (start, stop, and parity) to or from the serial data
- Independently controlled transmit, receive, line status
- Programmable baud rate generator allows division of clock by 1 to $(2^{16} - 1)$ and generates an internal 16X clock
- 5-, 6-, 7-, or 8-bit characters
- Even, odd, or no parity detection
- 1, 1½, or 2 stop bit generation
- Baud rate generation
- False start bit detection
- 64-byte Transmit FIFO
- 64-byte Receive FIFO
- Complete status reporting capability
- Internal diagnostic capabilities include:
 - Break
 - Parity
 - Overrun
 - Framing error simulation
- Fully prioritized interrupt system controls

3.12.3 UART Operation

The format of a UART data frame is shown in Figure 31.

Figure 31. UART Data Frame



Each data frame is between 7 bits and 12 bits long depending on the size of data programmed, if parity is enabled and if two stop bits is selected. The frame begins with a start bit that is represented by a high to low transition. Next, either 5 to 8 bits of data are transmitted, beginning with the least significant bit. An optional parity bit follows, which is set if even parity is enabled and an odd number of ones exist within the data byte, or if odd parity is enabled and the data byte contains an even number of ones. The data frame ends with one, one and a half, or two stop bits as programmed by the user, which is represented by one or two successive bit periods of a logic one.

3.12.3.1 UART FIFO OPERATION

The UART has one transmit FIFO and one receive FIFO. The transmit FIFO is 64 bytes deep and eight bits wide. The receive FIFO is 64 bytes deep and 11 bits wide.

3.12.3.1.1 UART FIFO Interrupt Mode Operation – Receiver Interrupt

When the Receive FIFO and receiver interrupts are enabled (UART_FCR[0]=1 and UART_IER[0]=1), receiver interrupts occur as follows:

- The receive data available interrupt is invoked when the FIFO has reached its programmed trigger level. The interrupt is cleared when the FIFO drops below the programmed trigger level.
- The UART_IIR receive data available indication also occurs when the FIFO trigger level is reached, and like the interrupt, the bits are cleared when the FIFO drops below the trigger level.
- The receiver line status interrupt (UART_IIR = C6H), as before, has the highest priority. The receiver data available interrupt (UART_IIR=C4H) is lower. The line status interrupt occurs only when the character at the top of the FIFO has errors.
- The data ready bit (DR in UART_LSR register) is set to 1 as soon as a character is transferred from the shift register to the Receive FIFO. This bit is reset to 0 when the FIFO is empty.

Character Time-out Interrupt

When the receiver FIFO and receiver time-out interrupt are enabled, a character time-out interrupt occurs when all of the following conditions exist:

- At least one character is in the FIFO.
- The last received character was longer than four continuous character times ago (if two stop bits are programmed the second one is included in this time delay).
- The most recent processor read of the FIFO was longer than four continuous character times ago.

The maximum time between a received character and a time-out interrupt is 160 ms at 300 baud with a 12-bit receive character (i.e., 1 start, 8 data, 1 parity, and 2 stop bits).

When a time-out interrupt occurs, it is cleared and the timer is reset when the processor reads one character from the receiver FIFO. If a time-out interrupt has not occurred, the time-out timer is reset after a new character is received or after the processor reads the receiver FIFO.

Transmit Interrupt

When the transmitter FIFO and transmitter interrupt are enabled (UART_FCR[0]=1, UART_IER[1]=1), transmit interrupts occur as follows:

- The Transmit Data Request interrupt occurs when the transmit FIFO is half empty or more than half empty. The interrupt is cleared as soon as the Transmit Holding register is written (1 to 64 characters may be written to the transmit FIFO while servicing the interrupt) or the IIR is read.

3.12.3.1.2 FIFO Polled Mode Operation

With the FIFOs enabled (TRFIFOE bit of UART_FCR set to 1), setting UART_IER[4:0] to all 0s puts the serial port in the FIFO polled mode of operation. Since the receiver and the transmitter are controlled separately, either one or both can be in the polled mode of operation. In this mode, software checks receiver and transmitter status via the UART_LSR. As stated in the register description:

- UART_LSR[0] is set as long as there is one byte in the receiver FIFO.
- UART_LSR[1] through UART_LSR[4] specify which error(s) has occurred for the character at the top of the FIFO. Character error status is handled the same way as interrupt mode. The UART_IIR is not affected since UART_IER[2] = 0.
- UART_LSR[5] indicates when the transmitter FIFO needs data.
- UART_LSR[6] indicates that both the transmitter FIFO and shift register are empty.
- UART_LSR[7] indicates whether there are any errors in the receiver FIFO.

3.12.4 Baud Rate Generator

The baud rate generator is a programmable block and generates a clock used in the transmit block. The output frequency of the baud rate generator is 16X the baud rate; baud rate is calculated as:

$$\text{BaudRate} = \text{APB Clock} / (16 \times \text{Divisor})$$

The Divisor ranges from 1 to $2^{16} - 1$. For example, for an APB clock of 1 MHz and a baud rate of 300 bps, the divisor is 209.

3.12.5 General Purpose I/O (GPIO)

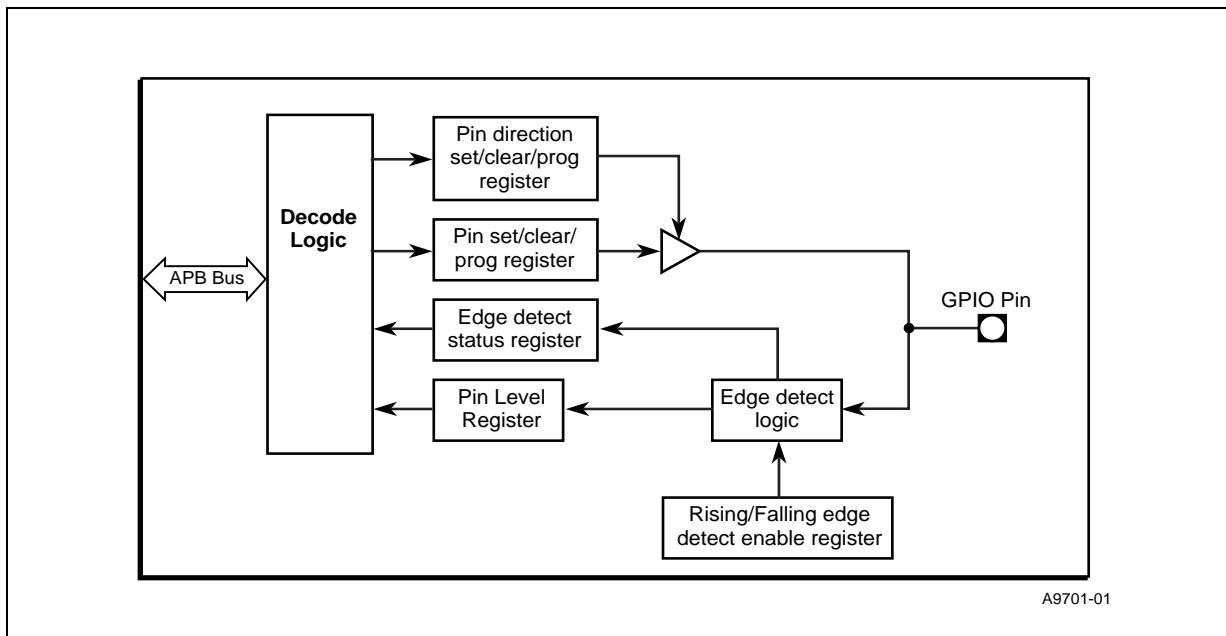
The IXP2800 Network Processor has eight General Purpose Input/Output (GPIO) port pins for use in generating and capturing application-specific input and output signals. Each pin is programmable as an input or output or as an interrupt signal sourcing from an external device. The GPIO can be used with appropriate software in I2C application.

Each GPIO pin can be configured as an input or an output by programming the corresponding GPIO pin direction register. When programmed as an input, the current state of the GPIO can be read through the corresponding GPIO pin level register. The register can be read at any time and can be used to confirm the state of the pin when it is configured as an output. In addition, each GPIO pin can be programmed to detect a rising or a falling edge by setting the corresponding GPIO rising/falling edge detect registers.

When configured as an output, the pin can be controlled by writing to the GPIO set register to write a 1 and by writing to the GPIO clear register to write a 0. These registers can be written regardless of whether the pin is configured as an input or an output.

Each of the GPIO pins is designed the same and instantiated to the number of GPIO port pins. [Figure 32](#) shows a GPIO functional diagram. The GPIO pin as seen can be programmed based on the configuration registers.

Figure 32. GPIO Functional Diagram





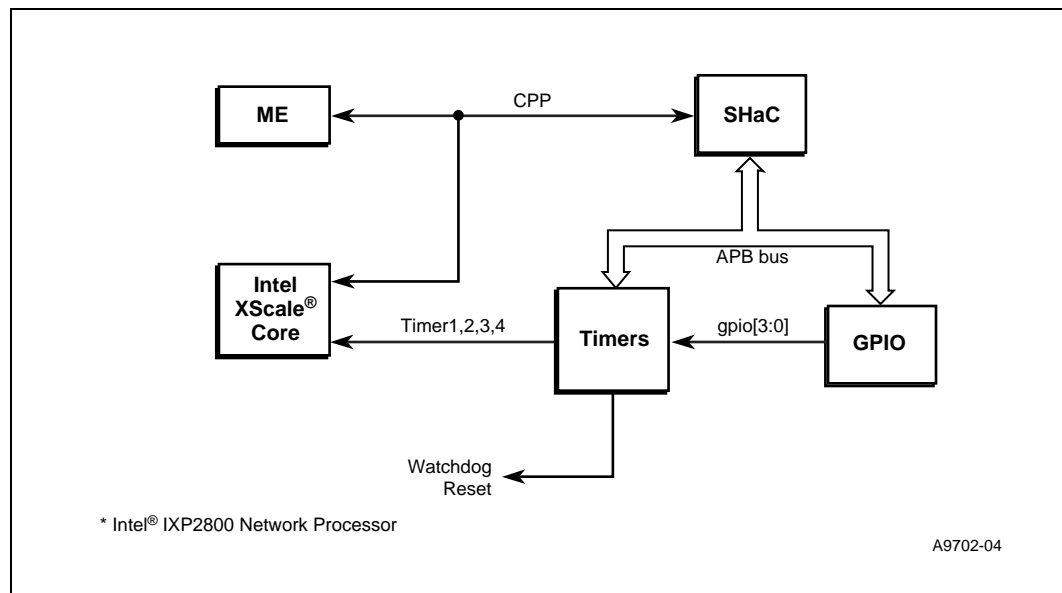
3.12.6 Timers

The IXP2800 Network Processor supports four timers. These timers are clocked by the Advanced Peripheral/Bus Clock (APB-CLK), which runs at 50 MHz to produce the PLPL_APB_CLK, PLPL_APB_CLK/16, or PLPL_APB_CLK/256 signals. The counters are loaded with an initial value, count down to 0, and raise an interrupt (if interrupts are not masked).

In addition, timer 4 can be used as a watchdog timer when the watchdog enable bits are configured to 1. When used as a watchdog timer, and when a count of 0 is encountered, it will initiate the reset sequence.

Figure 33 shows the timer control unit interfacing with other functional blocks.

Figure 33. Timer Control Unit Interfacing Diagram



3.12.6.1 Timer Operation

Each timer consists of a 32-bit counter.

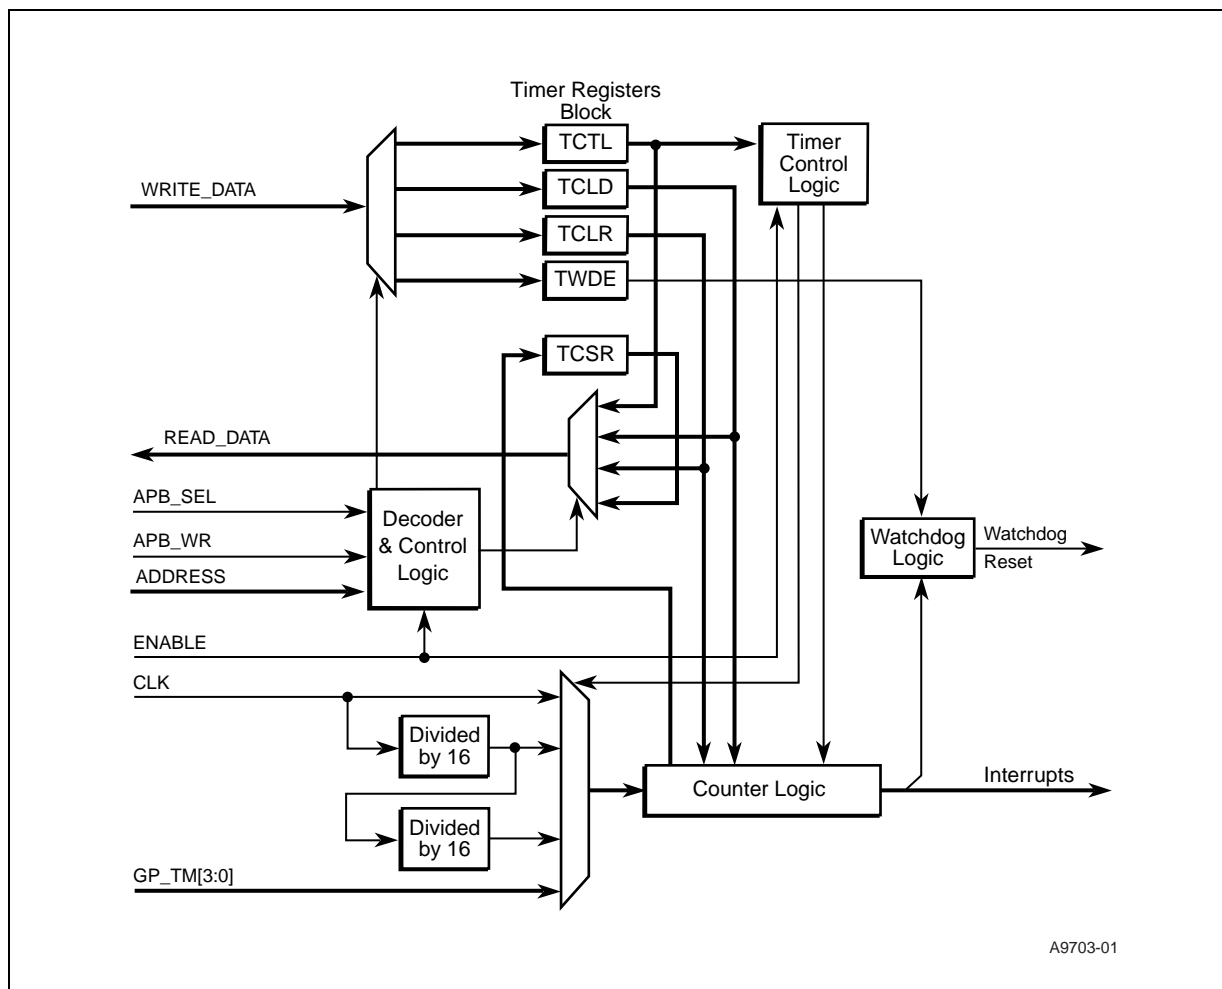
By default, the timer counter load register (TCLD) is set to 0xFFFFFFFF. The timer will count down from 0xFFFFFFFF to 0x00000000, then wrap back to 0xFFFFFFFF and continue to decrement if the TCLD is not programmed to any value. If a different value is programmed in the TCLD, then the counter will load this value every time it counts down to 0.

An interrupt is issued to the Intel XScale® core whenever the counter reaches 0. The interrupt signals can be enabled or disabled by the IRQEnable/FIQEnable registers. The interrupt remains asserted until it is cleared by writing a 1 to the corresponding timer clear register (TCLR).

The counter can be advanced by the clock, clock divided by 16, clock divided by 256, and the GPIO signals. The clock rate is controlled by the TCTL value programmed into the TCTL registers. There are four gpio signals, GPIO[3:0] that correspond to Timer 1, 2, 3, and 4, respectively. These signal are synchronized within the timer-clock domain before driving the counter.

Figure 34 shows the Timer Internal logic.

Figure 34. Timer Internal Logic Diagram



3.12.7 Slowport Unit

The IXP2800 Network Processor Slowport Unit supports basic PROM access and 8-, 16-, and 32-bit microprocessor device access. It allows a master, (Intel XScale® core or Microengine), to do a read/ write data transfer to these slave devices.

The address bus and data bus are multiplexed to reduce the pin count. In addition, the address bus is also compressed from A[25:0] down to A[7:0] and shifted out with three clock cycles. Therefore, an external set of buffers is needed for address storage and latch.

The access can be asynchronous. Insertion of delay cycles is possible for both setup and hold data. A programmable timing control mechanism is provided for this purpose. There are two types of interfaces supported in the Slowport Unit:

- Flash memory interface
- Microprocessor interface.



The Flash memory interface is used for the PROM device. The microprocessor interface can be used for SONET/SDH Framer microprocessor access.

There are two ports in the Slowport unit. The first is dedicated to the flash memory device while the second to the microprocessor device.

3.12.7.1 PROM Device Support

For all the Flash Memory access, only 8-bit devices are supported. APB write transactions are assumed to be 8-bits wide, and correspond to one write cycle at the flash interface. The extended APB read transactions are assumed to be 32-bits wide, and correspond to four read cycles at the flash memory interface for all the flash memory data read. However, for the flash register read inside the flash memory, like the flash status register, the returned data is one byte and placed in the lower order byte location. In this case, only one external transaction cycle is involved.

To accomplish this, a register (SP_FRM) is installed to allow to configure between 8-bit read mode and 32-bit read mode. By default, it goes to 32-bit read mode. For the 8-bit read mode, one read cycle is involved. No packing process is needed. The data will be directly placed onto the lower order byte, [7:0] and passed to APB. For the 32-bit read mode, it needs four read cycles. All 4 bytes are packed into a 32-bit data and passed to the APB. 16-bit mode is not supported for read.

Write always accesses the flash with one 8-bit cycle. Therefore, no unpacking process is needed. The PROM devices supported are listed in [Figure 54](#):

Table 54. 8-Bit Flash Memory Device Density

Vendor	Part Number	Size
Intel	28F128J3A	16 MB
Intel	28F640J3A	8 MB
Intel	28F320J3A	4 MB

3.12.7.2 Microprocessor Interface Support for the Framer

The Slowport Unit also supports a microprocessor interface with Framer components. Some supported devices are listed in [Table 55](#):

Table 55. SONET/SDH Devices (Sheet 1 of 2)

Vendor	Part Number	Microprocessor Interface	SP_PCR register Setting	DW Setting in SP_ADC register
PMC-Sierra*	PM3386	16 bits	0x3	0x1
PMC-Sierra*	PM5345	8 bits	0x2	0x0
PMC-Sierra*	PM5346	8 bits	0x2	0x0
PMC-Sierra*	PM5347	8 bits	0x2	0x0
PMC-Sierra*	PM5348	8 bits	0x2	0x0
PMC-Sierra*	PM5349	8 bits	0x2	0x0
PMC-Sierra*	PM5350	8 bits	0x2	0x0
PMC-Sierra*	PM5351	8 bits	0x2	0x0
PMC-Sierra*	PM5352	8 bits	0x2	0x0

Table 55. SONET/SDH Devices (Sheet 2 of 2)

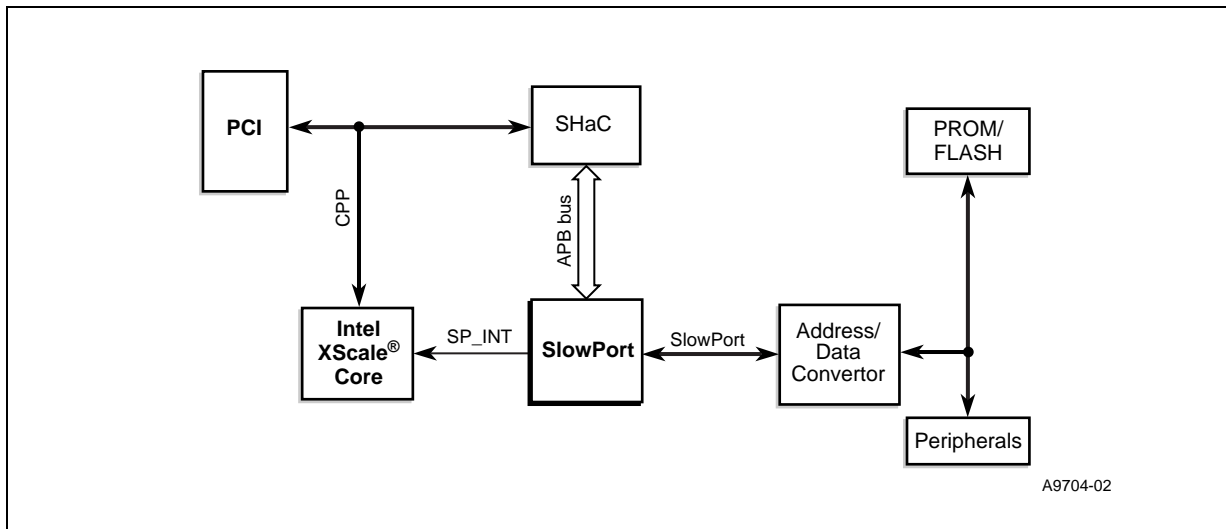
Vendor	Part Number	Microprocessor Interface	SP_PCR register Setting	DW Setting in SP_ADC register
PMC-Sierra*	PM5355	8 bits	0x2	0x0
PMC-Sierra*	PM5356	8 bits	0x2	0x0
PMC-Sierra*	PM5357	8 bits	0x2	0x0
PMC-Sierra*	PM5358	16 bits	0x2	0x1
PMC-Sierra*	PM5381	16 bits	0x2	0x1
PMC-Sierra*	PM5382	8 bits	0x2	0x0
PMC-Sierra*	PM5386	16 bits	0x2	0x1
AMCC*	S4801 (AMAZON)	8 bits	0x0	0x0
AMCC*	S4803 (YUKON)	8 bits	0x0	0x0
AMCC*	S4804 (RHINE)	8/16 bits	0x0/0x3	0x0/0x1
Intel	IXF6012 (Volga)	16 bits	0x3/0x4 ¹	0x1
Intel	IXF6048 (Amazon-A)	16 bits	0x3/0x4 ¹	0x1
Intel	Centaur	—	0x3/0x4 ¹	—
Intel	IXF6501	—	0x3/0x4 ¹	—
Intel	Ben Nevis	32 bits	0x3/0x4 ¹	0x2
Lucent*	TDAT042G5	16 bits	0x1/	0x1
Lucent*	TDAT04622	16 bits	0x1	0x1
Lucent*	TDAT021G2	16 bits	0x1	0x1

1. Usually there are two different protocols, Intel or Motorola*, of microprocessor interface in the Intel framer; the setting in the PCR should match with protocols activated in the framer.

3.12.7.3 Slowport Unit Interfaces

Figure 35 shows the Slowport unit interface diagram.

Figure 35. Slowport Unit Interface Diagram






3.12.7.4 Address Space

The total address space is defined as 64 Mbytes, which is further divided into two segments of 32 Mbytes each. Two devices can be connect to this bus. If these peripheral devices have a density of 256 Mbits (32 Mbytes) each, all the address space is going to be filled like a contiguous address space. However, if a small capacity device is used (like a device of 4, 8, or 16 MBytes), there will be a memory hole left in between these two devices. [Figure 36](#) is a 4-Mbyte memory example. Trying to read the space in between, you will get the repeating value for each 4-Mbyte location

Figure 36. Address Space Hole Diagram



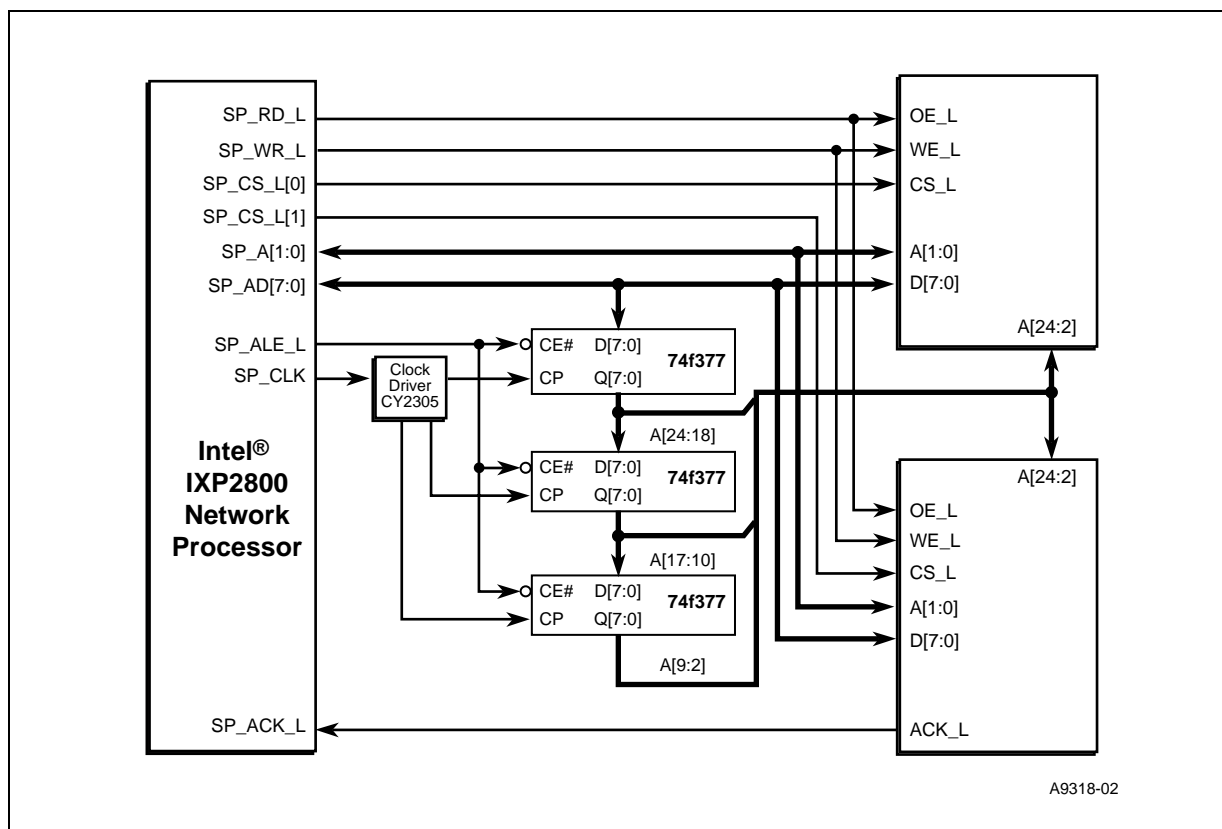
3.12.7.5 Slowport Interfacing Topology

[Figure 37](#) demonstrates one of the topologies used to connect to an 8-bit device. From the diagram, we can observe that address is shifted out eight bits at a time and buffered into three 74F377 or equivalent tri-state buffer devices in three consecutive clock cycles. These buffers also output separately to form a 25-bit wide address bus to address the 8-bit devices. The data are expected to be driven out after the address has been placed into the buffers.

There are two devices shown in [Figure 37](#). The top one is the fix-timed device, while the lower one, self-timing device. For the self-timing device, the access latency depends on the SP_ACK_L responded back from this device.

Three extra signals, SP_CP, SP_OE_L, and SP_DIR, are added to pack and unpack the data when a 16-bit or 32-bit device is hooked up to Slowport. They are used for special application only as described below.

Figure 37. Slowport Example Application Topology



3.12.7.6 Slowport 8-Bit Device Bus Protocols

The write/read transfer protocols are discussed in the following sections. The burst transfers are going to be broken down into single mode transfer. For each single write/read transaction, it can be either fixed-timed transaction or self-timing transaction. The fixed-timed transaction has the response fixed in a certain period, that can be controlled by the timing control registers.

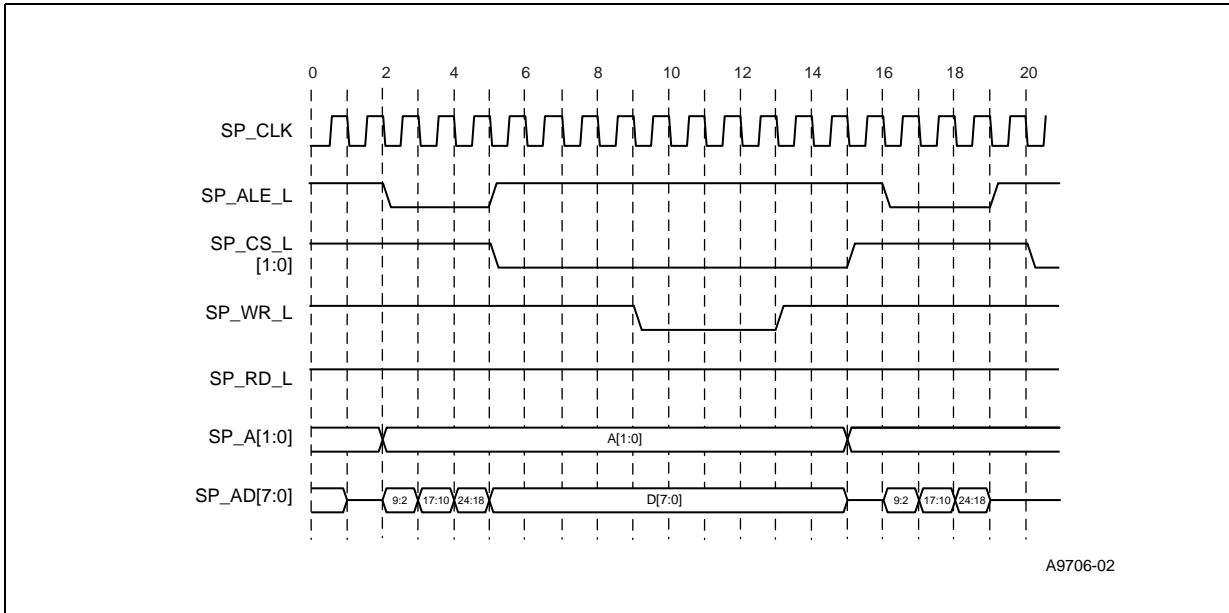
For the self-timing transaction, the response timing is dictated by the peripheral device. Hence, wait states can be inserted during the transaction. All the back-to-back transactions are intervened with one clock cycle. The Slowport clock, SP_CLK, shown in the following waveform diagrams, is generated by dividing the PLPL_APB_CLK. The divisor used is specified in the clock control register, SP_CCR.



3.12.7.6.1 Mode 0 Single Write Transfer for Fixed-Timed Device

Figure 38, shows the single write transfer for a fixed-timed device with the CSR programmed to a value of setup=4, pulse width=0, and hold=2, followed by another read transfer.

Figure 38. Mode 0 Single Write Transfer for a Fixed-Timed Device



The transaction is initiated with SP_ALE_L asserted. It latches the address from the SP_AD[7:0] bus into the external buffer, using three clock cycles. After that, it should deassert the SP_ALE_L to disable latching the address into the buffers.

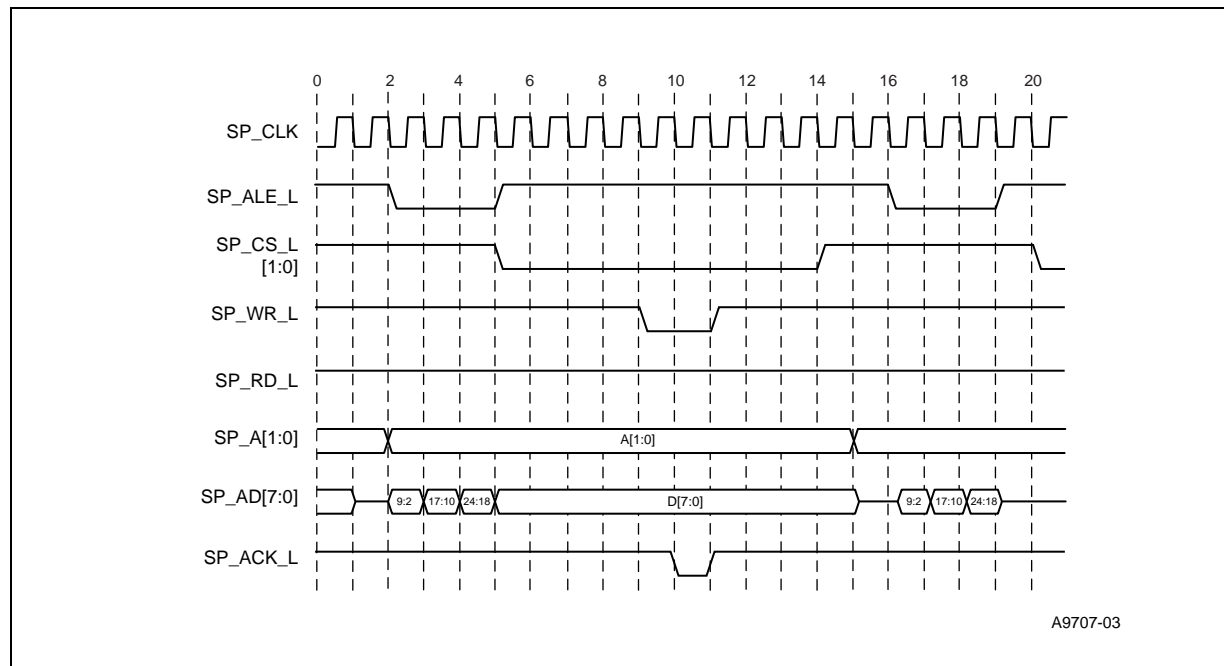
The SP_A[1:0] signals span the whole transaction cycle.

For the write, it drives the data onto the SP_AD[7:0]. Meanwhile, it asserts the SP_CS_L[1:0] signals. Depending on the timing control setup parameter, for this case, the SP_WR_L is not asserted until four clock cycles have elapsed. The SP_CS_L[1:0] signals are deasserted two clocks after the SP_WR_L is deasserted.

3.12.7.6.2 Mode 0 Single Write Transfer for Self-Timing Device

Figure 39 depicts the single write transfer for a self-timing device with the CSR programmed to setup=4, pulse width=0, and hold=3. Similarly, a read transaction is attached behind.

Figure 39. Mode 0 Single Write Transfer for a Self-Timing Device



Similar to the single write for fixed-timed device, the ALE_L, CS_L[1:0], AD[7:0], and A[1:0] follow the same pattern, and the timing is controlled by the timing control register — except for the WR_L, which is terminated depending on the SP_ACK_L returned from the self-timing device.

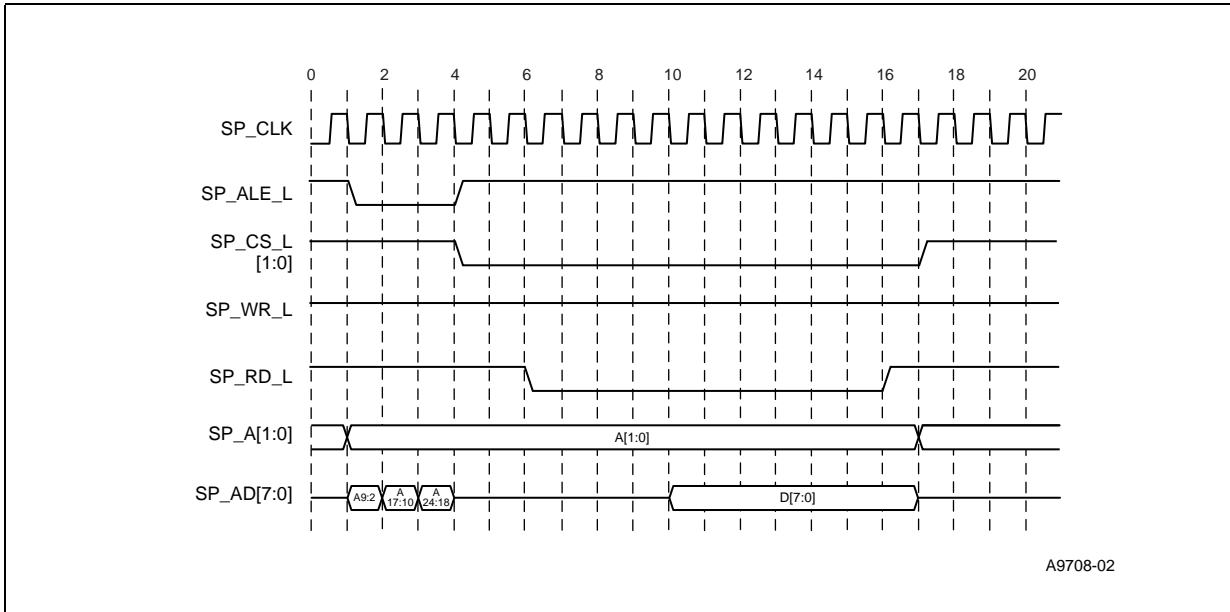
The time-out counter will be set to 255. If no SP_ACK_L responds back when the time-out counter reaches 0, the transaction is terminated with a time-out. An interrupt signal is issued to the bus master simultaneously with the time-out register update.



3.12.7.6.3 Mode 0 Single Read Transfer for Fixed-Timed Device

Figure 40 demonstrates the single read transfer issued to a fixed-timed PROM device followed by another write transaction. The CSR is assumed to be configured to the value $setup=2$, $pulse\ width=10$, and $hold=1$.

Figure 40. Mode 0 Single Read Transfer for Fixed-Timed Device

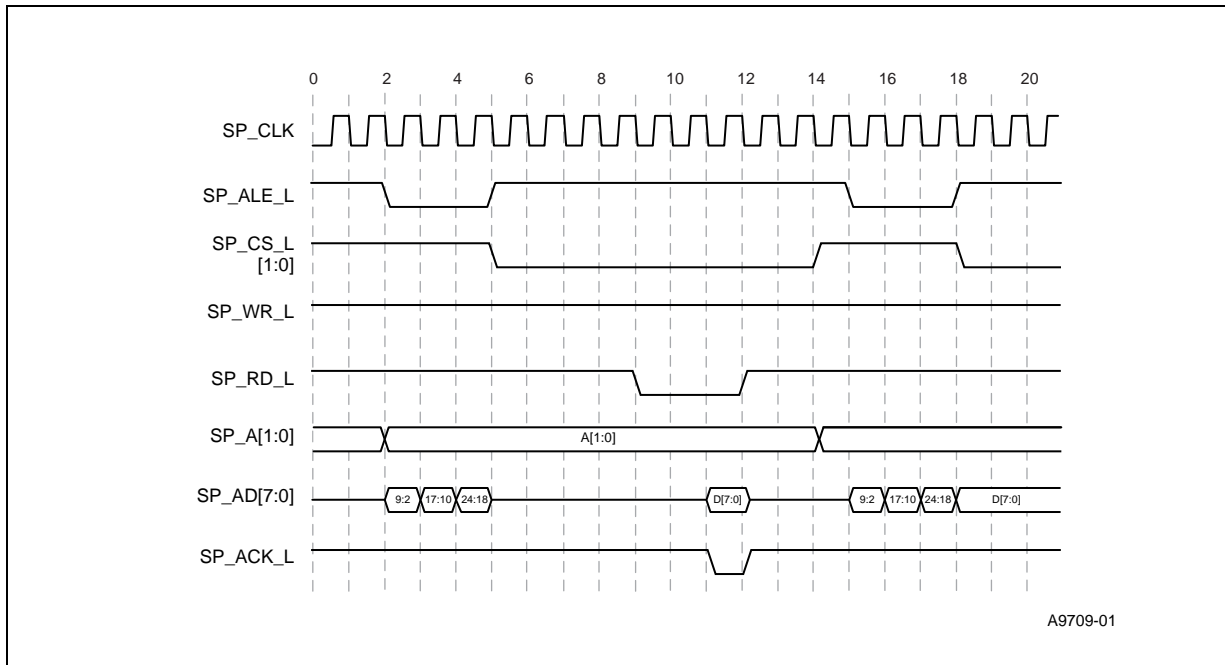


The address is loaded onto the external buffer in three clock cycles with the ALE_L asserted. Then, a clock cycle is inserted to tri-state all the AD[7:0] signals. The CS_L[1:0] signals come asserted on the fourth clock cycle. Now, the values stored in the timing control registers take effect. The RD_L is asserted after two clock cycles. It keeps asserted for ten clock cycles. The CS_L[1:0] should be de-asserted one clock cycle after RD_L is de-asserted. The data will be valid at clock cycle 16 as shown in the diagram. Since the hold delay has two cycles, and the transaction is terminated at clock cycle 16.

3.12.7.6.4 Single Read Transfer for a Self-Timing Device

Figure 41 demonstrates the single read transfer issued to a self-timing PROM device followed by another write transaction. The CSR assumed to be programmed to the value of setup=4, pulse width=0, and hold=2.

Figure 41. Mode 0 Single Read Transfer for a Self-Timing Device



The only difference for self-timed mode is in the SP_ACK_L signal. It has a dominant effect on the length of the transaction cycle or it overrides the value in the timing control register. A time-out counter is set to 256. The SP_ACK_L should arrive before the time-out counter counts down to 0. Similarly to the single write for self-timing device, an interrupt is launched for the time-out event and the time-out register is updated. In this case, the data will be sampled at clock cycle 12.

3.12.7.7 SONET/SDH Microprocessor Access Support

To support the SONET/SDH Microprocessor Interface, extra logic is added into this unit. Here we consider three SONET/SDH available components, including the Lucent* TDATA042G5, PMC-Sierra* PM5351, Intel, and AMCC* SONET/SDH devices.

However, because these microprocessor interfaces are not standardized, we treat them separately and a configuration register is installed to activate the bus to work with different interface protocol at a time. Extra pins are also added to accomplish this task.

A microprocessor interface type register is used to provide these kinds of solutions. The user is allowed to configure the interface to the following four different modes. The pin functionality and the interface protocol will be changed accordingly. By default, it activates the mode 0 with 8-bit generic PROM device support as mentioned above.



3.12.7.7.1 Mode 1: 16-Bit Microprocessor Interface Support with 16-Bit Address Lines

The address size control register is programmed to 16-bit address space for this case. This mode is designated for the devices with the similar protocol with the Lucent* TDAT042G5 SONET/SDH device.

16-Bit Microprocessor Interfacing Topology with 16-Bit address lines

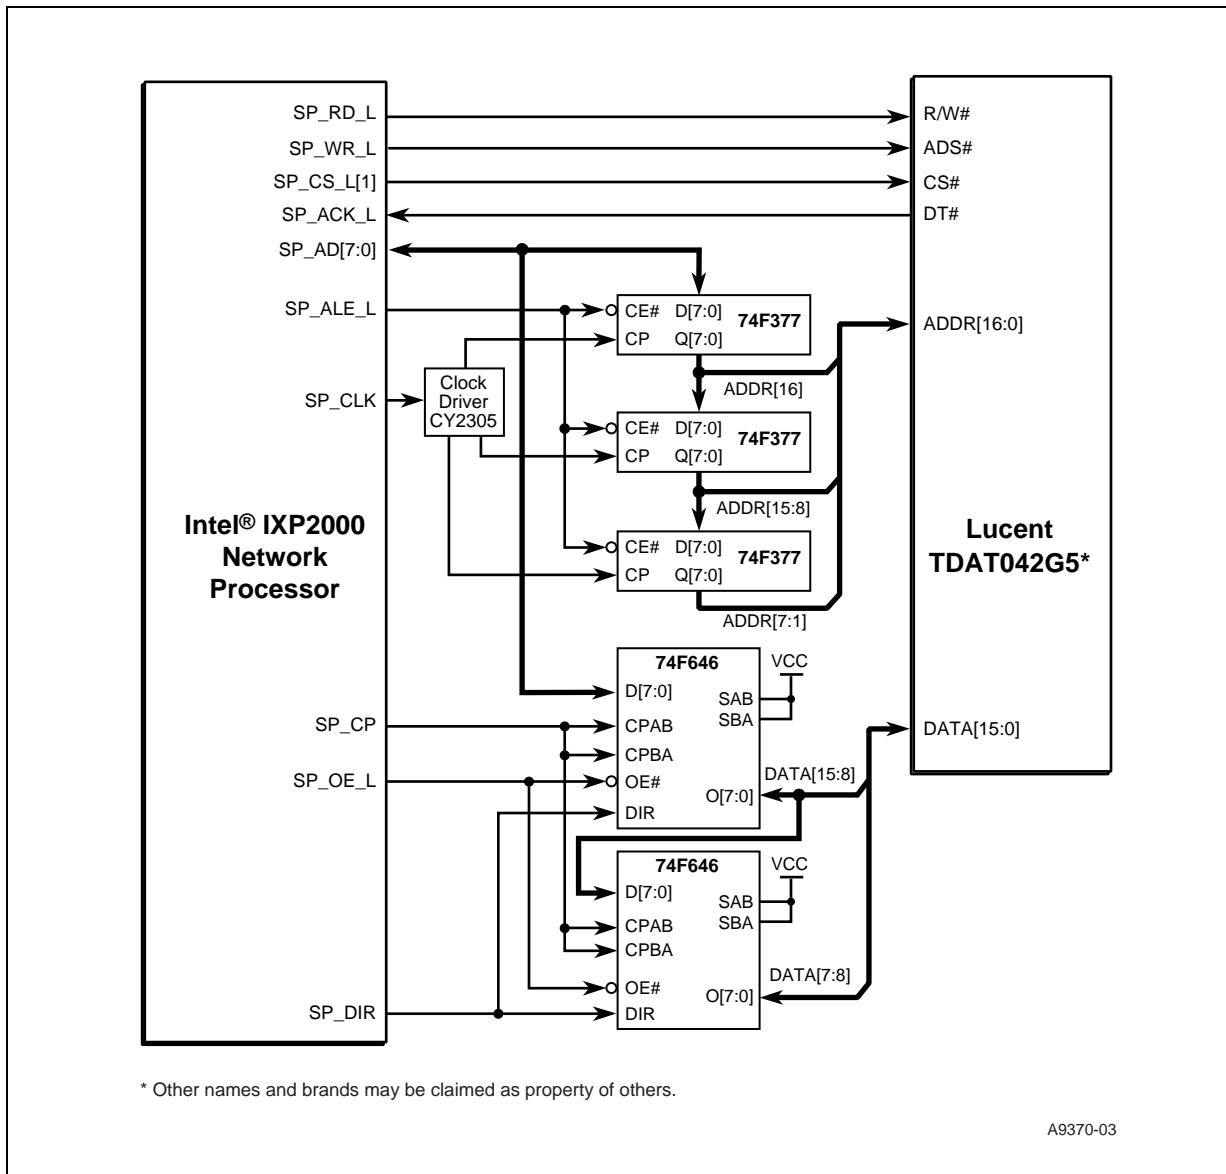
Figure 42 shows a solution for the 16-bit microprocessor interface. This solution bridges the Lucent* TDAT042G5 SONET/SDH 16-bit interface. From Figure 42, we observe that the control pins SP_RD_L and SP_WR_L are converted to R/W and ADS. The CS and DT are still compactible with SP_CS_L[1] and SP_ACK_L protocol.

Extra pins are added to accomplish the task of multiplexing and demultiplexing the data bus. The total pin count is 18.

During the write cycle, 8-bit data are stacked into 16-bit data. They are first shifted into two tri-state buffers, 74F646 or equivalent by SP_CP, using two consecutive clock cycle; then the SP_CS_L is used for output of the 16-bit data, which is shared with the CS.

During the read cycle, the 16-bit data are unpacked into 8-bit data by SP_CP. Two 74F646 or equivalent tri-state buffers are used. First, the 16-bit data are stored into these buffers. Then they are shifted out by SP_DIR, using two consecutive clock cycles.

Figure 42. An Interface Topology with Lucent® TDAT042G5 SONET/SDH

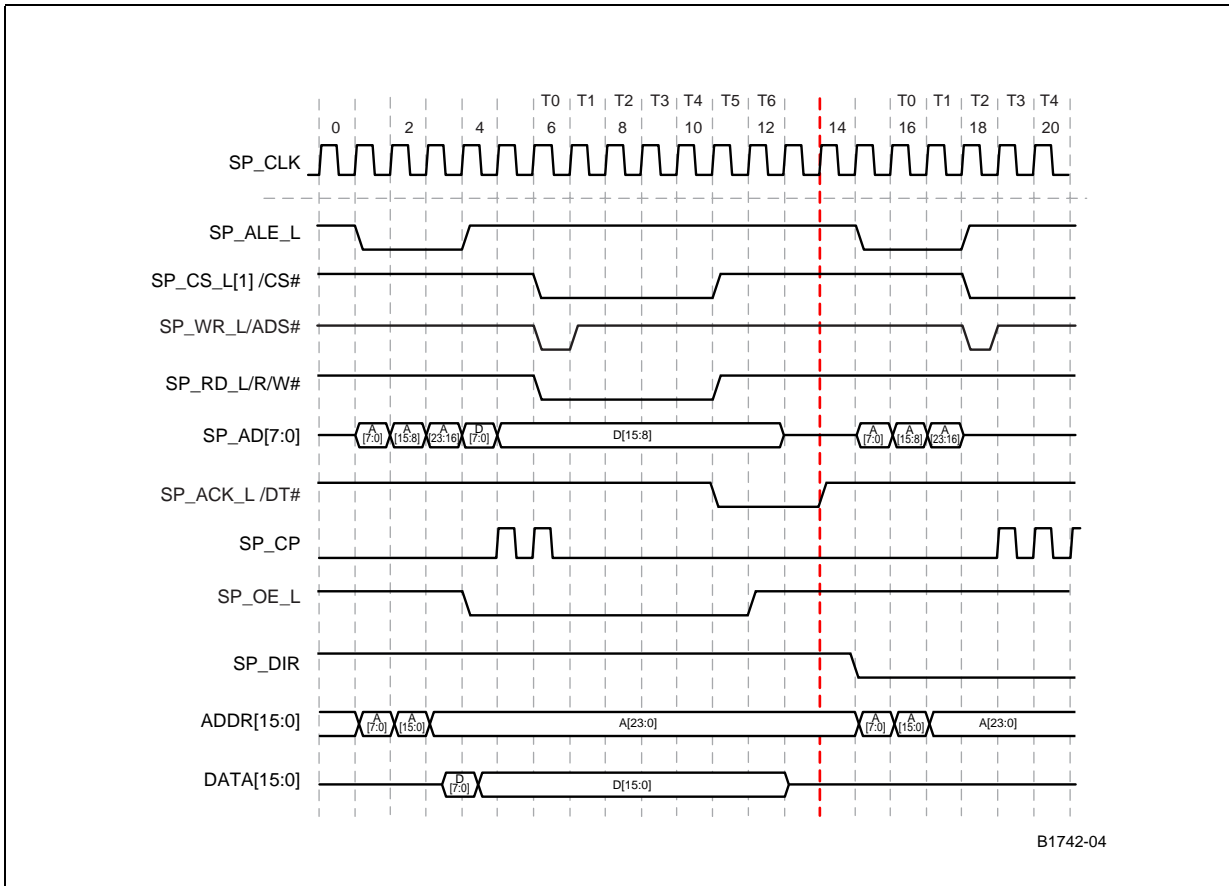


16-Bit Microprocessor Write Interface Protocol

Figure 43 uses the Lucent* TDAT042G5 device. In this case, the user should program the P_PCR register to mode 1 and also program the write timing control register to setup=7, pulse width=5, and hold=1, which represent seven clock cycles for \overline{CS} , five clock cycles for DT delay, and one clock cycle for \overline{ADS} . They are intervened with two idle cycles.

From Figure 43, we observe that there are a total of twelve clock cycles used for write access, (i.e., 240 ns), not including an intervened turnaround cycle after the write transaction. The throughput is 8.3 Mbytes per second.

Figure 43. Mode 1 Single Write Transfer for Lucent* TDAT042G5 Device (B0)



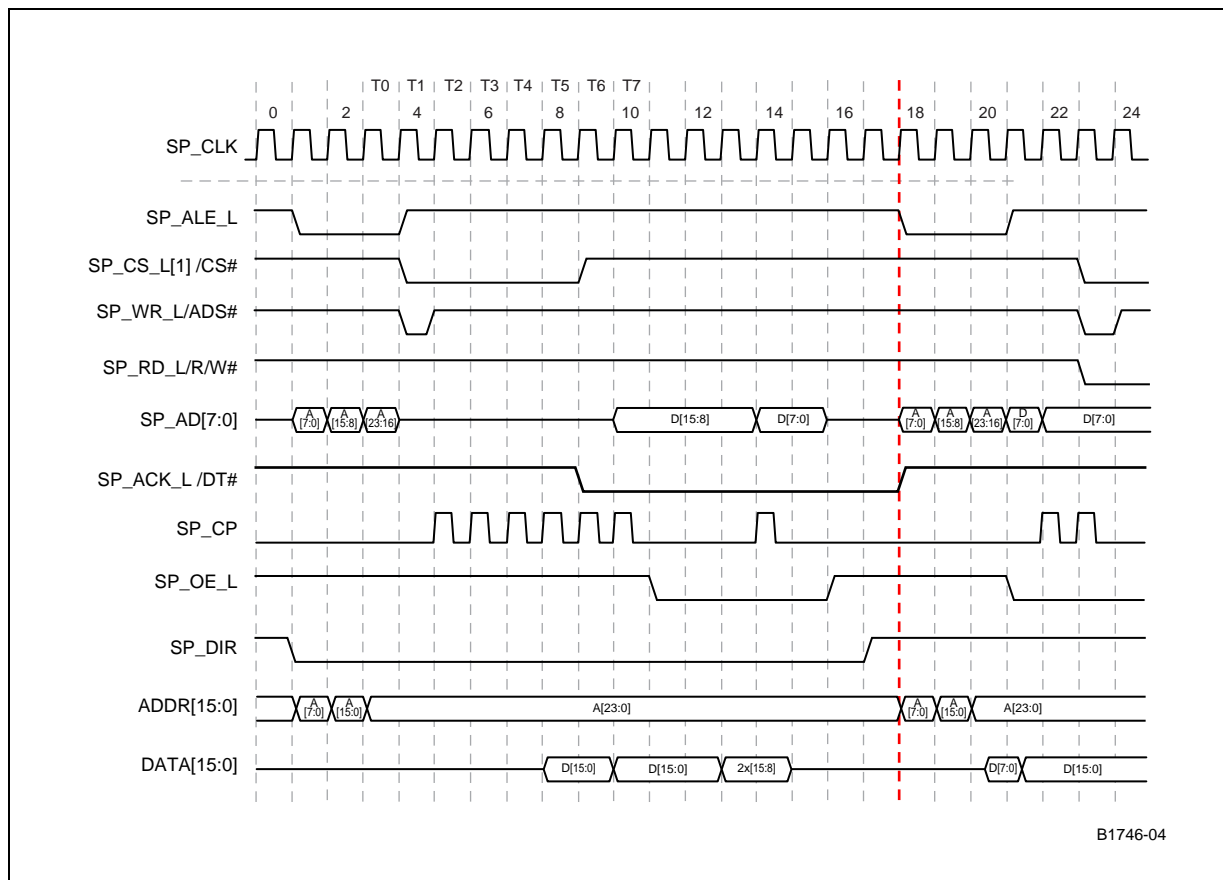


16-Bit Microprocessor Read Interface Protocol

Figure 44, likewise depicts a single read transaction launched from the IXP2800 Network Processor to the Lucent* TDAT042G5 device followed by a single read transaction. However, in this case the read timing control register has to be programmed to setup=0, pulse width=7, and hold=0.

In Figure 44, we can count twelve clock cycles used for the read transaction in total, (i.e., 240 ns) for a clock cycle of 10 ns, excluding a turnaround cycle after that. It has the throughput of 7.7 Mbytes per second.

Figure 44. Mode 1 Single Read Transfer for Lucent* TDAT042G5 Device (B0)





3.12.7.7.2 Mode 2: Interface with 8 Data Bits and 11 Address Bits

This application is designed for the PMC-Sierra* PM5351 S/UNI-TETRA* device. For the PMC-Sierra* PM5351, the address space is programmed to 11 bits; otherwise, other address space should be specified.

8-Bit PMC-Sierra* PM5351 S/UNI-TETRA* Interfacing Topology

Figure 45 displays one of the topologies used to connect to the Slowport with the PMC-Sierra* PM5351 S/UNI-TETRA* device.

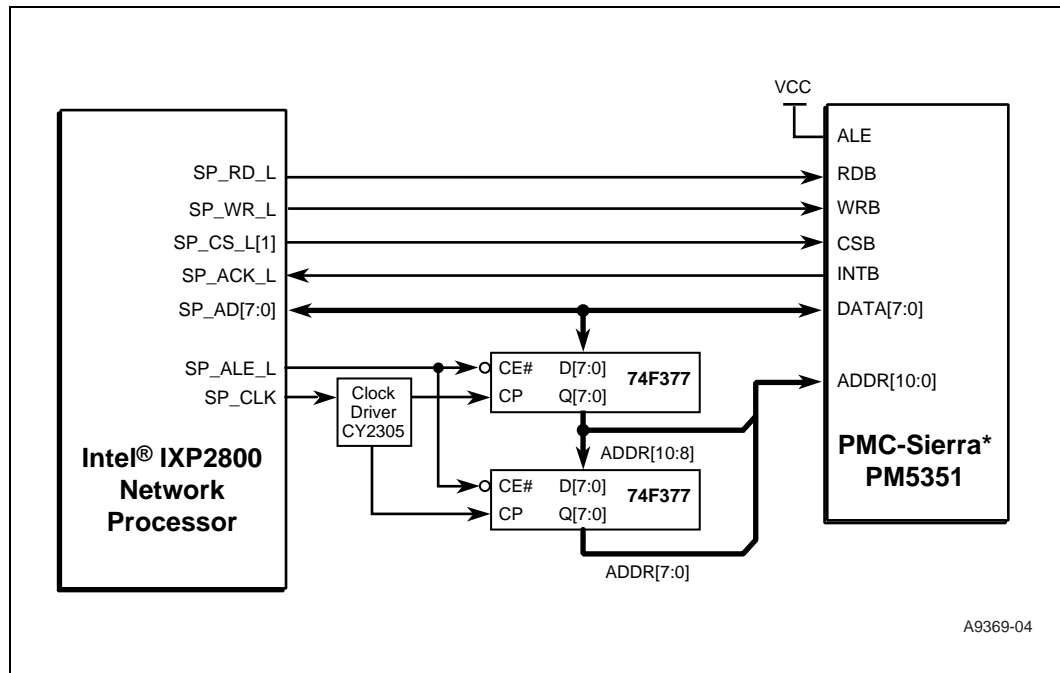
From Figure 45, because the protocols are very close to the generic Slowport protocol, the pin counts and the functionality is quite compatible. We do not need to use any more pins in this case. The only difference is in the INTB signal, which will be connected to the SP_ACK_L. Therefore, the SP_ACK_L needs to be converted to an interrupt signal.

Also because the address contains only 11bits, two 74F377 or equivalent buffers are needed.

The AS field in the SP_ADC register should be programmed to a 16-bit addressing space with the upper five address bits unconnected.

The timing controls are similar to the generic case.

Figure 45. An Interface Topology with PMC-Sierra* PM5351 S/UNI-TETRA*






PMC-Sierra* PM5351 S/UNI-TETRA* Write Interface Protocol

Figure 46 depicts a single write transaction launched from the IXP2800 to the PMC-Sierra* PM5351 device followed by single read transaction.

The write transaction for the PMC-Sierra* component has six clock cycles or a 120-ns access time for a 50-MHz Slowport clock. In this case, no intervening cycle is added after the transaction. The I/O throughput is 8.3 Mbytes per second. The SP_PCR should be programmed to mode 2 and the fields of SU, PW, and HD in the SP_WTC2 should be set to 1, 2, and 1, respectively. Here, SU, PW, and HD represent the SP_CS_L[1] pulse width, the SP_WR_L pulse width, and the SP_CP pulse width, respectively.

Figure 46. Mode 2 Single Write Transfer for PMC-Sierra* PM5351 Device (B0)




PMC-Sierra* PM5351 S/UNI-TETRA* Read Interface Protocol

Figure 47, depicts a single read transaction launched from the IXP2800 Network Processor to the PMC-Sierra* PM5351 device, followed by a single write transaction.

In this case, there are ten clock cycles of access time, or 200 ns in total, with three turnaround cycles attached at the back. The I/O throughput is 11.2 Mbytes per second.

Figure 47. Mode 2 Single Read Transfer for PMC-Sierra* PM5351 Device (B0)



3.12.7.7.3 Mode 3: Support for the Intel and AMCC* 2488 Mbps SONET/SDH Microprocessor Interface

The user has to configure the address bus to 10 bits.

Mode 3 Interfacing Topology

Figure 48 demonstrates one of the topologies used to connect the Slowport to the Intel and AMCC* 2488-Mbps SONET/SDH device. Similar to the Lucent* TDAT042G5 interface, the address and the data need demultiplexing. Totally, it requires four buffers to accomplish this task.

The SP_RD_L, SP_WR_L, and SP_CS_L[1] entirely match the RDB, WRB, and CSB pins in the Intel and AMCC* component. However, the INT has to be connected to the SP_ACK_L as the PMC-Sierra* Interface does. The ALE pin shares the SP_CP signal. If the timing does not meet specification, then ALE can be tied high as shown in Figure 49. It employs the same method as Lucent*'s TDAT042G5's topology to pack and unpack the data between the IXP2800 Slowport interface and the Intel and AMCC* microprocessor interface.

For a write, SP_CP loads the data onto the 74F646 (or equivalent) tri-state buffers, using two clock cycles. To reduce the pin count, the 16-bit data is latched with the same pin (SP_CS_L[1]), assuming that a turnaround cycle is inserted between the transaction cycles.

For a read, data are shifted out of two 74F646 or equivalent tri-state buffers by SP_CP, using two consecutive clock cycles.

Figure 48. An Interface Topology with Intel / AMCC* SONET/SDH Device

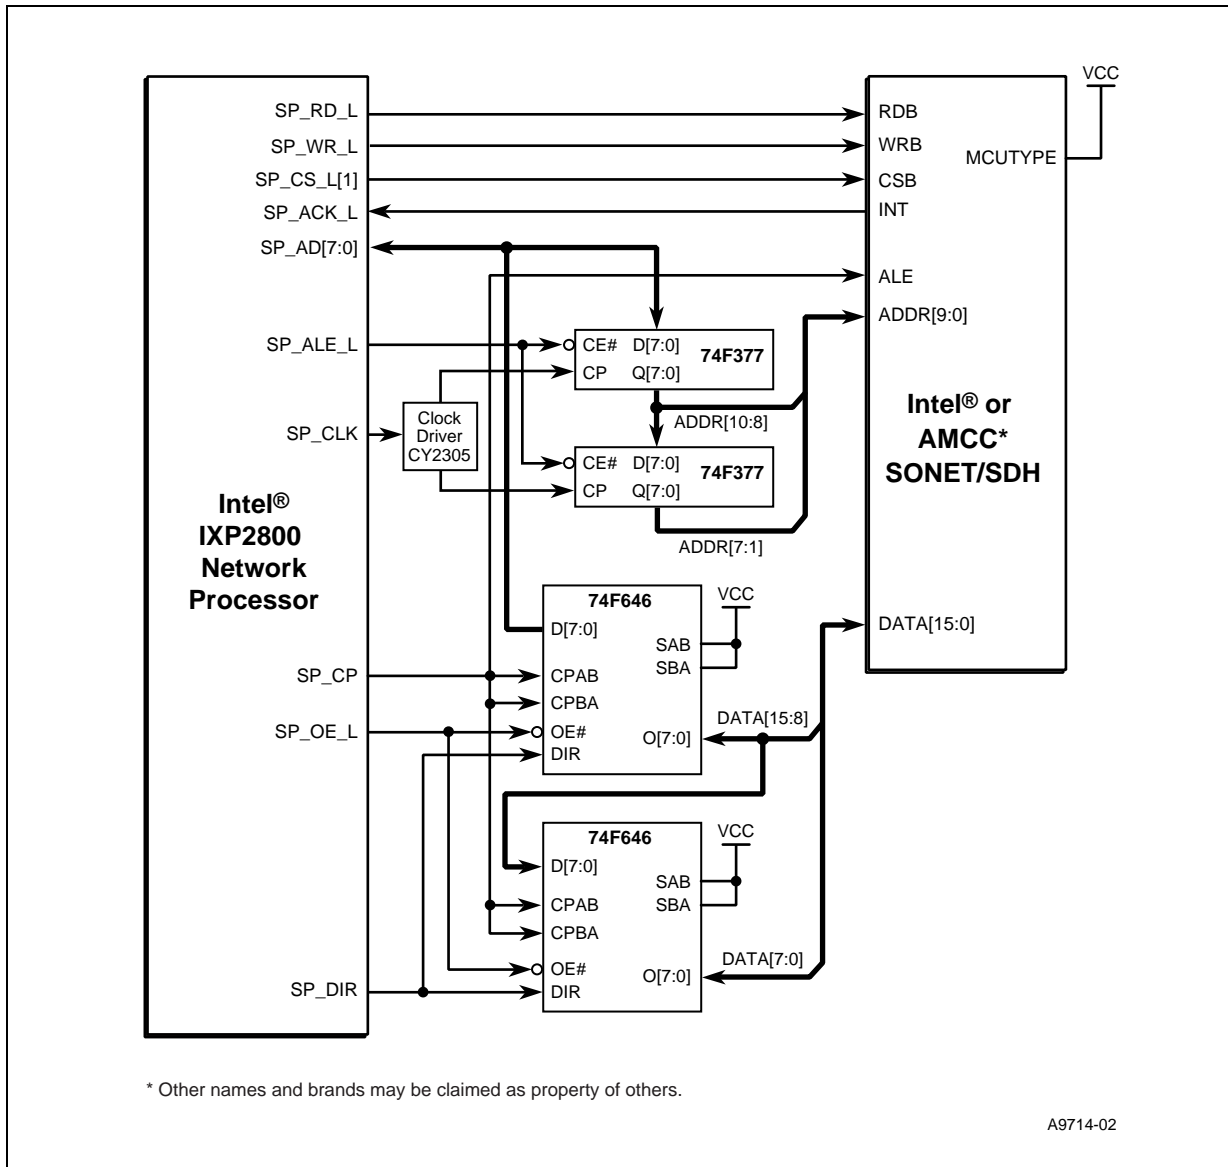
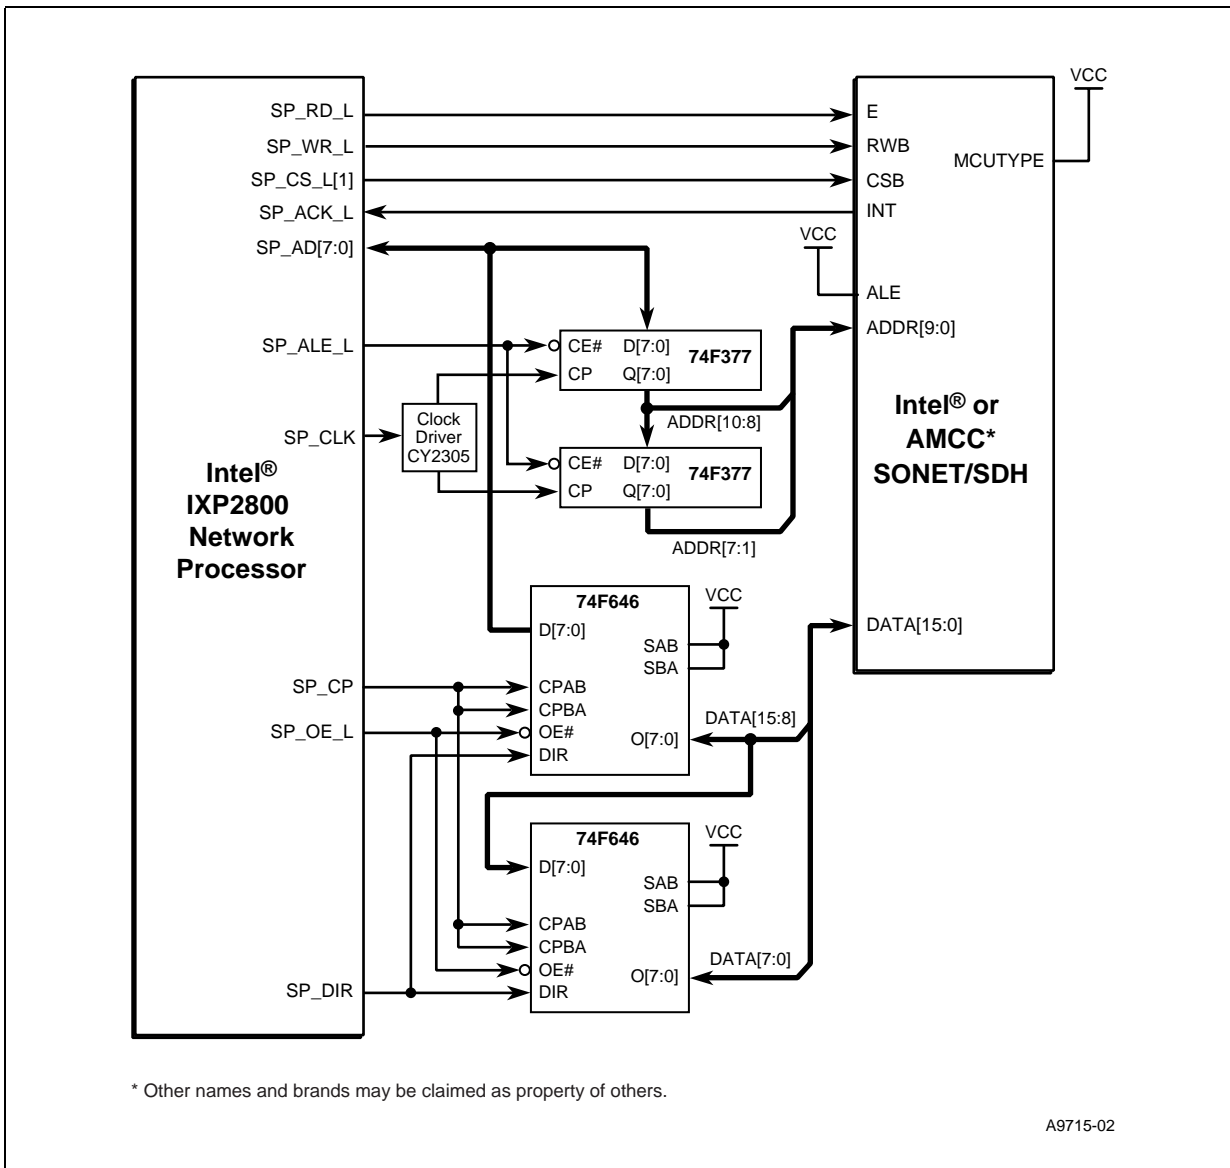


Figure 49. Mode 3 Second Interface Topology with Intel / AMCC* SONET/SDH Device



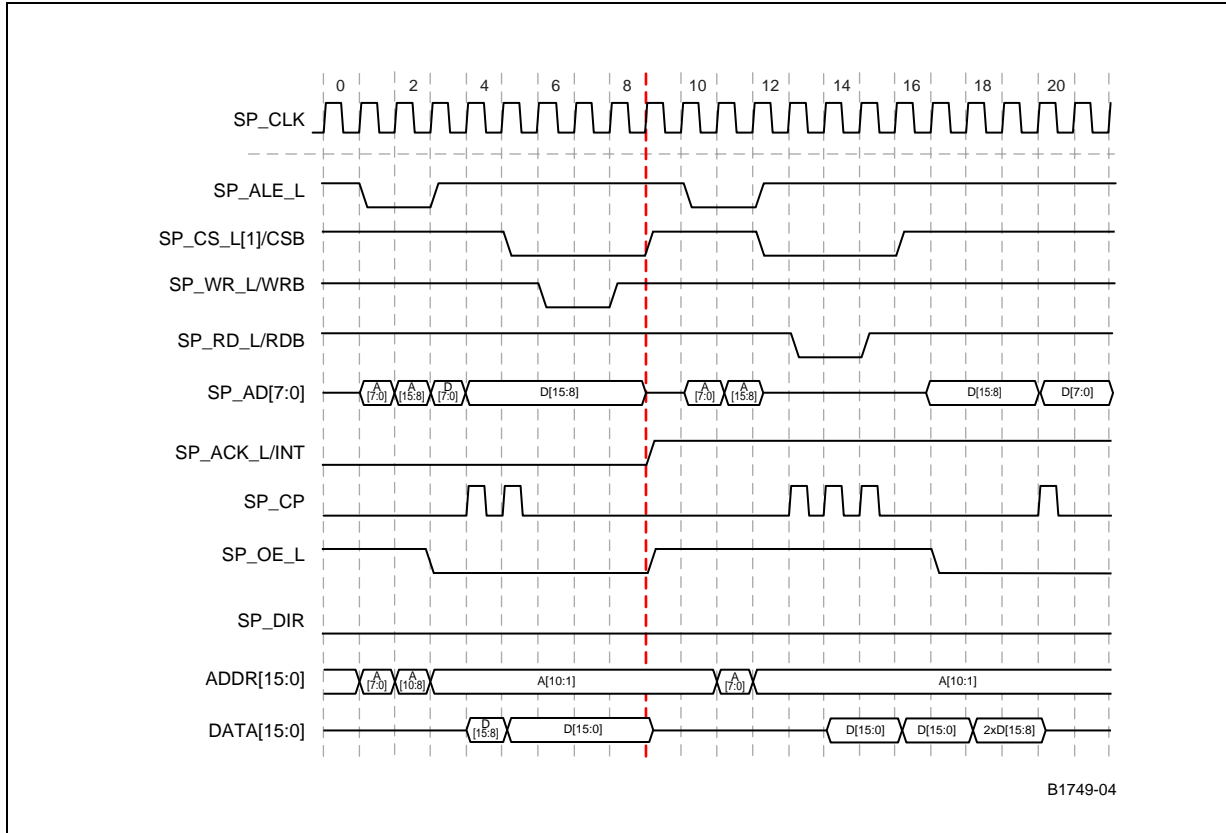


Mode 3 Write Interface Protocol

Figure 50 depicts a single write transaction launched from the IXP2800 Network Processor to the Intel and AMCC* SONET/SDH device, followed by two consecutive reads.

Compared with the Lucent* TDAT042G5, this device has a shorter access time, about eight clock cycles (i.e., 160 ns). In this case, an intervening cycle may not be needed for the write transactions. Therefore, the throughput is about 12.5 Mbytes per second.

Figure 50. Mode 3 Single Write Transfer Followed by Read (B0)

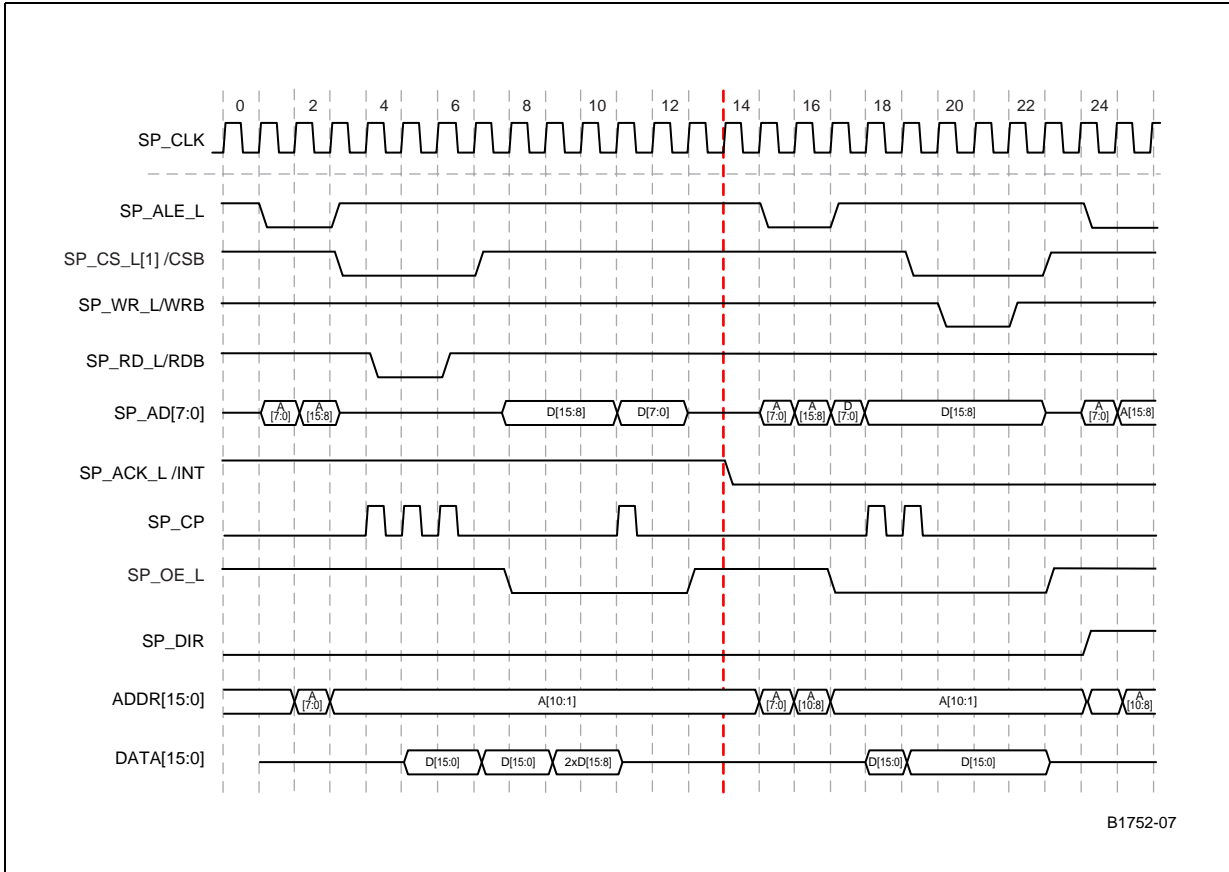


Mode 3 Read Interface Protocol

Figure 51 depicts a single read transaction launched from the IXP2800 to the Intel and AMCC* SONET/SDH device, followed by two consecutive writes.

Similarly, the access time is much better than the Lucent* TDAT042G5. The access time is eight clock cycles or 160 ns for a 50-MHz Slowport clock. Here, there are three intervening cycles between transactions. Therefore, the throughput is 11.1 Mbytes per second.

Figure 51. Mode 3 Single Read Transfer Followed by Write (B0)



Mode 4 Interfacing Topology

Figure 52 demonstrates one of the topologies used to connect Slowport to the Intel and AMCC* SONET/SDH device.

Similar to the Lucent* TDAT042G5 interface, the address and the data need demultiplexing. It requires a total of six buffers.

The RD_L, WR_L, and CS_L[1] entirely match the E, RWB, and CSB pins respectively, in the Intel framer configured to Motorola* mode. However, the INT has to be connected to the SP_ACK_L as the PMC-Sierra* Interface does. The ALE pin can share the SP_CP. However, if it does not meet the timing, the ALE pin can be tied high as shown in Figure 53.

It employs the same way to pack and unpack the data between the IXP2800 Network Processor Slowport interface and the Intel and AMCC* microprocessor interface.

For a write, W2B loads the data onto the 74F646 or equivalent tri-state buffers, using two clock cycles. To reduce the pin count, the 16-bit data are latched with the same pin (CS_L[1]), assuming that a turnaround cycle is inserted between the transaction cycles.

For a read, data are pipelined out of two 74F646 or equivalent tri-state buffers by B2S, using two consecutive clock cycles.

Figure 52. An Interface Topology with Intel / AMCC* SONET/SDH Device in Motorola* Mode

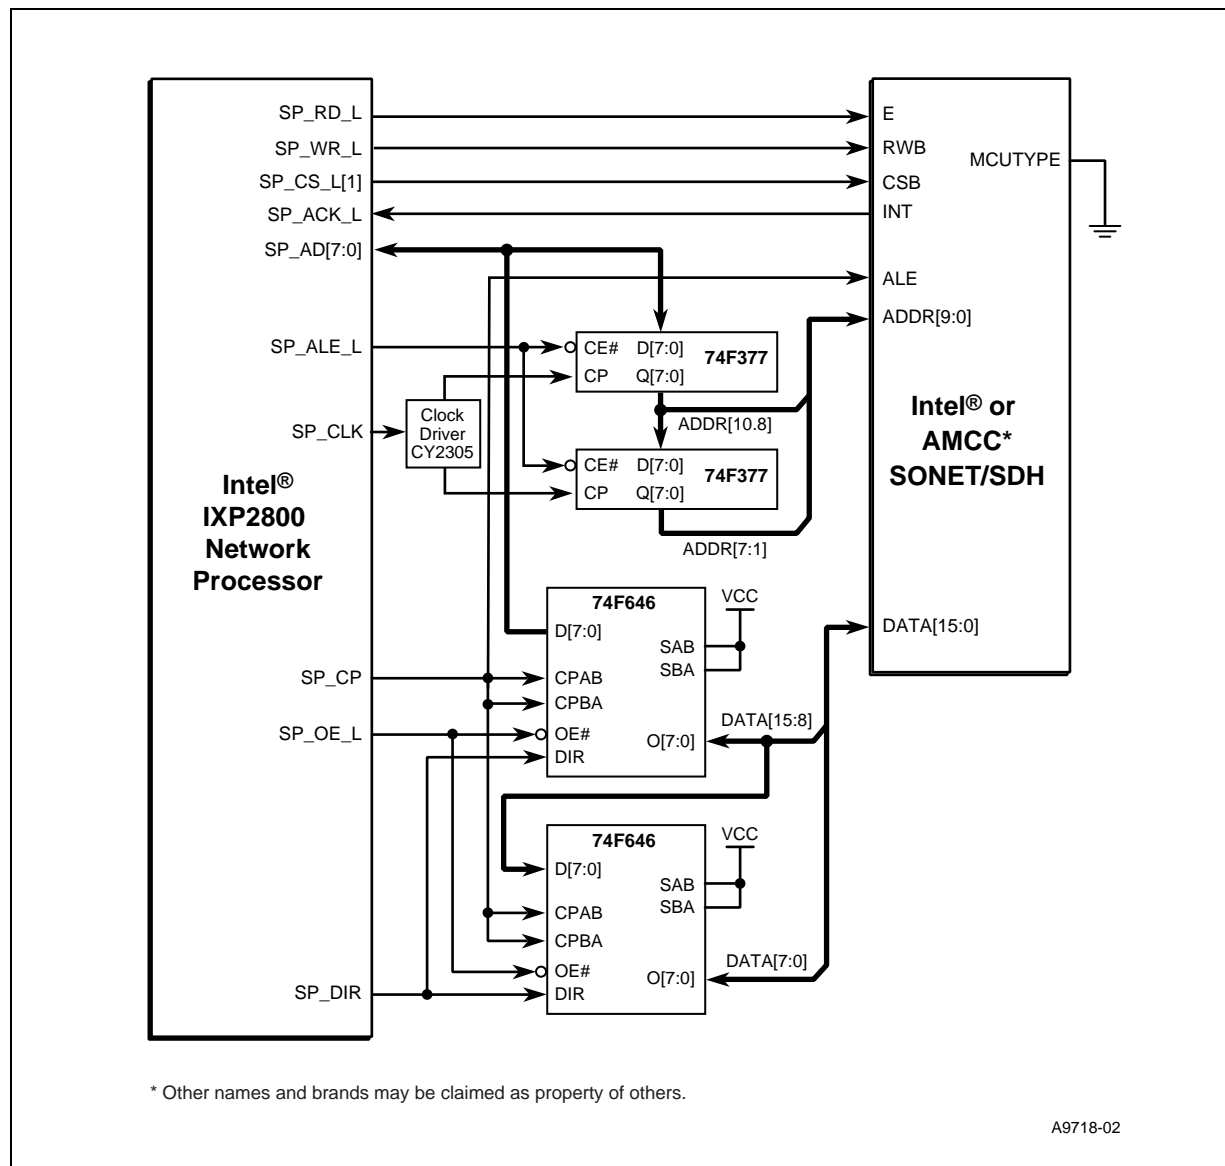
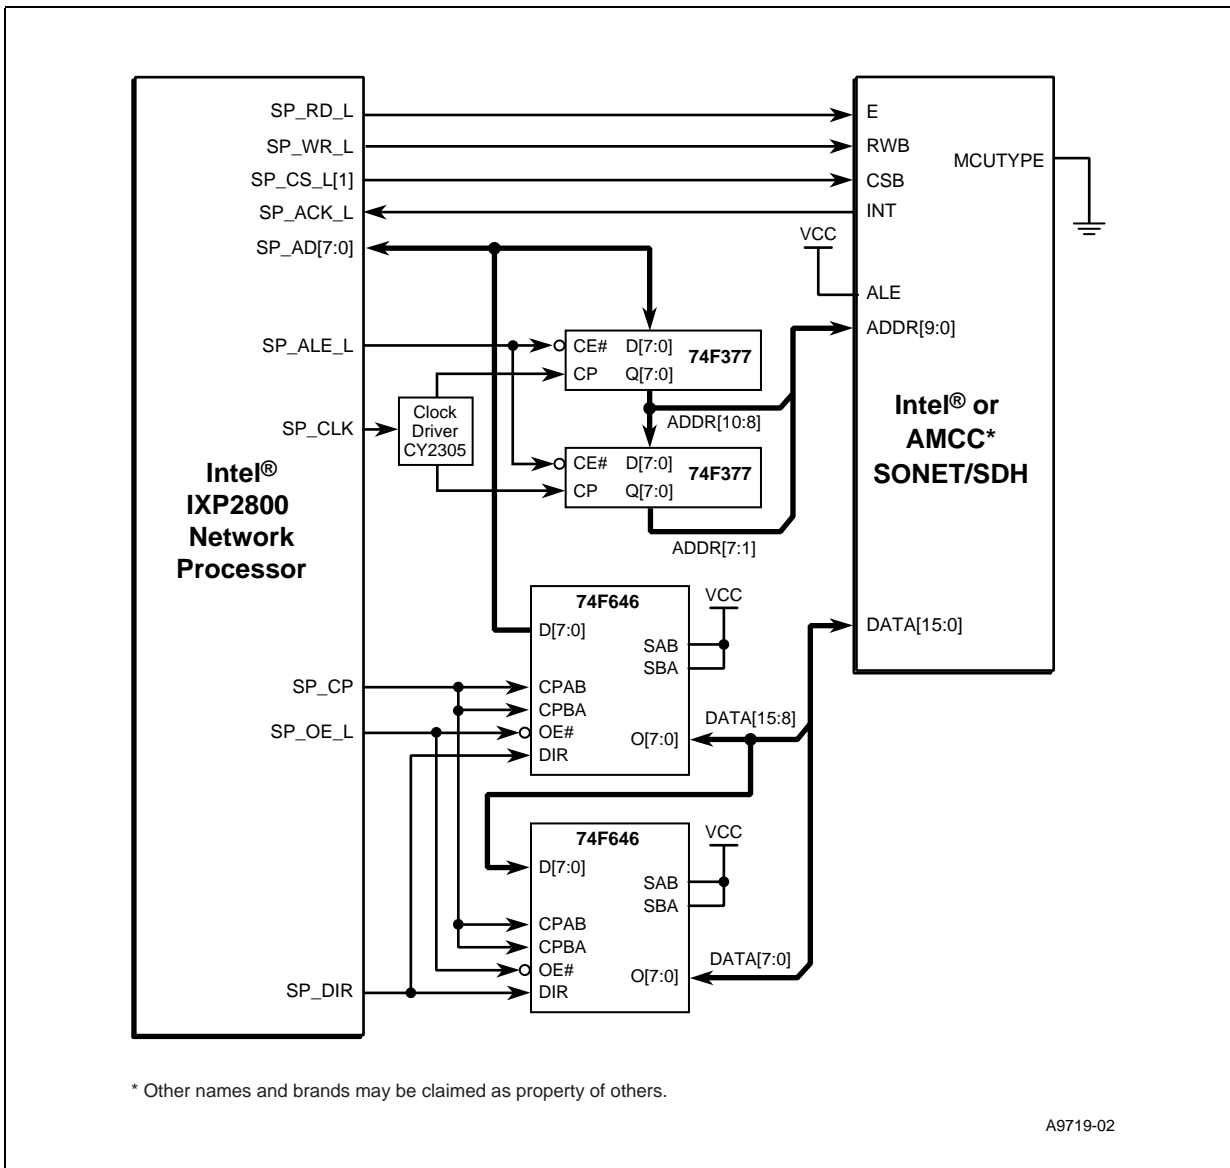


Figure 53. Second Interface Topology with Intel / AMCC* SONET/SDH Device



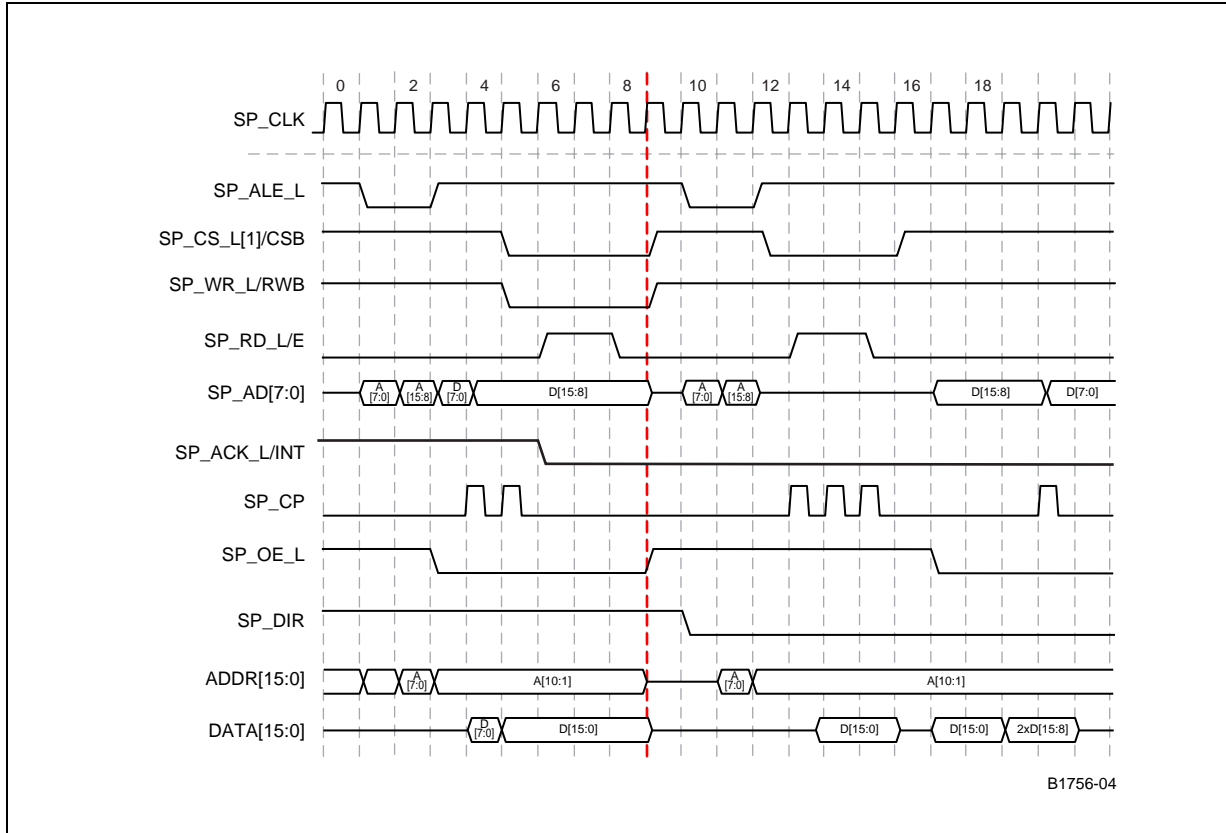


Mode 4 Write Interface Protocol

Figure 54 depicts a single write transaction launched from the IXP2800 Network Processor to the Intel and AMCC* SONET/SDH device, followed by two consecutive reads.

Compared with the Lucent* TDAT042G5, this device has a shorter access time, about eight clock cycles (i.e., 120 ns). In this case, an intervened cycle may not be needed; therefore, the throughput is about 12.5 Mbytes per second.

Figure 54. Mode 4 Single Write Transfer (B0)

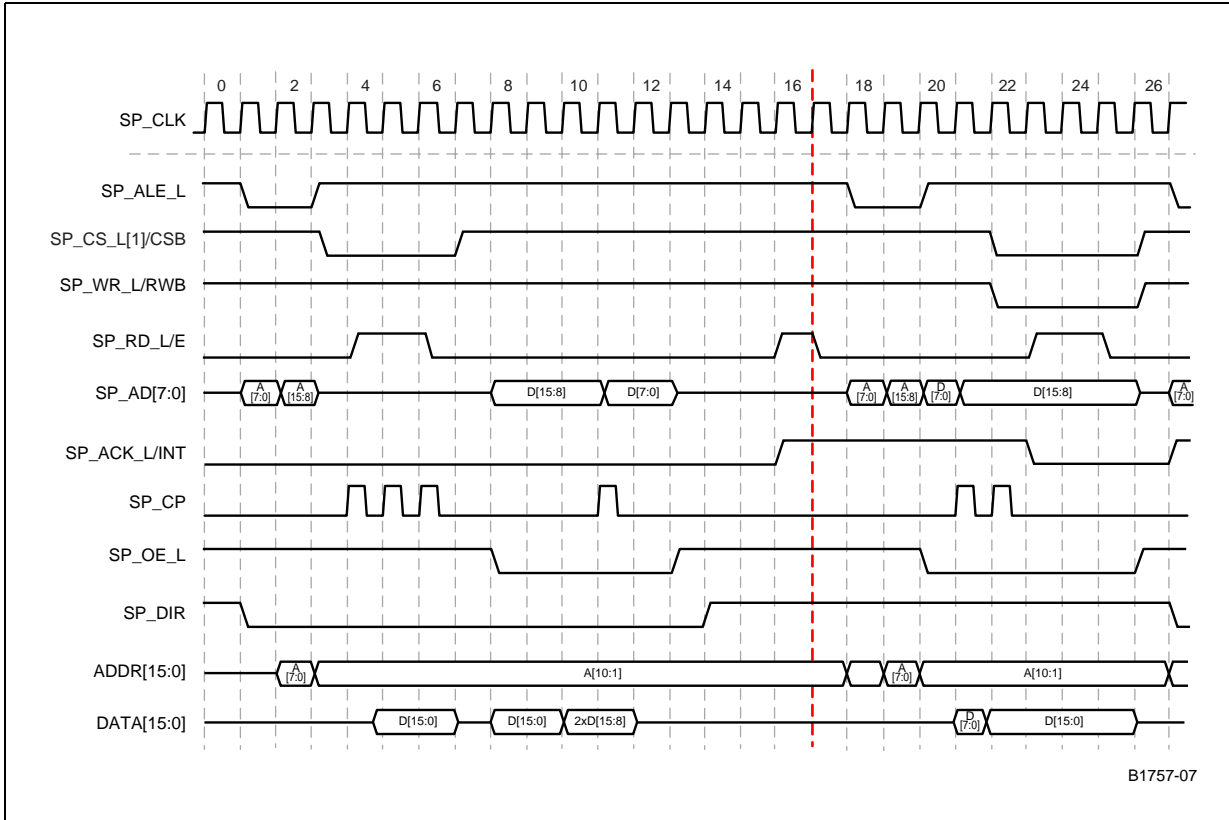


Mode 4 Read Interface Protocol

Figure 55 shows a single read transaction launched from the IXP2800 Network Processor to the Intel and AMCC* SONET/SDH device, followed by two consecutive writes.

Similarly, the access time is much better than the Lucent* TDAT042G5. The access time is about eight clock cycles or 160 ns. Here, we need an intervened cycle at the back. Therefore, the throughput is 11.2 Mbytes per second.

Figure 55. Mode 4 Single Read Transfer (B0)





This section defines the Network Processor Microengine (ME). This is the second version of the Microengine, and is often referred to as the MEv2 (Microengine Version 2).

4.1 Overview

The following sections describe the programmer's view of the Microengine. The block diagram in [Figure 56](#) is used in the description. Note that this block diagram is simplified for clarity, not all interface signals are shown, and some blocks and connectivity have been omitted to make the diagram more readable. This block diagram does not show any pipeline stages, rather it shows the logical flow of information.

The Microengine provides support for software controlled multi-threaded operation. Given the disparity in processor cycle times versus external memory times, a single thread of execution will often block waiting for external memory operations to complete. Having multiple threads available allows for threads to interleave operation—there is often at least one thread ready to run while others are blocked.



4.1.1 Control Store

The Control Store is a static RAM that holds the program that the Microengine executes. It holds 8192 instructions, each of which is 40 bits wide. It is initialized by an external device that writes to Ustore_Addr and Ustore_Data Local CSRs.

The Control Store can optionally be protected by parity against soft errors. The parity protection is optional, so that it can be disabled for implementations that don't need or want to pay the cost for it. Parity checking is enabled by CTX_Enable[Control Store Parity Enable]. A parity error on an instruction read will halt the Microengine and assert an output signal that can be used as an interrupt.

4.1.2 Contexts

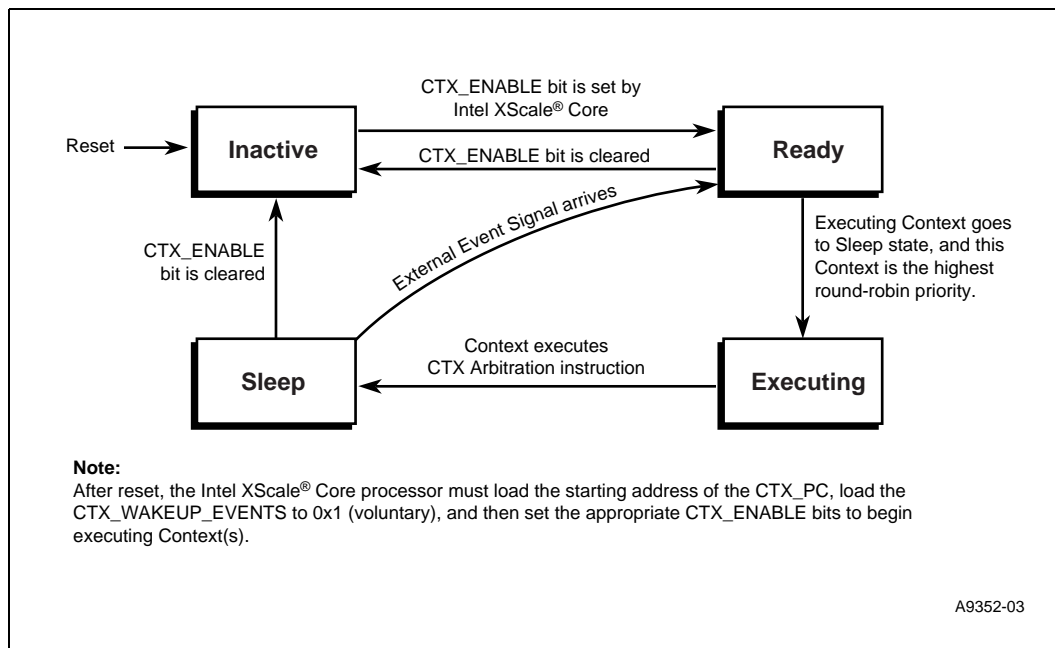
There are eight hardware Contexts available in the Microengine. To allow for efficient context swapping, each Context has its own register set, Program Counter, and Context specific Local registers. Having a separate copy per Context eliminates the need to move Context specific information to/from shared memory and Microengine registers for each Context swap. Fast context swapping allows a Context to do computation while other Contexts wait for IO (typically external memory accesses) to complete or for a signal from another Context or hardware unit. *Note:* a context swap is similar to a taken branch in timing.

Each of the eight Contexts is always in one of four states.

1. Inactive — Some applications may not require all eight contexts. A Context is in the Inactive state when its CTX_Enable CSR enable bit is a 0.
2. Executing — A Context is in Executing state when its context number is in Active_CTX_Status CSR. The executing Context's PC is used to fetch instructions from the Control Store. A Context will stay in this state until it executes an instruction that causes it to go to Sleep state (there is no hardware interrupt or preemption; Context swapping is completely under software control). At most one Context can be in Executing state at any time.
3. Ready — In this state, a Context is ready to execute, but is not because a different Context is executing. When the Executing Context goes to Sleep state, the Microengine's context arbiter selects the next Context to go to the Executing state from among all the Contexts in the Ready state. The arbitration is round robin.
4. Sleep — Context is waiting for external event(s) specified in the CTX_#_Wakeup_Events CSR to occur (typically, but not limited to, an IO access). In this state the Context does not arbitrate to enter the Executing state.

The state diagram in [Figure 57](#) illustrates the Context state transitions. Each of the eight Contexts will be in one of these states. At most one Context can be in Executing state at a time; any number of Contexts can be in any of the other states.

Figure 57. Context State Transition Diagram



The Microengine is in Idle state whenever no Context is running (all Contexts are in either Inactive or Sleep states). This state is entered:

1. After reset (because CTX_Enable Local CSR is clear, putting all Contexts into Inactive states).
2. When a context swap is executed, but no context is ready to wakeup.
3. When a `ctx_arb[bpt]` instruction is executed by the Microengine (this is a special case of condition 2 above, since the `ctx_arb[bpt]` clears CTX_Enable, putting all Contexts into Inactive states).

The Microengine provides the following functionality during Idle state:

1. The Microengine continuously checks if a Context is in Ready state. If so, a new Context begins to execute. If no Context is Ready, the Microengine remains in the Idle state.
2. Only the ALU instructions are supported. They are used for debug via special hardware defined in number 3 below.
3. A write to the Ustore_Addr Local CSR with the Ustore_Addr[ECS] bit set, causing the Microengine to repeatedly execute the instruction pointed by the address specified in the Ustore_Addr CSR. Only the ALU instructions are supported in this mode. Also, the result of the execution is written to the ALU_Out Local CSR rather than a destination register.
4. A write to the Ustore_Addr Local CSR with the Ustore_Addr[ECS] bit set, followed by a write to the Ustore_Data Local CSR loads an instruction into the Control Store. After the Control Store is loaded, execution proceeds as described in number 3 above. Note that the write to Ustore_Data causes Ustore_Addr to increment, so it must be written back to the address of the desired instruction.



4.1.3 Datapath Registers

As shown in the block diagram in [Figure 56](#), each Microengine contains four types of 32-bit datapath registers:

- 256 General Purpose registers
- 512 Transfer registers
- 128 Next Neighbor registers
- 640 32-bit words of Local Memory

4.1.3.1 General-Purpose Registers (GPRs)

GPRs are used for general programming purposes. They are read and written exclusively under program control. GPRs, when used as a source in an instruction, supply operands to the execution datapath. When used as a destination in an instruction, they are written with the result of the execution datapath. The specific GPRs selected are encoded in the instruction.

The GPRs are physically and logically contained in two banks, GPR A and GPR B, defined in [Table 57](#).

Note: The Microengine registers are defined in the *IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

4.1.3.2 Transfer Registers

There are four types of transfer (abbreviated as Xfer) registers used for transferring data to and from the Microengine and locations external to the Microengine (DRAMs, SRAMs, etc.).

- S_TRANSFER_IN
- S_TRANSFER_OUT
- D_TRANSFER_IN
- D_TRANSFER_OUT

Transfer_In registers, when used as a source in an instruction, supply operands to the execution datapath. The specific register selected is either encoded in the instruction or selected indirectly using T_Index. Transfer_In registers are written by external units based on the Push_ID input to the Microengine.

Transfer_Out registers, when used as a destination in an instruction, are written with the result from the execution datapath. The specific register selected is encoded in the instruction, or selected indirectly via T_Index. Transfer_Out registers supply data to external units based on the Pull_ID input to the Microengine.

As shown in [Figure 56](#), the S_TRANSFER_IN and D_TRANSFER_IN registers connect to both the S_Push and D_Push buses via a multiplexor internal to the Microengine. Additionally, the S_TRANSFER_OUT and D_TRANSFER_OUT Transfer registers connect to both the S_Pull and D_Pull buses. This feature enables a programmer to use the either type of transfer register regardless of the source or destination of the transfer.



Typically, the external units access the Transfer registers in response to commands sent by the Microengines. The commands are sent in response to instructions executed by the Microengine (for example, the command instructs a SRAM controller to read from external SRAM, and place the data into a S_TRANSFER_IN register). However, it is possible for an external unit to access a given Microengine's Transfer registers either autonomously, or under control of a different Microengine, or the Intel XScale® core, etc. The Microengine interface signals controlling writing/reading of the Transfer_In/Transfer_Out registers are independent of the operation of the rest of the Microengine.

4.1.3.3 Next Neighbor Registers

A new feature added for the Microengine Version 2 are 128 Next Neighbor registers that provide a dedicated datapath for transferring data from the previous/next neighbor Microengine. Next Neighbor registers, when used as a source in an instruction, supply operands to the execution datapath. They are written in two different ways: (1) by an external entity, typically, but not limited to, another adjacent Microengine, or (2) by the same Microengine they are in, as controlled by CTX_Enable[NN_Mode]. The specific register is selected in one of two ways: (1) Context-relative, the register number is encoded in the instruction, or (2) as a Ring, selected via NN_Get and NN_Put CSR registers.

When CTX_Enable[NN_Mode] is '0' – When Next Neighbor is used as a destination in an instruction, the instruction result data is sent out of the Microengine, typically to another, adjacent Microengine.

When CTX_Enable[NN_Mode] is '1' – When Next Neighbor is used as a destination in an instruction, the instruction result data is written to the selected Next Neighbor register in the Microengine. Note that there is a 5-instruction latency until the newly written data can be read. The data is not sent out of the Microengine as it would be when CTX_Enable[NN_Mode] is '0'.

Table 56. Next Neighbor Write as a Function of CTX_Enable[NN_Mode]

NN_Mode	Where the Write Goes	
	External?	NN Register in this Microengine?
0	Yes	No
1	No	Yes

4.1.3.4 Local Memory

Local Memory is addressable storage located in the Microengine, organized as 640 32-bit words. Local Memory is read and written exclusively under program control. Local Memory supplies operands to the execution datapath as a source, and receives results as a destination.

The specific Local Memory location selected is based on the value in one of the Local Memory_Addr registers, which are written by local_CSR_wr instructions. There are two LM_Addr registers per Context and a working copy of each. When a Context goes to the Sleep state, the value of the working copies is put into the Context's copy of LM_Addr. When the Context goes to the Executing state, the value in its copy of LM_Addr is put into the working copies. The choice of LM_Addr_0 or LM_Addr_1 is selected in the instruction.



It is also possible to make use of both or one LM_Addrs as global by setting CTX_Enable[LM_Addr_0_Global] and/or CTX_Enable[LM_Addr_1_Global]. When used globally, all Contexts use the working copy of LM_Addr in place of their own Context specific one; the Context specific ones are unused.

4.1.4 Addressing Modes

GPRs can be accessed in two different addressing modes: Context-Relative and Absolute. Some instructions can specify either mode; other instructions can specify only Context-Relative mode.

- Transfer and Next Neighbor registers can be accessed in Context-Relative and Indexed modes.
- Local Memory is accessed in Indexed mode.
- The addressing mode in use is encoded directly into each instruction, for each source and destination specifier.

4.1.4.1 Context-Relative Addressing Mode

The GPRs are logically subdivided into equal regions such that each Context has exclusive access to one of the regions. The number of regions (four or eight) is configured in the CTX_Enable CSR. Thus, a Context-Relative register name is actually associated with multiple different physical registers. The actual register to be accessed is determined by the Context making the access request (the Context number is concatenated with the register number specified in the instruction — see [Table 57](#)). Context-Relative addressing is a powerful feature that enables eight different contexts to share the same microcode, yet maintain separate data.

[Table 57](#) shows how the Context number is used in selecting the register number in relative mode. The register number in [Table 57](#) is the Absolute GPR address, or Transfer or Next Neighbor Index number to use to access the specific Context-Relative register. For example, with eight active Contexts, Context-Relative Register 0 for Context 2 is Absolute Register Number 32.

Table 57. Registers Used by Contexts in Context-Relative Addressing Mode

Number of Active Contexts	Active Context Number	GPR Absolute Register Numbers		S_Transfer or Neighbor Index Number	D_Transfer Index Number
		A Port	B Port		
8	0	0 – 15	0 – 15	0 – 15	0 – 15
	1	16 – 31	16 – 31	16 – 31	16 – 31
	2	32 – 47	32 – 47	32 – 47	32 – 47
	3	48 – 63	48 – 63	48 – 63	48 – 63
	4	64 – 79	64 – 79	64 – 79	64 – 79
	5	80 – 95	80 – 95	80 – 95	80 – 95
	6	96 – 111	96 – 111	96 – 111	96 – 111
	7	112 – 127	112 – 127	112 – 127	112 – 127
4	0	0 – 31	0 – 31	0 – 31	0 – 31
	2	32 – 63	32 – 63	32 – 63	32 – 63
	4	64 – 95	64 – 95	64 – 95	64 – 95
	6	96 – 127	96 – 127	96 – 127	96 – 127



4.1.4.2 Absolute Addressing Mode

With Absolute addressing, any GPR can be read or written by any one of the eight Contexts in a Microengine. Absolute addressing enables register data to be shared among all of the Contexts, e.g., for global variables or for parameter passing. All 256 GPRs can be read by Absolute address.

4.1.4.3 Indexed Addressing Mode

With Indexed addressing, any Transfer or Next Neighbor register can be read or written by any one of the eight Contexts in an Microengine. Indexed addressing enables register data to be shared among all of the Contexts. For indexed addressing the register number comes from the T_Index register for Transfer registers or NN_Put and NN_Get registers (for Next Neighbor registers).

4.2 Local CSRs

Local Control and Status registers (CSRs) are external to the Execution Datapath, and hold specific purpose information. They can be read and written by special instructions (local_csr_rd and local_csr_wr) and are typically accessed less frequently than datapath registers. Because Local CSRs are not built in the datapath, there is a write to use delay of either three or four cycles, and a read to consume penalty of one cycle.

4.3 Execution Datapath

The Execution Datapath can take one or two operands, perform an operation, and optionally write back a result. The sources and destinations can be GPRs, Transfer registers, Next Neighbor registers, and Local Memory. The operations are shifts, addition, subtraction, logicals, multiplication, byte-align, and “find first bit set”.

4.3.1 Byte Align

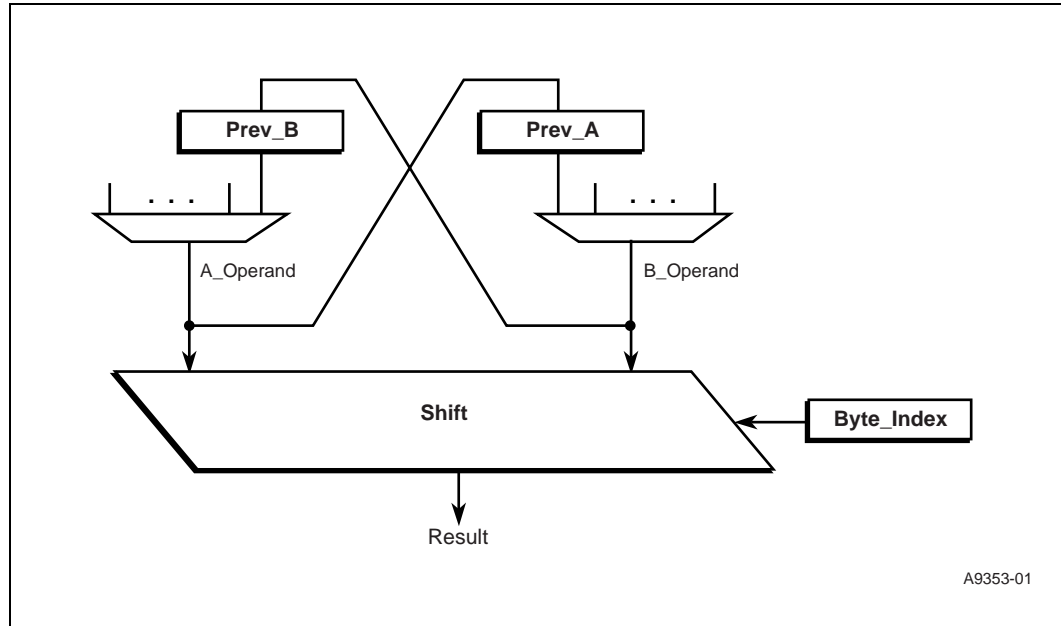
The datapath provides a mechanism to move data from source register(s) to any destination register(s) with byte aligning. Byte aligning takes four consecutive bytes from two concatenated values (eight bytes), starting at any of four byte boundaries (0, 1, 2, 3), and based on the endian type (which is defined in the instruction opcode), as shown in Table 58. The four bytes are taken from two concatenated values. Four bytes are always supplied from a temporary register that always holds the A or B operand from the previous cycle, and the other four bytes from the B or A operand of the Byte Align instruction. The operation is described below using the block diagram Figure 58. The alignment is controlled by the two LSBs of the Byte_Index Local CSR.

Table 58. Align Value and Shift Amount

Align Value (in Byte_Index[1:0])	Right Shift Amount (Number of Bits in Decimal)	
	Little-Endian	Big-Endian
0	0	32
1	8	24
2	16	16
3	24	8



Figure 58. Byte Align Block Diagram



Example 24 shows an align sequence of instructions and the value of the various operands. Table 59 shows the data in the registers for this example. The value in Byte_Index[1:0] CSR (which controls the shift amount) for this example is 2.

Table 59. Register Contents for Example 23

Register	Byte 3 [31:24]	Byte 2 [23:16]	Byte 1 [15:8]	Byte 0 [7:0]
0	0	1	2	3
1	4	5	6	7
2	8	9	A	B
3	C	D	E	F

Example 24. Big-Endian Align

Instruction	Prev B	A Operand	B Operand	Result
Byte_align_be[--, r0]	--	--	0123	--
Byte_align_be[dest1, r1]	0123	0123	4567	2345
Byte_align_be[dest2, r2]	4567	4567	89AB	6789
Byte_align_be[dest3, r3]	89AB	89AB	CDEF	ABCD

NOTE: A Operand comes from Prev_B register during byte_align_be instructions.



Example 25 shows another sequence of instructions and the value of the various operands. Table 60, shows the data in the registers for this example.

The value in Byte_Index[1:0] CSR (which controls the shift amount) for this example is 2.

Table 60. Register Contents for Example 24

Register	Byte 3 [31:24]	Byte 2 [23:16]	Byte 1 [15:8]	Byte 0 [7:0]
0	3	2	1	0
1	7	6	5	4
2	B	A	9	8
3	F	E	D	C

Example 25. Little-Endian Align

Instruction	A Operand	B Operand	Prev A	Result
Byte_align_le[--, r0]	3210	--	--	--
Byte_align_le[dest1, r1]	7654	3210	3210	5432
Byte_align_le[dest2, r2]	BA98	7654	7654	9876
Byte_align_le[dest3, r3]	FEDC	BA98	BA98	DCBA
NOTE: B Operand comes from Prev_A register during byte_align_le instructions.				

As the examples show, byte aligning “n” words takes “n+1” cycles due to the first instruction needed to start the operation.

Another mode of operation is to use the T_Index register with post-increment, to select the source registers. T_Index operation is described later in this chapter.

4.3.2 CAM

The block diagram in Figure 59 is used to explain the CAM operation.

The CAM has 16 entries. Each entry stores a 32-bit value, which can be compared against a source operand by instruction: CAM_Lookup[dest_reg, source_reg].

All entries are compared in parallel, and the result of the lookup is a 9-bit value that is written into the specified destination register in bits 11:3, with all other bits of the register set to 0 (the choice of bits 11:3 is explained below). The result can also optionally be written into either of the LM_Addr registers (see below in this section for details).

The 9-bit result consists of four State bits (dest_reg[11:8]), concatenated with a 1-bit Hit/Miss indication (dest_reg[7]), concatenated with 4-bit entry number (dest_reg[6:3]). All other bits of dest_reg are written with 0. Possible results of the lookup are:

- miss (0) — lookup value is not in CAM, entry number is Least Recently Used entry (which can be used as a suggested entry to replace), and State bits are 0000.
- hit (1) — lookup value is in CAM, entry number is entry that has matched; State bits are the value from the entry that has matched.



Note: The State bits are data associated with the entry. State bits are only used by software. There is no implication of ownership of the entry by any Context. The State bits hardware function is:

- the value is set by software (when the entry is loaded or changed in an already-loaded entry).
- its value is read out on a lookup that hits, and used as part of the status written into the destination register.
- its value can be read out separately (normally only used for diagnostic or debug).

The LRU (Least Recently Used) Logic maintains a time-ordered list of CAM entry usage. When an entry is loaded, or matches on a lookup, it is marked as MRU (Most Recently Used). Note that a lookup that misses does *not* modify the LRU list. The CAM is loaded by instruction:


`CAM_Write[entry_reg, source_reg, state_value].`

The value in the register specified by `source_reg` is put into the Tag field of the entry specified by `entry_reg`. The value for the State bits of the entry is specified in the instruction as `state_value`.

The value in the State bits for an entry can be written, without modifying the Tag, by instruction: `CAM_Write_State[entry_reg, state_value].`

Note: `CAM_Write_State` does *not* modify the LRU list.

Figure 59. CAM Block Diagram





One possible way to use the result of a lookup is to dispatch to the proper code using instruction:

```
jump[register, label#],defer [3]
```

where the register holds the result of the lookup. The State bits can be used to differentiate cases where the data associated with the CAM entry is in flight, or is pending a change, etc. Because the lookup result was loaded into bits[11:3] of the destination register, the jump destinations are spaced eight instructions apart. This is a balance between giving enough space for many applications to complete their task without having to jump to another region, versus consuming too much Control Store. Another way to use the lookup result is to branch on just the hit miss bit, and use the entry number as a base pointer into a block of Local Memory.

When enabled, the CAM lookup result is loaded into Local_Addr as follows:

```
LM_Addr[5:0] = 0 ([1:0] are read-only bits)
LM_Addr[9:6] = lookup result [6:3] (entry number)
LM_Addr[11:10] = constant specified in instruction
```

This function is useful when the CAM is used as a cache, and each entry is associated with a block of data in Local Memory. Note that the latency from when CAM_Lookup executes until the LM_Addr is loaded is the same as when LM_Addr is written by a Local_CSR_Wr instruction.

The Tag and State bits for a given entry can be read by instructions:

```
CAM_Read_Tag[dest_reg, entry_reg]
CAM_Read_State[dest_reg, entry_reg]
```

The Tag value and State bits value for the specified entry is written into the destination register, respectively for the two instructions (the State bits are placed into bits [11:8] of dest_reg, with all other bits 0). Reading the tag is useful in the case where an entry needs to be evicted to make room for a new value — the lookup of the new value results in a miss, with the LRU entry number returned as a result of the miss. The CAM_Read_Tag instruction can then be used to find the value that was stored in that entry. An alternative would be to keep the tag value in a GPR. These two instructions can also be used by debug and diagnostic software. Neither of these modify the state of the LRU pointer.

Note: The following rules must be adhered to when using the CAM.

- CAM is not reset by Microengine reset. Software must either do a CAM_clear prior to using the CAM to initialize the LRU and clear the tags to 0, or explicitly write all entries with CAM_write.
- No two tags can be written to have the same value. If this rule is violated, the result of a lookup that matches that value will be unpredictable, and LRU state is unpredictable.

The value, 0x00000000 can be used as a valid lookup value. However, note that the CAM_clear instruction puts 0x00000000 into all tags. To avoid violating rule 2 after doing CAM_clear, it is necessary to write all entries to unique values prior to doing a lookup of 0x00000000. An algorithm for debug software to determine the contents of the CAM is shown in [Example 26](#).



Example 26. Algorithm for Debug Software to Determine the Contents of the CAM

```

; First read each of the tag entries. Note that these reads
; don't modify the LRU list or any other CAM state.
tag[0] = CAM_Read_Tag(entry_0);
.....
tag[15] = CAM_Read_Tag(entry_15);
; Now read each of the state bits
state[0] = CAM_Read_State(entry_0);
...
state[15] = CAM_Read_State(entry_15);
; Knowing what tags are in the CAM makes it possible to
; create a value that is not in any tag, and will therefore
; miss on a lookup.

; Next loop through a sequence of 16 lookups, each of which will
; miss, to obtain the LRU values of the CAM.
for (i = 0; i < 16; i++)
  BEGIN_LOOP
    ; Do a lookup with a tag not present in the CAM. On a
    ; miss, the LRU entry will be returned. Since this lookup
    ; missed the LRU state is not modified.
    LRU[i] = CAM_Lookup(some_tag_not_in_cam);
    ; Now do a lookup using the tag of the LRU entry. This
    ; lookup will hit, which makes that entry MRU.
    ; This is necessary to allow the next lookup miss to
    ; see the next LRU entry.
    junk = CAM_Lookup(tag[LRU[i]]);
  END_LOOP
; Because all entries were hit in the same order as they were
; LRU, the LRU list is now back to where it started before the
; loop executed.
; LRU[0] through LRU[15] holds the LRU list.

```

The CAM can be cleared with `CAM_Clear` instruction. This instruction writes `0x00000000` simultaneously to all entries tag, clears all the state bits, and puts the LRU into an initial state (where entry 0 is LRU, ..., entry 15 is MRU).

4.4 CRC Unit

The CRC Unit operates in parallel with the Execution Datapath. It takes two operands, performs a CRC operation, and writes back a result. CRC-CCITT, CRC-32, CRC-10, CRC-5, and iSCSI polynomials are supported. One of the operands is the CRC_Remainder Local CSR, and the other is a GPR, Transfer_In register, Next Neighbor, or Local Memory, specified in the instruction and passed through the Execution Datapath to the CRC Unit. The instruction specifies the CRC operation type, whether to swap bytes and or bits, and the bytes of the operand to be included in the operation. The result of the CRC operation is written back into CRC_Remainder. The source operand can also be written into a destination register (however the byte/bit swapping and masking do not affect the destination register; they only affect the CRC computation). This allows moving data, for example, from `S_TRANSFER_IN` registers to `S_TRANSFER_OUT` registers at the same time as computing the CRC.



4.5 Event Signals

Event Signals are used to coordinate a program with completion of external events. For example, when a Microengine issues a command to an external unit to read data (which will be written into a Transfer_In register), the program must insure that it does not try to use the data until the external unit has written it. There is no hardware mechanism to flag that a register write is pending, and then prevent the program from using it. Instead the coordination is under software control, with hardware support.

When the program issues the command to the external event, it can request that the external unit supply an indication (called an Event Signal) that the command has been completed. There are 15 Event Signals per Context that can be used, and Local CSRs per Context to track which Event Signals are pending and which have been returned. The Event Signals can be used to move a Context from Sleep state to Ready state, or alternatively, the program can test and branch on the status of Event Signals.

Event Signals can be set in nine different ways.

1. When data is written into S_TRANSFER_IN registers (part of S_Push_ID input)
2. When data is written into D_TRANSFER_IN registers (part of D_Push_ID input)
3. When data is taken from S_TRANSFER_OUT registers (part of S_Pull_ID input)
4. When data is taken from D_TRANSFER_OUT registers (part of D_Pull_ID input)
5. On InterThread_Sig_In input
6. On NN_Sig_In input
7. On Prev_Sig_In input
8. On write to Same_ME_Signal Local CSR
9. By Internal Timer

Any or all Event Signals can be set by any of the above sources.

When a Context goes to the Sleep state (executes a `ctx_arb` instruction, or a Command instruction with `ctx_swap` token), it specifies which Event Signal(s) it requires to be put in the Ready state. The `ctx_arb` instruction also specifies whether the logical AND or logical OR of the Event Signal(s) is needed to put the Context into the Ready state.

When a Context Event Signals arrive, it goes to the Ready state, and then to the Executing state. In the case where the Event Signal is linked to moving data into or out of Transfer registers (numbers 1 through 4 in the list above), the code can safely use the Transfer register as the first instruction (for example, using a Transfer_In register as a source operand will get the new read data). The same is true when the Event Signal is tested for branches (`br_=signal` or `br_!signal` instructions).

The `ctx_arb` instruction, `CTX_Sig_Events`, and `CTX_Wakeup_#_Events` Local CSR descriptions provide details.



4.5.1 Microengine Endianness

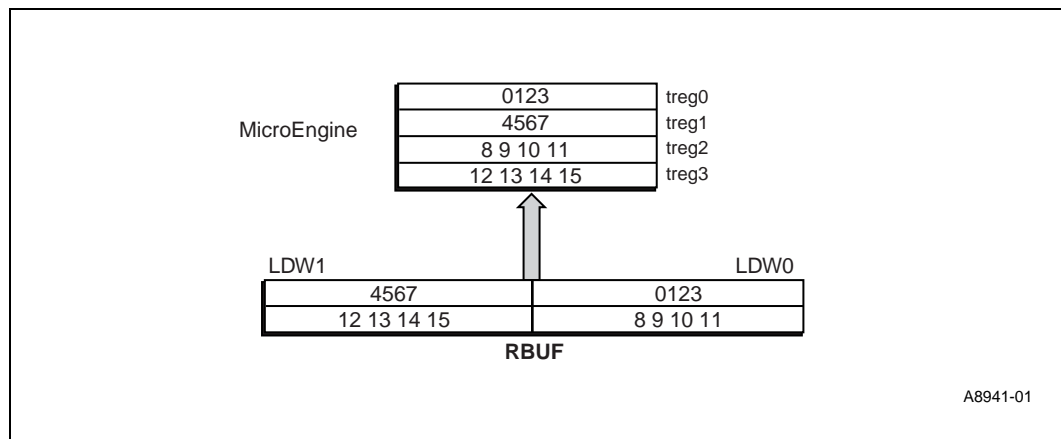
Microengine operation from an “endian” point of view can be divided into following categories:

- Read from RBUF (64 bits)
- Write to TBUF (64 bits)
- Read/write from/to SRAM
- Read/write from/to DRAM
- Read/write from/to SHaC and other CSRs
- Write to Hash

4.5.1.1 Read from RBUF (64 Bits)

Data in RBUF is arranged in LWBE order. Whenever the Microengine reads from RBUF, the low order longword (LDW0) is transferred into Microengine transfer register 0 (treg0), the high order longword (LDW1) is transferred into treg1, etc. This is explained in [Figure 60](#).

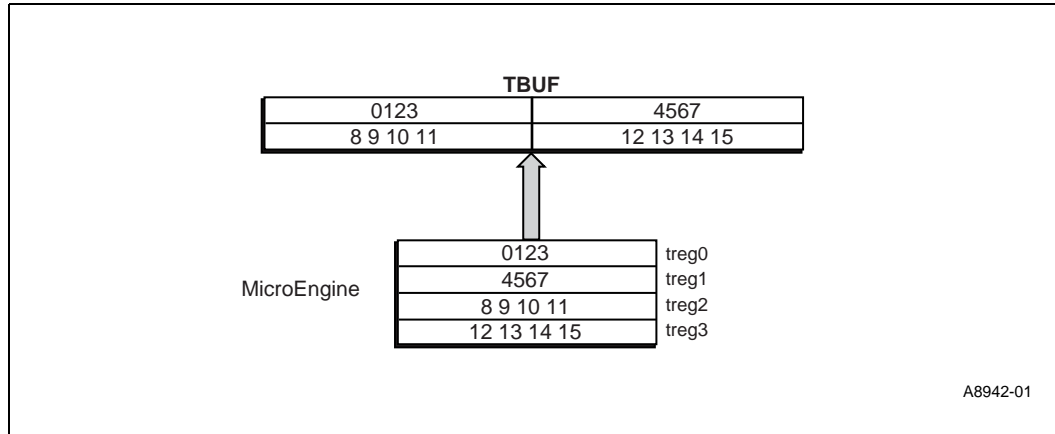
Figure 60. Read from RBUF (64 Bits)



4.5.1.2 Write to TBUF

Data in TBUF is arranged in LWBE order. When writing from the Microengine transfer registers to TBUF, treg0 goes into LDW0, treg1 goes into LDW1, etc. See [Figure 61](#).

Figure 61. Write to TBUF (64 Bits)



4.5.1.3 Read/Write from/to SRAM

Data inside SRAM is in big-endian order. While transferring data from SRAM to a Microengine, no endianness is involved and first read data goes into the first transfer register specified, the next read data into the second, etc.

4.5.1.4 Read/Write from/to DRAM

Data inside DRAM is in LWBE order. When a Microengine reads from DRAM, LDW0 goes into the first transfer register specified in the instruction, LDW1 goes into the next, and so on. While writing to DRAM, treg0 goes first, then followed by treg1, and both are combined in the DRAM controller as {LDW1, LDW0} and written as a 64-bit quantity into DRAM.

4.5.1.5 Read/Write from/to SHaC and Other CSRs

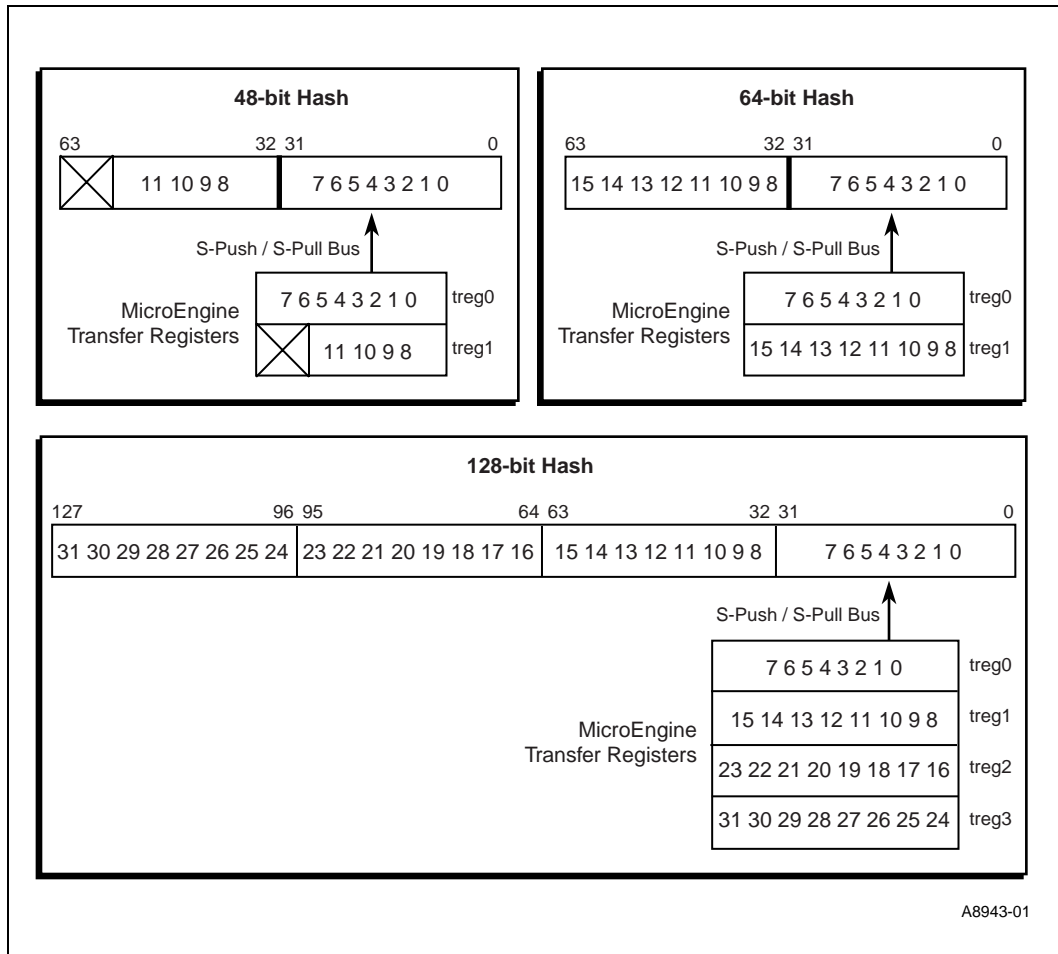
Read and write from SHaC and other CSRs happen as 32-bit operations only and are endian-independent. The low byte goes into the low byte of the transfer register and the high byte goes into the high byte of the transfer register.



4.5.1.6 Write to Hash Unit

Figure 62 explains 48-, 64-, and 128-bit hash operations. When the Microengine transfers a 48-bit hash operand to the hash unit, the operand resides in two transfer registers and is transferred, as shown in Figure 62. In the second longword transfer, only the lower half is valid. Hash unit concatenates the two longwords as shown in Figure 62. Similarly, 64-bit and 128-bit hash operand transfers from the Microengine to the hash unit happen as shown in Figure 62.

Figure 62. 48-, 64-, and 128-Bit Hash Operand Transfers



4.5.2 Media Access

Media operation can be divided in two parts:


- Read from RBUF (Section 4.5.2.1)
- Write to TBUF (Section 4.5.2.2)

4.5.2.1 Read from RBUF

To analyze the endianness on the media-receive interface and the way in which bytes are arranged inside RBUF, a brief introduction of how bytes are generated from the serial interface is provided here. Pipe A denotes the serial stream of data received at the serial interface (SERDES). Bit 0 of byte 0 comes first, followed by bit 1, etc. Pipe B converts this bit stream into byte stream byte 0 — byte 7, etc. So, byte 0 currently is the least significant byte received. In Pipe C, before being transmitted to the SPI-4 interface, these bytes are organized in 16-bit words in big-endian order where byte 0 is at B[15:8] and byte 1 is at B[7:0].

When the SPI-4 interface inside the IXP2800 received these 16-bit words, they are put into RBUF in LWBE order where longwords inside one RBUF entry are organized in little-endian order as shown in one RBUF element in Figure 63. In the least-significant-longword, byte 0 is at a higher address than byte 3 (therefore, big-endian). Similarly, in the most-significant-longword, byte 4 is at a higher address than byte 7 (therefore, big-endian). While transferring from RBUF to Microengine, the least significant longword from one RBUF element is transferred first, followed by the most significant longword into the Microengine transfer registers.

Figure 63. Bit, Byte, and Longword Organization in One RBUF Element



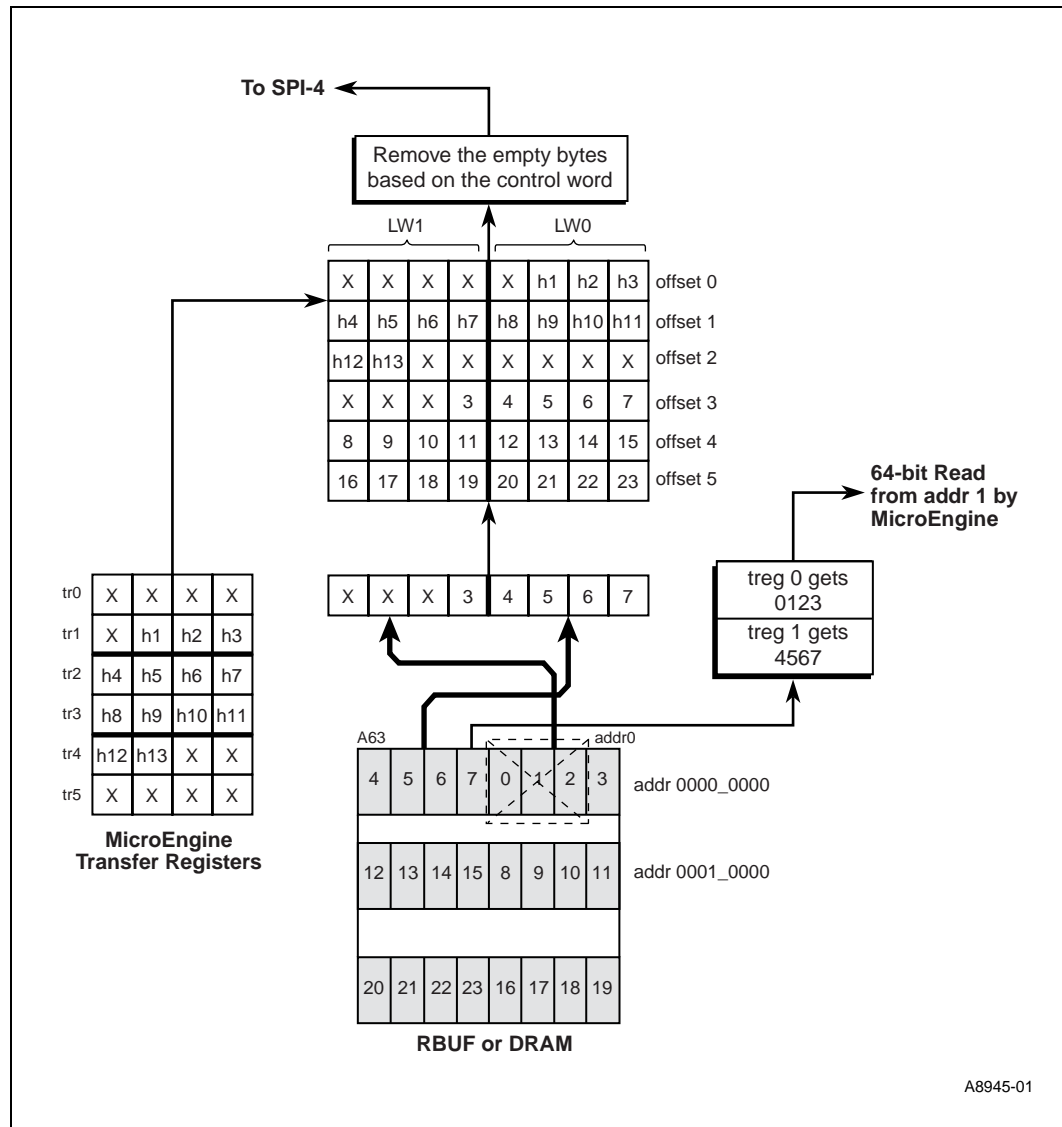
A9725-01



4.5.2.2 Write to TBUF

For writing to TBUF, the header comes from the Microengine and data comes from RBUF or DRAM. Since the Microengine to TBUF header transfer happened in 8-byte chunks, it is possible that the first longword that is inside tr0 may not contain any data if the valid header begins in transfer register tr1. Since data in tr0 goes to the LW1 location at offset 0 and data in tr2 goes to the LW0 location at offset 0, there are some invalid bytes at the beginning of the header, at offset 0. These invalid bytes are removed by the aligner on the way out of TBUF, based on the control word for this TBUF element. The data from tr2, tr3, ... tr6 is placed in TBUF, as shown in Figure 64 in big-endian order.

Figure 64. Write to TBUF



A8945-01




Since data in RBUF or DRAM is arranged in LWBE order, it is swapped on the way into the TBUF to make it truly big-endian, as shown in Figure 64. Again, the invalid bytes at the beginning of the payload that starts at offset 3 and at the end-of-header at offset 2 is removed by the aligner on the way out of TBUF.

4.5.2.3 TBUF to SPI-4 Transfer

Figure 65 shows how the MSF interface removes invalid bytes from TBUF data and transfers them onto the SPI-4 interface in 16-bit (2-byte) chunks.

Figure 65. MSF Interface



A8946-01

This section describes Rambus* DRAM operation.

5.1 Overview

The IXP2800 Network Processor has controllers for three Rambus* DRAM (RDRAM) channels. Either one, two, or three channels can be enabled. When more than one channel is enabled, the channels are interleaved (also known as striping) on 128-byte boundaries to provide balanced access to all populated channels. Interleaving is performed in hardware and is transparent to the programmer. The programmer views the DRAM memory space as a contiguous block of memory.

The total address space of two Gbytes is supported by the DRAM interface regardless of the number of channels that are enabled. The controllers support 64-, 128-, 256-, and 512-Mbyte, and 1-Gbyte devices; however, with interleaving, each of the channels must have the same number, size, and speed of RDRAMs populated. Each channel can be populated with up to 32 RDRAM devices. While each channel must have the same size and speed RDRAMs, it is possible for each individual channel to have different size and speed RDRAMs, as long as the total amount of memory is the same for all of the channels.

ECC (Error Correcting Code) is supported. Enabling ECC requires that x18 RDRAMs be used. If ECC is disabled, x16 RDRAMs can be used.

The Microengines (MEs), Intel XScale® core, and PCI (external Bus Masters and DMA Channels) have access to the DRAM memory space.

The controllers also automatically perform refresh as well as IO driver calibration to account for variations in operating conditions due to process, temperature, voltage, and board layout.

RDRAM Powerdown and nap modes are not supported.



5.2 Size Configuration

Each channel can be populated with 1 – 4 RDRAMs (Short Channel Mode). For supported loading configurations, refer to [Table 61](#). The RAM technology used determines the increment size and maximum memory per channel as shown in [Table 62](#).

Note: One or two channels can be left unpopulated if desired.

Table 61. RDRAM Loading

Bus Interface	Maximum Number of Loads	Trace Length (inches)
Short Channel: 400 and 533 MHz	4 devices per channel.	20 ¹
Long Channel: 400 MHz	2 RIMMs per channel – a maximum of 32 devices in both RIMMs.	20 ¹
Long Channel: 533 MHz	1 RIMM and 1 C-RIMM per channel – a maximum of 16 devices.	20 ¹

1. For termination, the DRAMs should be located as close as possible to the IXP2800 Network Processor.

Table 62. RDRAM Sizes

RDRAM Technology ¹	Increment Size	Maximum per Channel
64/72 MB	8 MBs	256 MB
128/144 MB	16 MB	512 MB
256/288 MB	32 MB	1 GB ²
512/576 MB	64 MB	2 GB ²

NOTES:

- The two numbers shown for each technology indicate x16 parts and x18 parts.
- The maximum memory that can be addressed across all channels is 2 Gbytes. This limitation is based on the partitioning of the 4-Gbyte address space (32-bit addresses). Therefore, if all three channels are used, each can be populated up to a maximum of 768 Mbytes. Two channels can be populated to a maximum of 1 Gbyte each. A single channel could be populated to a maximum of 1 Gbyte.

RDRAMs with 1 x 16 dependent banks, 2 x 16 dependent banks, and four independent banks are supported.

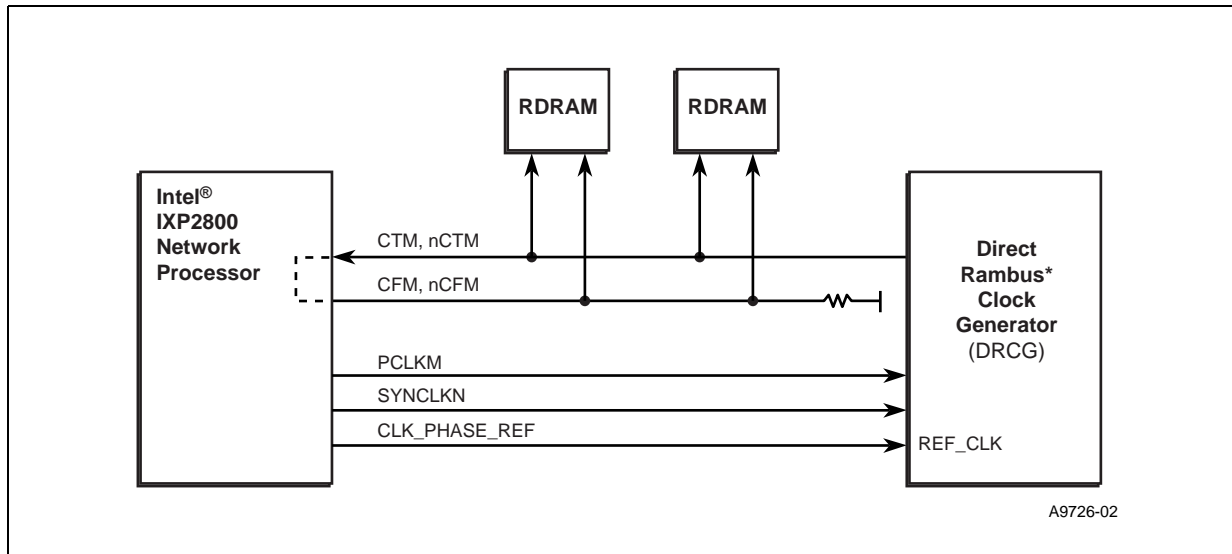


5.3 DRAM Clocking

Figure 66 shows the clock generation for one channel (this description is just for reference; for more information, refer to Rambus* design literature). The other channels use the same configuration.

Note: Refer to Section 10 for additional information on clocking.

Figure 66. Clock Configuration



The RDRAM Controller receives two clocks, both generated internal to the IXP2800 Network Processor.

The internal clock, is used to control all logic associated with communication with other on-chip Units. This clock is $\frac{1}{2}$ of the Microengine frequency, and is in the range of 500 – 700 MHz.

The other clock, the Rambus* Memory Controller (RMC) clock, is internally divided by two and brought out on the CLK_PHASE_REF pin, which is then used as the reference clock for the DRCG (see Figure 67 and Figure 68). The reason for this is that our internal RMC clock is derived from the Microengine clock (supported programmable divide range is from 8 – 15 for the A stepping and 6 – 15 for the B stepping) at a Microengine frequency of 1.4 GHz (the available RMC clock frequencies are 100 MHz and 127 MHz). In the RMC implementation, we have a fixed 1:1 clock relationship between the RMC clock and the SYNCLK ($\text{SYNCLK} = \text{Clock-to-Master(CTM)}/4$); to maintain this relationship, we provide the clock to the DRCG. The CTM is received by the DRAM controller which it drives back out as Clock-from-Master (CFM). Additionally, the controller creates PCLKM and SYNCLKN, which are also driven to the DRCG.

Figure 67. IXP2800 Clocking for RDRAM at 400 MHz

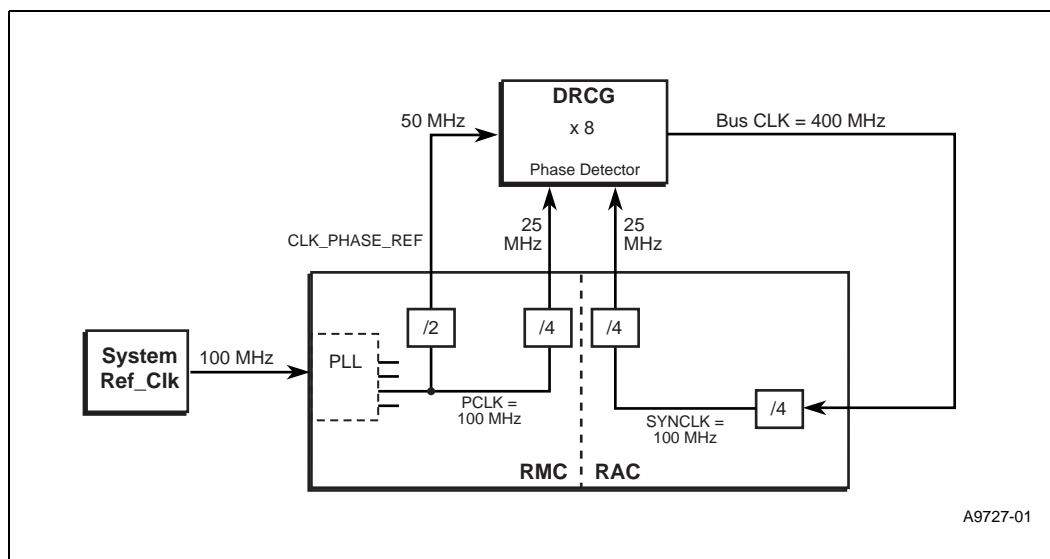
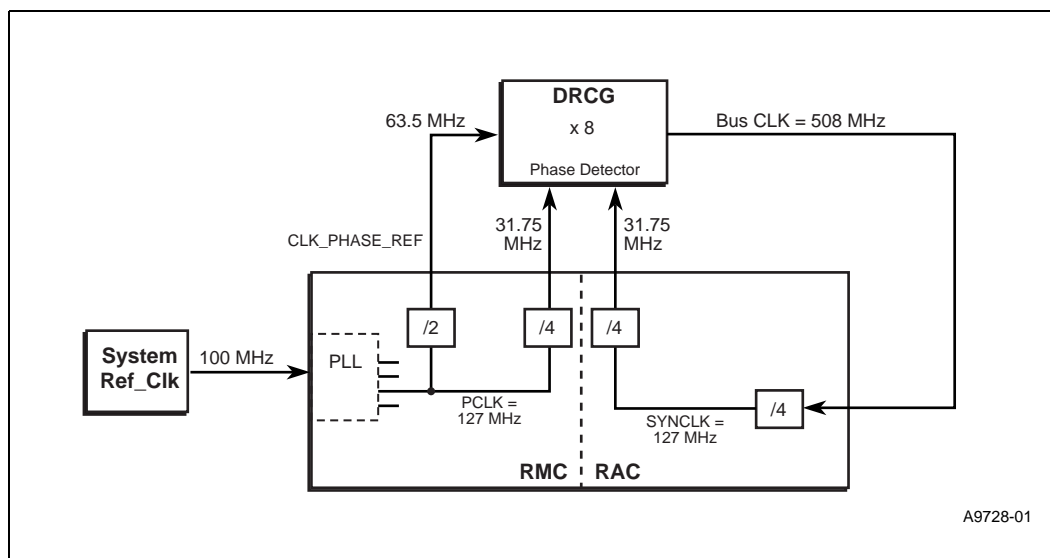


Figure 68. IXP2800 Clocking for RDRAM at 508 MHz



5.4 Bank Policy

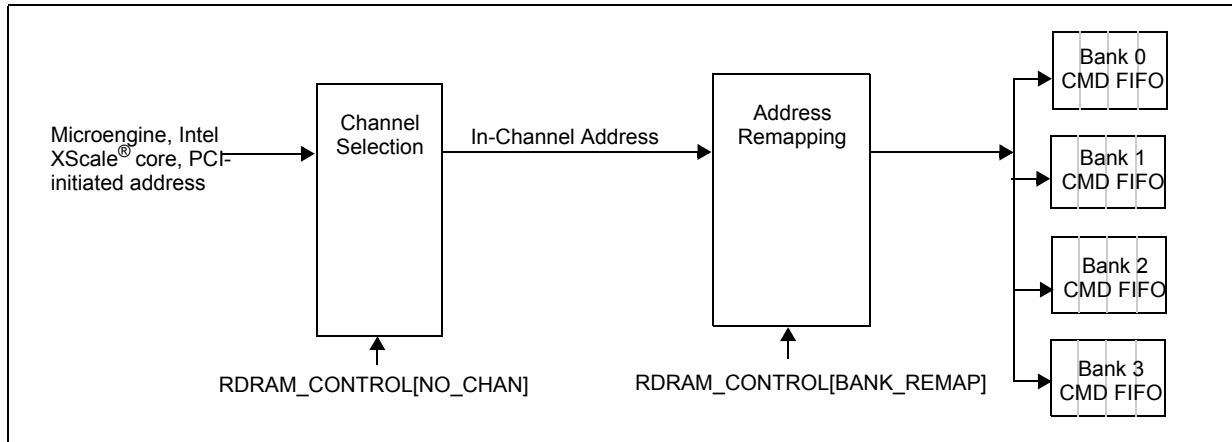
The RDRAM Controller uses a “closed bank” policy. Banks are activated long enough to do an access and then closed and precharged. They are not left open in anticipation of another access to the same page. This is unlike many CPU applications, where there is a high degree of locality. Since that locality does not exist in the typical applications in which the IXP2800 Network Processor uses RDRAM, the “closed bank” policy is used.



5.5 Interleaving

The RDRAM channels are interleaved on 128-byte boundaries in hardware to improve concurrency and bandwidth utilization. Contiguous addresses are directed to different channels by rearranging the physical address bits in a programmable manner described in [Section 5.5.1](#) through [Section 5.5.3](#) and then remapped as described in [Section 5.5.4](#). The block diagram in [Figure 69](#) illustrates the flow.

Figure 69. Address Mapping Flow



The mapping of addresses to channels is completely transparent to software. Software deals with physical addresses in RDRAM space; the mapping is done completely by hardware.

Note: Accessing an address above the amount of RDRAM populated will cause unpredictable results.

5.5.1 Three Channels Active (3-Way Interleave)

When all three channels are active, the interleave scheme selects the channel for each block, using modulo-3 reduction (address bits [31:7] are summed as modulo-3, and the remainder is the selected channel number). The algorithm ensures that adjacent blocks are mapped to different channels.

The address within the DRAM is then selected by rearranging the received address, as shown in [Table 63](#). In this case, the number of DRAMs on a channel must be either 1, 2, 4, 8, 16, or 32.

For Rev. B, the address within the DRAM is selected by adding the received address to the contents of one of the CSRs (K0 – K11), or 0, as shown in [Table 64](#). The values to load into K0 – K11 are a function of the amount of memory on the channel, and are specified in the *IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

For memory sizes of 32, 64, or 128 Mbytes, etc., the specified constants give the same remapping as was done in a previous revision.



Table 63. Address Rearrangement for 3-Way Interleave (Sheet 1 of 2)

When these bits of address are all "1"s... ¹	Shift 30:7 right by this many bits	Add this amount to shifted 30:7 (based on amount of memory on the channel) Address within channel is {30:7+table_value}, 6:0}							
		8 MB ³	16 MB	32 MB ³	64 MB	128 MB ³	256 MB	512 MB ³	1 GB
30:7	26	N/A ²	N/A	N/A	N/A	N/A	N/A	N/A	8388607
28:7	24	N/A	N/A	N/A	N/A	N/A	2097151	4194303	8388606
26:7	22	N/A	N/A	N/A	524287	1048575	2097150	4194300	8388600
24:7	20	N/A	131071	262143	524286	1048572	2097144	4194288	8388576
22:7	18	65535	131070	262140	524280	1048560	2097120	4194240	8388480
20:7	16	65532	131064	262128	524256	1048512	2097024	4194048	8388096
18:7	14	65520	131040	262080	524160	1048320	2096640	4193280	8386560
16:7	12	65472	130944	261888	523776	1047552	2095104	4190208	8380416
14:7	10	65280	130560	261120	522240	1044480	2088960	4177920	8355840
12:7	8	64512	129024	258048	516096	1032192	2064384	4128768	8257536
10:7	6	61440	122880	245760	491520	983040	1966080	3932160	7864320
8:7	4	49152	98304	196608	393216	786432	1572864	3145728	6291456
None	2	0	0	0	0	0	0	0	0

NOTES:

- This is a priority encoder; when multiple lines satisfy the condition, the line with the largest number of ones is used.
- N/A means not applicable.
- For these cases, the top 3 blocks (each block is 128 bytes) of the logical address space is not accessible. For example if each channel has 8 Mbytes, only (24 Mbytes - 384) total bytes are usable. This is an artifact of the remapping method.
- The numbers in the table are derived as follows:
For the first pair of ones (8:7) value is 3/4 the number of blocks. For each subsequent pair of ones, the value is the previous value, plus another 3/4 the remaining blocks.
 - [8:7]==11 - 3/4 * blocks
 - [10:7]==1111 - (3/4 + 3/16) * blocks
 - [12:7]==111111 - (3/4 + 3/16 + 3/64) * blocks
 - etc.



Table 64. Address Rearrangement for 3-Way Interleave (Sheet 2 of 2) (Rev B)

When these bits of address are all "1" ¹	Add the value in this CSR to the address
30:7	K11
28:7	K10
26:7	K9
24:7	K8
22:7	K7
20:7	K6
18:7	K5
16:7	K4
14:7	K3
12:7	K2
10:7	K1
8:7	K0
None	Value 0 added.
NOTES:	
1. This is a priority encoder; when multiple lines satisfy the condition, the line with the largest number of ones is used.	

5.5.2 Two Channels Active (2-Way Interleave)

It is possible to have only two channels populated for system cost and area savings. If only two channels are desired, then channels 0 and 1 should be populated and channel 2 should be left empty. In the Two Channel Mode, the address interleaving is designed with the goal of spreading adjacent accesses across the 2 channels.

When two channels are active, address bit 7 is used as the channel select. Addresses that have address 7 equal to 0 are mapped to channel 0 while those with address 7 equal to 1 are mapped to channel 1. The address within the channel is {[31:8], [6:0]}.

5.5.3 One Channel Active (No Interleave)

When only one channel is active, all accesses go to that channel. In this case, it is possible for an access to split across two DRAM banks (which could be in different RDRAMs).



5.5.4 Interleaving Across RDRAMs and Banks

In addition to interleaving across the different RDRAM channels, addresses are also interleaved across RDRAM chips and internal banks. This improves utilization since certain operations to different banks can be performed concurrently. The interleaving is done based on rearranging the remapped address derived from [Section 5.5.1](#), [Section 5.5.2](#), and [Section 5.5.3](#) as a function of the memory size as shown in [Table 65](#). The two MSBs of the rearranged address are used to select which Bank Command FIFO the command is place in. The rearranged address is also partitioned to choose RDRAM chip, bank within RDRAM, and page within bank.

Table 65. Address Bank Interleaving

Memory Size on Channel (MB) ³	Remapped Address Based on RDRAM_Control[Bank_Remap]			
	00	01	10	11
8	7:14, 22:15	9:14, 7:8, 22:15	11:14, 7:10, 22:15	13:14, 7:12, 22:15
16	7:14, 23:15	9:14, 7:8, 23:15	11:14, 7:10, 23:15	13:14, 7:12, 23:15
32	7:14, 24:15	9:14, 7:8, 24:15	11:14, 7:10, 24:15	13:14, 7:12, 24:15
64	7:14, 25:15	9:14, 7:8, 25:15	11:14, 7:10, 25:15	13:14, 7:12, 25:15
128	7:14, 26:15	9:14, 7:8, 26:15	11:14, 7:10, 26:15	13:14, 7:12, 26:15
256	7:14, 27:15	9:14, 7:8, 27:15	11:14, 7:10, 27:15	13:14, 7:12, 27:15
512	7:14, 28:15	9:14, 7:8, 28:15	11:14, 7:10, 28:15	13:14, 7:12, 28:15
1024	7:14, 29:15	9:14, 7:8, 29:15	11:14, 7:10, 29:15	13:14, 7:12, 29:15
Bits used to select Bank Command FIFO	7:8	9:10	11:12	13:14

NOTES:

1. Table shows device/bank sorting of the channel remapped block address, which is in address 31:7. LSBs of the address are always 6:0 (byte within the block), which are not remapped.
2. Unused MSBs of address have value of 0.
3. Size is programmed in RDRAM_Control[Size].

5.6 Parity and ECC

DRAM can be optionally protected by byte parity or by an 8-bit error detecting and correcting code (ECC). RDRAMn_Control[ECC] for each channel selects whether or not that channel should use Parity, ECC, or no protection. When parity or ECC is enabled x18 RDRAMs must be used with the extra bits connected to the dqa[8] and dqb[8] signals. Eight bits of ECC code cover eight bytes of data (aligned to an 8-byte boundary).

5.6.1 Parity and ECC Disabled

- On reads, the data is delivered to the originator of the read; no error is possible.
- Partial writes (writes of less than eight bytes) are done as masked writes.



5.6.2 Parity Enabled

On writes, odd byte parity is computed for each byte and written into the corresponding parity bit. Partial writes (writes of less than eight bytes) are done as masked writes.

On reads, odd byte parity is computed on each byte of data and compared to the corresponding parity bit. If there is an error RDRAMn_Error_Status_1[Uncorr_Err] bit is set, which can interrupt the Intel XScale® core if enabled. The Data Error signal will be asserted when the read data is delivered on D_Push_Data.

The address of the error, along with other information, is logged in RDRAMn_Error_Status_1[ADDR] and RDRAMn_Error_Status_2. Once the error bit is set, those registers are locked. That is, the information relating to subsequent errors is not loaded until the error status bit is cleared by the Intel XScale® core write.

5.6.3 ECC Enabled

On writes, eight ECC check bits are computed based on 64 bits of data, and are written into the check bits. Partial writes (writes of less than eight bytes) cause the channel controller to do a read-modify-write. Any single-bit error detected during the read portion is corrected prior to merging with the write data. An uncorrectable error detected during the read does not modify the data. Either type of error will set the appropriate error status bit, as described below.

On reads, the correct value for the check bits is computed from the data and is compared to the ECC check bits. If there is no error, data is delivered to the originator of the read, because it came from the RDRAMs. If there is a single-bit error, it is corrected before being delivered (the correction is done automatically, the reader is given the correct data). The error is also logged by setting the RDRAMn_Error_Status_1[Corr_Err] bit, which can interrupt the Intel XScale® core if enabled.

If there is an uncorrectable error, the RDRAMn_Error_Status_1[Uncorr_Err] bit is set, which can interrupt the Intel XScale® core if enabled. The Data Error signal is asserted when the read data is delivered on D_Push_Data, unless the token, Ignore Data Error, was asserted in the command. In that case, the RDRAM controller does not assert Data Error and does not assert a Signal (it will use 0xF, which is a null signal, in place of the requested signal number).

In both correctable and uncorrectable cases, the address of the error, along with other information, is logged in RDRAMn_Error_Status_1[ADDR] and RDRAMn_Error_Status_2. Once either of the error bits is set, those registers are locked. That is, the information relating to subsequent errors is not loaded until both error status bits are clear. That does not prevent the correction of single-bit errors, only the logging.

Note: When a single-bit error is corrected, the corrected data is not written back into memory (scrubbed) by hardware; this can be done by software if desired, because all of the information pertaining to the error is logged.



To avoid the detection of false ECC errors, the RDRAM ECC mode must be initialized using the procedure described below:

1. Ensure that parity/ECC is not enabled: program DRAM_CTRL[15:14] = 00
2. Write all zeros (0x00000000) to all the memory locations. By default, this initializes the memory with odd parity and in this case (data all 0), it coincides with ECC and does not require any read-modify-writes because ECC is not enabled.
3. Ensure that all of the writes are completed prior to enabling ECC. This is done by performing a read operation to 1000 locations.
4. Enable ECC mode: program DRAM_CTRL[15:14] accordingly.

5.6.4 ECC Calculation and Syndrome

The ECC check bits are calculated by forming parity checks on groups of data bits. The check bits are stored in memory during writes via the dq_a[8] and dq_b[8] signals. Note that memory initialization code must put good ECC into all of memory by writing each location before it can be read. Writing any arbitrary data into memory – for example 0, will accomplish this. This will take several milliseconds per Mbyte of memory.

On reads, the expected code is calculated from the data, and then compared to (XORed) the ECC that was read. The result of the comparison is called the syndrome. If the syndrome is equal to 0, then there was no error. There are eight syndromes that are calculated based on the read data and its corresponding ECC bit. When ECC is enabled, upon detecting a single-bit error, the syndrome is used to determine which bit needs to be flipped to correct the error.

5.7 Timing Configuration

Table 66 shows the example timing settings for RDRAMs of various speeds. The parameters are programmed in the RDRAM_Config CSRs (refer to the PRM for register descriptions)

Table 66. RDRAM Timing Parameter Settings

Parameter Name	-40-800	-45-800	-50-800	-45-711	-50-711	-45-600	-53-600
CfgTrcd	7	9	11	7	9	5	7
CfgTrasSyn	5	5	6	5	5	4	5
CfgTrp	8	8	10	8	8	6	8
CfgToffpSyn	4	4	4	4	4	4	4
CfgTrasrefSyn	5	5	6	5	5	4	5
CfgTprefSyn	2	2	2	2	2	2	2



5.8 Microengine Signals

Upon completion of a read or write, the RDRAM controller can signal a Microengine context, when enabled. It does so using the `sig_done` token; see [Example 27](#).

Example 27. RDRAM Controller Signaling a Microengine Context

```
dram [read,$xfer6,addr_a,0,1], sig_done_4
dram [read,$xfer7,addr_b,0,1], sig_done_6
ctx_arb[4, 5, 6, 7]
```

Because the RDRAM address space is interleaved, consecutive accesses can go to different RDRAM channels. There is no ordering guaranteed among different channels, so a separate signal is needed for each.

In addition, because accesses start at any address, and can specify up to 16 64-bit words (128 bytes), they can also split across two channels (refer to [Section 5.5](#)). The `ctx_arb` instruction must set two `Wakeup_Events` (an odd/even pair) per access. The RDRAM controllers coordinate as follows:

- If the access split across two channels, the channel handling the low part of the split delivers the even-numbered Event Signal, and the channel handling the upper part of the split delivers the odd-numbered Event Signal.
- If the access does not split, the channel delivers both Event Signals (by coordinating with the `D_Push` or `D_Pull` arbiter for read and writes respectively).
- In all cases, the channel delivers the Event Signal with the last data Push or Pull of a burst.

Using the above rules, the Microengine will be put into the Ready State (ready to resume executing) only when all accesses have completed.

5.9 Serial Port

The RDRAM chips are configured through a serial port, which consists of signals `D_SIO`, `D_CMD`, and `D_SCK`. Access to the serial port is via the `RDRAM_Serial_Command` and `RDRAM_Serial_Data` CSRs (refer to the *IXP2400 and IXP2800 Network Processor Programmer's Reference Manual* for the register descriptions).

All serial commands are initiated by a write to `RDRAM_Serial_Command`. Because the serial port is slow, `RDRAM_Serial_Command` has a Busy bit, which indicates that a serial port command is in progress. Software must test this bit before initiating a command. This ensures that software will not lose a command, while eliminating the need for a hardware FIFO for serial commands.

Serial *writes* are done by the following steps:

1. Read `RDRAM_Serial_Command`; test Busy bit until its a 0.
2. Write `RDRAM_Serial_Data`.
3. Write `RDRAM_Serial_Command` to start the write.

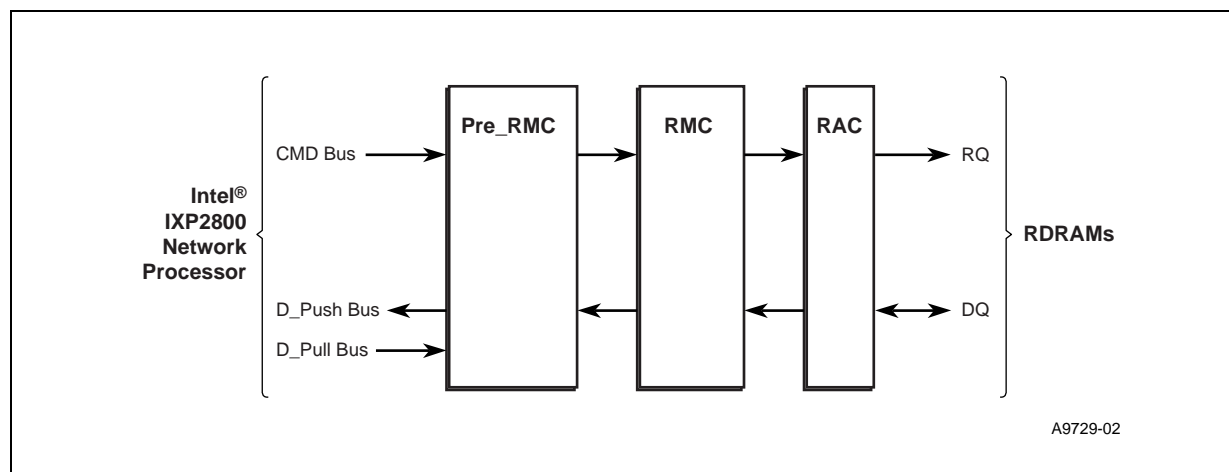
Serial *reads* are done by the following steps:

1. Read RDRAM_Serial_Command; test Busy bit until it is a 0.
2. Write RDRAM_Serial_Command to start the read.
3. Read RDRAM_Serial_Command; test Busy bit until it is a 0.
4. Read RDRAM_Serial_Data to collect the serial read data.

5.10 RDRAM Controller Block Diagram

The RDRAM controller consists of three pieces. Figure 70 is a simplified block diagram.

Figure 70. RDRAM Controller Block Diagram



Pre_RMC — has the queues for commands, data (both in and out), and interfaces to internal buses. It checks incoming commands and addresses to determine if they are targeted to the channel, and if so, enqueues them (if a command splits across two channels, the channel must enqueue the portion of the command that it owns). It sorts the enqueued commands to RDRAM banks, selects the command to be executed based on policy to get good bank utilization, and then hands off that command to RMC. It also arbitrates for refresh and calibration, which it requests RMC to perform. Pre_RMC also contains the ECC logic, and the CSRs that set size, timing, ECC, etc.

RMC — Rambus* Memory Controller, that handles the pin protocol. It controls all timing dependencies, pin turnaround, RAS-CAS, RAS-RAS, etc., including bank interactions. RMC handles all commands in the order that it receives them. RMC is based on the Rambus* RMC.

RAC — Rambus* ASIC Cell, a high-speed parallel-to-serial and serial-to-parallel interface. This is a hard macro that contains the I/O pads and drivers, DLL, and associated pin interface logic.

The following is a brief explanation of command operation:

Pre_RMC enqueues commands and sends them to RMC. It is responsible for initiating Pull operations to get Microengine/RBUF/Intel XScale® core/PCI data into the Pull_Data FIFO. A write is not eligible to go to RMC until Pre_RMC has all the data in the Pull Data FIFO.

Pre_RMC provides the Full signal to the Command Arbiter to inform it stop allowing RDRAM commands.



5.10.1 Commands

When a valid command is placed on the command bus, the control logic checks to see if the address matches the channel's address range, based on interleaving as described in [Section 5.5](#). The command, address, length, etc. are enqueued into the command Inlet FIFO.

If the command Inlet FIFO becomes full, the channel sends a signal to the command arbiter, which will prevent it from sending further DRAM commands. The full signal must be asserted while there is still enough room in the FIFOs to hold the worst case number of in-flight commands.

5.10.2 DRAM Write

When a write (or RBUF_RD, which does a DRAM write) command is at the head of the Command Inlet FIFO, it is moved to the proper Bank CMD FIFO, and the Pull_ID is sent to the Pull arbiter. This can only be done if there is room for the command in the Bank's CMD FIFO and for the pull data in the Bank's Pull Data FIFO (which must take into account all pull data in flight). If there is not enough room in the Bank's CMD FIFO or the Bank's Pull Data FIFO, the write command waits at the head of the Command Inlet FIFO. When the Pull_ID is sent to the Pull Arbiter, the Bank number is put into the PP (Pull in Progress) FIFO; this allows the channel to sort the Pull Data into the proper Bank Pull Data FIFO when it arrives.

The source of the Pull Data can be either RBUF, PCI, Microengine, or the Intel XScale® core, and is specified in the Pull_ID. When the source is RBUF or PCI, data will be supplied to the Pull Data FIFO, at 64 bits per cycle. When the source is Microengine or the Intel XScale® core, data will be supplied at 32 bits per cycle, justified to the low 32 bits of Pull Data. The Pull Arbiter must merge and pack data as required. In addition, the data must be aligned according to the start address, which has longword resolution; this is done in Pre_RMC.

The Length field of the command at the head of the Bank CMD FIFO is compared to the number of 64-bit words in the Bank Pull_Data FIFO. When the number of 64-bit words in Pull_Data FIFO is greater or equal to the length, the write arbitrates for the RMC. When it wins arbitration, it sends the address and command to RMC, which requests the write data from Pull_Data FIFO at the proper time to send it to the RDRAMs.

Note: The Microengine is signaled when the last data is pulled.

5.10.2.1 Masked Write

Masked writes (write of less than eight bytes) are done as either Read-Modify-Writes when ECC is enabled, or as Rambus*-masked writes (using COLM packets), when ECC is not enabled. In both cases, the masked write will modify seven or fewer bytes because the command bus limits a masked write to a ref_count of 1.

If a RMW is used, no commands from that Bank's CMD FIFO are started between the read and the write; other Bank commands can be done during that time.



5.10.3 DRAM Read

When a read (or TBUF_WR, which does a DRAM read) command is at the head of the Command Inlet FIFO, it is moved to the proper Bank CMD FIFO if there is room. If there is not enough room in the Bank's CMD FIFO, the read command waits at the head of the Command Inlet FIFO.

When a read command is at the head of the Bank CMD FIFO, and there is room for the read data in the Push Data FIFO (including all reads in flight at the RDRAM), it will arbitrate for RMC. When it wins arbitration, it sends the address and command to RMC. The Push_ID is put into the RP FIFO (Read in Progress) to coordinate it with read data from RMC.

When read data is returned from RMC, it is placed into the Push_Data FIFO. Each Push_Data is sent to the Push Arbiter with a Push_ID; the RDRAM controller increments the Push_ID for each data phase. If Push Arbiter asserts the full signal, Push Data is stopped and held in the Push Data skid FIFO. The Push Data is sent to the read destination under control of the Push Arbiter.

The destination of the Push Data can be either Intel XScale® core, PCI, TBUF, or Microengine, and is specified in the Push_ID. When the destination is TBUF or PCI, data is taken at 64 bits per cycle. When the destination is the Microengine or the Intel XScale® core, data is taken at 32 bits per cycle. The Push Arbiter justifies the data to the low 32 bits of Push Data. The Microengine is signaled when the last data is pushed.

5.10.4 CSR Write

When a CSR write command is at the head of the Command Inlet FIFO, it is moved to the CSR CMD register, and the Pull_ID is sent to the Pull arbiter. This can only be done if the CSR CMD register is not currently occupied. If it is, the CSR write command waits at the head of the Command Inlet FIFO.

When the Pull_ID is sent to the Pull Arbiter, a tag is put into the PP FIFO (Pull in Progress); this allows the channel to identify the Pull Data as CSR data when it arrives. When the CSR pull data arrives, it is put into the addressed CSR, and the CSR CMD register is marked as empty.

5.10.5 CSR Read

When a CSR read command is at the head of the Command Inlet FIFO, it is moved to the CSR CMD register. This can only be done if the CSR CMD register is not currently occupied. If it is, the CSR read command waits at the head of the Command Inlet FIFO.

On the first available cycle in which RDRAM data from RMC is not being put into the Push Data FIFO, the CSR data will be put into the Push Data FIFO. If it is convenient to guarantee a slot by putting a bubble on the RMC input, then that will be done.



5.10.6 Arbitration

The channel needs to arbitrate among several different operations at RMC. Arbitration rules are given here for those cases: from highest to lowest priority:

- Refresh RDRAM.
- Current calibrate RDRAM.
- Bank operations. When there are multiple bank operations ready, the rules are: (1) round robin among banks to avoid bank collisions, and (2) skip a bank to avoid DQ bus turnarounds. No bank can be skipped more than twice.

Commands are given to RMC in the order in which they will be executed.

5.10.7 Reference Ordering

Table 67 lists the ordering of reads and writes to the same address for DRAM. The definition of first and second is defined by the time the command is valid on the command bus.

Table 67. Ordering of Reads and Writes to the Same Address for DRAM

First Access	Second Access	Ordering Rules
Read	Read	None. If there are no side-effects on reads, both readers get the same data.
Read	Write	Reader must get the pre-modified data. This is <i>not</i> enforced in hardware. The write instruction must not be executed until after the Microengine receives the signal of read completion (i.e., program must use <code>sig_done</code> on the read).
Write	Read	Reader must get the post-modified data. This is <i>not</i> enforced in hardware. The read instruction must not be executed until after the Microengine receives the signal of write completion (i.e., program must use <code>sig_done</code> token on the write instruction and wait for the signal before executing the read instruction).
Write	Write	The hardware guarantees that the writes complete in the order in which they are issued.

5.11 DRAM Push/Pull Arbiter

The DRAM Push/Pull Arbiter contains the push and pull arbiters for the D-Cluster (DRAM Cluster). Both the PUSH and PULL data buses have multiple masters and multiple targets. The DRAM Push/Pull Arbiter determines which master gets to drive the data bus for a given transaction and makes sure that the data is delivered correctly.

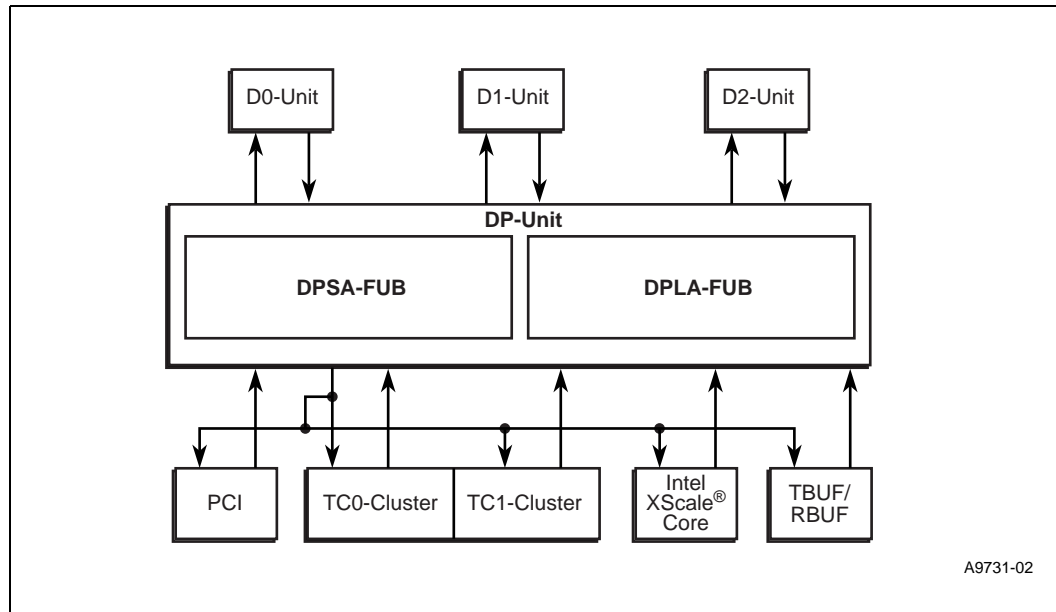
This unit has the following features:

- Up to three DX Unit (DRAM Unit) masters.
- 64-bit wide push and pull data buses.
- Round-robin arbitration scheme.
- Peak delivery of 64 bits per cycle.
- Supports third-party data transfers; the Microengine's can command data movements between the MSF (Media) and either the DX Units or the CR Units.

- Supports chaining for burst DRAM push operations to tell the arbiter to grant consecutive push requests.
- Supports data error bit handling and delivery.

Figure 71 shows the functional blocks for the DRAM Push/Pull Arbiter.

Figure 71. DRAM Push/Pull Arbiter Functional Blocks



5.11.1 Arbiter Push/Pull Operation

Within the arbiter there are two functional units: the push arbiter and the pull arbiter. Push and pull always refer to the way data is flowing from the bus master, i.e., a Microengine makes a read request, the DRAM channel does the read, and then “pushes” the data back to the Microengine.

For a push transaction, a push master drives the command and data to the DRAM push arbiter (DPSA) and into a dedicated request FIFO. When that command is at the head of the FIFO, and it is either the requesting unit’s turn to go based on the round-robin arbitration policy, or there are no other requesters, then the arbiter will “grant” the request. This grant means that the arbiter delivers the push data to the correct target with all the correct handshakes and retires the request (a data transaction is always eight bytes).

The DRAM pull arbiter (DPLA) is slightly different because it functions on bursts of data transactions instead of single transactions. For a pull transaction, a pull master drives a command to the pull arbiter and into a dedicated request FIFO. When the command gets to the head of the FIFO it is evaluated, as was done for the push arbiter. The difference is that each command may reference bursts of data movements (always in multiples of eight bytes). The arbiter grants the command, and keeps it granted until it increments through all of the data movements required by the command. As the data is read from its source, the command is modified to address the next data address, and a handshake to the requesting unit is driven when the data is valid.



5.11.2 DRAM Push Arbiter Description

The general data flow for a push operation is as shown in [Table 68](#). The DRAM Push Arbiter functional blocks are shown in [Figure 72](#).

Table 68. DRAM Push Arbiter Operation

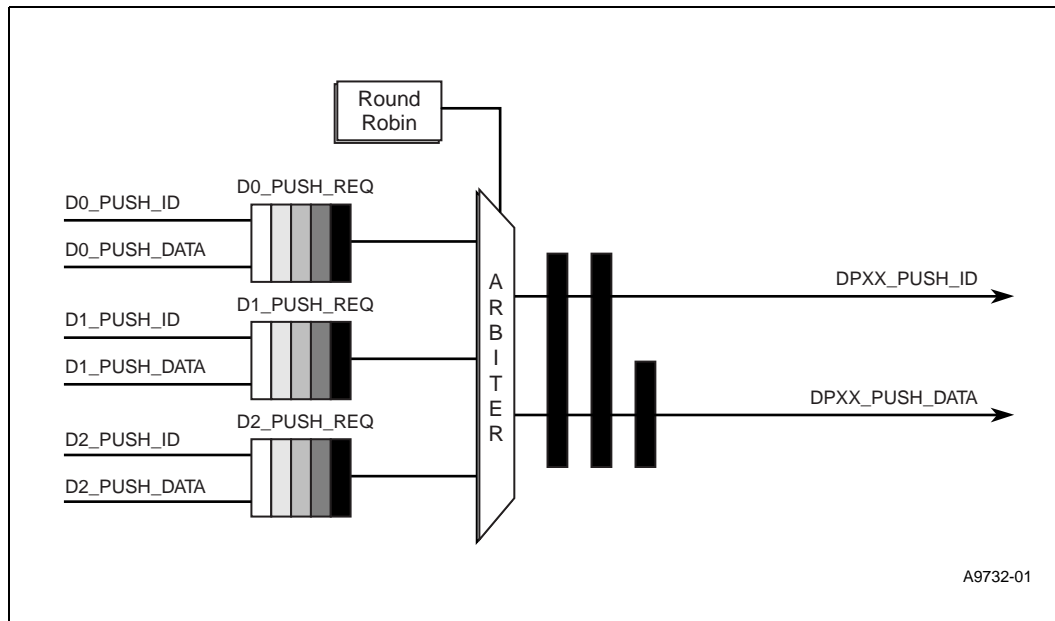
Push Bus Master/Requestor	Data Source	Data Destination
IXP2800 Network Processor		
D0 Unit D1 Unit D2 Unit	Current Master	TC0 Cluster (ME 0 – 7) TC1 Cluster (ME 10 – 17) Intel XScale® core PCI Unit MSF Unit

The push arbiter takes push requests from any requestors. Each requestor has a dedicated request FIFO. A request comes in the form of a PUSH_ID, and is accompanied by the data to be pushed, a data error bit, and a chain bit. All of this information is enqueued in the correct FIFO for each request, i.e., for each eight bytes of data. The push arbiter must drive a full signal to the requestor if the FIFO reaches a predefined “full” level to apply backpressure and stop requests from coming. The FIFO is 64 entries deep and goes full at 40 entries. The long skid allows for burst reads in flight to finish before stalling the DRAM controller. If the FIFO is not full, the push arbiter can enqueue a new request from each requestor on every cycle.

The push arbiter monitors the heads of each FIFO, and does a round robin arbitration between any available requestors. If the chain bit is asserted, it indicates that once the head request of a queue is granted, the arbiter should continue to grant that queue until the chain bit de-asserts. It is expected that the requestor will assert the chain bit for no longer than a full burst length. The push arbiter must also take special notice of requests destined for the receive buffer in the Media Switch Fabric (MSF). Finally, the push arbiter must manage the delivery of data at different rates, depending on how wide the bus is going into a given target.

The Microengines, PCI, and the Intel XScale® core all have 32-bit data buses. For these targets, the push arbiter takes two clock cycles to deliver 64 bits of data by first delivering bits 31:0 in the first cycle, and then putting bits 63:32 onto the low 32 bits of the PUSH_DATA in the second cycle.

Figure 72. DRAM Push Arbiter Functional Blocks



The DRM Push Arbiter boundary conditions are:

- Make sure each of the push_request queues assert the full signal and back pressure the requesting unit.
- Maintain 100% bus utilization, i.e., no holes.

5.12 DRAM Pull Arbiter Description

The general data flow for a push operation is as shown in Table 69. The DRAM Pull Arbiter functional blocks are shown in Figure 73.

Table 69. DPLA Description

Pull Bus Master/Requestor	Data Source	Data Destination
IXP2800 Network Processor		
D0 Unit D1 Unit D2 Unit	TC0 Cluster (Microengine 0 – 7) TC1 Cluster (Microengine 8 – 15) Intel XScale® core PCI Unit MSF Unit	Current Master

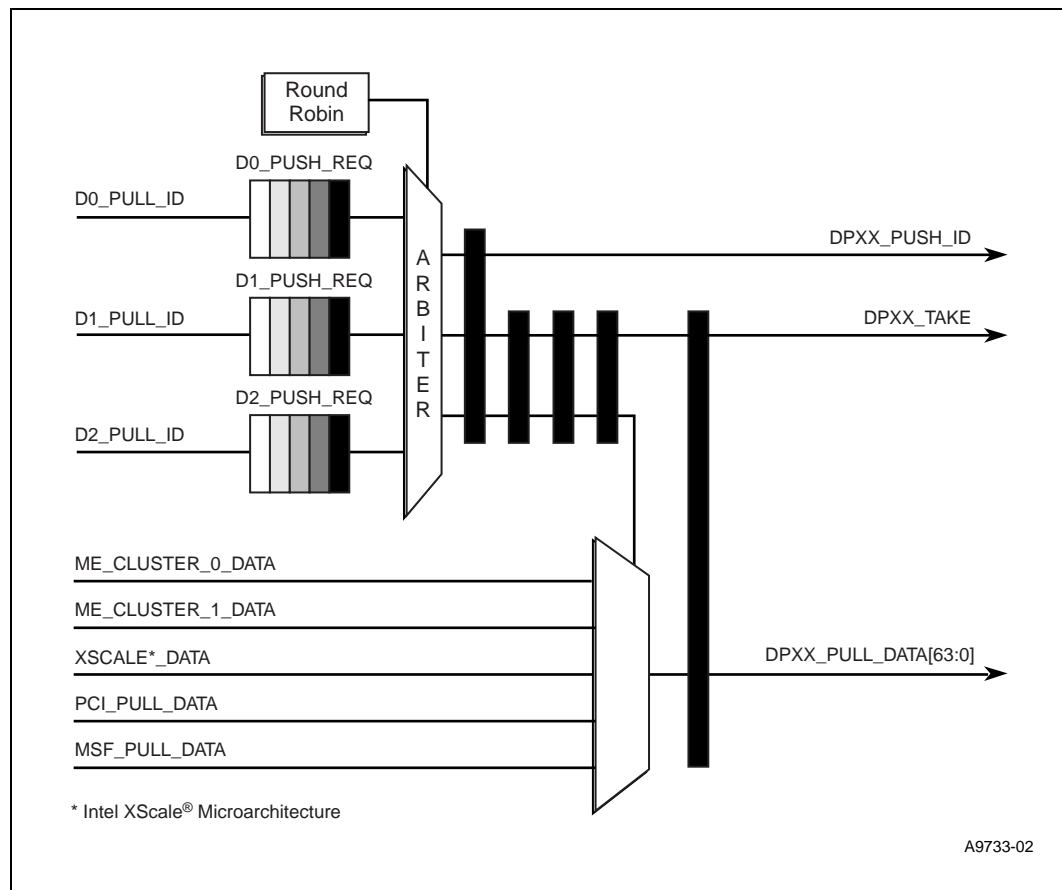
The pull arbiter is very similar to the push arbiter, except that it gathers the data from a data source ID and delivers it to the requesting unit where it is written to DRAM memory.



When a requestor gets a pull command on the CMD_BUS, the requestor sends the command to the pull arbiter. This is enqueued into a requestor-dedicated FIFO. The pull request FIFOs are much smaller than the push request FIFOs because pull requests can request up to 128 bytes of data. It is eight entries deep and asserts full when it has six entries to account for in-flight requests.

The pull arbiter monitors the heads of each of the three FIFOs. A round robin arbitration scheme is applied to all valid requests. When a request is granted, the request is completed regardless of how many data transfers are required. Therefore, one request can take as many as 16 – 32 DRAM cycles. The push data bus can only use 32 bits when delivering data to the Microengines, PCI, and the Intel XScale® core. For these data sources, it takes two cycles to pull every eight bytes requested; otherwise, it takes only one cycle per eight bytes. On four byte cycles, data is delivered as pulled.

Figure 73. DRAM Pull Arbiter Functional Blocks







SRAM Interface

6

6.1 Overview

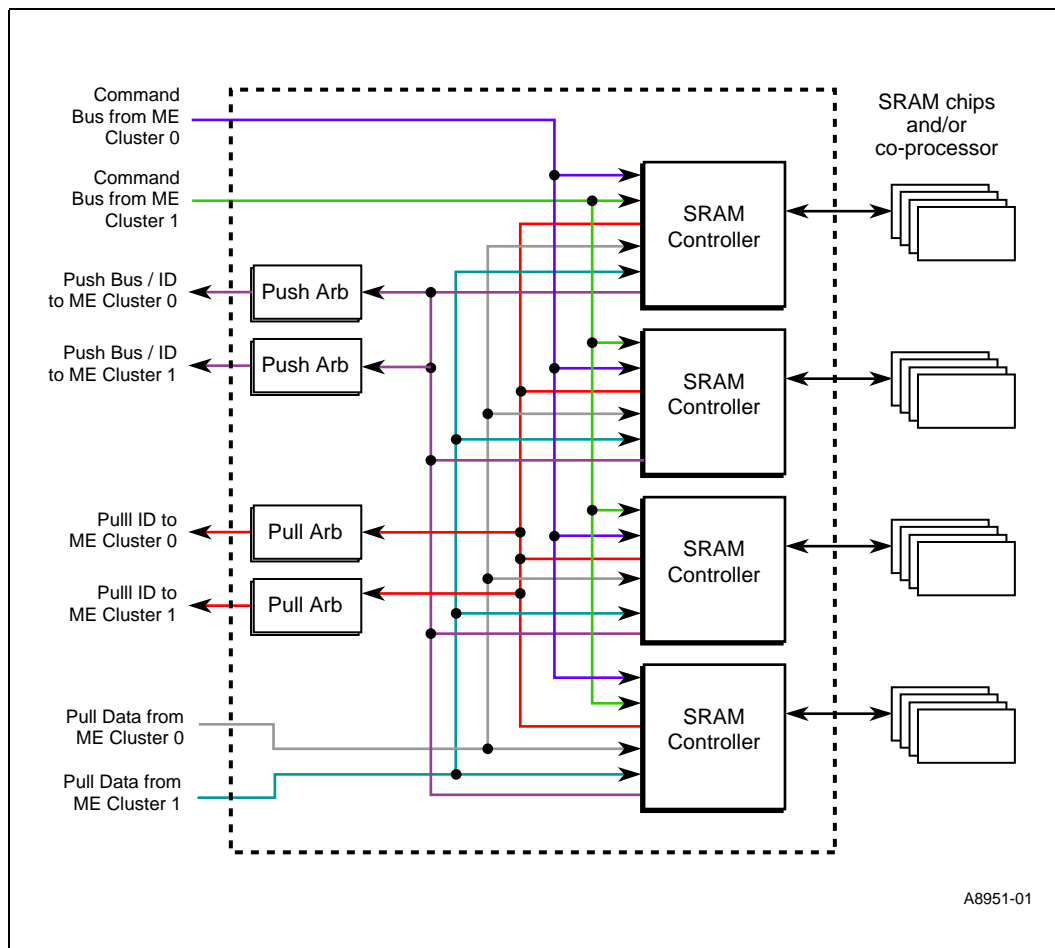
The IXP2800 Network Processor contains four independent SRAM controllers. SRAM controllers support pipelined QDR synchronous static RAM (SRAM) and a coprocessor that adheres to QDR signaling. Any or all controllers can be left unpopulated if the application does not need them.

Reads and writes to SRAM are generated by Microengines (MEs), the Intel XScale® core, and PCI Bus masters. They are connected to the controllers through Command Buses and Push and Pull Buses. Each of the SRAM controllers takes commands from the command bus and enqueues them. The commands are de-queued according to priority, and successive accesses to the SRAMs are performed.

Each SRAM controller receives commands using two Command Buses, one of which may be tied off as inactive, depending on the chip implementation. The SRAM Controller can enqueue a command from each Command Bus in each cycle. Data movement between the SRAM controllers and the Microengines is through the S_Push bus and S_Pull bus.

The overall structure of the SRAM controllers is shown in [Figure 74](#).

Figure 74. SRAM Controller/Chassis Block Diagram



6.2 SRAM Interface Configurations

Memory is logically four bytes (one longword) wide while physically, the data pins are two bytes wide and double-clocked. Byte parity is supported. Each of the four bytes has a parity bit, which is written when the byte is written and checked when the longword is read. There are byte-enables that select the bytes to write, for lengths of less than a longword.

The QDR controller implements a big-endian ordering scheme at the interface pins. For write operations, bytes 0/1, (data bits [31:16]), and associated parity and byte-enables are written on the rising edge of the K clock while bytes 2/3, (data bits [15:0]), and associated parity and byte-enables are written on the rising edge of the K_n clock. For read operations, bytes 0/1, (data bits [31:16]), and associated parity and byte-enables are captured on the rising edge of CIN0 clock while bytes 2/3, (data bits [15:0]), and associated parity and byte-enables are captured on the rising edge of CIN0_n clock.



In general, QDR and QDR II bursts of two SRAMs are supported at speeds up to 233 MHz. As other (larger) QDR SRAMs are introduced, they will also be supported.

The SRAM controller can also be configured to interface to an external coprocessor that adheres to the QDR or QDR II electrical and functional specification.

6.2.1 Internal Interface

Each SRAM channel receives commands through the command bus mechanism and transfers data to and from the Microengines, the Intel XScale® core, and PCI, using SRAM push and SRAM pull buses.

6.2.2 Number of Channels

The IXP2800 Network Processor supports four channels.

6.2.3 Coprocessor and/or SRAMs Attached to a Channel

Each channel supports the attachment of QDR SRAMs, a co-processor, or both, depending on the module level signal integrity and loading.

6.3 SRAM Controller Configurations

There are enough address pins (24) to support up to 64 Mbytes of SRAM. The SRAM controllers can directly generate multiple port enables (up to five pairs) to allow for depth expansion. Two pairs of pins are dedicated for port enables. Smaller RAMs use fewer address signals than the number provided to accommodate the largest RAMs, so some address pins (23:18) are configurable as either address or port-enable based on CSR SRAM_Control[Port_Control] as shown in Table 70.

Note: All of the SRAMs on a given channel must be the same size.

Note: Table 70 shows the capability of the logic — 1, 2, or 4 loads are supported as shown in the table, but this is subject to change.

Table 70. SRAM Controller Configurations

SRAM Configuration	SRAM Size	Addresses Needed to Index SRAM	Addresses Used as Port Enables	Total Number of Port Select Pairs Available
512K x 18	1 Mbyte	17:0	23:22, 21:20	4
1M x 18	2 Mbytes	18:0	23:22, 21:20	4
2M x 18	4 Mbytes	19:0	23:22, 21:20	4
4M x 18	8 Mbytes	20:0	23:22	3
8M x 18	16 Mbytes	21:0	23:22	3
16M x 18	32 Mbytes	22:0	None	2
32M x 18	64 Mbytes	23:0	None	1



Each channel can be expanded in depth according to the number of port enables available. If external decoding is used, then the number of SRAMs is not limited by the number of port enables generated by the SRAM controller.

Note: External decoding may require external pipeline registers to account for the decode time, depending on the desired frequency.

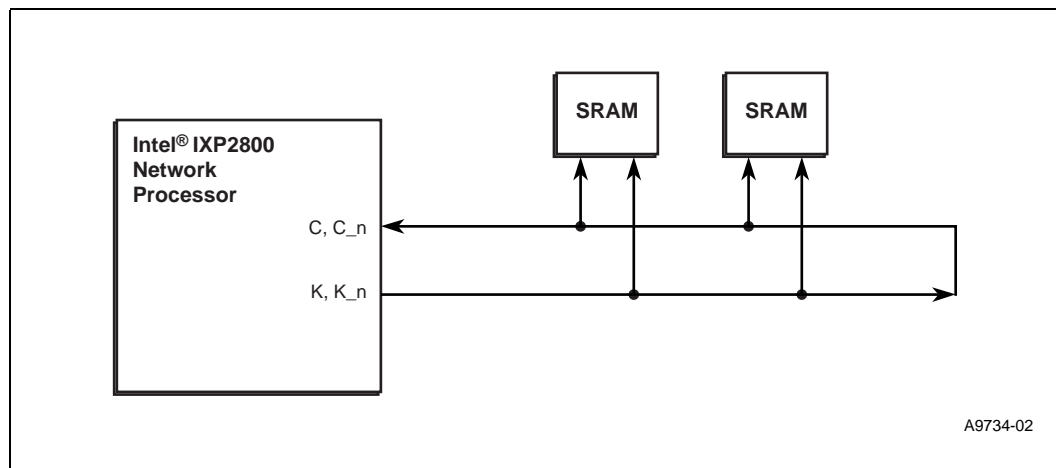
Maximum SRAM system sizes are shown in Table 71. Shaded entries require external decoding, because they use more port-enables than the SRAM controller can directly supply.

Table 71. Total Memory per Channel

SRAM Size	Number of SRAMs on Channel							
	1	2	3	4	5	6	7	8
512K x 18	1 MB	2 MB	3 MB	4 MB	5 MB	6 MB	7 MB	8 MB
1M x 18	2 MB	4 MB	6 MB	8 MB	10 MB	12 MB	14 MB	16 MB
2M x 18	4 MB	8 MB	12 MB	16 MB	20 MB	24 MB	28 MB	32 MB
4M x 18	8 MB	16 MB	24 MB	32 MB	64 MB	NA	NA	NA
8M x 18	16 MB	32 MB	48 MB	64 MB	NA	NA	NA	NA
16M x 18	32 MB	64 MB	NA	NA	NA	NA	NA	NA
32M x 18	64 MB	NA	NA	NA	NA	NA	NA	NA

Figure 75 shows how the SRAM clocks on a channel are connected. For receiving data from the SRAMs, the clock path and data path are matched to meet hold time requirements.

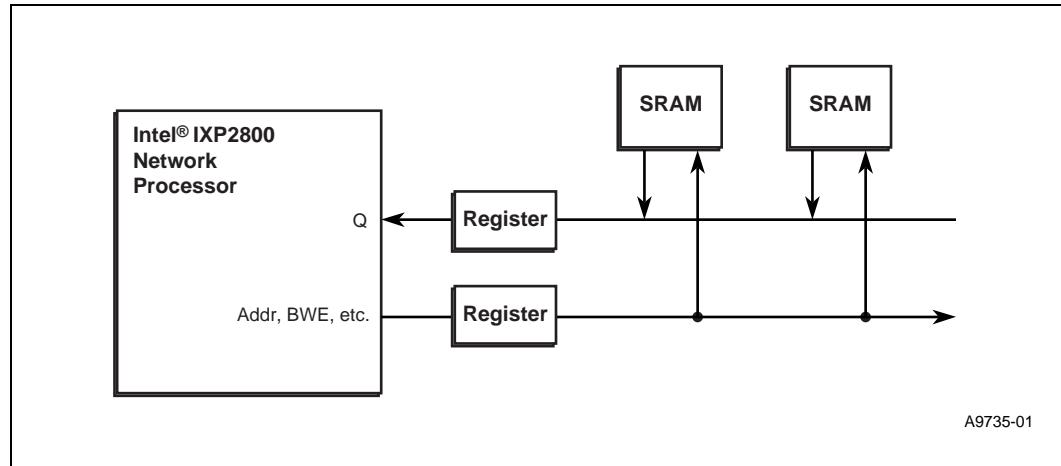
Figure 75. SRAM Clock Connection on a Channel



It is also possible to pipeline the SRAM signals with external registers. This is useful for the case when there is considerable loading on the address and data signals, which would slow down the cycle time. The pipeline stages make it possible to keep the cycle time fast by fanning out the address, byte write, and data signals. The RAM read data may also be put through a pipeline register, depending on configuration. External decoding of port selects can also be done to expand the number of SRAMs supported. Figure 76 is a block diagram that shows the concept of external pipelining.

A side-effect of the pipeline registers is to add latency to reads, and the SRAM controller must account for that delay by waiting extra cycles (relative to no external pipeline registers) before it registers the read data. The number of extra pipeline delays is programmed in SRAM_Control[Pipeline].

Figure 76. External Pipeline Registers Block Diagram



6.4 Command Overview

This section will give an overview of the SRAM commands and their operation. The details will be given later in the document. Memory reference ordering will be specified along with the detailed command operation.

6.4.1 Basic Read/Write Commands

The basic read and write commands will transfer from 1 – 16 longwords of data to or from the QDR SRAM external to the IXP2800 Network Processor.

For a read command, the SRAM is read and the data placed on the Push bus, one longword at a time. The command source (for example, the Microengine) is signaled that the command is complete during the last data phase of the push bus transfer.

For a write command, the data is first pulled from the source, then written to the SRAM in consecutive SRAM cycles. The command source is signaled that the command is complete during the last data phase of the pull bus transfer.

If a read operation stalls due to the pull-data FIFO filling, any concurrent write operation that is in progress to the same address is temporarily stopped. This technique results in atomic data reads.

6.4.2 Atomic Operations

The SRAM Controller does read-modify-writes for the atomic operations, and the pre-modified data can be returned if desired. Other (non-atomic) readers and writers can access the addressed location between the read and write portions of the read-modify-write. Table 72 describes the atomic operations supported by the SRAM Controller.



Table 72. Atomic Operations

Instruction	Pull Operand	Value Written to SRAM
Set_bits	Optional ¹	SRAM_Read_Data or Pull_Data
Clear_bits	Optional	SRAM_Read_Data and not Pull_Data
Increment	No	SRAM_Read_Data + 0x00000001
Decrement	No	SRAM_Read_Data - 0x00000001
Add	Optional	SRAM_Read_Data + Pull_Data
Swap	Optional	Pull_Data

1. There are two versions of the Set, Clear, Add, and Swap instructions. One version pulls operand data from the Microengine transfer registers, and the second version passes the operand data to the SRAM unit as part of the command.

Up to two Microengine signals are assigned to each read-modify-write reference. Microcode should always tag the read-modify-write reference with an even-numbered signal. If the operation requires a pull, the requested signal is sent on the pull. If the pre-modified data is to be returned to the Microengine, then the Microengine is sent (requested signal OR 1) when that data is pushed.

In [Example 28](#), the version of Test_and_Set requires both a pull and a push:

Example 28. SRAM Test_and_Set with Pull Data

```
immed [$xfer0, 0x1]
SRAM[test_and_set, $xfer0, test_address, 0, 1], sig_done_2
// SIGNAL_2 is set when $xfer0 is pulled from this ME. SIGNAL_3 is
// set when $xfer0 is written with the test value. Sleep until both
// SIGNALS have arrived.
CTX_ARB[sigal_2, sigal_3]
```

In [Example 29](#), the version of Test_and_Set does not require a pull, but does issue a push. A signal is generated when the push is complete.

Example 29. SRAM Test_and_Set with Optional No-Pull Data

```
#define no_pull_mode_bit 24
#define byte_mask_override_bit 20
#define no_pull_data_bit 12
#define upper_part_bit 21

// This constant can be created once at init time
ALU[no_pull_constant, --, b, 0x3, <<no_pull_mode_bit]
ALU[no_pull_constant, no_pull_constant, or, 1, <<byte_mask_override_bit]

// Here is a no_pull test_and_add
ALU[temp, no_pull_constant, or, 0xfe, <<no_pull_data_bit]
ALU[temp, temp, or, 0x5, <<upper_part_bit]
SRAM[test_and_add, $x0, addra, 0], indirect_ref, sig_done[sig2]
CTX_ARB[sig2]
```

In [Example 30](#), an Increment operation does not require a pull:

Example 30. SRAM Increment without Pull Data

```
SRAM [incr, $xfer0, incr_address, 0, 1], signal_2
// SIGNAL_2 is set when $xfer0 is written with the pre-increment value.
CTX_ARB[sigal_2]
```



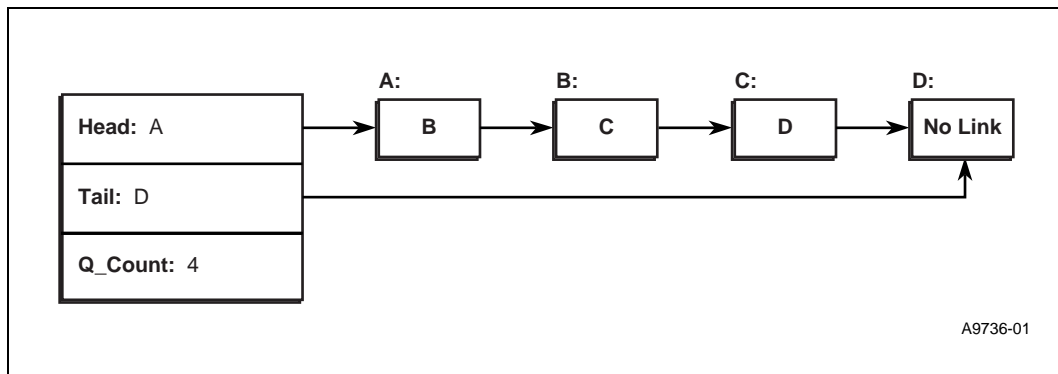
6.4.3 Queue Data Structure Commands

The ability to enqueue and dequeue data buffers at a fast rate is key to meeting chip performance goals. This is a difficult problem as it involves dependent memory references that must be turned around very quickly. The SRAM controller includes a data structure (called the Q_array) and associated control logic to perform efficient enqueue and dequeue operations. Optionally, this hardware (or a portion of it) can be used to implement rings and journals.

A queue is an ordered list of data buffers stored at non-contiguous addresses. The first buffer added to the queue will be the first buffer removed from the queue. Queue entries are joined together by creating links from one data buffer to the next. This hardware implementation supports only a forward link. A queue is described by a pointer to its first entry (called the head) and a pointer to its last entry (the tail). In addition, there is a count of the number of items currently on the queue. This triplet (head, tail, and count) is referred to as the queue descriptor. In the IXP2800 chips, the queue descriptor is stored in that order — head first, then tail, then count. The longword alignment of the head addresses for all queue descriptors must be a power of two. For example, when there are no extra parameters on the queue descriptor, there will be one unused longword per queue descriptor.

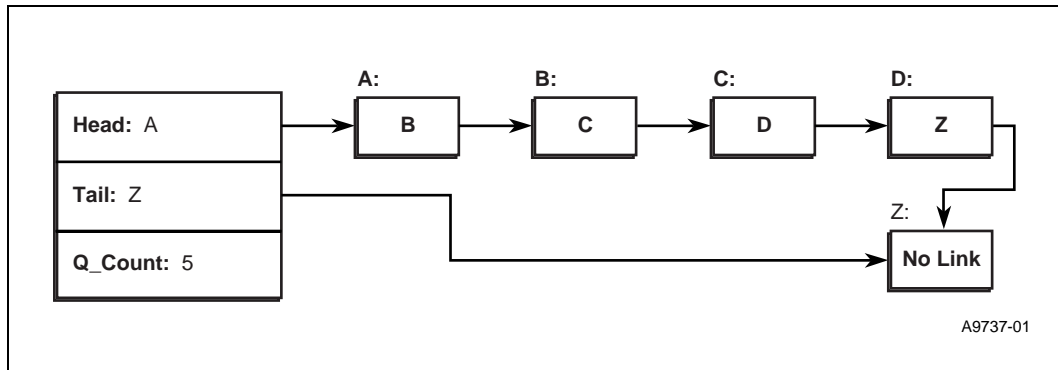
Figure 77 shows a queue descriptor and queue links for a queue containing four entries.

Figure 77. Queue Descriptor with Four Links



There are two different versions of the enqueue command, `ENQ_tail_and_link` and `ENQ_tail`. `ENQ_tail_and_link` enqueues one buffer at a time. In Figure 77, issuing an `ENQ_tail_and_link` to buffer link address z results in the queue shown in Figure 78.


Figure 78. Enqueueing One Buffer at a Time





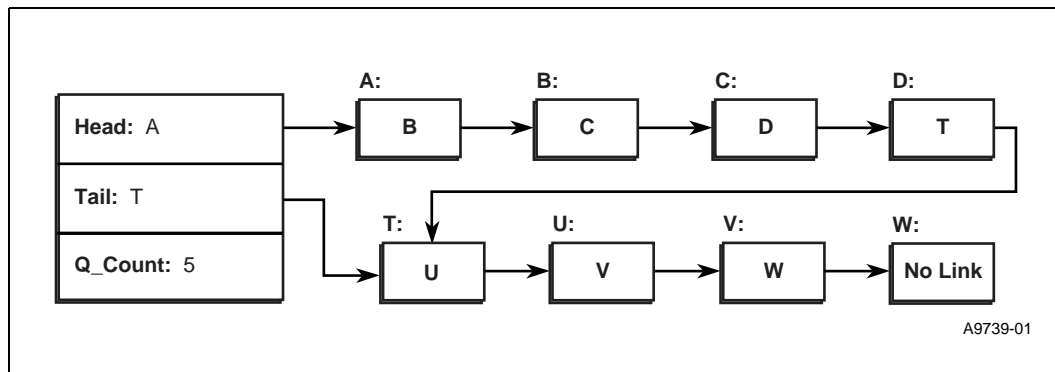
The ENQ_tail_and_link command followed by ENQ_tail enqueue a *previously* linked string of buffers. The string of buffers is used in the case where one packet is too large to fit in one buffer. Instead, it is divided among multiple buffers. Figure 79 is an example of a string of buffers.

Figure 79. Previously Linked String of Buffers



To enqueue the string of buffers in Figure 79 to the example queue in Figure 77, first issue ENQ_tail_and_link to address T; Figure 80 is the result.

Figure 80. First Step to Enqueue a String of Buffers to a Queue (ENQ_Tail_and_Link)

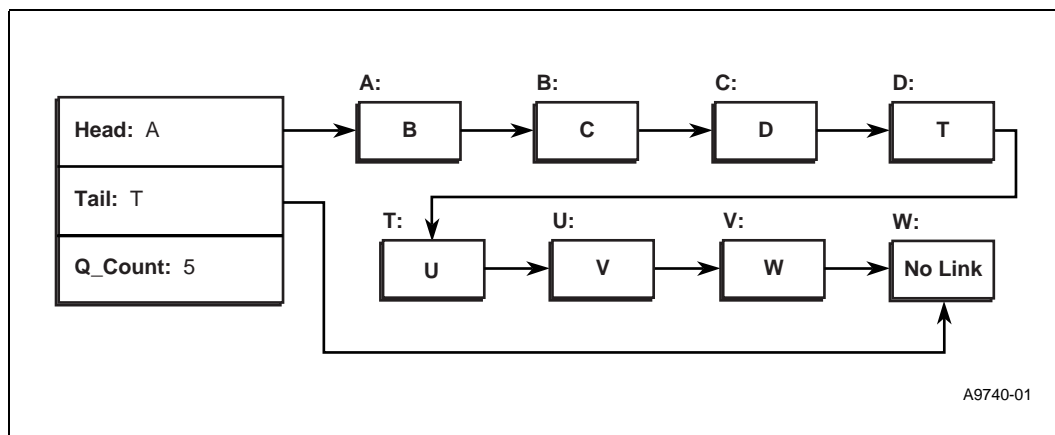


The second step is to issue an ENQ_tail to address W. This will fix the Tail to point to the last buffer of the string.

Note: Q_count is not changed by ENQ_tail because the string of buffers represents one packet.

Figure 81 is the final queue state.

Figure 81. Second Step to Enqueue a String of Buffers to a Queue (ENQ_Tail)





There are two different modes for the dequeue command. One mode removes an entire buffer from the queue. The second mode removes a piece of the buffer (referred to as a cell). The mode (cell dequeue or buffer dequeue) is selectable on a buffer-by-buffer basis by setting the `cell_count` bits (<30:24>) in the link longword.

A ring is an ordered list of data words stored in a fixed block of contiguous addresses. A ring is described by a head pointer and a tail pointer. Data is written, using the `put` command, to a ring at the address contained in the tail pointer and the tail pointer is incremented. Data is read, using the `get` command, from a ring at the address contained in the head pointer and the head pointer is incremented. Whenever either pointer reaches the end of the ring, the pointer is wrapped back to the address of the start of the ring.

A journal is similar to a ring. It is generally used for debugging. Journal commands only write to the data structure. New data overwrites the oldest data. Microcode can choose to tag the journal data with the Microengine number and CTX number of the journal writer.

The `Q_array` to support queuing, rings and journals contains 64 registers per SRAM channel. For a design with a large number of queues, the queue descriptors cannot all be stored on the chip, and thus a subset of the queue descriptors (16) is cached in the `Q_array`. (To implement the cache, 16 contiguous `Q_array` registers must be allocated.) The cache tag (the mapping of queue number to `Q_array` registers) for the `Q_array` is maintained by microcode in the CAM of a Microengine. The writeback and load of the cached registers in the `Q_array` is under the control of that microcode.

Note: The size of the `Q_array` does not set a limit on the number of queues supported.

For other queues (free buffer pools, for example), rings, and journals, the information does not need to be subsetted and thus can be loaded into the `Q_array` at initialization time and left there to be updated solely by the SRAM controller.

The sum total of the cached queue descriptors plus the number of rings, journals, and static queues must be less than or equal to 64 for a given SRAM channel.

The fields and sizes of the `Q_array` registers are shown in [Table 73](#) and [Table 74](#). All addresses are of type *longword*, and are 32 bits in length.

Table 73. Queue Format

Name	Longword Number	Bit Number ¹	Definition
EOP	0	31	End of Packet — decrement <code>Q_count</code> on dequeue
SOP	0	30	Start of Packet — used by the programmer
Cell Count	0	29:24	Number of cells in the buffer
Head	0	23:0	Head pointer
Tail	1	23:0	Tail pointer
<code>Q_count</code>	2	23:0	Number of packets on the queue or number of buffers on the queue
<code>SW_Private</code>	2	31:24	Ignored by hardware, returned to Microengine
Head Valid	N/A		Cached head pointer valid — maintained by hardware
Tail Valid	N/A		Cached tail pointer valid — maintained by hardware

1. Bits 31:24 of longword number 2 are available for use by microcode.



Table 74. Ring/Journal Format

Name	Longword Number	Bit Number	Definition
Ring Size	0	31:29	See Table 75 for size encoding.
Head	0	23:0	Get pointer
Tail	1	23:0	Put pointer
Ring Count	2	23:0	Number of longwords on the ring

Note: For a Ring or Journal, Head and Tail must be initialized to the same address.

Journals/Rings can be configured to be one of eight sizes, as shown in Table 75.

Table 75. Ring Size Encoding

Ring Size Encoding	Size of Journal/Ring Area	Head/Tail Field Base	Head and Tail Field Increment
000	512 longwords	23:9	8:0
001	1K	23:10	9:0
010	2K	23:11	10:0
011	4K	23:12	11:0
100	8K	23:13	12:0
101	16K	23:14	13:0
110	32K	23:15	14:0
111	64K	23:16	15:0

The following sections contain pseudo-code to describe the operation of the various queue and ring instructions.

Note: For these examples, NIL is the value 0.

6.4.3.1 Read_Q_Descriptor Commands

These commands are used to bring the queue descriptor data from QDR SRAM memory into the Q_array. Only portions of the Q_descriptor are read with each variant of the command, to minimize QDR SRAM bandwidth utilization. It is assumed that microcode has previously evicted the Q_descriptor data for the entry prior to overwriting the entry data with the new Q_descriptor data. Refer to the section, “SRAM (Read Queue Descriptor)”, in the *IXP2400 and IXP2800 Network Processor Programmer’s Reference Manual*, for more information.

6.4.3.2 Write_Q_Descriptor Commands

The write_Q_descriptor commands are used to evict an entry in the Q_array and return its contents to QDR SRAM memory. Only the valid fields of the Q_descriptor are written, to minimize QDR SRAM bandwidth utilization. Refer to the section, “SRAM (Write Queue Descriptor)”, in the *IXP2400 and IXP2800 Network Processor Programmer’s Reference Manual*, for more information.



6.4.3.3 ENQ and DEQ Commands

These commands add or remove elements from the queue structure while updating the `Q_array` registers. Refer to the sections, “SRAM (Enqueue)” and “SRAM (Dequeue)”, in the *IXP2400 and IXP2800 Network Processor Programmer’s Reference Manual*, for more information.

6.4.4 Ring Data Structure Commands

The ring structure commands use the `Q_array` registers to hold the head tail and count data for a ring data structure, which is a fixed-size array of data with insert and remove pointers. Refer to the section, “SRAM (Ring Operations)” in the *IXP2400 and IXP2800 Network Processor Programmer’s Reference Manual*, for more information.

6.4.5 Journaling Commands

Journaling commands use the `Q_array` registers to index into an array of memory in the QDR SRAM that will be periodically written with information to help debug applications running on the IXP2400 and IXP2800 processors. Once the array has been completely written once, subsequent journal writes overwrite the previously written data — only the most recent data will be present in the data structure. Refer to the section, “SRAM (Journal Operations)”, in the *IXP2400 and IXP2800 Network Processor Programmer’s Reference Manual*, for more information.

6.4.6 CSR Accesses

CSR accesses will write or read CSRs within each controller. The upper address bits will determine which channel will respond, while the CSR address within a channel are given in the lower address bits.

6.5 Parity

SRAM can be optionally protected by byte parity. Even parity is used — the combination of eight data bits and the corresponding parity bit will have an even number of ‘1s’. The SRAM controller generates parity on all SRAM writes. When parity is enabled (`SRAM_Control[Par_Enable]`), the SRAM controller checks for correct parity on all reads.

Upon detection of a parity error on a read, or the read portion of an atomic read-modify-write, the SRAM controller records the address of the location with bad parity in `SRAM_Parity[Address]` and sets the appropriate `SRAM_Parity[Error]` bit(s). Those bit(s) interrupt the Intel XScale® core when enabled in `IRQ_Enable[SRAM_Parity]` or `FIQ_Enable[SRAM_Parity]`.

The *Data Error* signal in the `Push_CMD` is asserted when the data to be read is delivered (unless the token `Ignore Data Error` was asserted in the command; in that case, the SRAM controller does not assert *Data Error*). When *Data Error* is asserted, the Push Arbiter suppresses the Microengine signal if the read was originated by a Microengine (it uses `0x0`, which is a null signal, in place of the requested signal number).



Note: If incorrect parity is detected on the read portion of an atomic read-modify-write, the incorrect parity is preserved after the write (that is, the byte(s) with bad parity during the read will have incorrect parity written during the write).

When parity is used, the Intel XScale® core software must initialize the SRAMs by:

1. Enabling parity (write a 1 to SRAM_Control[Par_Enable]).
2. Writing to every SRAM address.

SRAM should not be read prior to doing the above initialization; otherwise, parity errors are likely to be recorded.

6.6 Address Map

Each SRAM channel occupies a 1-Gbyte region of addresses. Channel 0 starts at 0, Channel 1 at 1 Gbyte, etc. Each SRAM controller receives commands from the command buses. It compares the target ID to the SRAM target ID, and address bits 31:30 to the channel number. If they both match, then the controller processes the command. See [Table 76](#).

Table 76. Address Map

Start Address	End Address	Responder
0x0000 0000	0x3FFF FFFF	Channel 0
0x4000 0000	0x7FFF FFFF	Channel 1
0x8000 0000	0xBFFF FFFF	Channel 2
0xc000 0000	0xFFFF FFFF	Channel 3

Note: If an access addresses a non-existent address within an SRAM controller’s address space, the results are unpredictable. For example the result of accessing address 0x0100 0000 when there is only one Mbyte of SRAM populated on the channel, produces unpredictable results.

For SRAM (memory or CSR) references from the Intel XScale® core, the channel select is in address bits 29:28. The Gasket shifts those bits to 31:30 to match addresses generated by the Microengines. Thus, the SRAM channel select logic is the same whether the command source is a Microengine or the Intel XScale® core.

The same channel start and end addresses are used both for SRAM memory and CSR references. CSR references are distinguished from memory references through the CSR encoding in the command field.

Note: Reads and writes to undefined CSR addresses yield unpredictable results.

The IXP2800 addresses are byte addresses. The fundamental unit of operation of the SRAM controller is longword access, so the SRAM controller ignores the two low-order address bits in *all* cases and utilizes the byte mask field on memory address space writes to determine the bytes to write into the SRAM. Any combination of the four bytes can be masked.

The operation of byte writes with a length other than 1 are unpredictable. That is, microcode should not use a ref_count greater than one longword when a byte_mask is active. CSRs are not byte-writable.



6.7 Reference Ordering

This section describes the ordering between accesses to any one SRAM controller. Various mechanisms are used to guarantee order — for example, references that always go to the same FIFOs remain in order. There is a CAM associated with *write addresses* that is used to order reads behind writes. Lastly, several counter pairs are used to implement “fences”. The input counter is tagged to a command and the command is not permitted to execute until the output counter matches the fence tag. All of this will be discussed in more detail in this section.

6.7.1 Reference Order Tables

Table 77 shows the architectural guarantees of order of accesses to the *same* SRAM address between a reference of any given type (shown in the column labels) and a subsequent reference of any given type (shown in the row labels). The definition of first and second is defined by the time the command is valid on the command bus. Verification requires testing only the order rules shown in Table 77 and Table 78. Note that a blank entry means no order is enforced.

Table 77. Address Reference Order

1 st ref → 2 nd ref ↓	Memory Read	CSR Read	Memory Write	CSR Write	Atomics	Queue / Ring / Q_Descr Commands
Memory Read			Order		Order	
CSR Read				Order		
Memory Write			Order		Order	
CSR Write				Order		
Atomics			Order		Order	
Queue / Ring / Q_Descr Commands						See Table 78.

Table 78 shows the architectural guarantees of order to access to the *same* SRAM Q_array entry between a reference of any given type (shown in the column labels) and a subsequent reference of any given type (shown in the row labels). The terms first and second are defined with reference to the time the command is valid on the command bus. The same caveats that apply to Table 77 apply to Table 78.



Table 78. Q_array Entry Reference Order

1 st ref → 2 nd ref ↓	Read_Q_Descr head, tail	Read_Q_Descr other	Write_Q_Descr	Enqueue	Dequeue	Put	Get	Journal
Read_Q_Descr head,tail			Order ¹					
Read_Q_Descr other	Order							
Write_Q_Descr ²								
Enqueue	Order	Order		Order	Order ³			
Dequeue	Order	Order		Order ³	Order			
Put						Order		
Get							Order	
Journal								Order

1. The order of Read_Q_Descr_head/tail after Write_Q_Descr to the same element will be guaranteed only if it is to a different descriptor SRAM address. The order of Read_Q_Descr_head/tail after Write_Q_Descr to the same element with the same descriptor SRAM address is not guaranteed and should be handled by the Microengine code.
2. Write_Q_Descr reference order is not guaranteed after any of the other references. The Queue array hardware assumes that the Microengine managing the cached entries will flush an element ONLY when it becomes the LRU in the Microengine CAM. Using this scheme, the time between the last use of this element and the write reference is sufficient to guarantee the order.
3. Order between Enqueue references and Dequeue references are guaranteed only when the Queue is empty or near empty.

6.7.2 Microcode Restrictions to Maintain Ordering

The microcode programmer must ensure order where the program flow requires order and where the architecture does not guarantee that order. One mechanism that can be used to do this is signaling. For example, if the microcode needs to update several locations in a table, a location in SRAM can be used to *lock* access to the table. [Example 31](#) is the microcode for this table update.

Example 31. Table Update Microcode

```

IMMED [$xfer0, 1]
SRAM [write, $xfer0, flag_address, 0, 1, ctx_swap [SIG_DONE_2]
; At this point, the write to flag_address has passed the point of coherency. Do
the table updates.
SRAM [write, $xfer1, table_base, offset1, 2] , sig_done [SIG_DONE_3]
SRAM [write, $xfer3, table_base, offset2, 2] , sig_done [SIG_DONE_4]
CTX_ARB [SIG_DONE_3, SIG_DONE_4]
; At this point, the table writes have passed the point of coherency. Clear the
flag to allow access by other threads.
IMMED [$xfer0, 0]
SRAM [write, $xfer0, flag_address, 0, 1, ctx_swap [SIG_DONE_2]

```

Other microcode rules:


- All access to atomic variables should be through read-modify-write instructions.
- If the flow must know that a write is completed (actually in the SRAM itself), follow the write with a read to the same address. The write is guaranteed to be complete when the read data has been returned to the Microengine.
- With the exception of initialization, never do *write* commands to the first three longwords of a *queue_descriptor* data structure (these are the longwords that hold head, tail, and count). All accesses to this data must be through the Q commands.
- To initialize the *Q_array* registers, perform a memory write of at least three longwords, followed by a memory read to the same address (to guarantee that the write completed). Then, for each entry in the *Q_array*, perform a *read_q_descriptor_head* followed by a *read_q_descriptor_other* using the address of the same three longwords.

6.8 Coprocessor Mode

Each SRAM controller may interface to an external coprocessor through its standard QDR interface. This interface allows for the cohabitation of both SRAM devices and coprocessors operating on the same bus, and the coprocessor behaves as a memory-mapped device on the SRAM bus. Figure 82 is a simplified block diagram of the SRAM controller. Figure 82 shows the connection to a coprocessor through a standard QDR interface.

Note: Most coprocessors do not need a large number of address bits — connect as many bits of A_n as required by the coprocessor.

Figure 82. Connection to a Coprocessor Through Standard QDR Interface



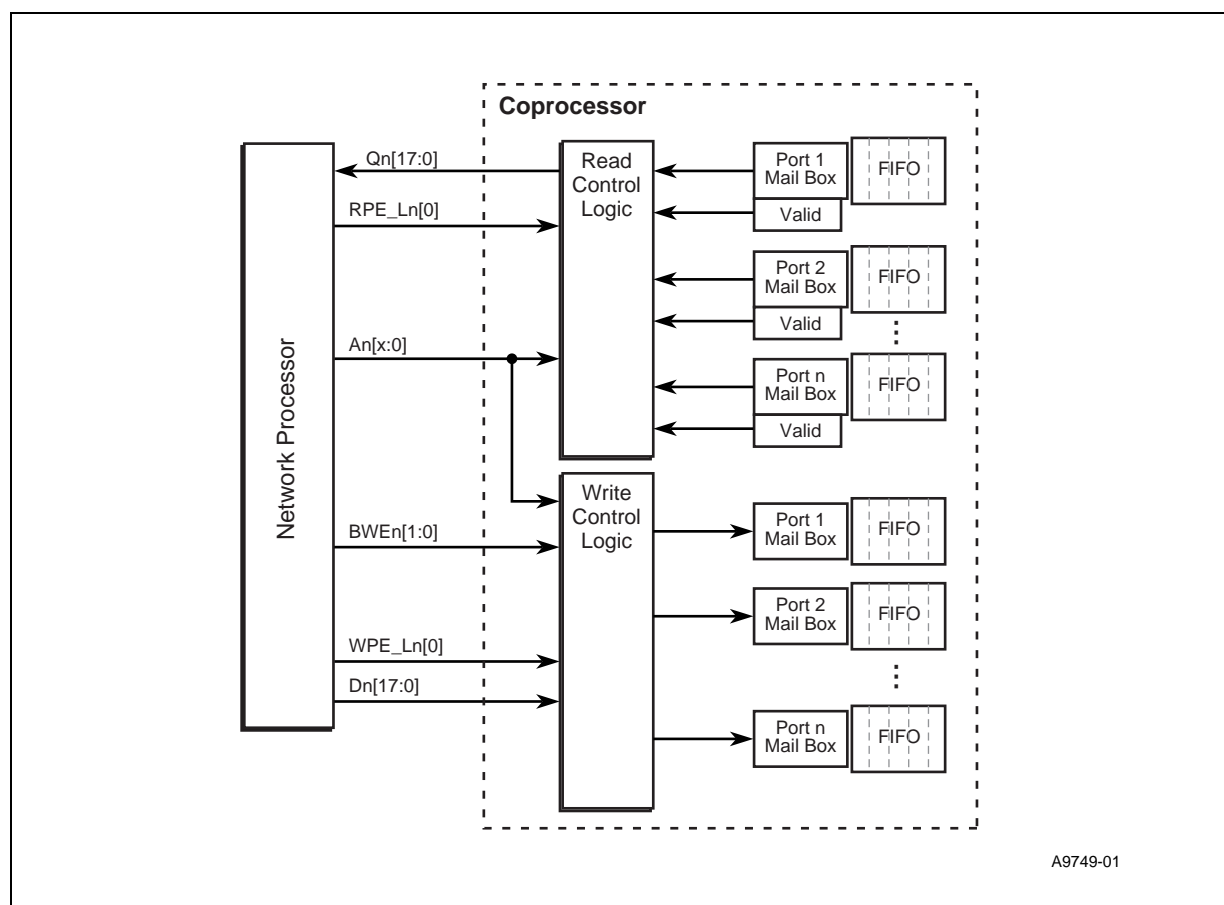
The external coprocessor interface is based on FIFO communication.

A thread can send parameters to the coprocessor by doing a normal SRAM write instruction:
`sram[write, $sram_xfer_reg, src1, src2, ref_count], optional_token`

The number of parameters (longwords) passed is specified by `ref_count`. The address can be used to support multiple coprocessor FIFO ports. The coprocessor performs some operation using the parameters, and then will later pass back some number of results values (the number of parameters and results will be known by the coprocessor designers). The time between the input parameter and return values is not fixed; it is determined by the amount of time the coprocessor requires to do its processing and can be variable. When the coprocessor is ready to return the results, it signals back to the SRAM controller through a mailbox-valid bit that the data in the read FIFO is valid. A thread can get the return values by doing a normal SRAM read instruction:
`sram[read, $sram_xfer_reg, src1, src2, ref_count], optional_token`

Figure 83 shows the coprocessor with 1-to-n memory-mapped FIFO ports.

Figure 83. Coprocessor with Memory Mapped FIFO Ports



If the read instruction executes before the return values are ready, the coprocessor signals data-invalid through the mailbox register on the read data bus ($Q_n[17:0]$). Signaling a thread upon pushing its read data works exactly as in a normal SRAM read.



There can be multiple operations in progress in the coprocessor. The SRAM controller sends parameters to the coprocessor in response to each SRAM write instruction without waiting for return results of previous writes. If the coprocessor is capable of re-ordering operations — i.e., returning the results for a given operation before returning the results of an earlier arriving operation — Microengine code must manage matching results to operations. Tagging the operation by putting a sequence value into the parameters, and having the coprocessor copy that value into the results is one way to accomplish this requirement.

Flow control is under the Network Processor's Microengine control. A Microengine thread accessing a coprocessor port maintains a count of the number of entries in that coprocessor's write-FIFO port. Each time an entry is written to that coprocessor port, the count is incremented. When a valid entry is read from that coprocessor read-port, the count is decremented by the thread.





SHaC — Unit Expansion

7

This section covers the operation of the Scratchpad, Hash Unit, and CSRs (SHaC).

7.1 Overview

The SHaC unit is a multifunction block containing Scratchpad memory and logic blocks used to perform hashing operations and interface with the Intel XScale® core peripherals and control status registers (CSRs) through the Advanced Peripheral Bus (APB) and CSR buses, respectively. The SHaC also houses the global registers, as well as Reset logic.

The SHaC unit has the following features:

- Communication to Intel XScale® core peripherals, such as GPIOs and timers, through the APB.
- Creation of hash indices of 48-, 64-, or 128-bit widths.
- Communication ring used by Microengines for interprocess communication.
- Third-option memory storage usable by Intel XScale® core and Microengines.
- CSR bus interface to permit fast writes to CSRs, as well as standard read and writes.
- Push/Pull Reflector to transfer data from the Pull bus to the Push bus.

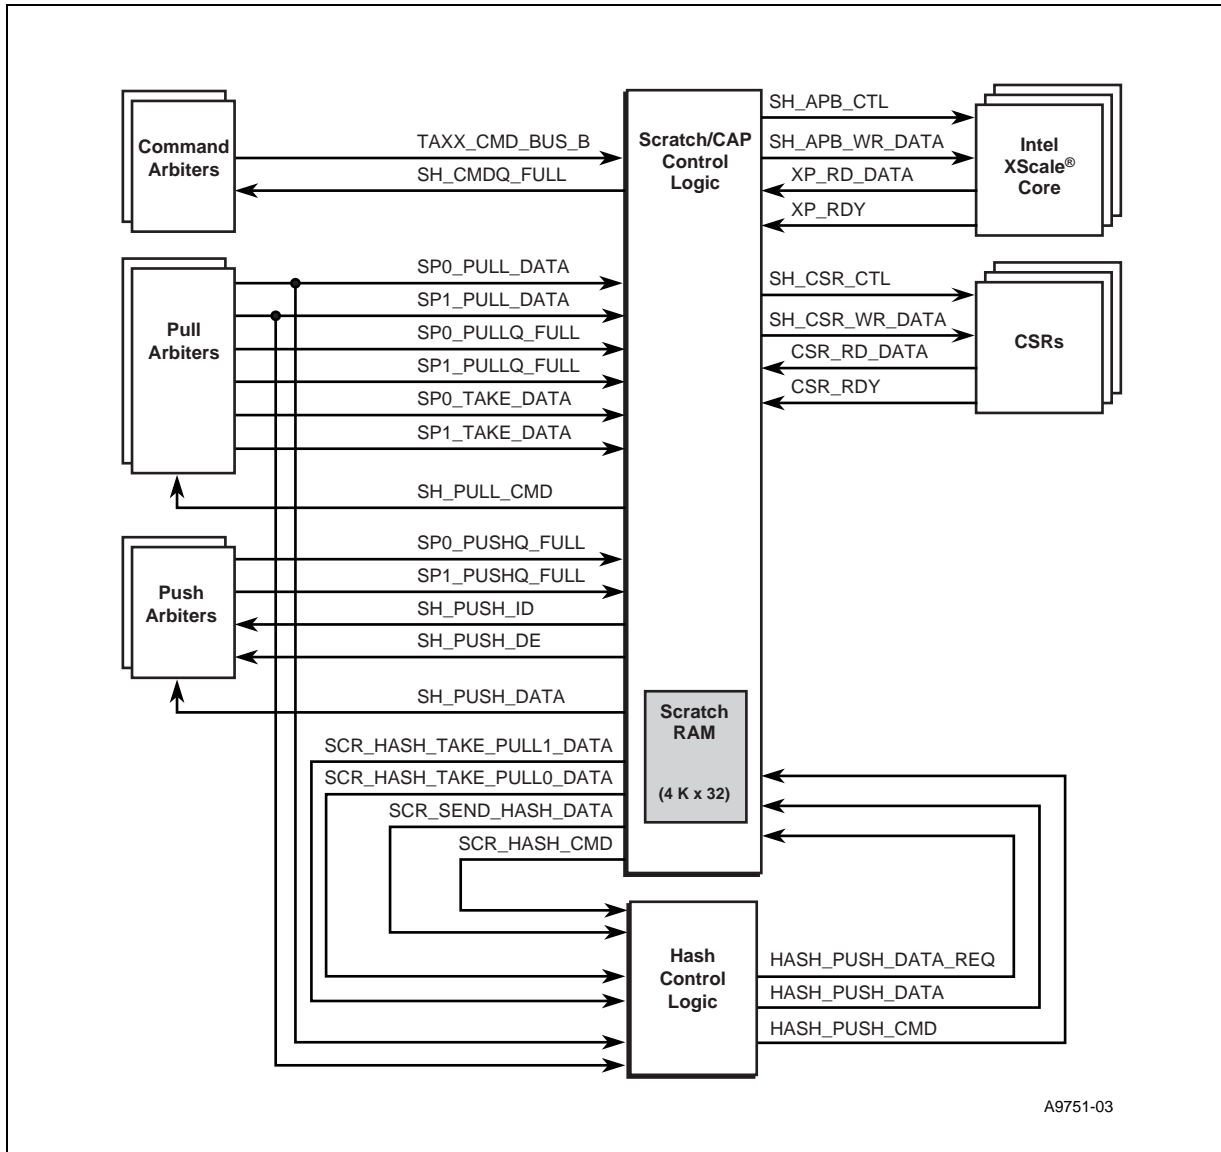
The CSR and ARM* Advanced Peripheral Bus (APB) bus interfaces are controlled by the Scratchpad state machine and will be addressed in the Scratchpad design detail section. (See [Section 7.1.2](#).)

Note: Detailed information about CSRs is contained in the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

7.1.1 SHaC Unit Block Diagram

The SHaC unit contains two functional units: the Scratchpad and Hash Unit. Each will be described in greater detail in the following sections. The CAP and APB interfaces are described as part of the Scratchpad description.

Figure 84. SHaC Top Level Diagram



A9751-03



7.1.2 Scratchpad

7.1.2.1 Scratchpad Description

The SHaC Unit contains a 16-Kbyte Scratchpad memory, organized as 4K 32-bit words, that is accessible by the Intel XScale® core and Microengines. The Scratchpad connects to the internal Command, S_Push and S_Pull, CSR, and APB buses, as shown in [Figure 85](#).

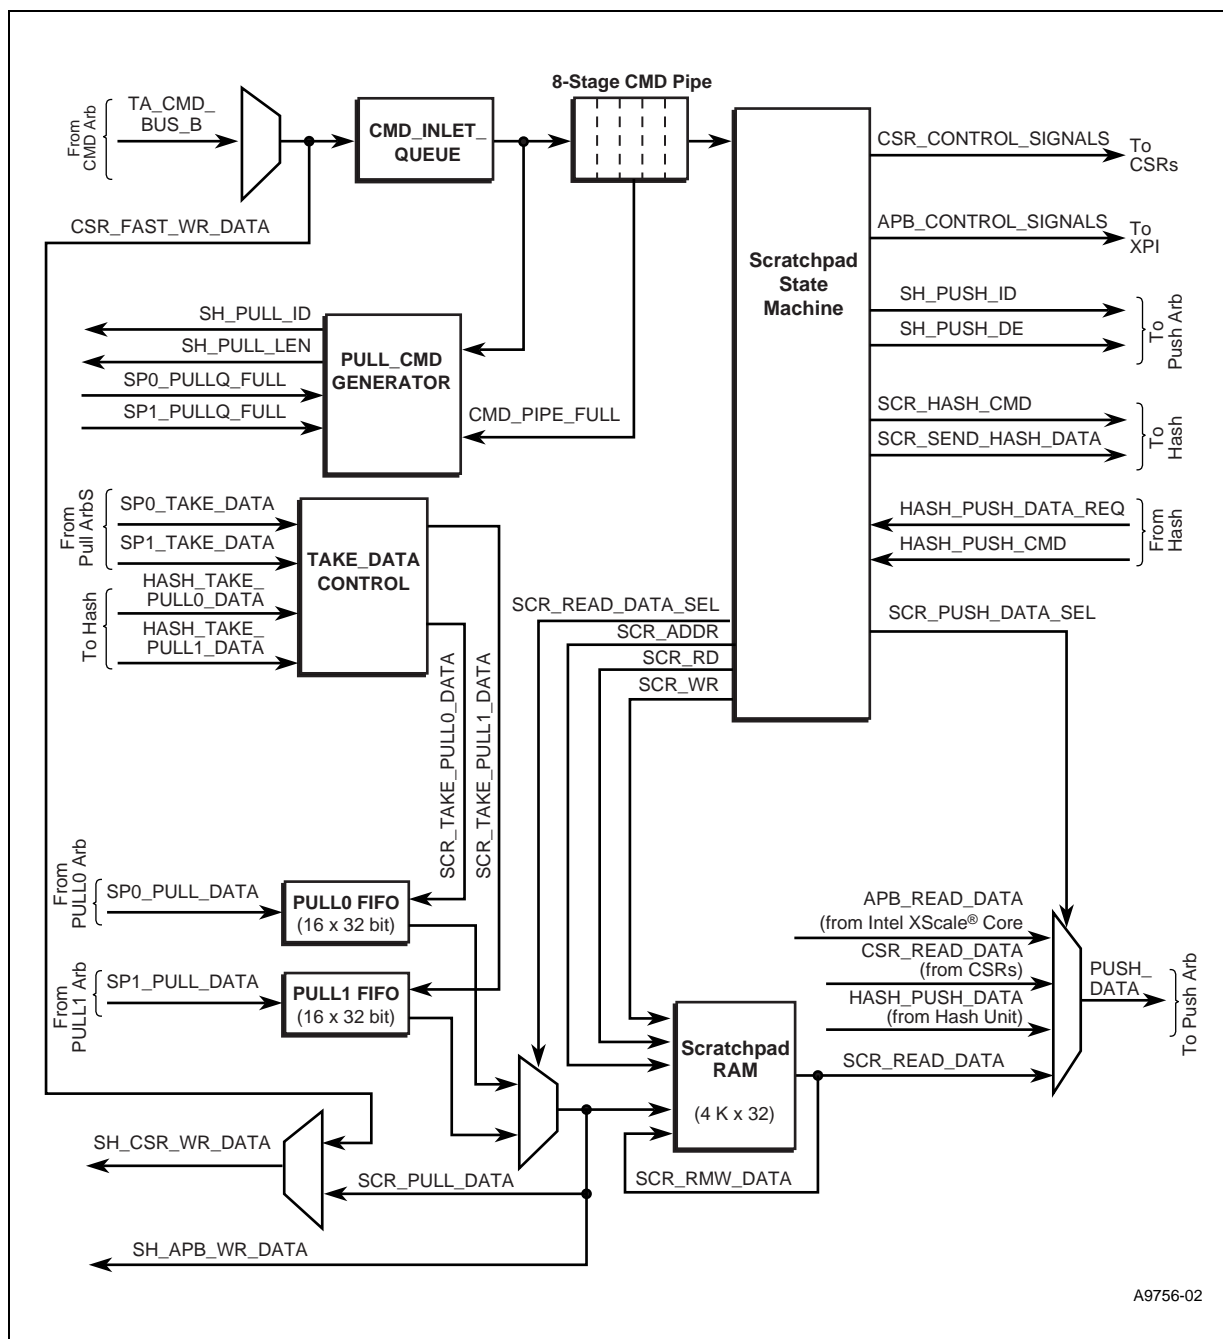
The Scratchpad memory provides the following operations:

- Normal reads and writes. 1 — 16 longwords (32 bits) can be read/written with a single command. Note that Scratchpad is not byte-writable. Each write must write *all four bytes*.
- Atomic read-modify-write operations, bit-set, bit-clear, increment, decrement, add, subtract, and swap. The Read-Modify-Write (RMW) operations can also optionally return the premodified data.
- 16 Hardware Assisted Rings for interprocess communication.¹
- Standard support of APB peripherals such as UART, Timers, and GPIOs through the ARM* Advanced Peripheral Bus (APB).
- Fast write and standard read and write operations to CSRs through the CSR Bus. For a fast write, the write data is supplied with the command, rather than pulled from the source.
- Push/Pull Reflector Mode that supports reading from a device on the pull bus and writing the data to a device on the push bus (reflecting the data from one bus to the other). A typical implementation of this mode is to allow a Microengine to read or write the transfer registers or CSRs in another Microengine. Note that the Push/Pull Reflector Mode only connects to a single Push/Pull bus. If a chassis implements more than one Push/Pull bus, it can only connect one specific bus to the CAP.

Scratchpad memory is provided as a third memory resource (in addition to SRAM and DRAM) that is shared by the Microengines and Intel XScale® core. The Microengines and Intel XScale® core can distribute memory accesses between these three types of memory resources to provide a greater number of memory accesses occurring in parallel.

1. A ring is a FIFO that uses a head and tail pointer to store/read information in Scratchpad memory.

Figure 85. Scratchpad Block Diagram



A9756-02



7.1.2.2 Scratchpad Interface

Note: The Scratchpad command and S_Push and S_Pull bus interfaces actually are shared with the Hash Unit. Only one command, to either of those units, can be accepted per cycle.

The CSR and APB buses are described in detail in the following sections.

7.1.2.2.1 Command Interface

The Scratchpad accepts commands from the Command Bus and can accept one command every cycle.

For Push/Pull reflector write and read commands, the command bus is rearranged before being sent to the Scratchpad state machine to allow a single state (REFLECT_PP) to be used to handle both commands.

7.1.2.2.2 Push/Pull Interface

The Scratchpad has the capability to interface to either one or two pairs of push/pull (PP) bus pairs. The interface from the Scratchpad to the PP bus pair is through the Push/Pull Arbiters. Each PP bus has a separate Push and Pull arbiter through which access to the Push bus and Pull bus, respectively, is regulated. Refer to the SRAM Push Arbiter and SRAM Pull Arbiter chapters for more information. When the Scratchpad is used in a chip that only utilizes one pair of PP buses, the other interface is unused.

7.1.2.2.3 CSR Bus Interface

The CSR Bus provides fast write and standard read and write operations from the Scratchpad to the CSRs in the CSR block.

7.1.2.2.4 Advanced Peripherals Bus Interface (APB)

The Advanced Peripheral Bus (APB) is part of the Advanced Microcontroller Bus Architecture (AMBA) hierarchy of buses that are optimized for minimal power consumption and reduced design complexity.

Note: The SHaC Unit uses a modified APB interface in which the APB peripheral is required to generate an acknowledge signal (APB_RDY_H) during read operations. This is done to indicate that valid data is on the bus. The addition of the acknowledge signal is an enhancement added specifically for the IXP2800 Network Processor architecture. (For more details refer to the *ARM* AMBA Specification 1.6.1.3.*)

7.1.2.3 Scratchpad Block Level Diagram

Scratchpad Command Overview

This section describes the operations performed for each Scratchpad command. Command order is preserved because all commands go through a single command inlet FIFO.

When a valid command is placed on the command bus, the control logic checks the instruction field for the Scratchpad or CAP ID. The command, address, length, etc., are enqueued into the Command Inlet FIFO. If the command requires pull data, signals are generated and immediately sent to the Pull Arbiter. The command is pushed from the Inlet FIFO to the command pipe where it is serviced according to the command type.



If the Command Inlet FIFO becomes full, the Scratchpad controller sends a full signal to the command arbiter that prevents it from sending further Scratchpad commands.

7.1.2.3.1 Scratchpad Commands

The basic read and write commands transfer from 1 – 16 longwords of data to/from the Scratchpad.

Reads

When a read command is at the head of the Command queue, the Push Arbiter is checked to see if it has enough room for the data. If so, the Scratchpad RAM is read, and the data is sent to the Push Arbiter one 32-bit word at a time (the Push_ID is updated for each word pushed). The Push Data is sent to the specified destination.

The read data is placed on the S_Push bus, one 32-bit word at a time. If the master is a Microengine, it is signaled that the command is complete during the last phase of the push bus transfer. Other masters (Intel XScale® core and PCI) must count the number of data pushes to know when the transfer is complete.

Writes

When a write command is at the head of the Command Inlet FIFO, signals are sent to the Pull Arbiter. If there is room in the queue, the command is sent to the Command pipe.

When a write command is at the head of the Command pipe, the command waits for a signal from the Pull Data FIFO, indicating that the data to be written is valid. Once the first longword is received, the data is written on consecutive cycles to the Scratchpad RAM until the burst (up to 16 longwords) is completed.

If the master is a Microengine, it is signaled that the command is complete during the last pull bus transfer. Other masters (Intel XScale® core and PCI) must count the number of data pulls to know when the transfer is complete.

Atomic Operations

The Scratchpad supports the following atomic operations.

- bit set
- bit clear
- increment
- decrement
- add
- subtract
- swap

The Scratchpad does read-modify-writes for the atomic operations, and the pre-modified data also can be returned, if desired. The atomic operations operate on a single longword. There is one cycle between the read and write while the modification is done. In that cycle, no operation is done, so an access cycle is lost.

When a read-modify-write command requiring pull data from a source is at the head of the Command Inlet FIFO, a signal is generated and sent to the Pull Arbiter (if there is room).

When the RMW command reaches the head of the Command pipe, the Scratchpad reads the memory location in the RAM. If the source requests the pre-modified data (Token[0] set), it is sent to the Push Arbiter at the time of the read. If the RMW requires pull data, the command waits for the data to be placed into the Pull Data FIFO before performing the operation; otherwise the operation is performed immediately. Once the operation has been performed, the modified data is written back to the Scratchpad RAM.

Up to two Microengine signals are assigned to each read-modify-write reference. Microcode should always tag the read-modify-write reference with an even-numbered signal. If the operation requires a pull, then the requested signal is sent on the pull. If the read data is to be returned to the Microengine, then the Microengine is sent (requested signal OR 1) when that data is pushed.

For all atomic operations, whether or not the read data is returned, is determined by Command bus Token[0].

Note: The Intel XScale® core can do atomic commands using aliased addresses in Scratchpad. An Intel XScale® core Store instruction to an atomic command address will do the RMW without returning the read data, and an Intel XScale® core Swap instruction (SWP) to an atomic command address will do the RMW and return the read data to Intel XScale® core.

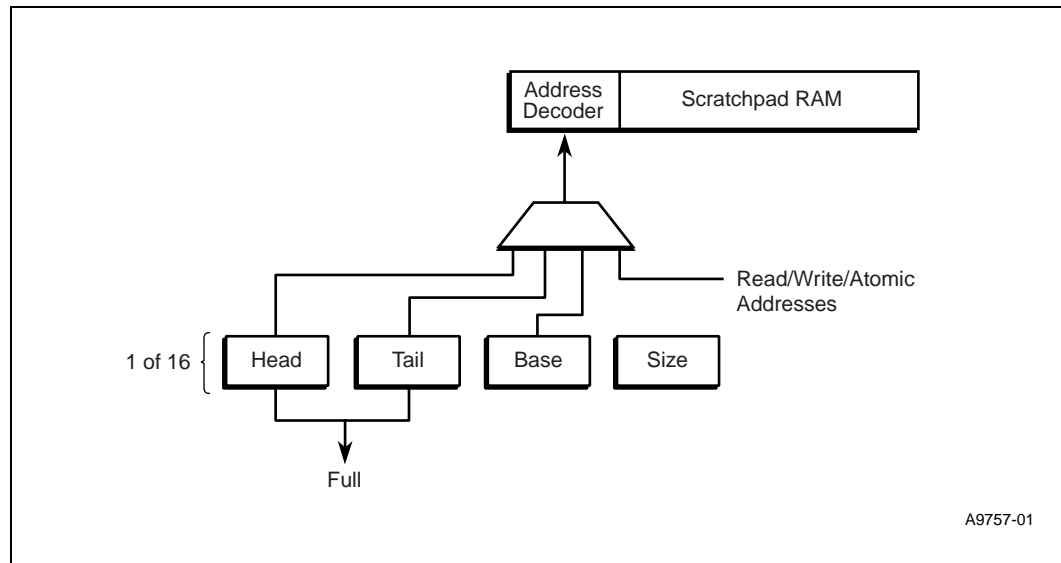
7.1.2.3.2 Ring Commands

The Scratchpad provides 16 Rings used for interprocess communication. The rings provide two operations.

- Get(ring, length)
- Put(ring, length)

Ring is the number of the ring (0 — 15) to get from or put to, and length specifies the number of longwords to transfer. A logical view of one of the rings is shown in [Figure 86](#).

Figure 86. Ring Communication Logic Diagram





Head, Tail, Base, and Size are registers in the Scratchpad Unit. Head and Tail point to the actual ring data, which is stored in the Scratchpad RAM. For each ring in use, a region of Scratchpad RAM must be reserved for the ring data. The reservation is by software convention. The hardware does not prevent other accesses to the region of Scratchpad used by the ring. Also, the regions of Scratchpad memory allocated to different rings must not overlap.

Head points to the next address to be read on a get, and Tail points to the next address to be written on a put. The size of each ring is selectable from the following choices: 128, 256, 512, or 1,024 32-bit words. The size is specified in the Ring_Base register.

Note: The above rule stating that rings must not overlap implies that many configurations are not legal. For example, programming five rings to a size of 1024 words would exceed the total size of Scratchpad memory, and therefore is not legal.

Note: The region of Scratchpad used for a ring is naturally aligned to its size.

Each ring asserts an output signal that is used as a state input to the Microengines. The software configures whether the Scratchpad asserts the signal if a ring becomes empty or if the ring is nearly full.

If configured to assert status when the rings are nearly full, Microengines must test the input state (by doing Branch on Input Signal) before putting data onto a ring. There is a lag in time from a put instruction executing to the Full signal being updated to reflect that put. To be guaranteed that a put does not overflow the ring, there is a limit on the number of Contexts and the number of 32-bit words per write, based on the size of the ring, as shown in Table 79. Each Context should test the Full signal, then do the put if not Full, and then wait until the Context has been signaled that the data has been pulled, before testing the Full signal again.

Table 79. Ring Full Signal Use – Number of Contexts and Length versus Ring Size

Number of Contexts	Ring Size			
	128	256	512	1024
1	16	16	16	16
2	16	16	16	16
4	8	16	16	16
8	4	12	16	16
16	2	6	14	16
24	1	4	9	16
32	1	3	7	15
40	Illegal	2	5	12
48	Illegal	2	4	10
64	Illegal	1	3	7
128	Illegal	Illegal	1	3

NOTE:
 1. Number in each table entry is the largest length that should be put. 16 is the largest length that a single put instruction can generate.
 2. Illegal - With that number of Contexts, even a length of 1 could cause the ring to overflow.



The ring commands operate as outlined in the pseudo-code in [Example 32](#). The operations are atomic, meaning that multi-word “Gets” and “Puts” do all the reads and writes, with no other intervening Scratchpad accesses.

Example 32. Ring Command Pseudo-Code

```

GET Command
Get(ring, length)
If count[ring] >= length //enough data in the ring?
ME <-- Scratchpad[head[ring]] // each data phase
head[ring] += length % ringSize
count[ring] -= length
else ME <--nil // 1 data phase signals read off empty list
NOTE: The Microengine signal is delivered with last data. In the case of nil, the signal is delivered with the 1 data phase.
PUT Command
Before issuing a PUT command, it is the responsibility of the Microengine thread issuing the command to make sure the Ring has enough room.
Put(ring, length)
SRAM[tail[ring]] <-- ME pull data // each data phase
tail[ring] += length % ringSize
Count[ring] += length
  
```

Table 80. Head/Tail, Base, and Full Threshold – by Ring Size

Size (Number of 32-Bit Words)	Base Address	Head/Tail Offset	Full Threshold (Entries)
128	13:9	8:2	32
256	13:10	9:2	64
512	13:11	10:2	128
1024	13:12	11:2	256

NOTE: Note that bits [1:0] of the address are assumed to be 00.

Prior to using the Scratchpad rings, software must initialize the Ring registers (by CSR writes). The Base address of the ring must be written, and also the size field that determines the number of 32-bit words for the Ring.

Note: Detailed information about CSRs is provided in the *Intel® IXP2400 and IXP2800 Network Processor Programmer’s Reference Manual*.

Writes

For an APB or CAP CSR write, the Scratchpad arbitrates for the S_Pull_Bus, pulls the write data from the source identified in the instruction (either a Microengine transfer register or an Intel XScale® core write buffer), and puts it into one of the Pull Data FIFOs. It then drives the address and writes data onto the appropriate bus. CAP CSRs locally decode the address to match their own. The Scratchpad generates a separate APB device select signal for each peripheral device (up to 15 devices). If the write is to an APB CSR, the control logic maintains valid signaling until the APB_RDY_H signal is returned (the APB RDY signal is an extension to the APB bus specification, specifically added for the Network Processor). Upon receiving the APB_RDY_H signal, the APB select signal is deasserted and the state machine returns to the idle state between commands. The CAP CSR bus does not support a similar acknowledge signal on writes since the Fast Write functionality requires that a write operation be retired on each cycle.



For writes using the Reflector mode, Scratchpad arbitrates for the S_Pull_Bus, pulls the write data from the source identified in the instruction (either a Microengine transfer register or an Intel XScale® core write buffer), and puts it into one of the Pull Data FIFOs (same as for APB and CAP CSR writes). The data is then removed from the Pull Data FIFO and sent to the Push Arbiter.

For CSR Fast Writes, the command bypasses the Inlet Command FIFO and is acted on at first opportunity. The CSR control logic has an arbiter that gives highest priority to fast writes. If an APB write is in progress when a fast write arrives, both write operations will complete simultaneously. For a CSR fast write, the Scratchpad extracts the write data from the command, instead of pulling the data from a source over the Pull bus. It then drives the address and writes data to all CSRs on the CAP CSR bus, using the same method used for the CAP CSR write.

The Scratchpad unit supports CAP write operations with burst counts greater than 1, except for fast writes, which only support a burst count of 1. Burst support is required primarily for Reflector mode and software must ensure that burst is performed to a non-contiguous set of registers. CAP looks at the length field on the command bus and breaks each count into a separate APB write cycle, incrementing the CSR number for each bus access.

Reads

For an APB read, the Scratchpad drives the address, write, select, and enable signals, and then waits for the acknowledge signal (APB_RDY_H) from the APB device. For a CAP CSR read, the address is driven, which controls a tree of multiplexers to select the appropriate CSR. CAP then waits for the acknowledge signal (CAP_CSR_RD_RDY).

Note: The CSR bus can support an acknowledge signal since the read operations occur on a separate read bus and will not interfere with Fast Write operations. In both cases, when the data is returned, the data is sent to the Push Arbiter and the Push Arbiter pushes the data to the destination.

For reads using the Reflector mode, the write data is pulled from the source identified in ADDRESS (either a Microengine transfer register or an Intel XScale® core write buffer), and put into one of the Scratchpad Pull Data FIFOs. The data is then sent to the Push Arbiter. The arbiter then moves the data to the destination specified in the command. Note that this is the same as a Reflector mode write, except that the source and destination are identified using opposite fields.

The Scratchpad performs one read operation at a time. In other words, CAP does not begin an APB read until a CSR read has completed, or vice versa. This simplifies the design by ensuring that, when lengths are greater than 1, the data is sent to the Push Arbiter in a contiguous order and not interleaved with data from a read on the other bus.

Signal Done

CAP can provide a signal to a Microengine upon completion of a command. For APB and CAP CSR operations, CAP signals the Microengine using the same method as any other target. For Reflector mode reads and writes, CAP uses the TOKEN field of the Command to determine whether to signal the command initiator, the Microengine that is the target of the reflection, both, or neither.



7.1.2.3.3 Clocks and Reset

Clock generation and distribution is handled outside of CAP and is dependent on the specific chip implementation. Separate clock rates are required for CAP CSRs/Push/Pull Buses and ARB since APB devices tend to run slower. CAP provides reset signals for the CAP CSR block and APB devices. These resets are based on the system reset signal and synchronized to the appropriate bus clock.

Table 81 shows the Intel XScale® core and Microengine instructions used to access devices on these buses and it shows the buses that are used during the operation. For example, to read an APB peripheral such as a UART CSR, a Microengine would execute a csr[read] instruction and the Intel XScale® core would execute a Load (ld) instruction. Data is then moved between the CSR and the Intel XScale® core/Microengine by first reading the CSR via the APB bus and then writing the result to the Intel XScale® core/Microengine via the Push Bus.

Table 81. Intel XScale® Core and Microengine Instructions

Accessing	Read Operation	Write Operation
APB Peripheral	Access Method: Microengine: csr[read] Intel XScale® core: ld	Access Method: Microengine: csr[write] Intel XScale® core: st
	Bus Usages: Read source: APB bus Write dest: Push bus	Bus Usages: Read source: Pull Bus Write dest: APB bus
CAP CSR	Access Method: Microengine: csr[read] Intel XScale® core: ld	Access Method: Microengine: csr[write], fast_wr Intel XScale® core: st
	Bus Usages: Read source: CSR bus Write dest: Push bus	Bus Usages: csr[write] and st Read source: Pull Bus Write dest: CSR bus fast_wr Write dest: CSR bus
Microengine CSR or Xfer register (Reflector Mode)	Access Method: Microengine: csr[read] Intel XScale® core: ld	Access Method: Microengine: csr[write] Intel XScale® core: st
	Bus Usages: Read source: Pull bus (Address) Write dest: Push bus(PP_ID)	Bus Usages: Reads: Pull Bus (PP_ID) Write dest: Push bus (Address)


7.1.2.3.4 Reset Registers

The reset registers reside in the SHaC. For more information on chip reset, refer to Section 10, “Clocks and Reset”. Strapping pins are used to select the reset count (currently 140 cycles after deassert). Options for reset count will be 64 (default), 128, 512, and 2048.

7.1.3 Hash Unit

The SHaC unit contains a Hash Unit that can take 48-, 64-, or 128-bit data and produce a 48-, 64-, or a 128-bit hash index, respectively. The Hash Unit is accessible by the Microengines and the Intel XScale® core. Figure 87 is a block diagram of the Hash Unit.

Figure 87. Hash Unit Block Diagram





7.1.3.1 Hashing Operation

Up to three hash indexes (see [Example 33](#)) can be created by using one Microengine instruction.

Example 33. Microengine Hash Instructions

```

hash1_48[$xfer], optional_token
hash2_48[$xfer], optional_token
hash3_48[$xfer], optional_token

hash1_64[$xfer], optional_token
hash2_64[$xfer], optional_token
hash3_64[$xfer], optional_token

hash1_128[$xfer], optional_token
hash2_128[$xfer], optional_token
hash3_128[$xfer], optional_token

Where:

$xfer                The beginning of a contiguous set of registers that supply the data used
                        to create the hash input and receive the hash index upon completion of
                        the hash operation.

optional_token      sig_done, ctx_swap, defer [1]
    
```

A Microengine initiates a hash operation by writing a contiguous set of SRAM Transfer registers and then executing the hash instruction. The SRAM Transfer registers can be specified as either Context-Relative or Indirect; Indirect allows any of the SRAM Transfer registers to be used. Two SRAM Transfer registers are required to create hash indexes for 48-bit and 64-bit, and four SRAM Transfer registers to create 128-bit hash indexes, as shown in [Table 82](#). In the case of the 48-bit hash, the Hash Unit ignores the upper two bytes of the second Transfer register.

Table 82. S_Transfer Registers Hash Operands (Sheet 1 of 2)

Register		Address
48-Bit Hash Operations		
Don't care	hash 3[47:32]	\$xfer n+5
	hash 3 [31:0]	\$xfer n+4
Don't care	hash 2[47:32]	\$xfer n+3
	hash 2 [31:0]	\$xfer n+2
Don't care	hash 1[47:32]	\$xfer n+1
	hash 1 [31:0]	\$xfer n
64-Bit Hash Operations		
	hash 3 [63:32]	\$xfer n+5
	hash 3 [31:0]	\$xfer n+4
	hash 2 [63:32]	\$xfer n+3
	hash 2 [31:0]	\$xfer n+2
	hash 1 [63:32]	\$xfer n+1
	hash 1 [31:0]	\$xfer n



Table 82. S_Transfer Registers Hash Operands (Sheet 2 of 2)

Register	Address
128-Bit Hash Operations	
hash 3 [127:96]	\$xfer n+11
hash 3 [95:64]	\$xfer n+10
hash 3 [63:32]	\$xfer n+9
hash 3 [31:0]	\$xfer n+8
hash 2 [127:96]	\$xfer n+7
hash 2 [95:64]	\$xfer n+6
hash 2 [63:32]	\$xfer n+5
hash 2 [31:0]	\$xfer n+4
hash 1 [127:96]	\$xfer n+3
hash 1 [64:95]	\$xfer n+2
hash 1 [63:32]	\$xfer n+1
hash 1 [31:0]	\$xfer n

The Intel XScale® core initiates a hash operation by writing a set of memory-mapped Hash Operand registers (which are built into the Intel XScale® core gasket) with the data to be used to generate the hash index. There are separate registers for 48-, 64-, and 128-bit hashes. Only one hash operation of each type can be done at a time. Writing to the last register in each group informs the gasket logic that it has all of the operands for that operation, and it will then arbitrate for the Command bus to send the command to the Hash Unit.

Note: Detailed information about CSRs is contained in the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

For Microengine-generated commands and those generated by the Intel XScale® core, the command enters the Command Inlet FIFO. As with the Scratchpad write and RMW operations, signals are generated and sent to the Pull Arbiter. The Hash unit Pull Data FIFO allows the data for up to three hash operations to be read into the Hash Unit in a single burst. When the command is serviced, the first data to be hashed enters the hash array while the next two wait in the FIFO.

The Hash Unit uses a hard-wired polynomial algorithm and a programmable hash multiplier to create hash indexes. Three separate multipliers are supported — one each, for 48-, 64-, and 128-bit hash operations. The multiplier is programmed through the registers, HASH_MULTIPLIER_64_1, HASH_MULTIPLIER_64_2, HASH_MULTIPLIER_48_1, HASH_MULTIPLIER_48_2, HASH_MULTIPLIER_128_1, HASH_MULTIPLIER_128_2, HASH_MULTIPLIER_128_3, and HASH_MULTIPLIER_128_4.

The multiplicand is shifted into the hash array 16 bits at a time. The hash array performs a 1's-complement multiply and polynomial divide, calculated by using the multiplier and 16 bits of the multiplicand. The result is placed into an output register and is also fed back into the array. This process is repeated three times for a 48-bit hash (16 bits x 3 = 48), four times for a 64-bit hash (16 bits x 4 = 64), and eight times for a 128-bit hash (16 x 8 = 128). After an entire multiplicand has been passed through the hash array, the resulting hash index is placed into a two-stage output pipeline and the next hash is immediately started.

The Hash Unit shares the Scratchpad's Push Data FIFO. After each hash index is completed, the index is placed into a three-stage output pipe and the Hash Unit sends a PUSH_DATA_REQ to the Scratchpad to indicate that it has a valid hash index to put into the Push Data FIFO for transfer. The Scratchpad issues a SEND_HASH_DATA signal, transfers the hash index to the Push Data FIFO, and sends the data to the Arbiter.

For hash operations initiated by the Intel XScale® core, the core reads the results from its memory-mapped Hash Result registers. The addresses of Hash Results are the same as the Hash Operand registers. Because of queuing delays at the Hash Unit, the time to complete an operation is not fixed. The Intel XScale® core can do one of two operations to get the hash results:

- Poll the Hash Done register. This register is cleared when the Hash Operand registers are written. Bit [0] of the Hash Done register is set when the Hash Result registers get the result from the Hash Unit (when the last word of the result is returned). The Intel XScale® core software can poll on Hash Done, and read Hash Result when Hash Done equals 0x00000001.
- Read Hash Result directly. The gasket logic acknowledges the read only when the result is valid. The Intel XScale® core stalls if the result is not valid when the read happens.

The number of clock cycles required to perform a single hash operation equals: two or four cycles through the input buffers, three, four, or eight cycles through the hash array, and two or four cycles through the output buffers. With the pipeline characteristics of the Hash Unit, performance is improved if multiple hash operations are initiated with a single instruction, rather than with separate hash instructions for each hash operation.

7.1.3.2 Hash Algorithm

The hashing algorithm allows flexibility and uniqueness since it can be programmed to provide different results for a given input. The algorithm uses binary polynomial multiplication and division under modulo-2 addition. The input to the algorithm is a 48-, 64-, or 128-bit value.

The data used to generate the hash index is considered to represent the coefficients of an order-47, order-63, or order-127 polynomial in x . The input polynomial (designated as $A(x)$) has the form:

Equation 1. $A_{48}(x) = a_0 + a_1x + a_2x^2 + \dots + a_{46}x^{46} + a_{47}x^{47}$ (48-bit hash operation)

Equation 2. $A_{64}(x) = a_0 + a_1x + a_2x^2 + \dots + a_{62}x^{62} + a_{63}x^{63}$ (64-bit hash operation)

Equation 3. $A_{128}(x) = a_0 + a_1x + a_2x^2 + \dots + a_{126}x^{126} + a_{127}x^{127}$ (128-bit hash operation)

This polynomial is multiplied by a programmable hash multiplier using a modulo-2 addition. The hash multiplier, $M(x)$ is stored in Hash Unit CSRs and represents the polynomial.

Equation 4. $M_{48}(x) = m_0 + m_1x + m_2x^2 + \dots + m_{46}x^{46} + m_{47}x^{47}$ (48-bit hash operation)

Equation 5. $M_{64}(x) = m_0 + m_1x + m_2x^2 + \dots + m_{62}x^{62} + m_{63}x^{63}$ (64-bit hash operation)

Equation 6. $M_{128}(x) = m_0 + m_1x + m_2x^2 + \dots + m_{126}x^{126} + m_{127}x^{127}$ (128-bit hash operation)

Since multiplication is performed using modulo-2 addition, the result is an order-94 polynomial, an order-126 polynomial, or an order-254 polynomial with coefficients that are also 1 or 0. This product is divided by a fixed generator polynomial given by:



Equation 7. $G_{48}(x) = 1 + x^{10} + x^{25} + x^{36} + x^{48}$ (48-bit hash operation)

Equation 8. $G_{64}(x) = 1 + x^{17} + x^{35} + x^{54} + x^{64}$ (64-bit hash operation)

Equation 9. $G_{128}(x) = 1 + x^{33} + x^{69} + x^{98} + x^{128}$ (128-bit hash operation)

The division results in a quotient $Q(x)$, a polynomial of order-46, order-62, or order-126, and a remainder $R(x)$, and a polynomial of order-47, order-63, or order-127. The operands are related by the equation:

Equation 10. $A(x)M(x) = Q(x)G(x) + R(x)$

The generator polynomial has the property of irreducibility. As a result, for a fixed multiplier $M(x)$, there is a unique remainder $R(x)$ for every input $A(x)$. The quotient $Q(x)$, can then be discarded, since input $A(x)$ can be derived from its corresponding remainder $R(x)$. A given bounded set of input values $A(x)$ — for example, 8K or 16K table entries — with bit weights of an arbitrary density function can be mapped one-to-one into a set of remainders $R(x)$ such that the bit weights of the resulting Hashed Arguments (a subset of all values of $R(x)$ polynomials) are all approximately equal.

In other words, there is a high likelihood that the low-order set of bits from the Hash Arguments are unique, so they can be used to build an index into the table. If the hash algorithm does not provide a uniform hash distribution for a given set of data, the programmable hash multiplier ($M(x)$) may be modified to provide better results.



Media and Switch Fabric Interface

8

8.1 Overview

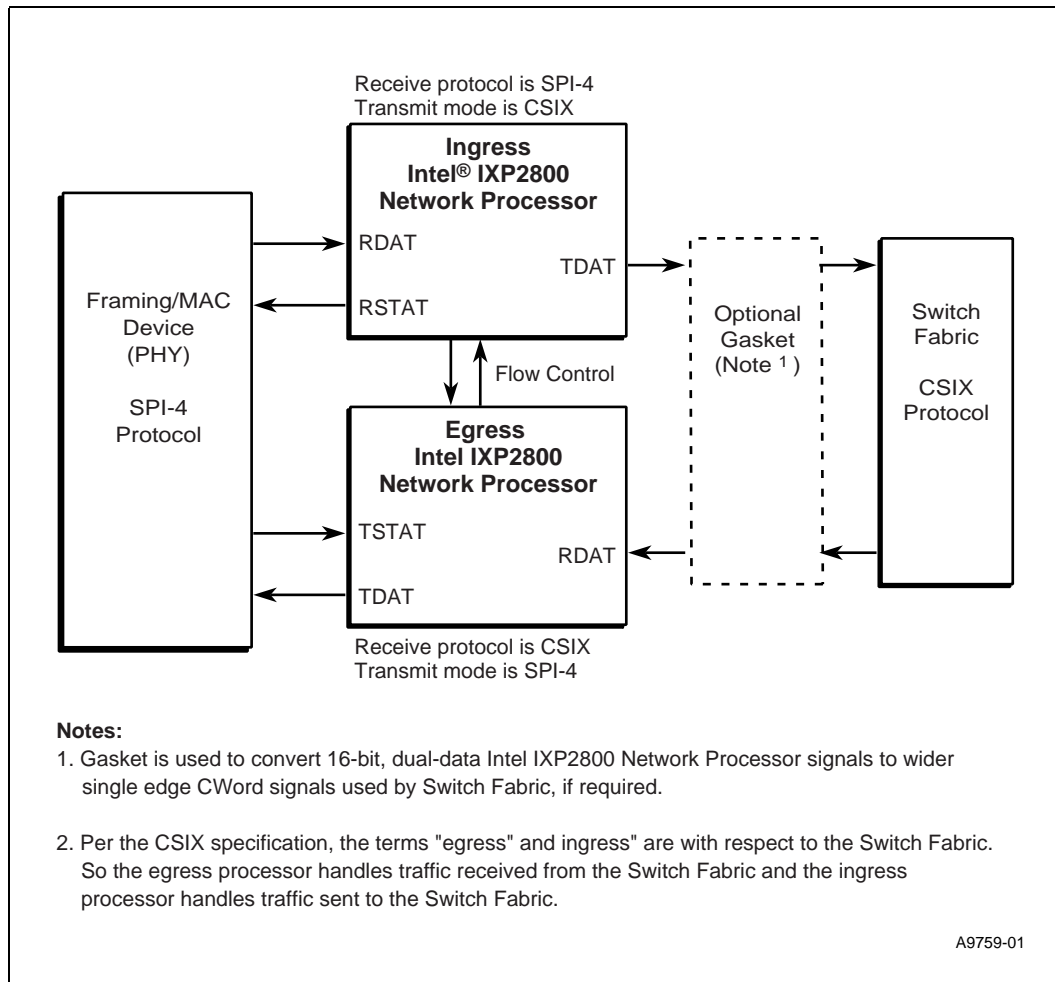
The Media and Switch Fabric (MSF) Interface connects the IXP2800 Network Processor to a physical layer device (PHY) and/or to a Switch Fabric. The MSF consists of separate receive and transmit interfaces, each of which can be separately configured for either SPI-4 Phase 2 (System Packet Interface), for PHY devices or for the CSIX-L1 protocol, for Switch Fabric Interfaces.

The receive and transmit ports are unidirectional and independent of each other. Each port has 16 data signals, a clock, a control signal, and a parity signal, all of which use LVDS (differential) signaling, and are sampled on both edges of the clock. There is also a flow control port consisting of a clock, data, and ready status bits, for communicating between two IXP2800 Network Processors, or a IXP2800 Network Processor and a Switch Fabric Interface; these are also LVDS, dual-edge data transfer.

Signal usage and the receive and transmit functions, are illustrated in [Figure 88](#), and described in the sections that follow.

Note: Detailed information about CSRs is contained in the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.


Figure 88. Example System Block Diagram



The use of some of the receive and transmit pins is based on protocol, SPI-4 or CSIX. For the LVDS pins, only the active high name is given (for LVDS, there are two pins per signal). The definitions of the pins can be found in the SPI-4 and CSIX specs, referenced below.

An alternate system configuration is shown in the block diagram in [Figure 89](#). In this case, a single IXP2800 Network Processor is used for both Ingress and Egress. The bit-rate supported would be less than in [Figure 88](#). A hypothetical Bus Converter chip, external to the IXP2800 Network Processor, is used. The block diagram in [Figure 89](#) is only an illustrative example.

Figure 89. Full-Duplex Block Diagram



8.1.1 SPI-4

SPI-4 is an interface for packet and cell transfer between a physical layer (PHY) device and a link layer device (the IXP2800 Network Processor), for aggregate bandwidths of OC-192 ATM and Packet over SONET/SDH (POS), as well as 10 Gb/s Ethernet applications.

The Optical Internetworking Forum (OIF), www.oiforum.com, controls the SPI-4 Implementation Agreement document.

SPI-4 has two types of transfers — Data when the RCTL signal is deasserted; Control when the RCTL signal is asserted. The Control Word format is shown in [Table 83](#) (this information is from the SPI-4 specification, shown here for convenience).

Table 83. SPI-4 Control Word Format

Bit Position	Label	Description
15	Type	Control Word Type. <ul style="list-style-type: none"> • 1—payload control word (payload transfer will immediately follow the control word). • 0—idle or training control word.
14:13	EOPS	End-of-Packet (EOP) Status. Set to the following values below according to the status of the immediately preceding payload transfer. <ul style="list-style-type: none"> • 00—Not an EOP. • 01—EOP Abort (application-specific error condition). • 10—EOP Normal termination, 2 bytes valid. • 11—EOP Normal termination, 1 byte valid. EOPS is valid in the first Control Word following a burst transfer. It is ignored and set to 00 otherwise.
12	SOP	Start-of-Packet (SOP). <ul style="list-style-type: none"> • Set to 1 if the payload transfer immediately following the Control Word corresponds to the start of a packet; set to 0 otherwise. • Set to 0 in all idle and training control words.
11:4	ADR	Port Address. 8-bit port address of the payload data transfer immediately following the Control Word. None of the addresses are reserved (all are available for payload transfer). <ul style="list-style-type: none"> • Set to all zeros in all idle Control Words. • Set to all ones in all training Control Words.
3:0	DIP-4	4-bit Diagonal Interleaved Parity. 4-bit odd parity computed over the current Control Word and the immediately preceding data words (if any) following the last Control Word.

Control words are inserted only between burst transfers; once a transfer has begun, data words are sent uninterrupted until either End of Packet or a multiple of 16 bytes is reached. The order of bytes within the SPI-4 data burst is shown in [Table 84](#).

The most significant bits of the bytes correspond to bits 15 and 7. On data transfers that do not end on an even byte-boundary, the unused byte on bits [7:0] is set to all zeros.

Table 84 shows the order of bytes on SPI-4; this example shows a 43-byte packet.

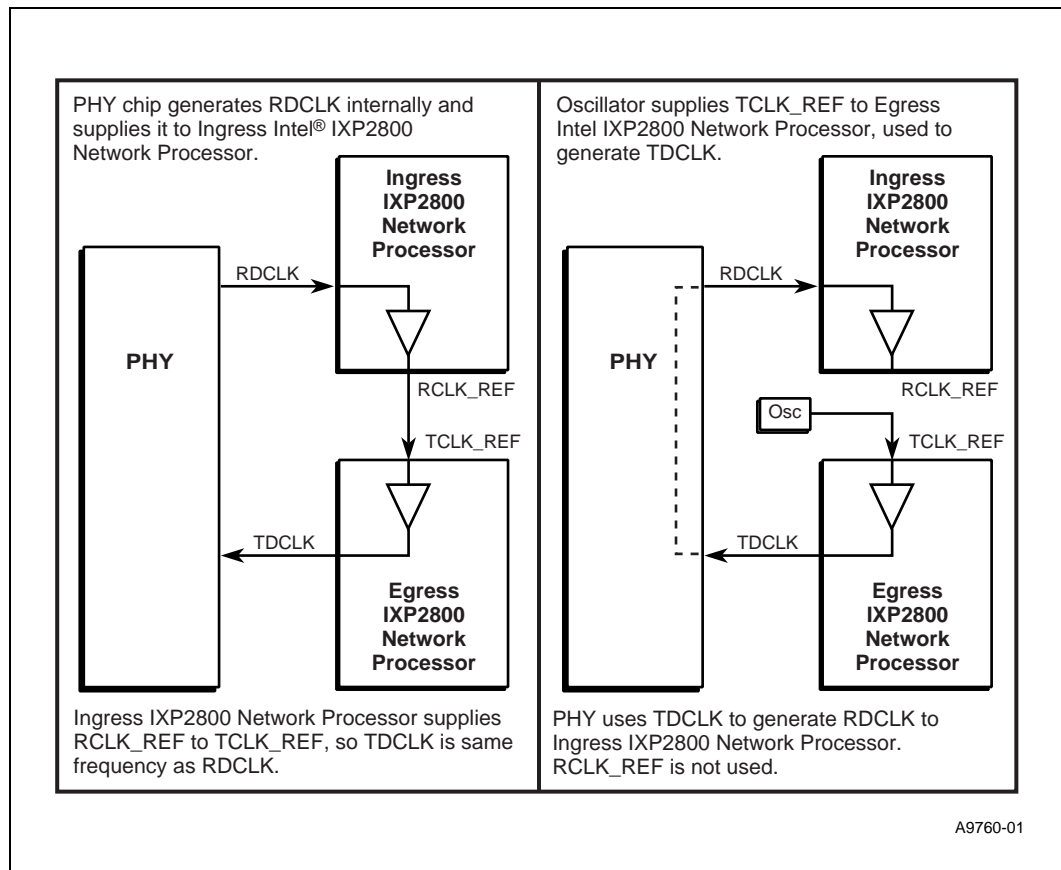
Table 84. Order of Bytes¹ within the SPI-4 Data Burst

	Bit 15	Bit 8	Bit 7	Bit 0
Data Word 1	Byte 1		Byte 2	
Data Word 2	Byte 3		Byte 4	
Data Word 3	Byte 5		Byte 5	
Data Word 4	Byte 7		Byte 6	
...	
...	
...	
Data Word 21	Byte 41		Byte 42	
Data Word 22	Byte 43 ²		00	

1. These bytes are valid *only* if EOP is set.
2. All transfers on the SPI-4 bus must be in multiples of 16 bytes if it is not associated with an End of Packet (EOP) transfer, to comply with the protocol. Hence, this 43-byte example would only be valid for an EOP transfer.

Figure 90 shows two ways in which the SPI-4 clocking can be done. Note that it is also possible to use an internally-supplied clock and leave **TCLK_REF** unused.

Figure 90. Receive and Transmit Clock Generation





8.1.2 CSIX

CSIX_L1 (Common Switch Interface) defines an interface between a Traffic Manager (TM) and a Switch Fabric (SF) for ATM, IP, MPLS, Ethernet, and similar data communications applications.

The Network Processor Forum (NPF) www.npforum.org, controls the CSIX_L1 specification.

The basic unit of information transferred between TMs and SFs is called a CFrame. There are a number of CFrame types defined as shown in [Table 85](#).

Table 85. CFrame Types

Type Encoding	CFrame Type
0	Idle
1	Unicast
2	Multicast Mask
3	Multicast ID
4	Multicast Binary Copy
5	Broadcast
6	Flow Control
7	Command and Status
8-F	CSIX Reserved

For transmission from the IXP2800 Network Processor, CFrames are constructed for transmission under Microengine software control, and written into the Transmit Buffer (TBUF).

On receive to the IXP2800 Network Processor, CFrames are either discarded, placed into Receive Buffer (RBUF), or placed into Flow Control Egress FIFO (FCEFIFO), according to mapping defined in the **CSIX_Type_Map** CSR. CFrames put into RBUF are passed to a Microengine to be parsed by software. CFrames put into FCEFIFO are sent to the Ingress IXP2800 Network Processor over the Flow Control bus. Link-level Flow Control information (CSIX Ready field) in the Base Header of all CFrames (including Idle) is handled by hardware.

8.1.3 CSIX/SPI-4 Interleave Mode

SPI-4 packets and CSIX CFrames are interleaved when the RBUF and TBUF are configured in 3-partition mode. When the protocol signal RPROT or TPROT is high, the data bus is transferring CSIX CFRAMEs or IDLE cycles. When protocol is low, the data bus is transferring SPI-4 packets or idle cycles. When operating in interleave mode, RPROT must be driven high (logic 1) for the entire CSIX CFRAME or low (logic 0) for the entire SPI-4 burst. When in 3-partition mode, the SPI-4 interval should be padded using SPI-4 idle cycles so that it ends on a 32-bit boundary or a complete RCLK or TCLK clock cycle. The actual SPI-4 data length can be any size. However, the SPI-4 interval, which includes the SPI-4 control words and payload data, must end on a 32-bit boundary.



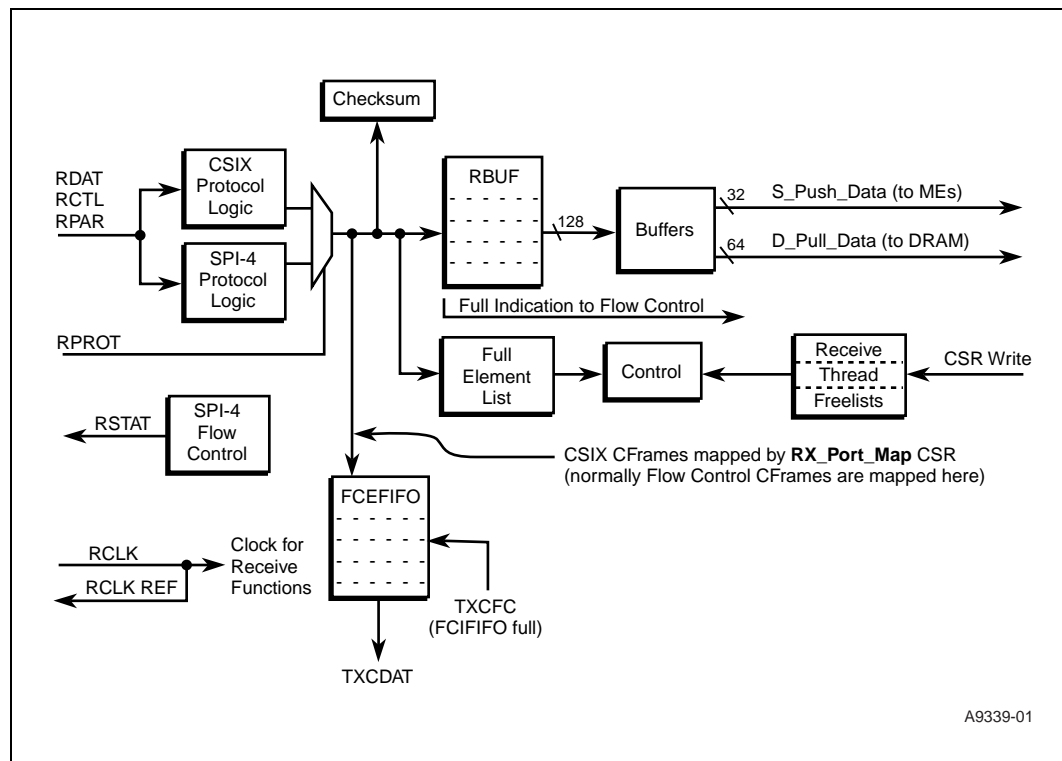
8.2 Receive

The receive section consists of:

- Receive Pins (Section 8.2.1)
- Checksum (Section 8.2.2)
- Receive Buffer (RBUF) (Section 8.2.2)
- Full Element List (Section 8.2.3)
- Rx_Thread_Freelist (Section 8.2.4)
- Flow Control Status (Section 8.2.7)

Figure 91 is a simplified block diagram of the receive section.

Figure 91. Simplified Receive Section Block Diagram





8.2.1 Receive Pins

The use of the receive pins is a function of RPROT input, as shown in [Table 86](#).

Table 86. Receive Pins Usage by Protocol

Name	Direction	SPI-4 Use	CSIX Use
RCLK	Input	RDCLK	TxCik
RDAT[15:0]	Input	RDAT[15:0]	TxData[15:0]
RCTL	Input	RCTL	TxSOF
RPAR	Input	Not Used	TxPar
RSCLK	Output	RSCLK	Not Used
RSTAT[1:0]	Output	RSTAT[1:0]	Not Used

In general, hardware does framing, parity checking, and flow control message handling. Interpretation of frame header and payload data is done by Microengine software.

The internal clock used is taken from the RCLK pin. **RCLK_Ref** output is a buffered version of the clock. It can be used to supply **TCLK_Ref** of the Egress IXP2800 Network Processor if desired.

The receive pins RDAT[15:0], RCTL, RPAR are sampled relative to RCLK. To work at high frequencies, each of those pins has de-skewing logic as described in [Section 8.6](#).

8.2.2 RBUF

RBUF is a RAM that holds received data. It stores received data in sub-blocks (referred to as elements), and is accessed by a Microengine or the Intel XScale® core reading the received information. Details of how RBUF elements are allocated and filled is based on the receive data protocol, and is described in [Section 8.2.2.1](#) – [Section 8.2.2.2](#). When data is received, the associated status is put into the **Full_Element_List** FIFO and subsequently sent to a Microengine for processing. **Full_Element_List** insures that received elements are sent to a Microengine in the order in which the data was received.

RBUF contains a total of eight Kbytes of data. [Table 87](#) shows the order in which received data is stored in RBUF. Each number represents a byte, in order of arrival from the receiver interface.

Table 87. Order in which Received Data Is Stored in RBUF

Data/Payload								Address Offset (Hex)
4	5	6	7	0	1	2	3	0
C	D	E	F	8	9	A	B	8
14	15	16	17	10	11	12	13	10

The mapping of elements to address offset in RBUF is based on the RBUF partition and element size, as programmed in the **MSF_Rx_Control** CSR. RBUF can be partitioned into one, two, or three partitions based on **MSF_Rx_Control[RBUF_Partition]**. The mapping of received data to partitions is shown in [Table 88](#).

Table 88. Mapping of Received Data to RBUF Partitions

Number of Partitions in Use	Receive Data Protocol	Data Use by Partition, Fraction of RBUF Used, Start Byte Offset (Hex)		
		Partition Number		
		0	1	2
1	SPI-4 only	SPI-4 All Byte 0	n/a	n/a
2	CSIX only	CSIX Data 3/4 of RBUF Byte 0	CSIX Control 1/4 of RBUF Byte 0x1800	n/a
3	Both SPI-4 and CSIX	CSIX Data 1/2 of RBUF Byte 0	SPI-4 3/8 of RBUF Byte 0x1000	CSIX Control 1/8 of RBUF Byte 0x1C00

The data in each partition is further broken up into elements, based on **MSF_Rx_Control[RBUF_Element_Size_#]** (n = 0, 1, 2). There are three choices of element size – 64, 128, or 256 bytes.

Table 89 shows the RBUF partition options. Note that the choice of element size is independent for each partition.

Table 89. Number of Elements per RBUF Partition

RBUF_Partition Field	RBUF_Element_Size_# Field	Partition Number		
		0	1	2
00 (1 partition)	00 (64 bytes)	128	Unused	Unused
	01 (128 bytes)	64		
	10 (256 bytes)	32		
01 (2 partitions)	00 (64 bytes)	96	32	Unused
	01 (128 bytes)	48	16	
	10 (256 bytes)	24	8	
10 (3 partitions)	00 (64 bytes)	64	48	16
	01 (128 bytes)	32	24	8
	10 (256 bytes)	16	12	4

The Microengine can read data from the RBUF to Microengine S_TRANSFER_IN registers using the `msf[read]` instruction, where the starting byte number is specified (which must be aligned to 4-byte units), and also the number of 32-bit words to read. The number in the instruction can be either the number of 32-bit words, or the number of 32-bit word pairs, using the single- and double-instruction modifiers, respectively. The data is pushed to the Microengine on the S_Push_Bus by RBUF control logic:

```
msf[read, $s_xfer_reg, src_op_1, src_op_2, ref_cnt], optional_token
```



The `src_op_1` and `src_op_2` operands are added together to form the address in RBUF (note that the base address of the RBUF is 0x2000). The `ref_cnt` operand is the number of 32-bit words or word pairs, that are pushed into two sequential `S_TRANSFER_IN` registers, starting with `$s_xfer_reg`.

Using the data in RBUF in [Table 87](#) above, reading eight bytes from offset 0 into transfer registers 0 and 1 would yield the result in [Example 34](#).

Example 34. Data from RBUF Moved to Microengine Transfer Registers

Transfer Register Number	Bit Number within Transfer Register							
	31	24	23	16	15	8	7	0
0	0		1		2		3	
1		4		5		6		7

Microengine can move data from RBUF to DRAM using the instruction:
`dram[rbuf_rd, --, src_op1, src_op2, ref_cnt], indirect_ref`

The `src_op_1` and `src_op_2` operands are added together to form the address in DRAM, so the `dram` instruction must use the `indirect_ref` modifier to specify the RBUF address (refer to the IXP2800 Network Processor chassis chapter for details). The `ref_cnt` operand is the number of 64-bit words that are read from RBUF.

Using the data in RBUF in [Table 87](#) above, reading 16 bytes from offset 0 in RBUF into DRAM would yield the result in [Example 35](#) in DRAM (addresses in DRAM must be aligned to 8-byte units. The data from lower-offset RBUF offsets goes into lower addresses in DRAM.)

Example 35. Data from RBUF Moved to DRAM

63	56	55	48	47	40	39	32	31	24	23	16	15	8	7	0
4		5		6		7		0		1		2		3	
C		D		E		F		8		9		A		B	

For both types of RBUF read, reading an element does not modify any RBUF data, and does not free the element, so buffered data can be read as many times as desired.

8.2.2.1 SPI-4

SPI-4 data is placed into RBUF as follows:

- At chip reset all elements are marked invalid (available).
- When a SPI-4 Control Word is received (i.e., when RCTL is asserted) it is placed in a temporary holding register. The Checksum accumulator is cleared. The subsequent action is based on the Type field.

— If Type is Idle or Training, the Control Word is discarded.

— If Type is not Idle or Training:

An available RBUF element is allocated by receive control logic. (If there is no available element, the data is discarded and `MSF_Interrupt_Status[RBUF_Overflow]` is set.) Note that this normally should not happen because, when the number of RBUF elements falls below a programmed limit, the flow control status is sent back to the PHY device (refer to

Section 8.2.7.1). The SPI-4 Control Word Type, EOPS, SOP, and ADR fields are placed into a temporary status register. The Byte_Count field of the element status is set to 0x0. As each Data Word is received, the data is written into the element, starting at offset 0x0 in the element, and Byte_Count is updated. Subsequent Data transfers are placed at higher offsets (0x2, 0x4, etc.). The 16-bit Checksum Accumulator is also updated with the 1's-complement addition of each byte pair. (Note that, if the data transfer has an odd number of bytes, a byte of zeros is appended as the more significant byte, before the checksum addition is done.)

- If a Control Word is received before the element is full — the element is marked valid. EOP for the element is taken from the value of the EOPS field (see [Table 83](#)) from the just-received Control Word. If the EOPS field from the just-received Control Word indicates that EOP is asserted, then the Byte_Count for the element is decremented by 0 or 1, according to the EOPS field (i.e., decrement by 0 if *two* bytes are valid, and by 1 if *one* byte is valid). If the EOPS field indicates Abort, Byte_Count is rounded up to the next multiple of 4. The temporary status register value is put into Full_Element_List.
- If the element becomes full before receipt of another Control Word — the element is marked as pre-valid. The eventual status is based on the next SPI-4 transfer(s).
- If the next transfer is a Data Word — the previous element is changed from pre-valid to valid. The EOP for the element is 0. The temporary status register value is put into Full_Element_List. Another available RBUF element is allocated, and the new data is written into it. The temporary status for the new element gets the same ADR field of the previous element, and SOP is set to 0. The Status word Byte_Count field is set to 0x2, and will be incremented as more Data Words arrive. The Checksum Accumulator is cleared.
- If the next transfer is a Control Word — the previous element is changed from pre-valid to valid. EOP for the element is taken from the value of the EOPS field from the just-received Control Word. If the EOPS field from the just-received Control Word indicates that EOP is asserted, then the Byte_Count for the element is decremented by 0 or 1, according to the EOPS field (i.e., decrement by 0 if *two* bytes valid, and by 1 if *one* byte is valid). The temporary status register value is put into Full_Element_List.

Data received from the bus is placed into the element lowest offset first, in big-endian order (i.e., with the first byte received in the most significant byte of the 32-bit word, etc.).



The status contains the following information:

3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
RPROT	Element						Byte Count						SOP	EOP	Err	Len Err	Par Err	Abort Err	Null	Type	ADR											

6	6	6	6	5	5	5	5	5	5	5	5	5	5	4	4	4	4	4	4	4	4	4	4	3	3	3	3	3	3	3	3
Reserved																Checksum															

The definitions of the fields are shown in [Table 90](#).

Table 90. RBUF SPIF-4 Status Definition

Field	Definition
RPROT	This bit is a 0, indicating that the Status is for SPI-4. It is derived from the RPROT input signal.
Element	The element number in the RBUF that holds the data. This is equal to the offset in RBUF of the first byte in the element, shifted right by six places
Byte_Count	Indicates the number of Data bytes, from 1 to 256, in the element (value 0x00 means 256). This field is derived from the number of data transfers that fill the element, and also the EOPS field of the Control Word that most recently succeeded the data transfer.
SOP	Indicates whether the element is the start of a packet. This field is taken from the SOP field of the Control Word that most recently preceded the data transfer for the first element allocated after a Control Word. For subsequent elements (i.e., if more than one element worth of data follow the Control Word) this value is 0.
EOP	Indicates whether the element is the end of a packet. This field is taken from the EOPS field of the Control Word that most recently succeeded the data transfer.
Err	Error. This is the logical OR of Par Err, Len Err, and Abort Err.
Len Err	A non-EOP burst occurred that was not a multiple of 16 bytes.
Par Err	Parity Error was detected in the DIP-4 parity field. See the description in Section 8.2.8.1 .
Abort Err	An EOP with Abort was received on bits [14:13] of the Control Word that most recently succeeded the data transfer.
Null	Null receive. If this bit is set, it means that the Rx_Thread_Freelist timeout expired before any more data was received, and that a null Receive Status Word is being pushed, to keep the receive pipeline flowing. The rest of the fields in the Receive Status Word must be ignored; there is no data or RBUF entry associated with a null Receive Status Word.
Type	This field is taken from the Type field of the Control Word that most recently preceded the data transfer.
ADR	The port number to which the data is directed. This field is taken from the ADR field of the Control Word that most recently preceded the data transfer.
Checksum	Checksum calculated over the Data Words in the element. This can be used for TCP.



8.2.2.2 CSIX

CSIX CFrames are placed into either RBUF or FCEFIFO as follows:

At chip reset, all RBUF elements are marked invalid (available) and FCEFIFO is empty.

When a Base Header is sent (i.e., when RxSof is asserted) it is placed in a temporary holding register. The Ready Field is extracted and held to be put into **FC_Egress_Status** CSR when (and if) the entire CFrame is received without error. The Type field is extracted and used to index into **CSIX_Type_Map** CSR to determine one of four actions.

- Discard (except for the Ready Field as described in [Section 8.2.7.2.1](#)).
- Place into RBUF Control CFrame partition.
- Place into RBUF Data CFrame partition.
- Place into FCEFIFO.

Note: Normally Idle CFrames (Type 0x0) will be discarded, Command and Status CFrames (Type 0x7) will be placed into Control Partition, Flow Control CFrames (Type 0x6) will be placed into FCEFIFO, and all others will be placed into Data Partition (see [Table 87](#)). The remapping done through the **CSIX_Type_Map** CSR allows for more flexibility in usage, if desired.

If the action is Discard, the CFrame is discarded (except for the Ready Field as described in [Section 8.2.7.2.1](#)). The Base Header, as well as Extension Header and Payload (if any) are discarded.

If the destination is FCEFIFO:

The Payload is placed into the FCEFIFO, to be sent to the Ingress IXP2800 Network Processor over the TXCDAT pins. If there is not enough room in FCEFIFO for the entire CFrame, based on the Payload Size in the Base Header, the entire CFrame is discarded and **MSF_Interrupt_Status[FCEFIFO_Overflow]** is set.

If the destination is RBUF (either Control or Data):

An available RBUF element of the corresponding type is allocated by receive control logic. If there is not an available element, the CFrame is discarded and **MSF_Interrupt_Status[RBUF_Overflow]** is set. Note that this normally should not happen because, when the number of RBUF elements falls below a programmed limit, backpressure is sent to the Switch Fabric. (Refer to [Section 8.2.7.2](#).) The Type, Payload Length, CR (CSIX Reserved), and P (Private) bits, and (subsequently arriving) Extension Header are placed into a temporary status register. As the Payload (including padding if any) is received, it is placed into the allocated RBUF element, starting at offset 0x0. (Note that it is more exact to state that the first four bytes after the Base Header are placed into the status register as Extension Header. For Flow Control CFrames, there is no Extension Header; the first four bytes are part of the Payload. They would be found in the Extension Header field of the Status — no bytes are lost.)

When all of the Payload data (including padding if any), as indicated by the Payload Length field, and Vertical Parity has been received, the element is marked valid. If another RxSof is received prior to receiving the entire Payload, the element is also marked valid, and the Length Error status bit is set. If the Payload Length field of the Base Header is greater than the element size (as configured in **MSF_Rx_Control[RBUF_Element_Size]**), then the Length Error bit in the status will be set, and all payload bytes above the element size will be discarded. The temporary status register value is put into **Full_Element_List**.



Note: In CSIX protocol, an RBUF element is allocated only on RxSof assertion. Therefore, the element size must be programmed based on the Switch Fabric usage. For example, if the switch never sends a payload greater than 128 bytes, then 128-byte elements can be selected. Otherwise, 256-byte elements must be selected.

Data received from the bus is placed into the element lowest-offset first in big-endian order (that is, with the first byte received in the most significant byte of the 32-bit word, etc.).

The status contains the following information:

3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
RPROT	Element						Payload Length						CR	P	Err	Len Err	HP Err	VP Err	Null	Reserved						Type						

6	6	6	6	5	5	5	5	5	5	5	5	5	5	5	4	4	4	4	4	4	4	4	4	4	3	3	3	3	3	3	3	3
3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	
Extension Header																																

The definitions of the fields are shown in [Table 91](#).

Table 91. RBUF CSIX Status Definition

Field	Definition
RPROT	This bit is a 1, indicating that the Status is for CSIX-L1. It is derived from the RPROT input signal.
Element	The element number in the RBUF that holds the data. This is equal to the offset in RBUF of the first byte in the element, shifted right by six places.
Payload Length	Payload Length Field from the CSIX Base Header. A value of 0x0 indicates 256 bytes.
CR	CR (CSIX Reserved) bit from the CSIX Base Header.
P	P (Private) bit from the CSIX Base Header.
Err	Error. This is the logical OR of VP Err, HP Err, and Len Err.
Len Err	Length Error; either amount of Payload received (before receipt of next Base Header) did not match value indicated in Base Header Payload Length field) or Payload Length field was greater than size of RBUF element.
HP Err	Horizontal Parity Error was detected on the CFrame. See description in Section 8.2.8.2.1 .
VP Err	Vertical Parity Error was detected on the CFrame. See description in Section 8.2.8.2.2 .
Null	Null receive. If this bit is set, it means that the Rx_Thread_Freelist timeout expired before any more data was received, and that a null Receive Status Word is being pushed, to keep the receive pipeline flowing. The rest of the fields in the Receive Status Word must be ignored; there is no data or RBUF entry associated with a null Receive Status Word.
Type	Type Field from the CSIX Base Header.
Extension Header	The Extension Header from the CFrame. The bytes are received in big-endian order; byte 0 is in bits 63:56, byte 1 is in bits 55:48, byte 2 is in bits 47:40, and byte 3 is in bits 39:32.



8.2.3 Full Element List

Receive control hardware maintains the Full Element List to hold the status of valid RBUF elements, in the order in which they were received. When an element is marked valid (as described in [Section 8.2.2.1](#) for SPI-4 and [Section 8.2.2.2](#) for CSIX), its status is added to the tail of the Full Element List. When a Microengine is notified of element arrival (by having the status written to its S_Transfer register; see [Section 8.2.4](#)), it is removed from the head of the Full Element List.

8.2.4 Rx_Thread_Freelist_#

Each **Rx_Thread_Freelist_#** is a FIFO that indicates Microengine Contexts that are awaiting an RBUF element to process. This allows the Contexts to indicate their ready-status prior to the reception of the data, as a way to eliminate latency. Each entry added to a Freelist also has an associated S_Transfer register and signal number. The receive logic maintains either one, two, or three separate lists based on **MSF_Rx_Control[RBUF_Partition]**, **MSF_Rx_Control[CSIX_Freelist]**, and **Rx_Port_Map** as shown in [Table 92](#).

Table 92. Rx_Thread_Freelist Use

Number of Partitions ¹	Use	CSIX_Freelist ²	Rx_Thread_Freelist_# Used		
			0	1	2
1	SPI-4 only	n/a	SPI-4 Ports equal to or below Rx_Port_Map	SPI-4 Ports above Rx_Port_Map	Not Used
2	CSIX only	0	CSIX Data	CSIX Control	Not Used
		1	CSIX Data and CSIX Control	Not Used	Not Used
3	Both SPI-4 and CSIX	0	CSIX Data	SPI-4	CSIX Control
		1	CSIX Data and CSIX Control	SPI-4	Not Used

1. Programmed in **MSF_Rx_Control[RBUF_Partition]**.
2. Programmed in **MSF_Rx_Control[CSIX_Freelist]**.

To be added as ready to receive an element, an Microengine does an `msf[write]` or `msf[fast_write]` to the **Rx_Thread_Freelist_#** address; the write data is the Microengine/Context/S_Transfer register number to add to the Freelist. Note that using the data (rather than the command bus ID) permits a Context to add either itself or other Contexts as ready.

When there is valid status at the head of the Full Element List, it will be pushed to a Microengine. The receive control logic pushes the status information (which includes the element number) to the Microengine in the head entry of **Rx_Thread_Freelist_#**, and sends an Event Signal to the Microengine. It then removes that entry from the **Rx_Thread_Freelist_#**, and removes the status from Full Element List. (Note that this implies the restriction — a Context waiting on status must not read the S_Transfer register until it has been signaled.) See [Section 8.2.6](#) for more information. In the event that **Rx_Thread_Freelist_#** is empty, the valid status will be held in Full Element List until an entry is put into **Rx_Thread_Freelist_#**.



8.2.5 Rx_Thread_Freelist_Timeout_#

Each **Rx_Thread_Freelist_#** has an associated countdown timer. If the timer expires and no new receive data is available yet, the receive logic will autopush a Null Receive Status Word to the next thread on the **Rx_Thread_Freelist_#**. A Null Receive Status Word has the “Null” bit set, and does not have any data or RBUF entry associated with it.

The **Rx_Thread_Freelist_#** timer is useful for certain applications. Its primary purpose is to keep the receive processing pipeline (implemented as microcode running on the Microengine) moving even when the line has gone idle. It is especially useful if the pipeline is structured to handle mpackets in groups, i.e., eight mpackets at a time.

If seven mpackets are received, the line goes idle, and the timeout triggers the autopush of a null Receive Status Word, filling the eighth slot and allowing the pipeline to advance. Another example is if one valid mpacket is received before the line goes idle for a long period; seven null Receive Status Words will be autopushed, allowing the pipeline to proceed.

Typically, the timeout interval is programmed to be slightly larger than the minimum arrival time of the incoming cells or packets. The timer is controlled by using the **Rx_Thread_Freelist_Timeout_#** CSR. The timer may be enabled or disabled, and the timeout value specified using this CSR.

The following rules define the operation of the **Rx_Thread_Freelist** timer.

1. Writing a non-zero value to the **Rx_Thread_Freelist_Timeout_#** CSR both resets the timer and enables it. Writing a zero value to this CSR resets the timer and disables it.
2. If the timer is disabled, then only valid (non-null) Receive Status Words are autopushed to the receive threads; null Receive Status Words are never pushed.
3. If the timer expires and the **Rx_Thread_Freelist_#** is non-empty, but there is no mpacket available, this will trigger the autopush of a null Receive Status Word.
4. If the timer expires and the **Rx_Thread_Freelist_#** is empty, the timer stays in the EXPIRED state and is not restarted. A null Receive Status Word cannot be autopushed, since the logic has no destination to push anything to.
5. An expired timer is reset and restarted if and only if an autopush, null or non-null, is performed.
6. Whenever there is a choice, the autopush of a non-null Receive Status Word takes precedence over a null Receive Status Word.

8.2.6 Receive Operation Summary

During receive processing, received CFrames, cells, and packets (which in this context are all called mpackets) are placed into the RBUF, and then, when marked valid, are immediately handed off to a Microengine for processing. Normally, by application design, some number of Microengine Contexts will be assigned to receive processing. Those Contexts will have their number added to the proper **Rx_Thread_Freelist_#** (via `msf[write]` or `msf[fast_write]`), and then will go to sleep to wait for arrival of an mpacket (or alternatively poll waiting for arrival of an mpacket).

When an mpacket becomes valid as described in [Section 8.2.2.1](#) for SPI-4 and [Section 8.2.2.2](#) for CSIX, receive control logic will autopush eight bytes of information for the element to the Microengine/Context/S_Transfer registers at the head of **Rx_Thread_Freelist_#**. The information pushed is (see [Table 90](#) and [Table 91](#) for detailed definitions):

- Status Word (SPI-4) or Header Status (CSIX) to Transfer register n (n is the Transfer register programmed to the **Rx_Thread_Freelist_#**)
- Checksum (SPI-4) or Extension Header (CSIX) to Transfer register n+1

To handle the case where the receive Contexts temporarily fall behind and **Rx_Thread_Freelist_#** is empty, all received element numbers are held in the Full Element List. In that case, as soon as an **Rx_Thread_Freelist_#** entry is entered, the status of the head element of Full Element List will be pushed to it.

The Microengine may read part of (or the entire) RBUF element to their S_Transfer registers (via `msf[read]` instruction) for header processing, etc., and may also move the element data to DRAM (via `dram[rbuf_rd]` instruction).

When a Context is done with an element, it does an `msf[write]` or `msf[fast_write]` to the **RBUF_Element_Done** address; the write data is the element number. This marks the element as free and available to be re-used. There is no restriction on the order in which elements are freed; Contexts can do different amounts of processing per element based on the contents of the element — therefore, elements can be returned in a different order than they were handed to Contexts.

The states that an RBUF element goes through are shown in [Figure 92](#).

Figure 92. RBUF Element State Diagram

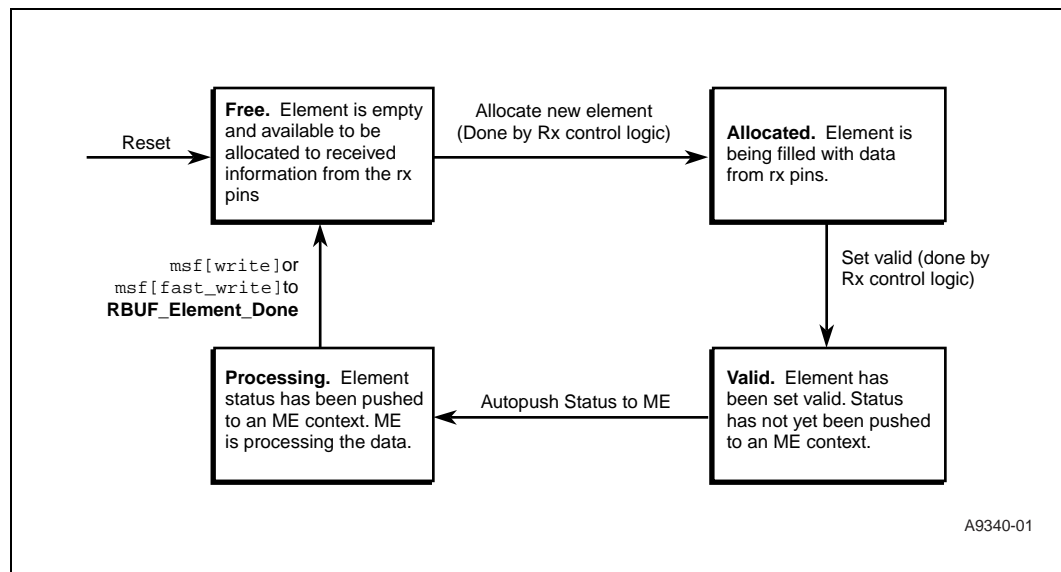


Table 93 summarizes the differences in RBUF operation between the SPI-4 and CSIX protocols.

Table 93. Summary of SPI-4 and CSIX RBUF Operations

Operation	SPI-4	CSIX
When is RBUF Element Allocated?	Upon receipt of Payload Control Word or when Element data section fills and more Data Words arrive. The Payload Control Word allocates an element for data that will be received subsequent to it.	Start of Frame and Base Header Type is mapped to RBUF (in the CSIX_Type_Map CSR).
How Much Data is Put into Element?	All Data Words received between two Payload Control Words, or number of bytes in the element, whichever is less.	Number of bytes specified in Payload Length field of Base Header.
How is RBUF Element Set Valid?	Upon receipt of Payload Control Word or when Element data section fills. The Payload Control Word validates the element holding data received prior to it.	All Payload is received (or if premature SOF, which will set an error bit in Element Status).
How is RBUF Element Handed to Microengine?	Element Status is pushed to Microengine at the head of the appropriate Rx_Thread_Freelist_# (based on the protocol). Status is pushed to two consecutive Transfer registers; bits[31:0] of Element Status to the first Transfer register and bits[63:32] to the next higher numbered Transfer register.	
How is RBUF Element returned to free list?	CSR write to RBUF_Element_Done .	

8.2.7 Receive Flow Control Status

Flow control is handled in hardware. There are specific functions for SPI-4 and CSIX.

8.2.7.1 SPI-4

SPI-4, FIFO status information is sent periodically over the RSTAT signals from the Link Layer device (which is the IXP2800 Network Processor) to the PHY device. (Note that TXCDAT pins can act as RSTAT based on the **MSF_Rx_Control[RSTAT_Select]** bit.) The information to be sent is based on the number of RBUF elements available to receive SPI-4.

The FIFO status of each port is encoded in a 2-bit data structure — code 0x3 is used for framing the data, and the other three codes are valid status values. The FIFO status words are sent according to a repeating calendar sequence. Each sequence begins with the framing code to indicate the start of a sequence, followed by the status codes, followed by a parity code covering the preceding frame. The length of the calendar is defined in **Rx_Calendar_Length**, which is a CSR field that is initialized with the length of the calendar, since in many cases fewer than 256 ports are in use.

When **TRAIN_DATA[RSTAT_En]** is disabled, RSTAT is held at 0x3.

The IXP2800 Network Processor transmits FIFO status only if **TRAIN_DATA[RSTAT_En]** is set. The logic sends “Satisfied,” “Hungry,” or “Starving” based on either the upper limit of the RBUF, a global override value set in **MSF_Rx_Control[RSTAT_OV_VALUE]**, or a port-specific override value set in **RX_PORT_CALENDAR_STATUS_#**. The choice is controlled by **MSF_RX_CONTROL[RX_Calendar_Mode]**.

When set to **Conservative_Value**, the status value sent for each port is the most conservative of:

- The RBUF upper limit
- **MSF_RX_CONTROL[RSTAT_OV_VALUE]**
- **RX_PORT_CALENDAR_STATUS_#**

“Satisfied” is more conservative than “Hungry,” which is more conservative than “Starving.”



When **MSF_RX_CONTROL[RX_Calendar_Mode]** is set to **Force_Override**, the value of **RX_PORT_CALENDAR_STATUS_#** is used to determine which status value is sent. If **RX_PORT_CALENDAR_STATUS_#** is set to 0x3, then the global status value set in **MSF_RX_CONTROL[RSTAT_OV_VALUE]** is sent; otherwise, the port-specific status value set in **RX_PORT_CALENDAR_STATUS_#** is sent.

The RBUF upper limit is based on the **MSF_RX_CONTROL** register and is defined in [Table 89](#). The upper limit is programmed in **HWM_Control[RBUF_S_HWM]**. Note that either RBUF partition 0 or partition 1 will be used for SPI-4 ([Table 88](#)).

8.2.7.2 CSIX

There are two types of CSIX flow control:

- Link-level
- Virtual Output Queue (VOQ)

Information received from the Switch Fabric by the Egress IXP2800 Network Processor, must be communicated to the Ingress IXP2800 Network Processor, which is sending data to the Switch Fabric.

8.2.7.2.1 Link-Level

Link-level flow control can be used to stop all transmission. Separate Link-level flow control is provided for Data CFrames and Control CFrames. CSIX protocol provides link-level flow control as follows. Every CFrame Base Header contains a Ready Field, which contains two bits; one for Control traffic (bit 6 of byte 1) and one for Data traffic (bit 7 of byte 1). The CSIX requirement for response is:

From the tick that the Ready Field leaves a component the maximum response time for a pause operation is defined as: $n \cdot T$, $n = C + L$ where:

- T is the clock period of the interface
- n is the maximum number of ticks for the response
- C is a constant for propagating the field within the “other” component (or chipset as the case may be) to the interface logic control the reverse direction data flow. C is defined to be 32 ticks.
- L is the maximum number of ticks to transport the maximum fabric CFrame size.

As each CFrame is received, the value of these bits is copied (by receive hardware) into the **FC_Egress_Status[SF_CReady]** and **FC_Egress_Status[SF_DReady]** respectively. The value of these two bits is sent from the Egress to the Ingress IXP2800 Network Processor on the TXCSR_B signal, and can be used to stop transmission to the Switch Fabric, as described in [Section 8.3.4.2](#). The TXCSR_B signal is described in [Section 8.5.1](#).

8.2.7.2.2 Virtual Output Queue

CSIX protocol provides Virtual Output Queue Flow Control via Flow Control CFrames. CFrames that were mapped to FCEFIFO (via the **CSIX_Type_Map** CSR) are parsed by the receive control logic and placed into FCEFIFO, which provides buffering while they are sent from the Egress IXP2800 Network Processor to the Ingress IXP2800 Network Processor over the TXCDAT signals (normally Flow Control CFrames would be mapped to FCEFIFO).

The entire CFrame is sent over TXCDAT, including the Base Header and Vertical Parity field. The 32-bit CWord is sent four bits at a time, most significant bits first. The CFrames are forwarded in a “cut-through” manner, meaning that the Egress IXP2800 Network Processor does not wait for the entire CFrame to be received before forwarding (each CWord can be forwarded as it is received).

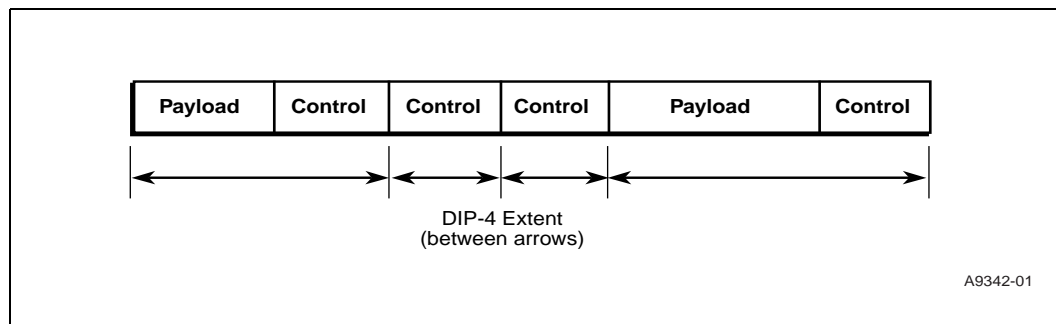
If FCEFIFO gets full, as defined by **HWM_Control[FCEFIFO_HWM]**, then the **FC_Egress_Status[TM_CReady]** bit will be deasserted (to inform the Ingress IXP2800 Network Processor to deassert Control Ready in CFrames sent to the Switch Fabric). [Section 8.3.4.2](#) describes how Flow Control information is used in the Ingress IXP2800 Network Processor.

8.2.8 Parity

8.2.8.1 SPI-4

The receive logic computes 4-bit Diagonal Interleaved Parity (DIP-4) as specified in the SPI-4 specification. The DIP-4 field received in a control word contains odd parity computed over the current Control Word and the immediately preceding data words (if any) following the last Control Word. [Figure 93](#) shows the extent of the DIP-4 codes.

Figure 93. Extent of DIP-4 Codes



There is a DIP-4 Error Flag and a 4-bit DIP-4 Accumulator register. After each Control Word is received, the Flag is conditionally reset (see Note below this paragraph) and the Accumulator register is cleared. As each Data Word (if any), and the first succeeding Control Word is received, DIP-4 parity is accumulated in the register, as defined in the SPI-4 spec. The accumulated parity is compared to the value received in the DIP-4 field of that first Control Word. If it does not match, the DIP-4 Error Flag is set. The value of the flag becomes the element status Par Err bit.

Note: An error in the DIP-4 code invalidates the transfers before and after the Control Word, since the control information is assumed to be in error. Therefore the DIP-4 Error Flag is not reset after a Control Word with bad DIP-4 parity. It is only reset after a Control Word with correct DIP-4 parity.



8.2.8.2 CSIX

8.2.8.2.1 Horizontal Parity

The receive logic computes Horizontal Parity on each 16 bits of each received Cword (there is a separate parity for data received on rising and falling edge of the clock).

There is an internal HP Error Flag. At the end of each CFrame, the flag is reset. As each 16 bits of each Cword is received, the expected odd-parity value is computed from the data, and compared to the value received on RxPar. If there is a mismatch, the flag is set. The value of the flag becomes the element status HP Err bit.

If the HP Error Flag is set:

- the **FC_Egress_Status[SF_CReady]** and **FC_Egress_Status[SF_DReady]** bits are cleared
- the **MSF_Interrupt_Status[HP_Error]** bit is set (which can interrupt the Intel XScale® core if enabled)

8.2.8.2.2 Vertical Parity

The receive logic computes Vertical Parity on CFrames.

There is a VP Error Flag and a 16-bit VP Accumulator register. At the end of each CFrame, the flag is reset and the register is cleared. As each Cword is received, odd parity is accumulated in the register as defined in the CSIX specification (16 bits of vertical parity are formed on 32 bits of received data by treating the data as words; i.e., bit 0 and bit 16 of the data are accumulated into parity bit 0, bit 1, and bit 17 of the data are accumulated into parity bit 1, etc.). After the entire CFrame has been received (including the Vertical Parity field; the two bytes following the Payload) the accumulated value should be 0xFFFF. If it is not, the VP Error Flag is set. The value of the flag becomes the element status VP Err bit.

Note: The Vertical Parity always follows the Payload, which may include padding to the CWord width if the Payload Length field is not an integral number of CWords. The CWord width is programmed in **MSF_Rx_Control[Rx_CWord_Size]**.

If the VP Error Flag is set:

- the **FC_Egress_Status[SF_CReady]** and **FC_Egress_Status[SF_DReady]** bits are cleared
- the **MSF_Interrupt_Status[VP_Error]** bit is set (which can interrupt the Intel XScale® core)

8.2.9 Error Cases

Receive errors are specific to the protocol, SPI-4 or CSIX. The element status, described in [Table 90](#) and [Table 91](#), has appropriate error bits defined. Also, there are some IXP2800 Network Processor specific error cases — for example, when an mpacket arrives with no free elements — that are logged in the **MSF_Interrupt_Status** register, which can interrupt the Intel XScale® core if enabled.

8.3 Transmit

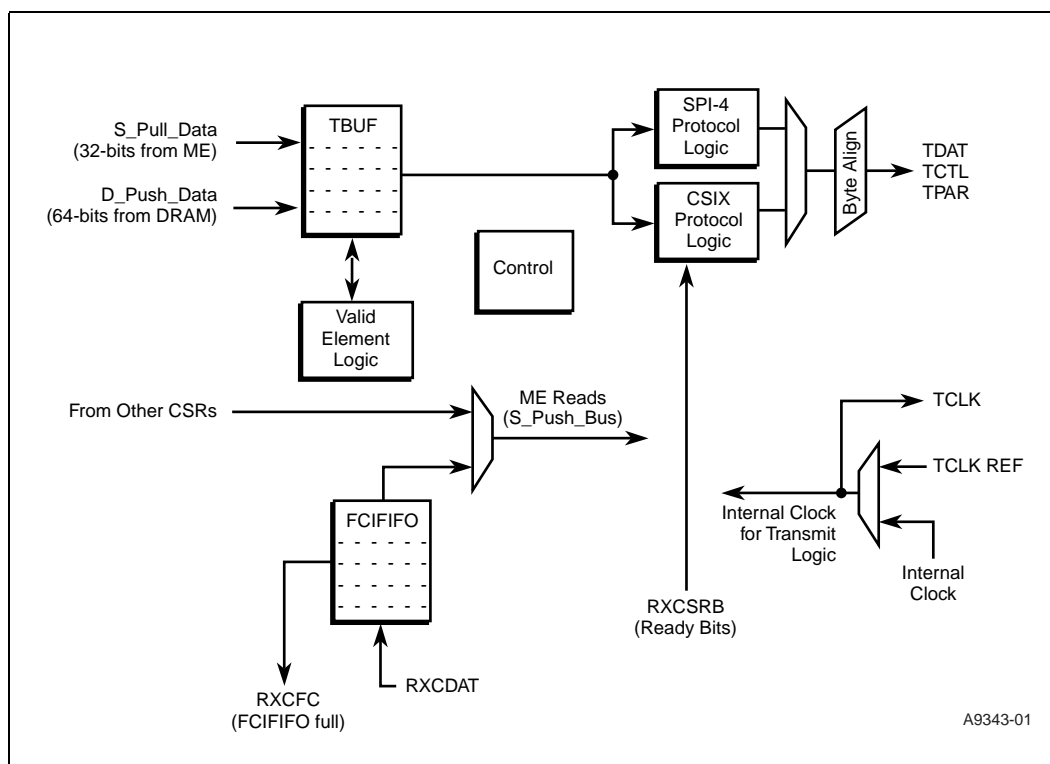
The transmit section consists of:

- Transmit Pins (Section 8.3.1)
- Transmit Buffer (Section 8.3.2)
- Byte Aligner (Section 8.3.2)

Each of these is described below.

Figure 94 is a simplified block diagram of the MSF transmit block.

Figure 94. Simplified Transmit Section Block Diagram



8.3.1 Transmit Pins

The use of the transmit pins is a function of the protocol (which is determined by TBUF partition in MSF_Tx_Control CSR) as shown in Table 94.

Table 94. Transmit Pins Usage by Protocol (Sheet 1 of 2)

Name	Direction	SPI-4 Use	CSIX Use
TCLK	Output	TDCLK	RxCIk
TDAT[15:0]	Output	TDAT[15:0]	RxData[15:0]
TCTL	Output	TCTL	RxSOF



Table 94. Transmit Pins Usage by Protocol (Sheet 2 of 2)

Name	Direction	SPI-4 Use	CSIX Use
TPAR	Output	Not Used	RTxPar
TSCLK	Input	TSCLK	Not Used
TSTAT[1:0]	Input	TSTAT[1:0]	Not Used

8.3.2 TBUF

The TBUF is a RAM that holds data and status to be transmitted. The data is written into sub-blocks referred to as elements, by Microengine or the Intel XScale® core. TBUF contains a total of 8 Kbytes of data, and associated control.

Table 95 shows the order in which data is written into TBUF. Each number represents a byte, in order of transmission onto the tx interface. Note that this is reversed on a 32-bit basis relative to RBUF — the swap of 4 low bytes and 4 high bytes is done in hardware to facilitate the transmission of bytes.

Table 95. Order in which Data is Transmitted from TBUF

Data/Payload								Address Offset (Hex)
0	1	2	3	4	5	6	7	0
8	9	A	B	C	D	E	F	8
10	11	12	13	14	15	16	17	10

The mapping of elements to address offset in TBUF is based on the TBUF partition and element size, as programmed in **MSF_Tx_Control** CSR. TBUF can be partitioned into one, two, or three partitions based on **MSF_Tx_Control[TBUF_Partition]**. The mapping of partitions to transmit data is shown in Table 96.

Table 96. Mapping of TBUF Partitions to Transmit Protocol

Number of Partitions in Use	Transmit Data Protocol	Data Use by Partition, Fraction of TBUF Used, Start Byte Offset (Hex)		
		Partition Number		
		0	1	2
1	SPI-4 only	SPI-4 All Byte 0	n/a	n/a
2	CSIX only	CSIX Data 3/4 of TBUF Byte 0	CSIX Control 1/4 of TBUF Byte 0x1800	n/a
3	Both SPI-4 and CSIX	CSIX Data 1/2 of TBUF Byte 0	SPI-4 3/8 of TBUF Byte 0x1000	CSIX Control 1/8 of TBUF Byte 0x1C00

The data in each segment is further broken up into elements, based on **MSF_Tx_Control[TBUF_Element_Size_#]** (n = 0,1,2). There are three choices of element size: 64, 128, or 256 bytes.



Table 97 shows the TBUF partition options. Note that the choice of element size is independent for each partition.

Table 97. Number of Elements per TBUF Partition

TBUF_Partition Field	TBUF_Element_Size_# Field	Partition Number		
		0	1	2
00 (1 partition)	00 (64 bytes)	128	Unused	Unused
	01 (128 bytes)	64		
	10 (256 bytes)	32		
01 (2 partitions)	00 (64 bytes)	96	32	Unused
	01 (128 bytes)	48	16	
	10 (256 bytes)	24	8	
10 (3 partitions)	00 (64 bytes)	64	48	16
	01 (128 bytes)	32	24	8
	10 (256 bytes)	16	12	4

The Microengine can write data from Microengine S_TRANSFER_OUT registers to the TBUF using the `msf[write]` instruction, where they specify the starting byte number (which must be aligned to four bytes), and number of 32-bit words to write. The number in the instruction can be either the number of 32-bit words, or number of 32-bit word pairs, using the single and double instruction modifiers, respectively. Data is pulled from the Microengine to TBUF via S_Pull_Bus. `msf[write, $s_xfer_reg, src_op_1, src_op_2, ref_cnt], optional_token`

The `src_op_1` and `src_op_2` operands are added together to form the address in TBUF (note that the base address of the TBUF is 0x2000). The `ref_cnt` operand is the number of 32-bit words or word pairs, which are pulled from sequential S_TRANSFER_OUT registers, starting with `$s_xfer_reg`.

The Microengine can move data from DRAM to TBUF using the instruction `dram[tbuf_wr, --, src_op1, src_op2, ref_cnt], indirect_ref`

The `src_op_1` and `src_op_2` operands are added together to form the address in DRAM, so the `dram` instruction must use indirect mode to specify the TBUF address. The `ref_cnt` operand is the number of 64-bit words that are written into TBUF.

Data is stored in big-endian order. The most significant (lowest numbered) byte of each 32-bit word is transmitted first.

All elements within a TBUF partition are transmitted in the order. Control information associated with the element (Section 98 and Section 99) defines which bytes are valid. The data from the TBUF will be shifted and byte-aligned to the TDAT pins as required. Four parameters are defined.

Prepend Offset — Number of the first byte to send. This is information that is prepended onto the payload, for example as a header. It need not start at offset 0 in the element.

Prepend Length — Number of bytes of prepended information. This can be 0 to 31 bytes. If it is 0, then the Prepend Offset must also be 0.



Payload Offset — Number of bytes to skip from the last 64-bit word of the Prepend to the start of Payload. The absolute byte number of the first byte of Payload in the element is:
 $((\text{Prepend Offset} + \text{Prepend Length} + 0x7) \&\& 0xF8) + \text{Payload Offset}$.

Payload Length — Number of bytes of Payload.

The sum of Prepend Length, Payload length, and any gaps in between them $((\text{prepend_offset} + \text{prepend_length} + 7) \&\& 0xF8) + \text{payload_offset} + \text{payload_length}$) must be no greater than the number of bytes in the element. Typically, the Prepend is computed by a Microengine and written into the TBUF by `msf[write]` and the Payload will be written by `dram[tbuf_wr]`. These two operations can be done in either order; the microcode is responsible for making sure the element is not marked valid to transmit until all data is in the element (see [Section 8.3.3](#)).

[Example 36](#) illustrates the usage of the parameters. The element in [Example 36](#) is shown as 8 bytes wide because the smallest unit that can be moved into the element is 8 bytes. In [Example 36](#), bytes to be transmitted are shown in black (the offsets are byte numbers); bytes in *gray* are written into TBUF (because the writes always write 8 bytes), but are not transmitted.

Prepend Offset = 6 (bytes 0x0 — 0x5 are not transmitted).

Prepend Length = 16 (bytes 0x6 — 0x15 are transmitted).

Payload Offset = 7 (bytes 0x16 — 0x1E are not transmitted). The Payload starts in the next 8-byte row (i.e., the next “empty” row above where the Prepend stops), even if there is room in the last row containing Prepend information. This is done because the TBUF does not have byte write capability, and therefore would not merge the `msf[write]` and `dram[tbuf_wr]`. The software computing the Payload Offset only needs to know how many bytes of the payload that were put into DRAM need to be removed.

Payload Length = 33 (bytes 0x1F through 0x3F are transmitted).

Example 36. TBUF Prepend and Payload

0	1	2	3	4	5	6	7
8	9	A	B	C	D	E	F
10	11	12	13	14	15	16	17
18	19	1A	1B	1C	1D	1E	1F
20	21	22	23	24	25	26	27
28	29	2A	2B	2C	2D	2E	2F
30	31	32	33	34	35	36	37
38	39	3A	3B	3C	3D	3E	3F

The transmit logic will send the valid bytes onto TDAT correctly aligned and with no gaps. The protocol transmitted, SPI-4 or CSIX (and the value of the TPROT output) are based on which partition of TBUF the data was placed (see [Table 95](#)).



8.3.2.1 SPI-4

For SPI-4, data is put into the data portion of the element, and information for the SPI-4 Control Word that will precede the data is put into the Element Control Word.

When the Element Control Word is written the information is (the data comes from two consecutive Transfer registers; bits [31:0] from the lower numbered and bits[63:32] from the higher numbered):

3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
Payload Length										Prepend Offset		Prepend Length					Payload Offset		Res	Skip	Abort	SOP	EOP	ADR						

6	6	6	6	5	5	5	5	5	5	5	5	5	5	4	4	4	4	4	4	4	4	4	3	3	3	3	3	3	3	3	
3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2
Res																															

The definitions of the fields are shown in [Table 98](#).

Table 98. TBUF SPI-4 Control Definition

Field	Definition
Payload Length	Indicates the number of Payload bytes, from 1 to 256, in the element. The value of 0x00 means 256 bytes. The sum of Prepend Length and Payload Length will be sent. That value will also control the EOPS field (1 or 2 bytes valid indicated) of the Control Word that will succeed the data transfer. Note 1.
Prepend Offset	Indicates the first valid byte of Prepend, from 0 to 7, as defined in Section 8.3.2 .
Prepend Length	Indicates the number of bytes in Prepend, from 0 to 31.
Payload Offset	Indicates the first valid byte of Payload, from 0 to 7, as defined in Section 8.3.2 .
Skip	Allows software to allocate a TBUF element and then not transmit any data from it. 0—transmit data according to other fields of Control Word. 1—free the element without transmitting any data.
Abort	Indicates if the element is the end of a packet that should be aborted. If this bit is set, the status code of EOP Abort will be sent in the EOPS field of the Control Word that will succeed the data transfer. Note 1.
SOP	Indicates if the element is the start of a packet. This field will be sent in the SOPC field of the Control Word that will precede the data transfer.
EOP	Indicates if the element is the end of a packet. This field will be sent in the EOPS field of the Control Word that will succeed the data transfer. Note 1.
ADR	The port number to which the data is directed. This field will be sent in the ADR field of the Control Word that will precede the data transfer.
NOTE:	
1. Normally EOPS is sent on the next Control Word (along with ADR and SOP) to start the next element. If there is no valid element pending at the end of sending the data, the transmit logic will insert an Idle Control Word with the EOPS information.	



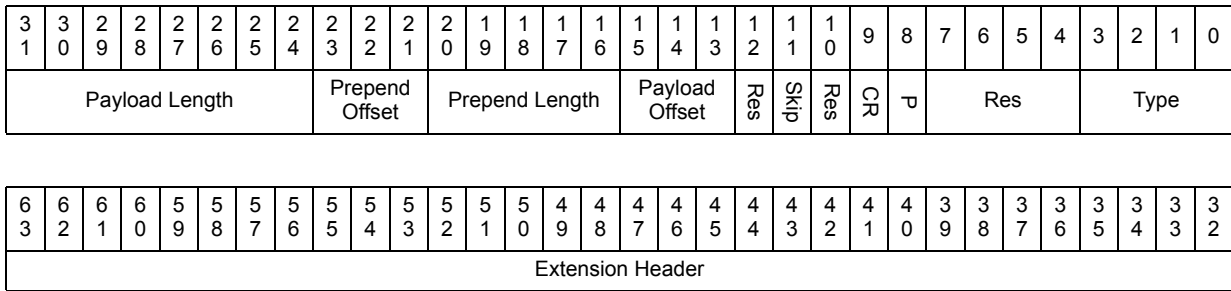
8.3.2.2 CSIX

For CSIX protocol, the TBUF should be set to two partitions in MSF_Tx_Control[TBUF_Partition], one for Data traffic and one for Control traffic.

Payload information is put into the Payload area of the element, and Base and Extension Header information is put into the Element Control Word.

Data is stored in big-endian order. The most significant byte of each 32-bit word is transmitted first.

When the Element Control Word is written the information is (note that the data comes from two consecutive Transfer registers; bits [31:0] from the lower numbered and bits[63:32] from the higher numbered):



The definitions of the fields are shown in Table 99.

Table 99. TBUF CSIX Control Definition

Field	Definition
Payload Length	Indicates the number of Payload bytes, from 1 to 256, in the element. The value of 0x00 means 256 bytes. The sum of Prepend Length and Payload Length will be sent, and also put into the CSIX Base Header Payload Length field. Note that this length does not include any padding that may be required. Padding is inserted by transmit hardware as needed.
Prepend Offset	Indicates the first valid byte of Prepend, from 0 to 7, as defined in Section 8.3.2.
Prepend Length	Indicates the number of bytes in Prepend, from 0 to 31.
Payload Offset	Indicates the first valid byte of Payload, from 0 to 7, as defined in Section 8.3.2.
Skip	Allows software to allocate a TBUF element and then not transmit any data from it. 0—transmit data according to other fields of Control Word 1—free the element without transmitting any data.
CR	CR (CSIX Reserved) bit to put into the CSIX Base Header.
P	P (Private) bit to put into the CSIX Base Header.
Type	Type Field to put into the CSIX Base Header. Idle type is <i>not</i> legal here.
Extension Header	The Extension Header to be sent with the CFrame. The bytes are sent in big-endian order; byte 0 is in bits 63:56, byte 1 is in bits 55:48, byte 2 is in bits 47:40, and byte 3 is in bits 39:32.



8.3.3 Transmit Operation Summary

During transmit processing data to be transmitted is placed into the TBUF under Microengine control, which allocates an element in software. The transmit hardware processes TBUF elements within a partition, in strict sequential order so the software can track the element to allocate next.

Microengines may write directly into an element by the `msf[write]` instruction, or have data from DRAM written into the element by the `dram[tbuf_wr]` instruction. Data can be merged into the element by doing both.

There is a Transmit Valid bits per element, which marks the element as ready to be transmitted. Microengines move all data into the element, by either or both of the `msf[write]` and `dram[tbuf_wr]` instructions to the TBUF. The Microengines also write the element Transmit Control Word with information about the element. The Microengines should use a single operation to perform the TCW write, i.e., a single `msf[write]` with a `ref_count` of 2. When all of the data movement is complete, the Microengine sets the element valid bit as shown in the following steps.

1. Move data into TBUF by either or both of `msf[write]` and `dram[tbuf_wr]` instructions to the TBUF.
2. Wait for 1 to complete.
3. Write Transmit Control Word at `TBUF_Element_Control_#` address. Using this address sets the Transmit Valid bit.

Note: When moving data from DRAM to TBUF using `dram[tbuf_wr]`, it is possible that there could be an uncorrectable error on the data read from DRAM (if ECC is enabled). In that case, the Microengine does not get an Event Signal, to prevent use of the corrupt data. The error is recorded in the DRAM controller (including the number of the Microengine that issued the TBUF Wr command — refer to the DRAM chapter for details), and will interrupt the Intel XScale® core, if enabled, so that it can take appropriate action. Such action is beyond the scope of this document. However, it must include recovering the TBUF element by setting it valid with the Skip bit set in the Control Word.

The transmit pipeline will be stalled since all TBUF elements must be transmitted in order; it will be un-stalled when the element is skipped.

8.3.3.1 SPI-4

Transmit control logic sends valid elements on the transmit pins in element order. First, a Control Word is sent — it is formed as shown in [Table 100](#). After the Control Word, the data is sent; the number of bytes to send is the total of Element Control Word Prepend Length field plus the Element Control Word Payload Length.

Table 100. Transmit SPI-4 Control Word

SPI-4 Control Word Field	Derived from:
Type	Type Bit of Element Control Word
EOPS	EOP Bit, Prepend Length, Payload Length of previous element's Element Control Word
SOP	SOP Bit of Element Control Word
ADR	ADR field of Element Control Word
DIP-4	Parity accumulated on previous element's data and this Control Word



If the next sequential element is not valid when its turn comes up:

1. Send an idle Control Word with SOP set to 0, and EOPS set to the values determined from the most recently sent element, ADR field 0x00, correct parity.
2. Until an element becomes valid, send idle Control Words with SOP set to 0, EOPS set to 00, ADR field 0x00, and correct parity.

Note: Sequential elements with same ADR are not “merged”, a Control Word is sent for each element.

Note: SPI-4 requires that all data transfers, except the last fragment (with EOP), be multiples of 16 bytes. It is up to the software loading the TBUF element to enforce this rule.

After an element has been sent on the transmit pins, the valid bit for that element is cleared. The **Tx_Sequence** register is incremented when the element has been transmitted; by also maintaining a sequence number of elements that have been allocated (in software), the microcode can determine how many elements are in-flight.

8.3.3.2 CSIX

Transmit control logic sends valid elements on the transmit pins in element order. Each element sends a single CFrame. First the Base Header is sent — it is formed as shown in [Table 101](#). Next, the Extension Header is sent. Finally, the data is sent; the number of bytes to send is the total of Element Control Word Prepend Length field plus the Element Control Word Payload Length, plus padding to fill the final CWord if required (the CWord Size is programmed in `MSF_Tx_Control[Tx_CWord_Size]`). Both Horizontal Parity and Vertical Parity are transmitted, as described in [Section 8.3.5.2.1](#) and [Section 8.3.5.2.2](#).

Note: When transmitting a Flow Control CFrame, the entire payload must be written into the TBUF entry. The extension header field of the Transmit Control Word is not used for Flow Control CFrames.

Table 101. Transmit CSIX Header

CSIX Header Field	Derived From
Type	Type field of Element Control Word
Data Ready	FC_Ingress_Status[TM_DReady]
Control Ready	FC_Ingress_Status[TM_CReady]
Payload Length	Element Control Word Prepend Length + Element Control Word Payload Length
P	P Bit of Element Control Word
CR	CR Bit of Element Control Word
Extension Header	Extension Header field of Element Control Word

Control elements and Data elements share use of the transmit pins. Each will alternately transmit a valid element, if present.

If the next sequential element is not valid when its turn comes up, or if transmission is disabled by `FC_Ingress_Status[SF_CReady]` or `FC_Ingress_Status[SF_DReady]`, then transmit logic will alternate sending Idle CFrames with Dead Cycles; it will continue to do so until a valid element is ready. Idle CFrames get the value for the Ready Field from `FC_Ingress_Status[TM_CReady]` and `FC_Ingress_Status[TM_DReady]`, the Payload Length is set to 0.

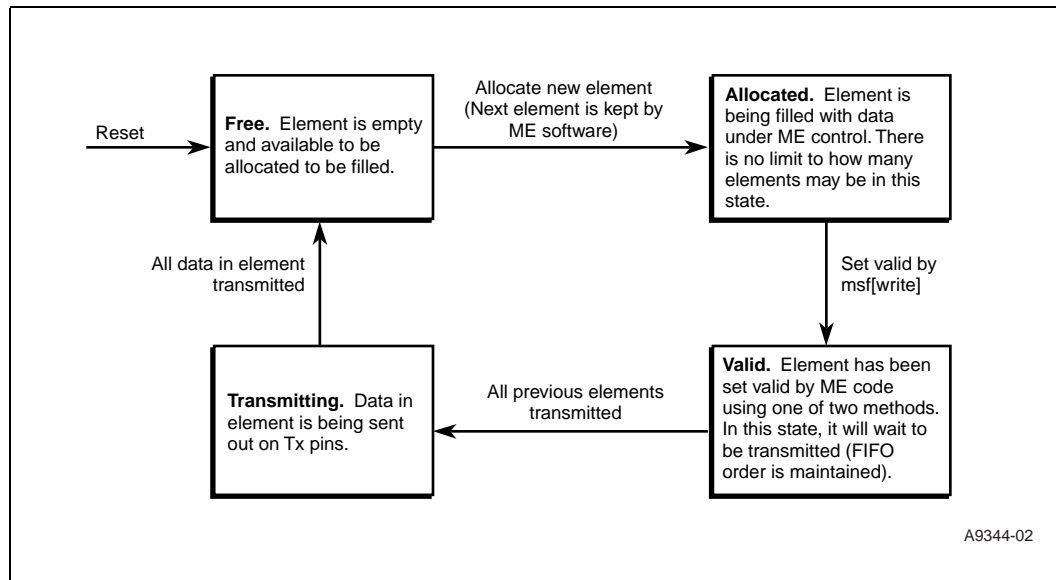
Note: A Dead Cycle is any cycle after the end of a CFrame, and prior to the start of another CFrame (i.e., SOF is not asserted). The end of a CFrame is defined as after the Vertical Parity has been transmitted. This in turn is found by counting the Payload Bytes specified in the Base Header and rounding up to CWord size.

After an element has been sent on the transmit pins, the valid bit for that element is cleared. The **Tx_Sequence** register is incremented when the element has been transmitted; by also maintaining a sequence number of elements that have been allocated (in software), the microcode can determine how many elements are in-flight.

8.3.3.3 Transmit Summary

The states that a TBUF element goes through (Free, Allocated, Transmitting, and Valid) are shown in Figure 95.

Figure 95. TBUF State Diagram



8.3.4 Transmit Flow Control Status

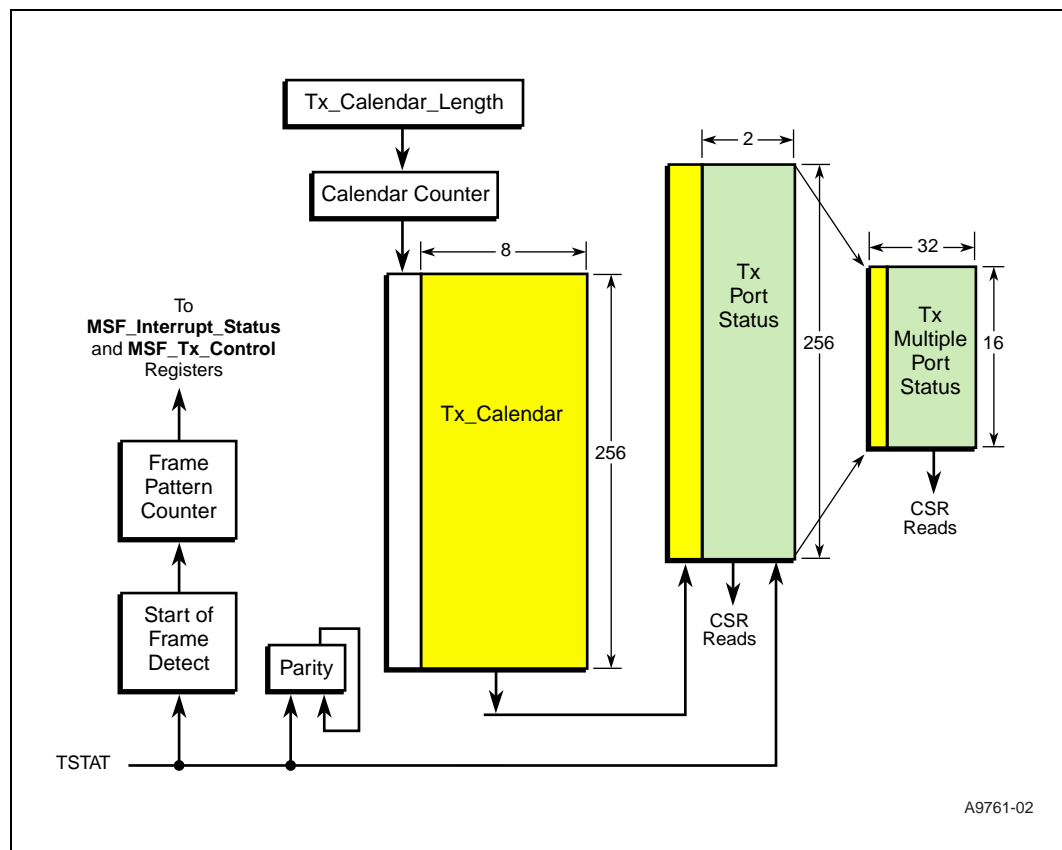
Transmit Flow Control is handled partly by hardware and partly by software. Information from the Egress IXP2800 Network Processor can be transmitted to the Ingress IXP2800 Network Processor (as described in Section 8.2.7 on Receive Flow Control); how it is used is described in the remainder of this section.

8.3.4.1 SPI-4

FIFO status information is sent periodically over the TSTAT signals from the PHY to the Link Layer device, which is the IXP2800 Network Processor. (The RXCDAT pins can act as TSTAT based on the **MSF_Tx_Control[TSTAT_Select]** bit.) The FIFO status of each port is encoded in a 2-bit data structure — code 0x3 is used for framing the data, and the other three codes are valid status values, which are interpreted by Microengine software.

The FIFO status words are received according to a repeating calendar sequence. Each sequence begins with the framing code to indicate the start of a sequence, followed by the status codes, followed by a DIP-2 parity code covering the preceding frame. The length of the calendar, as well as the port values, are defined in this section, and shown in [Figure 96](#).

Figure 96. Tx Calendar Block Diagram



Tx_Port_Status_# is a register file containing 256 registers, one for each of the SPI-4.2 ports. The port status is updated each time a new calendar status is received for each port, according to the mode programmed in **MSF_Tx_Control[Tx_Status_Update_Mode]**. The **Tx_Port_Status_#** register file holds the latest received status for each port, and can be read by CSR reads.

There are 16 **Tx_Multiple Port Status_#** registers. Each aggregates the status for each group of 16 ports. These registers provide an alternative method for reading the FIFO status of multiple ports with a single CSR read. For example, **Tx_Multiple Port Status_0** contains the 2-bit status for ports 0 – 16, and provides the same status as reading the individual registers, **Tx_Port_Status_0** through **Tx_Port_Status_15**.



The **TX_Port_Status_#** or the **TX_Multiple_Port_Status_#** registers must be read by the software to determine the status of each port and send data to them accordingly. The MSF hardware does not check these registers for port status before sending data out to a particular port.

The **MSF_Tx_Control[Tx_Status_Update_Mode]** field is used to select one of two methods for updating the port status. The first method updates the port status with the new status value, regardless of the value received. The second method updates the port status only when a value is received that is equal to or less than the current value.

Note: Detailed information about the status update modes is contained in the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

Reading a port status causes its value to be changed. This provides a way to avoid reading stale status bits. The **MSF_Tx_Control[Tx_Status_Read_Mode]** field is used to select the method for changing the bits after they are read.

Tx_Calendar is a RAM with 256 entries of eight bits each. It is initialized with the calendar information by software (the calendar is a list that indicates the sequence of port status that will be sent — the PHY and the IXP2800 Network Processor must be initialized with the same calendar). **Tx_Calendar_Length** is a CSR field that is initialized with the length of the calendar, since in many cases, not all 256 entries of **Tx_Calendar** are used.

When the start of a Status frame pattern is detected (by a value of 0x3 on TSTAT) the Calendar Counter is initialized to 0. On each data cycle, the Calendar Counter is used to index into **Tx_Calendar** to read a port number. The port number is used as an index to **Tx_Port_Status**, and the information received on TSTAT is put into that location in **Tx_Port_Status**. The count is incremented each cycle.

DIP-2 Parity is also accumulated on TSTAT. At the start of the frame, parity is cleared. When the count reaches **Tx_Calendar_Length**, the next value on TSTAT is used to compare to the accumulated parity. The control logic then looks for the next frame start. If the received parity does not match the expected value, the **MSF_Interrupt_Status[TSTAT_Par_Err]** bit is set, which can interrupt the Intel XScale® core if enabled.

Note: An internal status flag records whether or not the most recently received DIP-2 was correct. When that flag is set (indicating bad DIP-2 parity) all reads to **Tx_Port_Status** return a status of “Satisfied” instead of the value in the **Tx_Port_Status** RAM. The flag is re-loaded at the next parity sample; so the implication is that all ports will return “Satisfied” status for at least one calendar.

SPI-4 protocol uses a continuous stream of repeated frame patterns to indicate a disabled status link. The IXP2800 Network Processor flow control status block has a Frame Pattern Counter that counts up each time a frame pattern is received on TSTAT, and is cleared when any other pattern is received. When the Frame Pattern Counter reaches 32, **MSF_Interrupt_Status[Detect_No_Calendar]** is set and **Train_Data[Detect_No_Calendar]** is asserted (**MSF_Interrupt_Status[Detect_No_Calendar]** must be cleared by a write to the **MSF_Interrupt_Status** register; **Train_Data[Detect_No_Calendar]** will reflect the current status and will deassert when the frame pattern stops). The transmit logic will generate training sequence on transmit pins while both **Train_Data[Detect_No_Calendar]** and **Train_Data[Train_Enable_TSTAT]** are asserted.



8.3.4.2 CSIX

There are two types of CSIX flow control:

- Link-level
- Virtual Output Queue (VOQ)

8.3.4.2.1 Link-Level

The Link-level flow control function is done via hardware and consists of two parts:

1. Enable/disable transmission of valid TBUF elements.
2. Ready field to be sent in CFrames sent to the Switch Fabric.

As described in [Section 8.2.7](#), the Ready Field of received CFrames is placed into **FC_Egress_Status[SF_CReady]** and **FC_Egress_Status[SF_DReady]**. The value in those bits is sent to the Ingress IXP2800 Network Processor on TXCSR. In Full Duplex Mode, the information is received on RXCSR by the Ingress IXP2800 Network Processor and put into **FC_Ingress_Status[SF_CReady]** and **FC_Ingress_Status[SF_DReady]**. Those bits allow or stop transmission of Control and Data elements, respectively. When one of those bits transitions from allowing transmission to stopping transmission, the current CFrame in progress (if any) is completed, and the next CFrame of that type is prevented from starting.

As described in [Section 8.2.7](#), if the Egress IXP2800 Network Processor RBUF gets near full, or if the Egress IXP2800 Network Processor FCEFIFO gets near full, it will send that information on TXCSR. Those bits are put into **FC_Ingress_Status[TM_CReady]** and **FC_Ingress_Status[TM_DReady]**, and are used as the value in CFrame Base Header Control Ready and Data Ready, respectively.

8.3.4.2.2 Virtual Output Queue

The Virtual Output Queue flow control function is done by software, with hardware support.

As described in [Section 8.2.7](#), the CSIX Flow Control CFrames received on the Egress IXP2800 Network Processor are passed to the Ingress IXP2800 Network Processor over TXCDAT. The information is received on RXCDAT and placed into the FCIFIFO. A Microengine reads that information by `msf[read]`, and uses it to maintain per-VOQ information. The way in which that information is used is application-dependent and is done in software. The hardware mechanism is described in [Section 8.5.3](#).

8.3.5 Parity

8.3.5.1 SPI-4

DIP-4 parity is computed by Transmit hardware placed into the Control Word sent at the beginning of transmission of a TBUF element, and also on Idle Control Words sent when no TBUF element is valid. The value to place into the DIP-4 field is computed on the preceding Data Words (if any), and the current Control Word.



8.3.5.2 CSIX

8.3.5.2.1 Horizontal Parity

The transmit logic computes odd Horizontal Parity for each transmitted 16-bits of each Cword, and transmits it on TxPar.

8.3.5.2.2 Vertical Parity

The transmit logic computes Vertical Parity on CFrames. There is a 16-bit VP Accumulator register. At the beginning of each CFrame, the register is cleared. As each Cword is transmitted, odd parity is accumulated in the register as defined in the CSIX specification (16 bits of vertical parity are formed on 32 bits of transmitted data by treating the data as words; i.e., bit 0 and bit 16 of the data are accumulated into parity bit 0, bit 1, and bit 17 of the data are accumulated into parity bit 1, etc.). The accumulated value is transmitted in the Cword along with the last byte of Payload and any padding, if required.

8.4 RBUF and TBUF Summary

Table 102 summarizes and contrasts the RBUF and TBUF operations.

Table 102. Summary of RBUF and TBUF Operations (Sheet 1 of 2)

Operation	RBUF	TBUF
Allocate element	<p>SPI-4</p> <p>Hardware allocates an element upon receipt of a non-idle Control Word, or when a previous element becomes full and another Data Word arrives with no intervening Control Word. Any available element in the SPI-4 partition may be allocated, however, elements are guaranteed to be handed to threads in the order in which they arrive.</p> <p>CSIX</p> <p>Hardware allocates an element upon receipt of RxSof asserted. Any available element in the CSIX Control or CSIX Data partition may be allocated (according to the type), however, elements are guaranteed to be handed to threads in the order in which they arrive.</p>	<p>Microengine allocates an element. Because the elements are transmitted in FIFO order (within each TBUF partition), the Microengine can keep the number of the next element in software.</p>
Fill element	<p>SPI-4</p> <p>Hardware fills the element with Data Words.</p> <p>CSIX</p> <p>Hardware fills the element with Payload.</p>	<p>Microcode fills the element from DRAM using the <code>dram[tbuf_wr]</code> instruction and from Microengine registers using <code>msf[write]</code> instruction.</p>
Set element valid	<p>SPI-4</p> <p>Set valid by hardware when either it becomes full or when a Control Word is received.</p> <p>CSIX</p> <p>Set valid by hardware when the number of bytes in Payload Length have been received.</p>	<p>The element's Transmit Valid bit is set. This is done by a write to the TBUF_Element_Control_\$_# CSR (\$is A or B, # is the element number).</p>

Table 102. Summary of RBUF and TBUF Operations (Sheet 2 of 2)

Operation	RBUF	TBUF
Remove data from element	Microcode moves data from the element to DRAM using the <code>dram[rbuf_rd]</code> instruction and to Microengine registers using the <code>msf[read]</code> instruction.	Hardware transmits information from the element to the Tx pins. Transmission of elements is in FIFO order within each partition; that is an element will be transmitted only when all preceding elements in that partition have been transmitted. Choice of element to transmit among partitions is round-robin.
Return element to Free List	Microcode writes to <code>Rx_Element_Done</code> with the number of the element to free.	Microengine software uses the <code>TX_Sequence_n</code> CSRs to track elements that have been transmitted.

8.5 CSIX Flow Control Interface

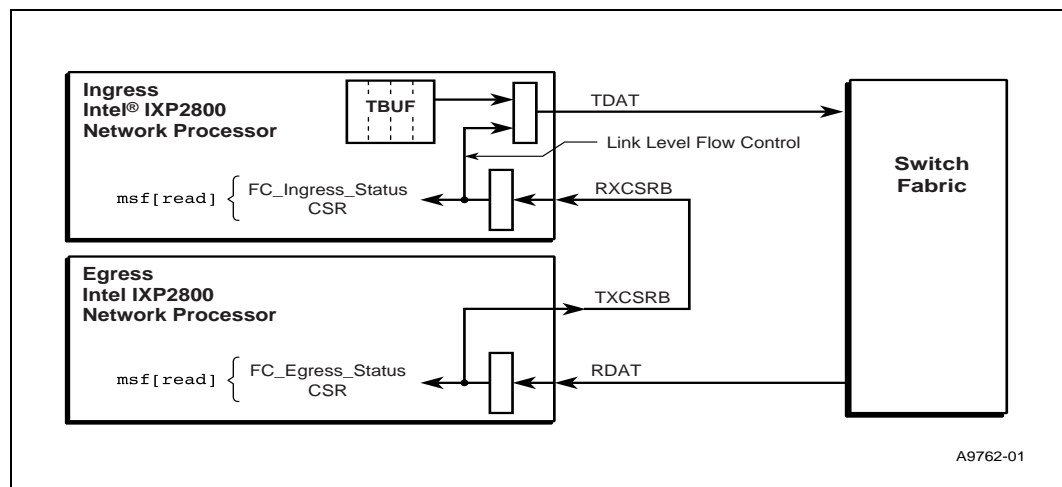
This section describes the Flow Control Interface. Section 8.2 and Section 8.3 of this chapter also contain descriptions of how those functions interact with Flow Control. There are two modes — Full Duplex, where flow control information goes from Egress IXP2800 Network Processor to the Ingress IXP2800 Network Processor, and Simplex mode, where the information from the Switch Fabric is sent directly to the Ingress IXP2800 Network Processor, and from the Egress IXP2800 Network Processor to the Switch Fabric.

8.5.1 TXCSRB and RXCSRB Signals

TXCSRB and RXCSRB are used only in Full Duplex mode. (See Figure 97.) They send information from the Egress to the Ingress IXP2800 Network Processor for two reasons:

1. Pass the CSIX Ready Field (link-level flow control) from the Switch Fabric to the Ingress IXP2800. The information is used by the Ingress IXP2800's transmit control logic to stop transmission of CFrames to the Switch Fabric.
2. Set the value of the Ready field sent from the Ingress IXP2800 to the Switch Fabric. This is to inform the Switch Fabric to stop transmitting CFrames to the Egress IXP2800, based on receive buffer resource availability in the Egress IXP2800.

Figure 97. CSIX Flow Control Interface — TXCSRB and RXCSRB





The information transmitted on TXCSRБ can be read in **FC_Egress_Status** CSR, and the information received on RXCSRБ can be read in **FC_Ingress_Status** CSR.

The TXCSRБ or RXCSRБ signals carry the Ready information in a serial stream. Four bits of data are carried in 10 clock phases, LSB first, as shown in [Table 103](#).

Table 103. SRB Definition by Clock Phase Number

Clock Cycle Number	Description	
	Source of bit on Egress IXP2800 Network Processor (TXCSRБ)	Use of bit on Ingress IXP2800 Network Processor (RXCSRБ)
0–5	Framing information. Data is 000001; this pattern allows the Ingress IXP2800 Network Processor to get synchronized to the serial stream regardless of the data values.	
6	Most recently received Control Ready from a CFrame Base Header. Also visible in FC_Egress_Status[SF_CReady] .	When 0—Stop sending Control CFrames to the Switch Fabric. When 1—OK to send Control CFrames to the Switch Fabric. Also visible in FC_Ingress_Status[SF_CReady] .
7	Most recently received Data Ready from a CFrame Base Header. Also visible in FC_Egress_Status[SF_DReady]	When 0—Stop sending Data CFrames to the Switch Fabric. When 1—OK to send Data CFrames to the Switch Fabric. Also visible in FC_Ingress_Status[SF_DReady] .
8	RBUF or FCEFIFO are above high water mark. Also visible in FC_Egress_Status[TM_CReady] .	Place this bit in the Control Ready bit of all outgoing CSIX Base Headers. Also visible in FC_Ingress_Status[TM_CReady] .
9	RBUF is above high water mark. Also visible in FC_Egress_Status[TM_DReady] .	Place this bit in the Data Ready bit of all outgoing CSIX Base Headers. Also visible in FC_Ingress_Status[TM_DReady] .

The Transmit Data Ready bit sent from Egress to Ingress IXP2800 Network Processor will be deasserted if the following condition is met.

- RBUF CSIX Data partition is full, based on **HWM_Control[RBUF_D_HWM]**.

The Transmit Control Ready bit sent from the Egress to the Ingress IXP2800 Network Processor will be deasserted if either of the following conditions is met.

- RBUF CSIX Control partition is full, based on **HWM_Control[RBUF_C_HWM]**.
- FCEFIFO full, based on **HWM_Control[FCEFIFO_HWM]**.

8.5.2 FCIFIFO and FCEFIFO Buffers

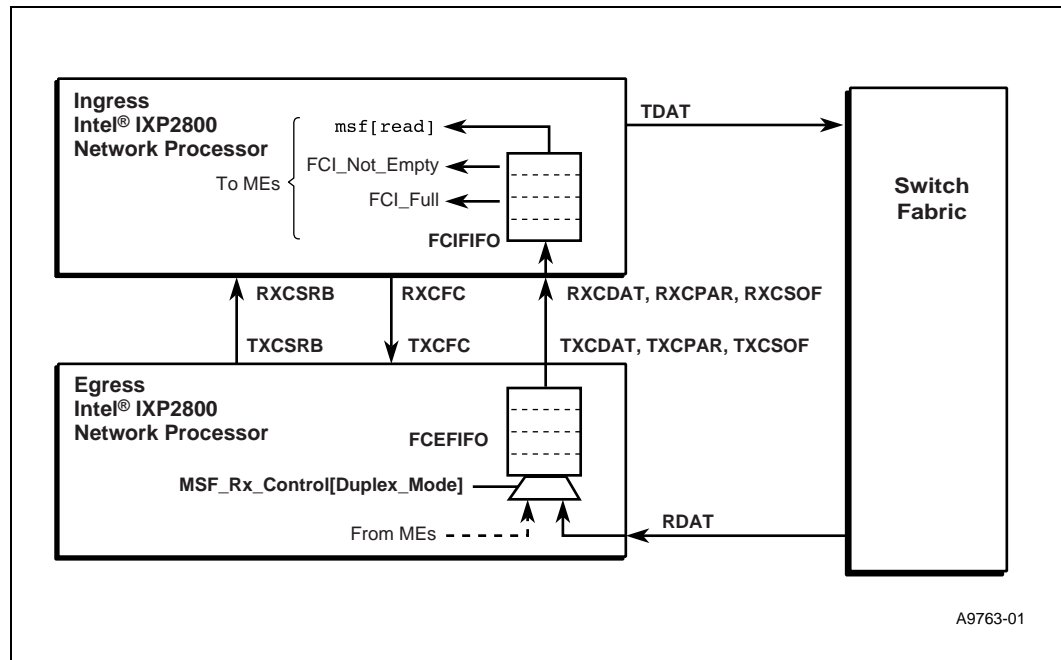
FCIFIFO and FCEFIFO are 1-Kbyte (256 entry x 32-bit) buffers for the flow control information. FCEFIFO holds data while it is being transmitted off of the Egress IXP2800 Network Processor. FCIFIFO holds data received into the Ingress IXP2800 Network Processor until Microengines can read it. There are two usage models for the FIFOs — Full Duplex Mode and Simplex Mode, selected by **MSF_Rx_Control[Duplex_Mode]**.

8.5.2.1 Full Duplex CSIX

In Full Duplex Mode, the information from the Switch Fabric is sent to the Egress IXP2800 Network Processor and must be communicated to the Ingress IXP2800 Network Processor via TXCSRB or RXCSRB. CSIX CFrames received from the Switch Fabric on the Egress IXP2800 Network Processor are put into FCEFIFO, based on the mapping in the **CSIX_Type_Map** CSR (normally they will be the Flow Control CFrames). The entire CFrame is put in, including the Base Header and Vertical Parity field.

The CFrames are forwarded in a “cut-through” manner, meaning the Egress IXP2800 Network Processor does not wait for the entire CFrame to be received before forwarding. The Egress processor will corrupt the Vertical Parity of the CFrame being forwarded if either a Horizontal or Vertical Parity is detected during receive, to inform the Ingress processor that an error occurred. The Ingress IXP2800 Network Processor checks both Horizontal Parity and Vertical Parity and will discard the entire CFrame if bad parity is detected. The signal protocol details of how the information is sent from the Egress IXP2800 Network Processor to the Ingress IXP2800 Network Processor is described in [Section 8.5.3](#). (See [Figure 98](#).)

Figure 98. CSIX Flow Control Interface — FCIFIFO and FCEFIFO in Full Duplex Mode



The Ingress IXP2800 Network Processor puts the CFrames into the FCIFIFO, including the Base Header and Vertical Parity fields. It does not make a CFrame visible in the FCIFIFO until the entire CFrame has been received without errors. If there is an error, the entire CFrame is discarded and **MSF_Interrupt_Status[FCIFIFO_Error]** is set.

CFrames in the FCIFIFO of the Ingress IXP2800 Network Processor are read by Microengines, which use them to keep current VOQ Flow Control information. (The application software determines how and where that information is stored and used.)

The FCIFIFO supplies two signals to Microengines, which can be tested using the BR_STATE instruction:

1. FCI_Not_Empty — indicates that there is at least one CWord in the FCIFIFO. This signal stays asserted until all CWords have been read. (Note that when FCIFIFO is empty, this signal will not assert until a full CFrame has been received into FCIFIFO; as that CFrame is removed by the Microengine, this signal will stay asserted until all CWords have been removed, including any subsequently received CFrames.)
2. FCI_Full — indicates that FCIFIFO is above the upper limit defined in **HWM_Control[FCIFIFO_Int_HWM]**.

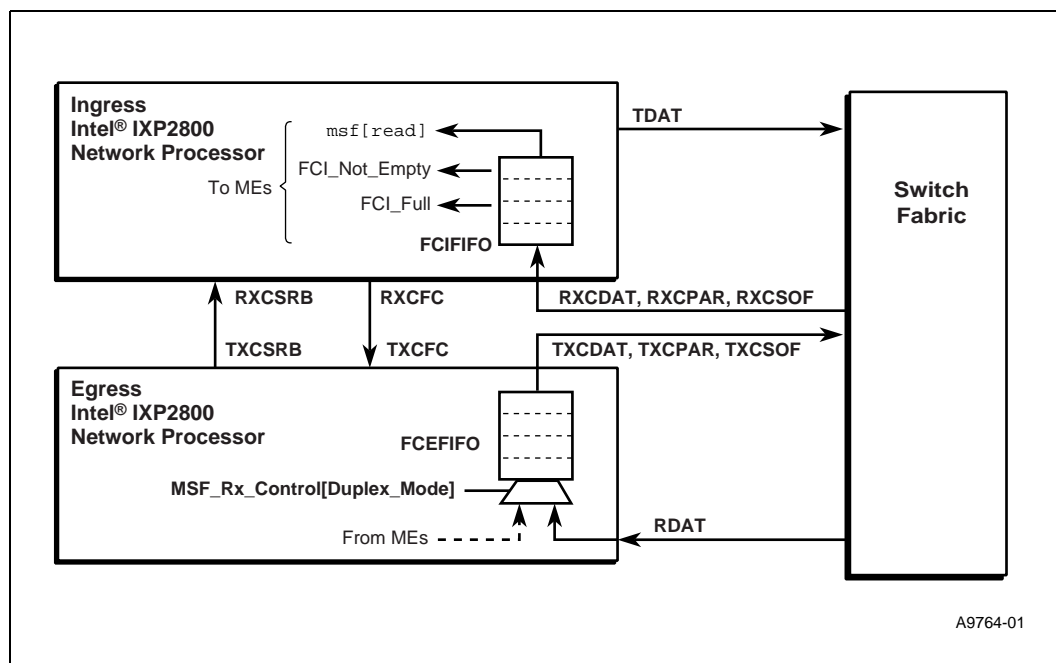
The Microengine that has been assigned to handle FCIFIFO must read the CFrame, 32 bits at a time, from the FCIFIFO by using the `msf[read]` instruction to the FCIFIFO address; the length of the read can be anywhere from 1 to 16. The FCIFIFO handler thread must examine the Base Header to determine how long the CFrame is and perform the necessary number of reads from the FCIFIFO to dequeue the entire CFrame. If a read is issued to FCIFIFO when it is empty, an Idle CFrame will be read back (0x0000FFFF). Note that when FCIFIFO is receiving a CFrame, it does not make it visible until the entire CFrame has been received without errors.

The nearly-full signal is based on the upper limit programmed into **HWM_Control[FCIFIFO_Int_HWM]**. When asserted, this means that higher priority needs to be given to draining the FCIFIFO to prevent flow control from being asserted to the Egress IXP2800 Network Processor (by assertion of RXCFC).

8.5.2.2 Simplex CSIX

In Simplex Mode, the Flow Control signals are connected directly to the Switch Fabric; flow control information is sent directly from the Egress IXP2800 Network Processor to the Switch Fabric, and directly from the Switch Fabric to the Ingress IXP2800 Network Processor. (See Figure 99.)

Figure 99. CSIX Flow Control Interface — FCIFIFO and FCEFIFO in Simplex Mode



The TXCSR and RXCSR pins are not used in Simplex Mode. The RXCF and TXCF pins are used for flow control in both Simplex and Duplex Modes. The Egress IXP2800 Network Processor uses the TXCSOF, TXCDAT, and TXCPAR pins to send CFrames to the Switch Fabric.

The Ingress IXP2800 Network Processor uses the RXCSOF, RXCDAT, and RXCPAR pins to receive CFrames from the Switch Fabric (the Switch Fabric is expected to send Flow Control CFrames on these pins instead of the RDAT pins in Simplex Mode). The **FC_Ingress_Status[SF_CReady]** and **FC_Ingress_Status[SF_DReady]** bits are set are from the “Ready bits” received in all incoming CFrames received on this interface. Transmit hardware in the Ingress IXP2800 Network Processor uses the **FC_Ingress_Status[SF_CReady]** and **FC_Ingress_Status[SF_DReady]** bits to flow control the data and control transmit on TDAT.

CFrames in the FCIFIFO of the Ingress IXP2800 Network Processor are read by Microengines, which use them to keep current VOQ Flow Control information (this is the same as for Full Duplex Mode). The FCI_Not_Empty and FCI_Full status flags, as described in [Section 8.5.2.1](#) let the Microengine know if the FCIFIFO has any CWords in it. When FCI_Full is asserted, **FC_Ingress_Status[TM_CReady]** will be deasserted; that bit is put into the Ready field of CFrames going to the Switch Fabric, to inform it to stop sending Control CFrames.

Flow Control CFrames to the Switch Fabric are put into FCEFIFO, instead of TBUF, as in the Full Duplex Mode case. In this mode, the Microengines create CFrames and write them into FCEFIFO using the `msf[write]` instruction to the FCEFIFO address; the length of the write can be from 1 – 16. The Microengine creating the CFrame must put a header (conforming to CSIX Base Header format) in front of the message, indicating to the hardware how many bytes to send.

The Microengine first tests if there is room in FCEFIFO by reading the **FC_Egress_Status[FCEFIFO_Full]** status bit. After the CFrame has been written to FCEFIFO, the Microengine writes to the **FCEFIFO_Validate** register, indicating that the CFrame should be sent out on TXCDAT; this prevents underflow by ensuring that the entire CFrame is in FCEFIFO before it can be transmitted. A validated CFrame at the head of FCEFIFO is started on TXCDAT if **FC_Egress_Status[SF_CReady]** is asserted, and held off, if it is deasserted. However, once started, the entire CFrame is sent, regardless of changes in **FC_Egress_Status[SF_CReady]**. The **FC_Egress_Status[SF_DReady]** is ignored in controlling FCEFIFO.

FC_Egress_Status[TM_CReady] and **FC_Egress_Status[TM_DReady]** are placed by hardware into the Base Header of outgoing CFrames. Horizontal and Vertical parity are created by hardware.

If there is no valid CFrame in FCEFIFO, or if **FC_Egress_Status[SF_CReady]** is deasserted, then idle CFrames are sent on TXCDAT. The idle CFrames also carry (in the Base Header Ready Field), both **FC_Egress_Status[TM_CReady]** and **FC_Egress_Status[TM_DReady]**. In all cases, the Switch Fabric must honor the “ready bits” to prevent overflowing RBUF.

Note: For simplex mode, there is a condition in which the Flow Control Bus may take too long to properly control incoming traffic on CSIX. This condition may occur when large packets are transmitted on the Flow Control Bus and small packets are transmitted on CSIX. For example, this condition may occur if the Switch Fabric’s CSIX Receive FIFO is full, and the FIFO wants to deassert the **x_RDY** bit, but a maximum-sized flow control CFrame just went out. The Flow Control Bus is a 4-bit wide LVDS interface that sends data on both the rising and falling edges of the clock. As such, it takes 260 clock cycles to transmit a maximum-sized CFrame, which consists of 256 bytes, plus a 4-byte base header/vertical parity (i.e, 260 bytes total). The interface does not see the transition of the **X_RDY** bit until this CFrame has been transmitted or until 260 cycles later.



8.5.3 TXCDAT/RXCDAT, TXCSOF/RXCSOF, TXCPAR/RXCPAR, and TXCFC/RXCFC Signals

TXCDAT and RXCDAT, along with TXCSOF/RXCSOF and TXCPAR/RXCPAR are used to send CSIX Flow Control information from the Egress IXP2800 Network Processor to the Ingress IXP2800 Network Processor.

The protocol is basically the same as CSIX-LI, but with only four data signals.

TXCSOF is asserted to indicate start of a new CFrame. The format is the same as any normal CFrame — Base Header, followed by Payload and Vertical Parity; the only difference is that each CWord is sent on TXCDAT in four cycles, with the most significant bits first. TXCPAR carries odd parity for each four bits of data. The transmit logic also creates valid Vertical Parity at the end of the CFrame, with one exception. If the Egress IXP2800 Network Processor detected an error on the CFrame, it will create bad Vertical parity so that the Ingress IXP2800 Network Processor will detect that and discard it.

The Egress IXP2800 Network Processor sends CFrames from FCEFIFO in cut-through manner. If there is no data in FCEFIFO, then the Egress IXP2800 Network Processor alternates sending Idle CFrames and Dead Cycles. (Note that FCIFIFO never enqueues Idle CFrames in either Full Duplex or Simplex Modes. The transmitted Idle CFrames are injected by the control state machine, not taken from the FCEFIFO.)

The Ingress IXP2800 Network Processor asserts RXCFC to indicate that FCIFIFO is full, as defined by **HWM_Control[FCIFIFO_Ext_HWM]**. The Egress IXP2800 Network Processor, upon receiving that signal asserted, will complete the current CFrame, and then transmit Idle CFrames until RXCFC deasserts. During that time, the Egress IXP2800 Network Processor can continue to buffer Flow Control CFrames in FCEFIFO; however, if that fills, the further CFrames mapped to FCEFIFO will be discarded.

Note: If there is no Switch Fabric present, this port could be used for interchip message communication. FC pins must connect between network processors as in Full Duplex Mode. Set **MSF_RX_CONTROL[DUPLICATION_MODE] = 0** and **MSF_TX_CONTROL[DUPLICATION_MODE] = 0** (Simplex) and **FC_STATUS_OVERRIDE=0x3ff**. Microengines write CFrames to the FCEFIFO CSR as in Simplex Mode. The RXCFC and TXCFC pins must be connected between network processors to provide flow control.

8.6 Deskew and Training

There are three methods of operation that can be used, based on the application requirements.

1. Static Alignment — the receiver latches all data and control signals at a fixed point in time, relative to clock.
2. Static Deskew — the receiver latches each data and control signal at a programmable point in time, relative to clock. The programming value for each signal is characterized for a given system design and loaded into deskew control registers at system boot time.
3. Dynamic Deskew — the transmitter periodically sends a training pattern that the receiver uses to automatically select the optimal timing point for each data and control signal. The timing values are loaded into the deskew control registers by the training hardware.



The IXP2800 Network Processor supports all three methods. There are three groups of high-speed pins to which this applies, as shown in Table 104, Table 105, and Table 106. The groups are defined by the clock signal that is used.

Table 104. Data Deskew Functions

Clock	Signals	IXP2800 Network Processor Operation
RCLK	RDAT	1. Sample point for each pin is programmed in Rx_Deskew . 2. Deskew values set automatically when training pattern (Section 8.6.1) is received and is enabled in Train_Data[Ignore_Training] .
	RCTL	
	RPAR	
	RPROT	
TCLK	TDAT	1. Send training pattern <ul style="list-style-type: none"> under software control (write to Train_Data[Continuous_Train] or Train_Data[Single_Train]) when TSTAT input has framing pattern for more than 32 cycles and enabled in Train_Data[Train_Enable].
	TCTL	
	TPAR	
	TPROT	

Table 105. Calendar Deskew Functions

Clock	Signals	IXP2800 Network Processor Operation
RSCLK	RSTAT	1. Used to indicate need for data training on receive pins by forcing to continual framing pattern (write to Train_Data[RSTAT_En]). 2. Send training pattern under software control (write to Train_Calendar[Continuous_Train] or Train_Calendar[Single_Train]).
TSCLK	TSTAT	1. Sample point for each pin is set in Rx_Deskew , either by manual programming or automatically. 2. Deskew values set automatically when training pattern (Section) is received and is enabled in Train_Calendar[Ignore_Training] . 3. Received continuous framing pattern can be used to initiate data training (Train_Data[Detect_No_Calendar]), and/or interrupt the Intel XScale® core.

Table 106. Flow Control Deskew Functions

Clock	Signals	IXP2800 Network Processor Operation
RXCCLK	RXCFSOF	1. Sample point for each pin is programmed in Rx_Deskew . 2. Deskew values set automatically when training pattern (Section 8.6.2) is received and is enabled in Train_Flow_Control[Ignore_Training] . Note 1, 2
	RXCDFAT	
	RXCPCAR	
	RXCSCRB	
TXCCLK	TXCFSOF	1. Send training pattern <ul style="list-style-type: none"> under software control (write to Train_Flow_Control[Continuous_Train] or Train_Flow_Control[Single_Train]) when TXCFC input has been asserted for more than 32 cycles and enabled in Train_Flow_Control[Train_Enable]. Notes 1, 2
	TXCDFAT	
	TXCPCAR	
	TXCSCRB	
NOTES: 1. TXCFC is not trained. RXCFC is driven out relative to RXCCLK; TXCFC is received relative to TXCCLK, but is treated as asynchronous. 2. RXCFC can be forced asserted by write to Train_Flow_Control[RXCFC_En] .		



8.6.1 Data Training Pattern

The data pin training sequence is shown in Table 107. This is a superset of SPI-4 training sequence, because it includes the TPAR/RPAR and TPROT/RPOT pins, which are not included in SPI-4.

Table 107. Data Training Sequence

Cycle (Note 4)	PROT	PAR	CTL	DATA																	
				15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
1 (Note 5)	0	x	1	0	x	x	0	0	0	0	0	0	0	0	0	0	0	a	b	c	d
2 to 11	0	1	1	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
12 to 21	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
20 α -18 to 20 α -9	0	1	1	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
20 α -8 to 20 α +1	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0

NOTES:

- In cycle 1, x and abcd depend on the contents of the interval after the last preceding control word. This is an Idle Control Word.
- α represents the number of repeats, as specified in SPI-4 specification. When the IXP2800 Network Processor is transmitting training sequences the value is in **Train_Data[Alpha]**.
- On receive, the IXP2800 Network Processor will do dynamic deskew when **Train_Data[Ignore_Training]** is 0, and RCTL = 1 and RDATA = 0x0FFF for three consecutive samples. Note that RPROT and RPAR are ignored when recognizing the start of training sequence.
- These are really phases (i.e., each edge of the clock is counted as one sample).
- This cycle is valid for SPI-4, it is not used in CSIX training.

8.6.2 Flow Control Training Pattern

This section defines training for the flow control pins (Table 108). These pins are normally used for CSIX flow control (Section 8.5), but can be programmed for use as SPI-4 Status Channel. The training pattern used is based on the usage.

The flow control pin training sequence when the pins are used for CSIX flow control is shown in Table 108.

Table 108. Flow Control Training Sequence

Cycle (Note 3)	XCOSF	XCDAT				XCPAR	XC SRB
		3	2	1	0		
1 to 10	1	1	1	0	0	0	0
11 to 20	0	0	0	1	1	1	1
20 α -19 to 20 α -10	1	1	1	0	0	0	0
20 α -9 to 20 α	0	0	0	1	1	1	1

NOTE:

- α represents the number of repeats, as specified in SPI-4 specification. When the IXP2800 Network Processor is transmitting training sequences the value is in **Train_Flow_Control[Alpha]**.
- On receive, the IXP2800 Network Processor will do dynamic deskew when **Train_Flow_Control[Ignore_Training]** is 0, and RXCOSF = 1, RXCDATA = 0xC, RXCPAR = 0, and RXCSR B = 0 for three consecutive samples.
- These are really phases (i.e., each edge of the clock is counted as one sample).



The training sequence when the pins are used for SPI-4 Status Channel is shown in [Table 109](#). This is compatible to SPI-4 training sequence.

Table 109. Calendar Training Sequence

Cycle (Note 3)	XCDAT	
	1	0
1 to 10	0	0
11 to 20	1	1
20 α -19 to 20 α -10	0	0
20 α -9 to 20 α	1	1

NOTE:

1. α represents the number of repeats, as specified in SPI-4 specification. When the IXP2800 Network Processor is transmitting training sequences the value is in **Train_Calendar[Alpha]**.
2. On receive, the IXP2800 Network Processor will do dynamic deskew when **Train_Calendar[Ignore_Training]** is 0, and **TCDAT**= 0x0 for ten consecutive samples.
3. These are really phases (i.e., each edge of the clock is counted as one sample).
4. Only XCDAT[1:0] are included in training.

8.6.3 Use of Dynamic Training

Dynamic training is done by cooperation of hardware and software as defined in this section.

The IXP2800 Network Processor will need training at reset or it loses training. Loss of training will typically be detected by parity errors on received data. [Table 110](#) lists the steps to initiate the training. SPI-4, CSIX Full Duplex, and CSIX Simplex cases follow similar but slightly different sequences. The SPI-4 protocol uses the calendar status pins, TSTAT/RSTAT (or RXCDAT/TXCDAT if those are used for calendar status), as an indicator that data training is required. For CSIX use, the IXP2800 Network Processor uses a proprietary method of in-band signaling using Idle CFrames and Dead Cycles to indicate the need for training.

Until the LVDS IOs are deskewed correctly, DIP-4 errors will occur. At startup, the receiver should request training followed by the transmitting device being sent training. The receiver should initially see received_training set and DIP-4 parity errors. The receiver should then clear the parity errors, wait for receive_training set and dip4_error cleared and check that all of the applicable RX_PHASEMON registers indicate no training errors. Then the LVDS IOs are properly trained.



Table 110. IXP2800 Network Processor Requires Data Training

Step	SPI-4 (IXP2800 Network Processor is Ingress Device)	CSIX (IXP2800 Network Processor is Egress Device)	
		Full Duplex	Simplex
1	Detect need for training (for example, reset or excessive parity errors).		
2	Force RSTAT (when using LVTTTL status channel) to continuous framing pattern (Write a 0 to Train_Data[RSTAT_En]), or force RXCDAT (when using LVDS status channel) to continuous training (Write a 1 to Train_Calendar [Continuous_Train]).	Force Transmission of Idle CFrames on Flow Control (Write a 1 to Train_Flow_Control [Force_FCIdle]).	Force Transmission of Dead Cycles on Flow Control (Write a 1 to Train_Flow_Control [Force_FCDead]).
3	Framer device detects RSTAT in continuous framing (when using LVTTTL status channel, or RXCDAT in continuous training (when using LVDS status channel).	Ingress IXP2800 Flow Control port detects Idle CFrames and sets Train_Flow_Control [Detect_FCIdle] .	Switch Fabric detects Dead Cycles on Flow Control.
4	Framer device transmits Training Sequence (IXP2800 receives on RDAT).	Ingress IXP2800 sends Dead Cycles on TDAT (if Train_Data [Dead_Enable_FCIdle] is set).	
5		Switch Fabric detects Dead Cycles on Data.	
6		Switch Fabric transmits Training Sequence on Data.	
7	When MSF_Interrupt_Status[Received_Training_Data] interrupt indicates training happened, and all of the applicable RX_PHASEMON registers indicate no training errors. Write MSF_Interrupt_Status[DIP4_ERR] to clear previous errors.		
	Write a 1 to Train_Data[RSTAT_En] or Write a 0 to Train_Calendar [Continuous_Train] .	Write a 0 to Train_Flow_Control [Force_FCIdle] .	Write a 0 to Train_Flow_Control [Force_FCDead] .

The second case is when the Switch Fabric or SPI-4 framing device indicates it needs Data training. [Table 111](#) lists that sequence.



Table 111. Switch Fabric or SPI-4 Framer Requires Data Training

Step	SPI-4	CSIX	
		Full Duplex	Simplex
1	Framer sends continuous framing code on IXP2800 calendar status pins TSTAT (when using LVTTTL status channel) or sends continuous training on IXP2800 calendar status pins RXCDAT (when using LVDS status channel).	Switch Fabric sends continuous Dead Cycles on Data.	Switch Fabric sends continuous Dead Cycles on Flow Control.
2	IXP2800 detects no calendar on TSTAT (when using LVTTTL status channel) or detects continuous training on RXCDAT (when using LVDS status channel), and sets Train_Data [Detect_No_Calendar] .	Egress IXP2800 detects Dead Cycles and sets Train_Data [Detect_CDead] .	Ingress IXP2800 detects Dead Cycles and sets Train_Flow_Control [Detect_FCDead] .
3	IXP2800 transmits Training Pattern (if Train_Data [Train_Enable_TDAT] is set).	Egress IXP2800 Flow Control port sends continuous Dead Cycles if Train_Flow_Control [TD_Enable_CDead] .	
4		Ingress IXP2800 Flow Control port detects continuous Dead Cycles and set Train_Flow_Control [Detect_FCDead] .	
5		Ingress IXP2800 transmits continuous Training Sequence on data if Train_Data [Train_EN_FCDead] .	
6	When Framer/Switch Fabric is trained it indicates that fact by reverting to normal operation.		
	Framer stops continuous framing code on calendar status pins.	Switch Fabric stops continuous Dead Cycles on Data.	Switch Fabric stops continuous Dead Cycles on Flow Control.

The IXP2800 Network Processor needs training at reset, or whenever it loses training. Loss of training is typically detected by parity errors on received flow control information.



Table 112 lists the steps to initiate the training. CSIX Full Duplex and CSIX Simplex cases follow similar, but slightly different sequences.

Table 112. IXP2800 Network Processor Requires Flow Control Training

Step	CSIX (IXP2800 Network Processor is Ingress Device)	
	Full Duplex	Simplex
1	Force TXCFC pin asserted (Write a 0 to Train_Flow_Control [RXCFC_En]).	Force Data pins to continuous Dead Cycles (Write a 1 to Train_Data[Force_CDead]).
2	Egress IXP2800 Network Processor Flow Control port detects RXCFC sustained assertion and sets Train_Flow_Control [Detect_TXCFC_Sustained] .	Switch Fabric detects Dead Cycles on Data.
3	Ingress IXP2800 Network Processor transmits Training Sequence on Flow Control pins (if Train_Flow_Control [Train_Enable_CFC] is set).	Switch Fabric transmits Training Sequence on Flow Control pins.
4	When MSF Interrupt_Status[Received_Training_FC] interrupt indicates training happened and all of the applicable RX_PHASEMON registers indicate no training errors, write CSR bits set in Step 1 to inactive value.	
	Write a 1 to Train_Flow_Control [RXCFC_En] .	Write a 1 to Train_Data[Force_CDead] .

The last case is when the Switch Fabric indicates it needs Flow Control training. Table 113 lists that sequence.

Table 113. Switch Fabric Requires Flow Control Training

Step	Simplex (IXP2800 Network Processor is Egress Device)
1	Switch Fabric sends continuous Dead Cycles on Data.
2	Egress IXP2800 Network Processor detects Dead Cycles and sets Train_Data [Detect_CDead] .
3	Egress IXP2800 Network Processor transmits Training Sequence on Flow Control pins (if Train_Flow_Control [Train_Enable_CDead] is set).
4	Switch Fabric, upon getting trained stops continuous Dead Cycles on Data.



8.7 CSIX Startup Sequence

This section defines the sequence required to startup the CSIX interface.

8.7.1 CSIX Full Duplex

8.7.1.1 Ingress IXP2800 Network Processor

1. On reset, FC_STATUS_OVERRIDE[Egress_Force_En] is set to force the Ingress IXP2800 to send Idle CFrames with low CReady and DReady bits to the Egress IXP2800 over TXCSR.B.
2. The Microengine or the Intel XScale® core writes a 1 to MSF_Rx_Control[RX_En_C] so that Idle CFrames can be received.
3. The Microengine or the Intel XScale® core polls on MSF_Interrupt_Status[Detected_CSIX_Idle] to see when the first Idle CFrame is received. The Intel XScale® core may use the Detected_CSIX_Idle Interrupt if MSF_Interrupt_Enable[Detected_CSIX_Idle] is set.
4. When the first Idle CFrame is received, the Microengine or the Intel XScale® core writes a 0 to FC_STATUS_OVERRIDE[Egress_Force_En] to deactivate SRB Override or writes 2'b11 to FC_STATUS_OVERRIDE[7:6] ([TM_CReady] and [TM_DReady]). This will inform the Egress IXP2800 that the Switch Fabric has sent an Idle CFrame and the Ingress IXP2800 has detected it.

8.7.1.2 Egress IXP2800 Network Processor

1. On reset, FC_STATUS_OVERRIDE[Ingress_Force_En] is set.
2. The Microengine or the Intel XScale® core writes a 1 to MSF_Tx_Control[Transmit_Idle] and MSF_Tx_Control[Transmit_Enable] so that Idle CFrames with low CReady and Dready bits are sent over TDAT.
3. The Microengine or the Intel XScale® core writes a 0 to FC_STATUS_OVERRIDE[Ingress_Force_En]. The Egress IXP2800 will then be sending Idle CFrames with CReady and DReady according to what is received on RXCSR.B from the Ingress IXP2800. If the Egress IXP2800 has not detected an Idle CFrame, low TM_CReady and TM_DReady bits will be transmitted over its TXCSR.B pin. If it has detected an Idle CFrame, the TM_CReady and TM_DReady bits are high. The TM_CReady and TM_DReady bits received on RXCSR.B by the Ingress IXP2800 are used in the Base Headers of CFrames transmitted over TDAT.
4. The Microengine or the Intel XScale® core polls on FC_Ingress_Status[TM_CReady] and FC_Ingress_Status[TM_DReady]. When they are seen active, the Microengine or the Intel XScale® core writes a 1 to MSF_Tx_Control[TX_En_CC] and MSF_Tx_Control[TX_En_CD]. Egress IXP2800 then resumes normal operation. Likewise, when the Switch Fabric recognizes Idle CFrames with “ready bits” high, it will assume normal operation.



8.7.1.3 Single IXP2800 Network Processor

1. The Microengine or the Intel XScale® core writes a 1 to MSF_Tx_Control[Transmit_Idle] and MSF_Tx_Control[Transmit_Enable] so that Idle CFrames with low CReady and DReady bits are sent over TxDAT.
2. The Microengine or the Intel XScale® core writes a 1 to MSF_Rx_Control[RX_En_C] so that Idle CFrames can be received.
3. The Microengine or the Intel XScale® core writes a 0 to FC_STATUS_OVERRIDE[Ingress_Force_En].
4. The Microengine or the Intel XScale® core polls on MSF_Interrupt_Status[Detected_CSIX_Idle] to see when the first Idle CFrame is received. The Intel XScale® core may use the Detected_CSIX_Idle Interrupt if MSF_Interrupt_Enable[Detected_CSIX_Idle] is set.
5. When the first Idle CFrame is received, the Microengine or the Intel XScale® core writes a 0 to FC_STATUS_OVERRIDE[Egress_Force_En] to deactivate SRB Override or writes 2'b11 to FC_STATUS_OVERRIDE[7:6] ([TM_CReady and TM_DReady]).
6. The Microengine or the Intel XScale® core writes a 1 to MSF_Tx_Control[TX_En_CC] and MSF_Tx_Control[TX_En_CD]. IXP2800 resumes normal operation.

8.7.2 CSIX Simplex

8.7.2.1 Ingress IXP2800 Network Processor

1. On reset, FC_STATUS_OVERRIDE[Egress_Force_En] is set to force Ingress IXP2800 to send Idle CFrames with low CReady and DReady bits to Switch Fabric over TXCDAT.
2. The Microengine or the Intel XScale® core writes a 1 to MSF_Rx_Control[RX_En_C] so that Idle CFrames can be received.
3. The Microengine or the Intel XScale® core polls on MSF_Interrupt_Status[Detected_CSIX_Idle] to see when the first Idle CFrame is received. The Intel XScale® core may use the Detected_CSIX_Idle Interrupt if MSF_Interrupt_Enable[Detected_CSIX_Idle] is set.
4. When the first Idle CFrame is received, the Microengine or the Intel XScale® core writes a 0 to FC_STATUS_OVERRIDE[Egress_Force_En]. Idle CFrames with “ready bits” high will be transmitted over TXCDAT. Ingress IXP2800 may resume normal operation.



8.7.2.2 Egress IXP2800 Network Processor

1. On reset, FC_STATUS_OVERRIDE[Ingress_Force_En] is set.
2. The Microengine or the Intel XScale® core writes a 1 to MSF_Tx_Control[Transmit_Idle] and MSF_Tx_Control[Transmit_Enable] so that Idle CFrames with low CReady and DReady bits are sent over TDAT.
3. The Microengine or the Intel XScale® core polls on MSF_Interrupt_Status[Detected_CSIX_FC_Idle] to see when the first Idle CFrame is received. The Intel XScale® core may use the Detected_CSIX_FC_Idle Interrupt if MSF_Interrupt_Enable[Detected_CSIX_FC_Idle] is set.
4. When the first Idle CFrame is received, the Microengine or the Intel XScale® core writes a 0 to FC_STATUS_OVERRIDE[Ingress_Force_En] to deactivate SRB Override.
5. The Microengine or the Intel XScale® core polls on FC_Ingress_Status[TM_CReady] and FC_Ingress_Status[TM_DReady]. When they are seen active, the Microengine or the Intel XScale® core writes a 1 to MSF_Tx_Control[TX_En_CC] and MSF_Tx_Control[TX_En_CD]. Egress IXP2800 then resumes normal operation. Likewise, when the Switch Fabric recognizes Idle CFrames with “ready bits” high, it will assume normal operation.

8.7.2.3 Single IXP2800 Network Processor

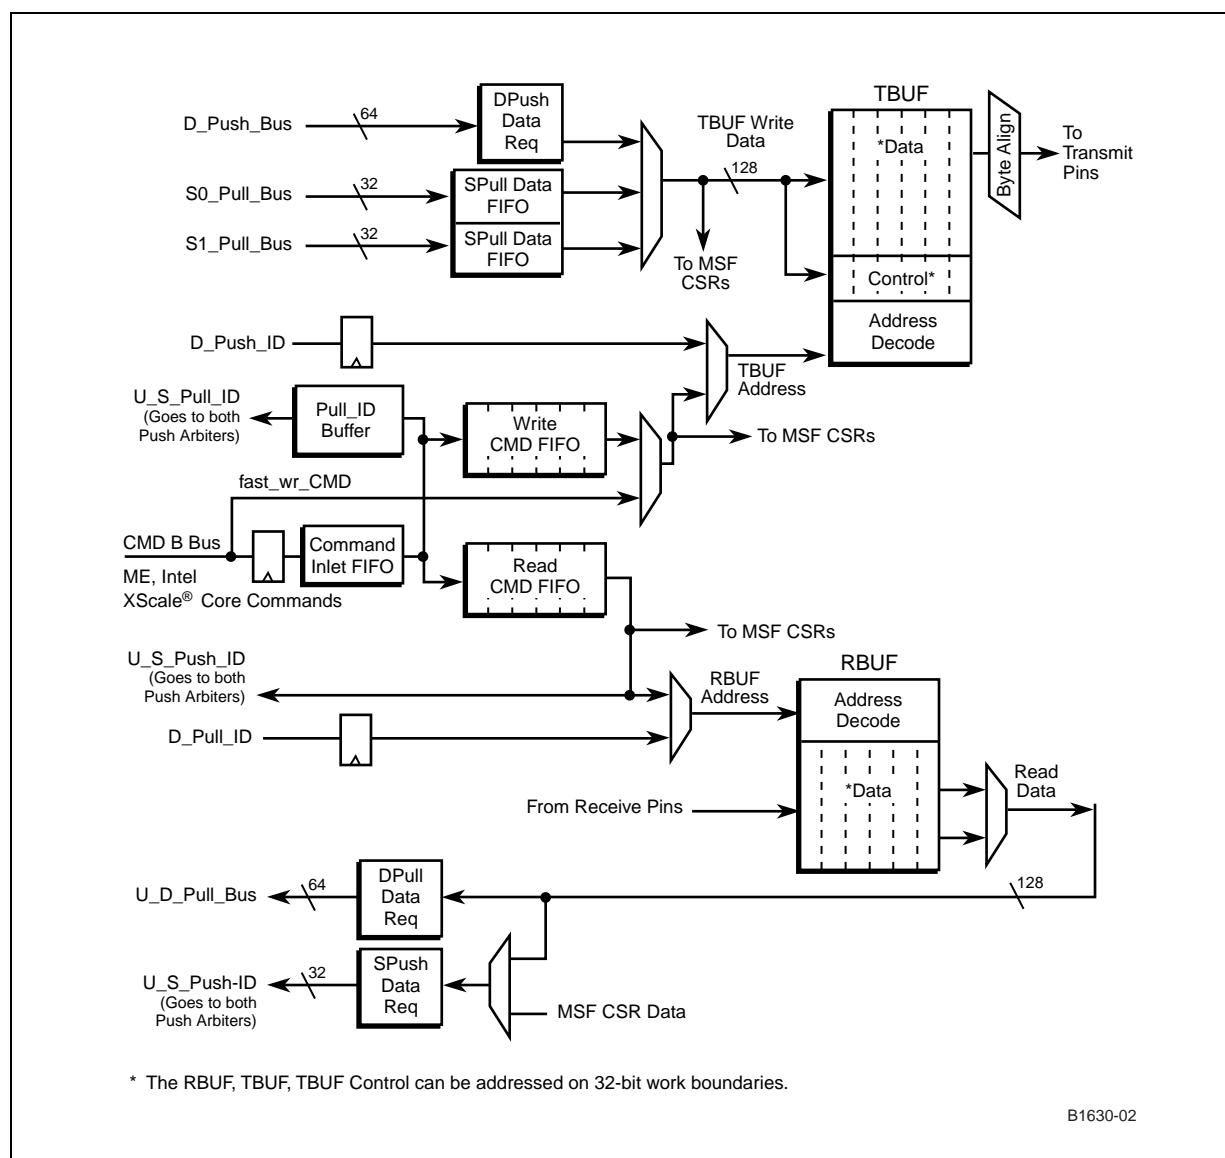
Both CSIX startup routines described above will be needed to complete the CSIX startup sequence. Using Simplex mode on a single IXP2800 with RDAT, TDAT and RXCDAT, TXCDAT using CSIX, there are essentially two independent CSIX receive and transmit buses.

8.8 Interface to Command and Push and Pull Buses

Figure 100 shows the interface of the MSF to the command and push and pull buses. Data transfers to and from the TBUF/RBUF are done in the following cases (refer to section):

- RBUF or MSF CSR to Microengine S_TRANSFER_IN Register for Instruction:
- Microengine S_TRANSFER_OUT Register to TBUF or MSF CSR for Instruction:
- Microengine to MSF CSR for Instruction:
- From RBUF to DRAM for Instruction:
- From RBUF to DRAM for Instruction:

Figure 100. MSF to Command and Push and Pull Buses Interface Block Diagram





8.8.1 RBUF or MSF CSR to Microengine S_TRANSFER_IN Register for Instruction:

```
msf[read, $s_xfer_reg, src_op_1, src_op_2, ref_cnt], optional_token
```

For transfers to a Microengine, the MSF acts as a target. Commands from Microengines and the Intel XScale® core are received on the command bus. The commands are checked to see if they are targeted to the MSF. If so, they are enqueued into the Command Inlet FIFO, and then moved to the Read Cmd FIFO. When the Command Inlet FIFO is nearly full, it asserts a signal to the command arbiters. The command arbiters prevent further commands to the MSF until after the full signal is asserted. The RBUF element or CSR specified in the address field of the command is read and the data is registered in the SPUSH_DATA register. The control logic then arbitrates for S_PUSH_BUS, and when granted, it drives the data.

8.8.2 Microengine S_TRANSFER_OUT Register to TBUF or MSF CSR for Instruction:

```
msf[write, $s_xfer_reg, src_op_1, src_op_2, ref_cnt], optional_token
```

For transfers from a Microengine, the MSF acts as a target. Commands from Microengines are received on the two command buses. The commands are checked to see if they are targeted to the MSF. If so, they are enqueued into the Command Inlet FIFO, and then moved to the Write Cmd FIFO. When the Command Inlet FIFO is nearly full, it asserts a signal to the command arbiters. The command arbiters prevent further commands to the MSF until after the full signal is asserted. The control logic then arbitrates for S_PULL_BUS, and when granted, it receives and registers the data from the Microengine into the S_PULL_DATA register. It then writes that data into the TBUF element or CSR specified in the address field of the command.

8.8.3 Microengine to MSF CSR for Instruction:

```
msf[fast_write, src_op_1, src_op_2]
```

For fast write transfers from the Microengine, the MSF acts as a target. Commands from Microengines are received on the two command buses. The commands are checked to see if they are targeted to the MSF. If so, they are enqueued into the Command Inlet FIFO, and then moved to the Write Cmd FIFO. When the Command Inlet FIFO is nearly full, it asserts a signal to the command arbiters. The command arbiters prevent further commands to the MSF until after the full signal is asserted. The control logic uses the address and data, both found in the address field of the command. It then writes the data into the CSR specified.

8.8.4 From RBUF to DRAM for Instruction:

```
dram[rbuf_rd, --, src_op1, src_op2, ref_cnt], indirect_ref
```

For the transfers to DRAM, the RBUF acts like a slave. The address of the data to be read is given in D_PULL_ID. The data is read from RBUF and registered in the D_PULL_DATA register. It is then multiplexed and driven to the DRAM channel on D_PULL_BUS.

8.8.5 From DRAM to TBUF for Instruction:

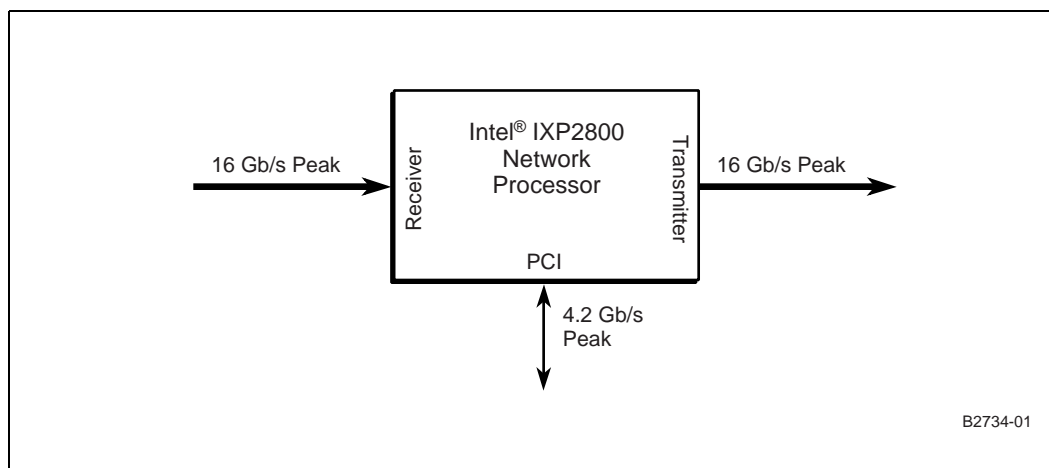
```
dram[tbuf_wr, --, src_op1, src_op2, ref_cnt], indirect_ref
```

For the transfers from DRAM, the TBUF acts like a slave. The address of the data to be written is given in D_PUSH_ID. The data is registered and assembled from D_PUSH_BUS, and then written into TBUF.

8.9 Receiver and Transmitter Interoperation with Framers and Switch Fabrics

The Intel® IXP2800 Network Processor can process data received at a peak rate of 16 Gb/s and transmit data at a peak rate of 16 Gb/s. In addition, data may be received and transmitted via the PCI bus at an aggregate peak rate of 4.2 Gb/s, as shown in Figure 101.

Figure 101. Basic I/O Capability of the Intel® IXP2800 Network Processor



The network processor's receiver and transmitter can be independently configured to support either an SPI-4.2 framer interface or a fabric interface consisting of DDR LVDS signaling and the CSIX-L1 protocol. The dynamic training sequence of SPI-4.2, used for de-skewing the signals, has been optionally incorporated into the fabric interface.

"SPI-4.2 is an interface for packet and cell transfer between a physical layer (PHY) device and a link layer device, for aggregate bandwidths of OC-192 ATM and Packet over SONET/SDH (POS), as well as 10 Gb/s Ethernet applications."¹

"CSIX-L1 is the Common Switch Interface. It defines a physical interface for transferring information between a traffic manager (Network Processor) and a switching fabric..."² The network processor adopts the protocol of CSIX-L1, but uses a DDR LVDS physical interface rather than an LVCMOS or HSTL physical interface.

1. "System Packet Interface Level 4 (SPI-4) Phase 2: OC-192 System Interface for Physical and Link Layer Devices," Implementation Agreement: OIF-SPI4-02.0, Optical Internetworking Forum
2. "CSIX-L1: Common Switch Interface Specification-L1," CSIX

SPI-4.2 supports up to 256 port addresses, with independent flow control for each. For data received by the PHY and passed to the link layer device, flow control is optional. The flow control mechanism is based upon independent pools of credits, corresponding to 16-byte blocks, for each port.

The CSIX-L1 protocol supports 4096 ports and 256 unicast classes of traffic. It supports various forms of multicast and 256 multicast queues of traffic. The protocol supports independent link-level flow control for data and control traffic and supports virtual output queue (VOQ) flow control for data traffic.

8.9.1 Receiver and Transmitter Configurations

The network processor receiver and transmitter *independently* support three different configurations:

- Simplex (SPI-4.2 or CSIX-L1 protocol), described in [Section 8.9.1.1](#).
- Hybrid simplex (transmitter only, SPI-4.2 data path, and CSIX-L1 protocol flow control), described in [Section 8.9.1.2](#).
- Dual Network Processor, full duplex (CSIX-L1 protocol), described in [Figure 8.9.1.3](#).

Additionally, the *combined* receiver and transmitter support a single Network Processor, full-duplex configuration using two different protocols:


- Multiplexed SPI-4.2 protocol, described in [Section 8.9.1.4](#).
- CSIX-L1 protocol, described in [Section 8.9.1.5](#).

In both the simplex and hybrid simplex configurations, the path receiving from a framer, fabric, or Network Processor is independent of the path transmitting to a framer, fabric, or Network Processor. In a full duplex configuration, the receiving path forwards CSIX-L1 control information for the transmit path and vice versa.

8.9.1.1 Simplex Configuration

In the simplex configuration, as shown in [Figure 102](#), the reverse path provides control information to the transmitter. This control information may include flow control information and requests for dynamic training sequences.

Figure 102. Simplex Configuration



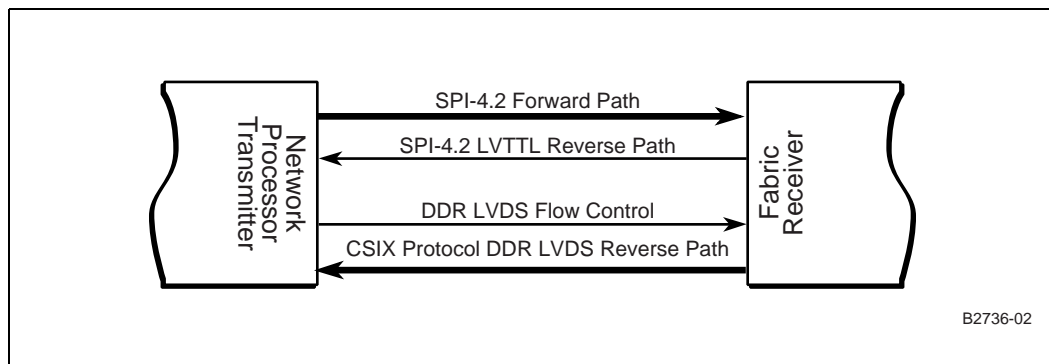
The SPI-4.2 mode of the simplex configuration supports an LVTTTL reverse path or status interface clocked at up to 125 MHz or a DDR LVDS reverse path or status interface clocked at up to 500 MHz. The SPI-4.2 mode status interface consists of a clock signal and two data signals.

The CSIX-L1 protocol mode of the simplex configuration supports a full-duplex implementation of the CSIX-L1 protocol, but no Data CFrames are transferred on the reverse path and the reverse path is a quarter of the width of the forward path. The CSIX-L1 protocol mode supports a DDR LVDS reverse path interface clocked at up to 500 MHz. The CSIX-L1 protocol mode reverse path control interface consists of a clock signal, four data signals, a parity signal, and a start-of-frame signal.

8.9.1.2 Hybrid Simplex Configuration

In the hybrid simplex configuration, data transfers and link-level flow control is supported via the SPI-4.2 modes of the receiver and transmitter, as shown in Figure 103. Only the LVTTTL SPI-4.2 status interface is supported in this configuration.

Figure 103. Hybrid Simplex Configuration



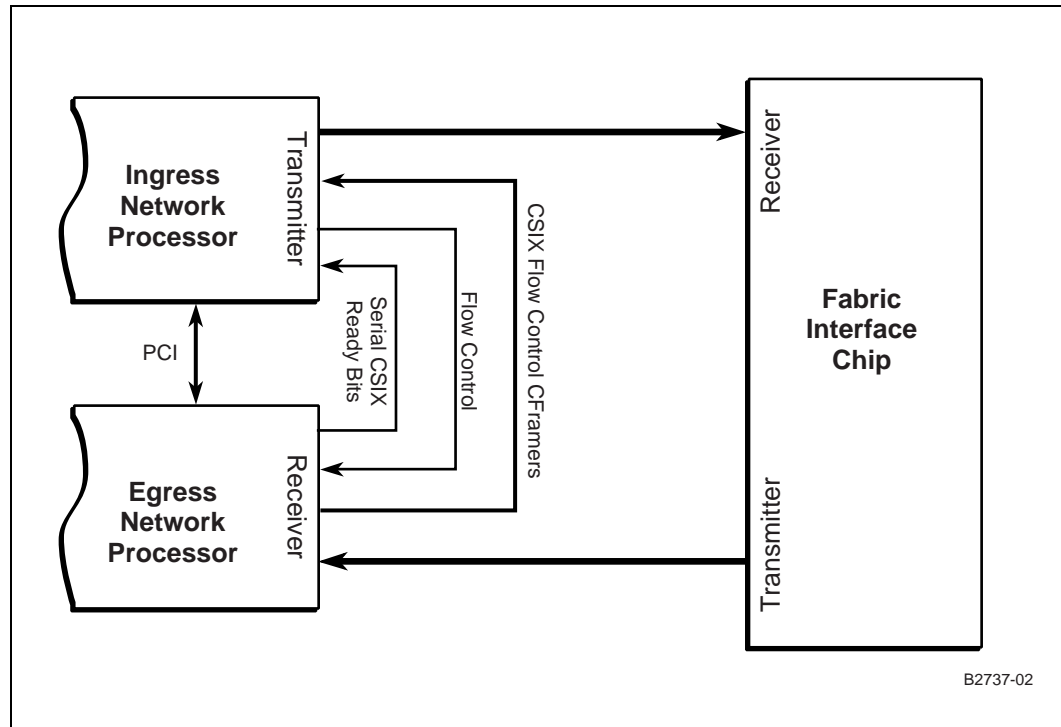
Virtual output queue flow control information (or other information) is delivered to the transmitter via the CSIX-L1 protocol via an interface similar to the reverse path of the CSIX-L1 protocol mode of the simplex configuration. Flow control for the CSIX-L1 CFrames is provided by an asynchronous LVDS signal back to the fabric and not by the “ready bits” of the CSIX-L1 protocol.

The hybrid simplex configuration for a fabric interface may be especially useful to implementers when an SPI-4.2 interface implementation is readily available. The CSIX-L1 protocol reverse path may not need to operate at a clock rate as aggressive as the SPI-4.2 interface and, as such, may be easier to implement than a full-rate data interface.

8.9.1.3 Dual Network Processor Full Duplex Configuration

In the dual Network Processor, full duplex configuration, an ingress Network Processor and an egress Network Processor are integrated to offer a single full duplex interface to a fabric, similar to the CSIX-L1 interface, as shown in Figure 104. This configuration provides an interface that is closest to the standard CSIX-L1 interface. It is easiest to bridge between this configuration and an actual CSIX-L1 interface.

Figure 104. Dual Network Processor, Full Duplex Configuration



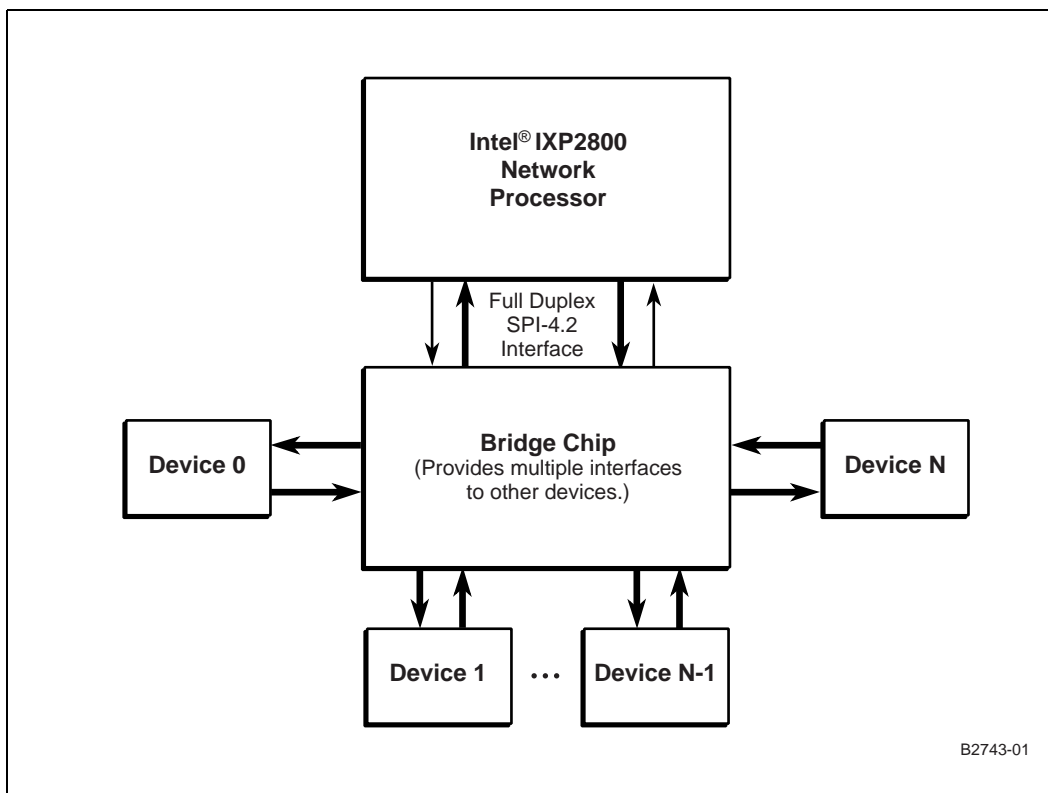
Flow control CFrames are forwarded by the egress Network Processor to the ingress Network Processor over a separate flow control interface. The bandwidth of this interface is a quarter of the primary interface offered to the fabric. A signal from ingress Network Processor to egress Network Processor provides flow control for this interface. (This interface is the same interface that was used in the hybrid simplex configuration.) A separate signal from egress Network Processor to ingress Network Processor provides the state of the CSIX-L1 “ready bits” that were received from the fabric, conveying the state of the fabric receiver, and those that should be sent to the fabric, conveying the state of the egress Network Processor receiver.

The PCI may be used to convey additional information between the egress Network Processor and ingress Network Processor.

8.9.1.4 Single Network Processor Full Duplex Configuration (SPI-4.2)

The single Network Processor, full duplex configuration (SPI-4.2 only) allows a single Network Processor to interface to multiple discrete devices, processing both the receiver and transmitter data for each, as shown in Figure 105 (where $N=255$). Up to 256 devices can be addressed by the SPI-4.2 implementation. The bridge chip implements the specific interfaces for each of those devices.

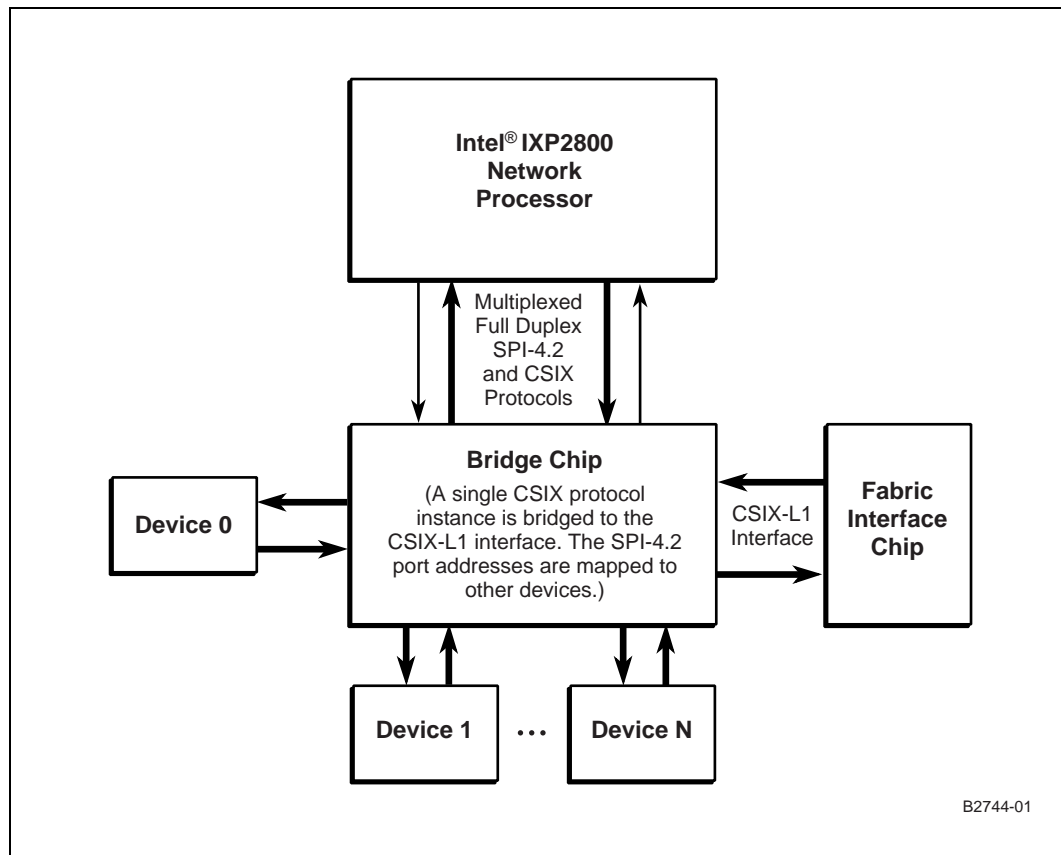
Figure 105. Single Network Processor, Full Duplex Configuration (SPI-4.2 Protocol)



8.9.1.5 Single Network Processor, Full Duplex Configuration (SPI-4.2 and CSIX-L1)

The Single Network Processor, Full Duplex Configuration (SPI-4.2 and CSIX-L1 Protocol) allows a single Network Processor to interface to a fabric via a CSIX-L1 interface and to multiple other discrete devices, as shown in Figure 106. The CSIX-L1 and SPI-4.2 protocols are multiplexed on the network processor receiver and transmitter interface. Independent processing and buffering resources are allocated to each protocol.

Figure 106. Single Network Processor, Full Duplex Configuration (SPI-4.2 and CSIX-L1 Protocols)



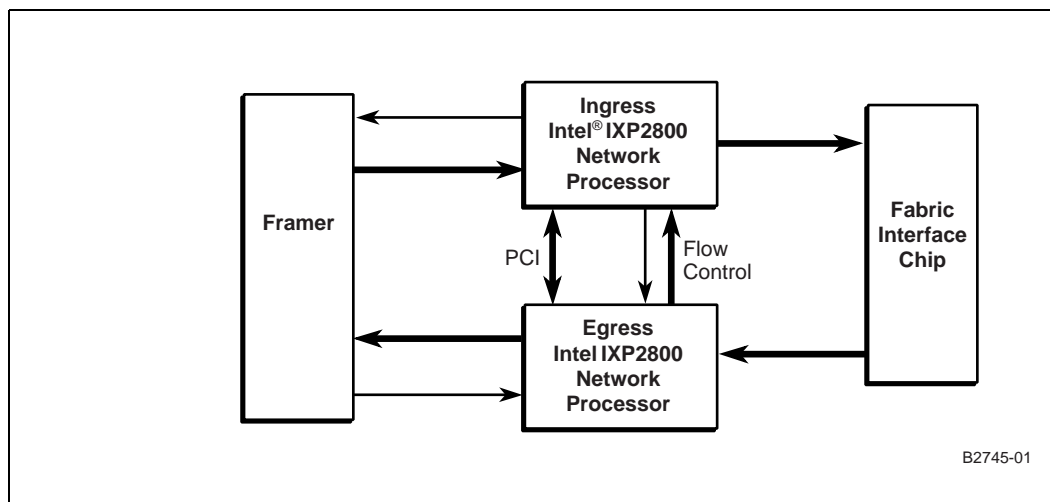
8.9.2 System Configurations

The receiver and transmitter configurations in the preceding Section 8.9.1 enable several system designs, as shown in Figure 107 through Figure 111.

8.9.2.1 Framer, Single Network Processor Ingress and Egress, and Fabric Interface Chip

Figure 107 illustrates the baseline system configuration consisting of the dual chip, full-duplex fabric configuration of network processors with a framer chip and a fabric interface chip

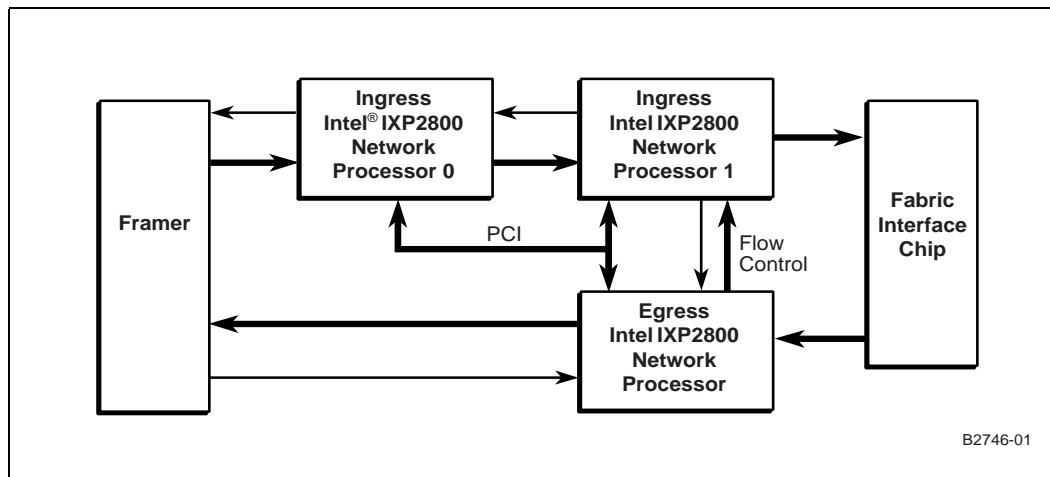
Figure 107. Framer, Single Network Processor Ingress and Egress, and Fabric Interface Chip



8.9.2.2 Framer, Dual Network Processor Ingress, Single Network Processor Egress, and Fabric Interface Chip

If additional processing capacity is required in the ingress path, an additional network processor can be added to the configuration, as shown in Figure 108. The configuration of the interface between the two ingress network processors can use either the SPI-4.2 or CSIX-L1 protocol.

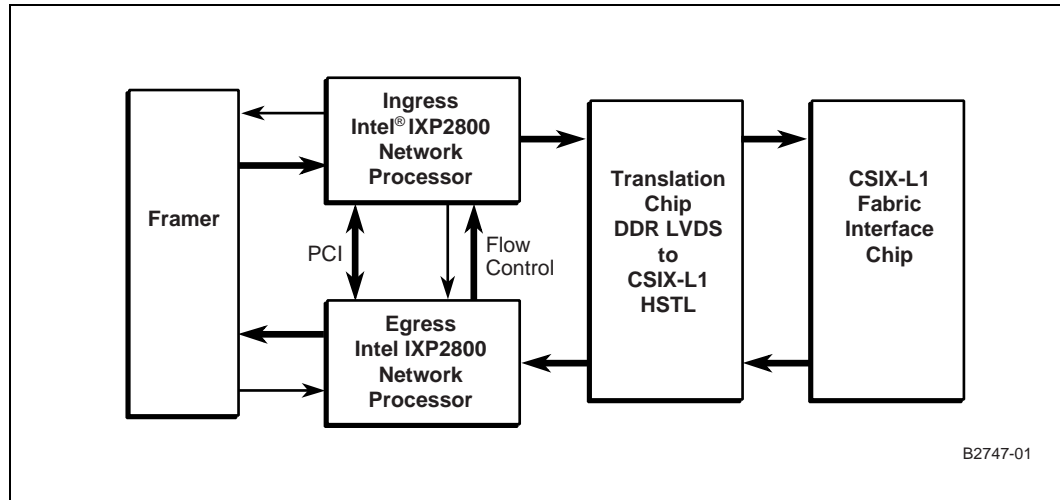
Figure 108. Framer, Dual Processor Ingress, Single Processor Egress, and Fabric Interface Chip



8.9.2.3 Framers, Single Network Processor Ingress and Egress, and CSIX-L1 Chips for Translation and Fabric Interface

To interface to existing standard CSIX-L1 fabric interface chips, a translation bridge can be employed, as shown in Figure 109. Translation between the network processor interface and standard CSIX-L1 is very simple by design.

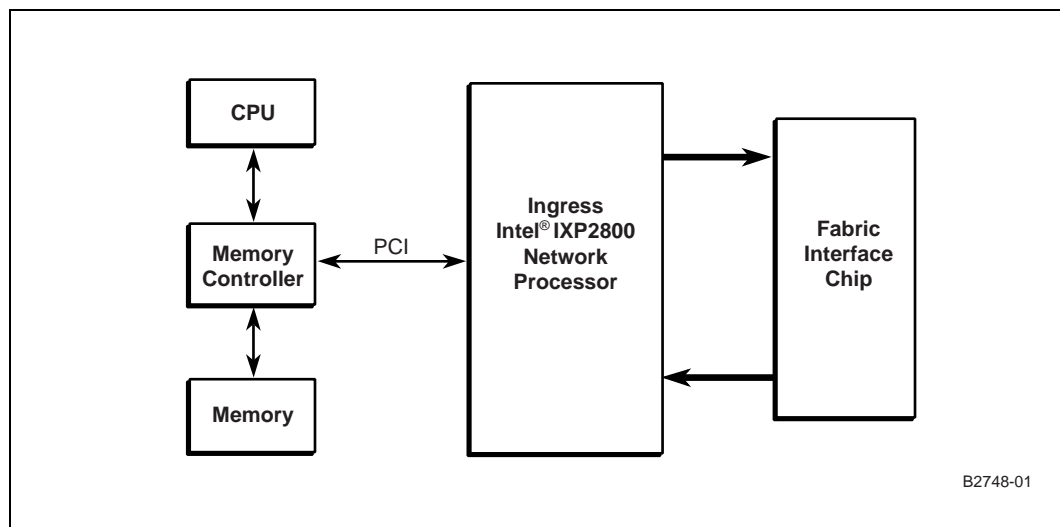
Figure 109. Framers, Single Network Processor Ingress, Single Network Processor Egress, CSIX-L1 Translation Chip and CSIX-L1 Fabric Interface Chip



8.9.2.4 CPU Complex, Network Processor, and Fabric Interface Chip

If a processor card requires access to the fabric, a single network processor can provide both ingress and egress access to the fabric for the processor via the PCI interface, as shown in Figure 110. In many cases the available aggregate peak bandwidth of 4.2 Gb/s is sufficient for the processor's capacity.

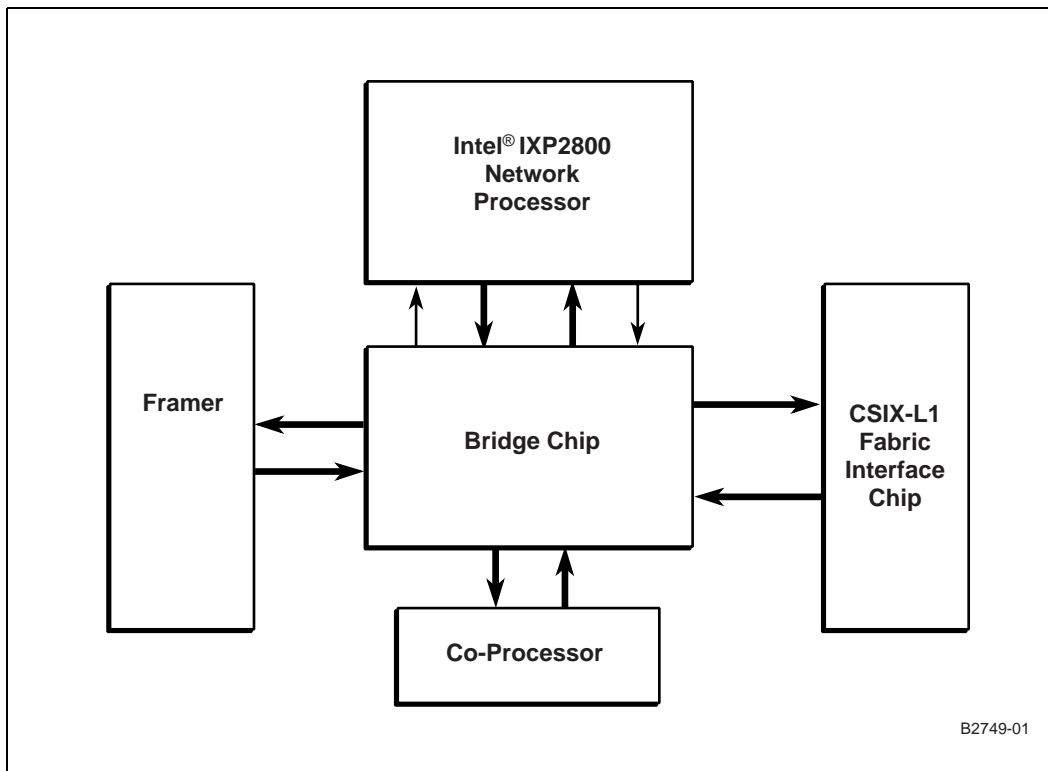
Figure 110. CPU Complex, Network Processor, and Fabric Interface Chips



8.9.2.5 Framer, Single Network Processor, Co-Processor, and Fabric Interface Chip

The network processor supports multiplexing the SPI-4.2 and CSIX-L1 protocols over its physical interface via a protocol signal. This capability enables using a bridge chip to allow a single network processor to support the ingress and egress paths between a framer and a fabric, provided the aggregate system bandwidth does not exceed the capabilities of that single network processor, as shown in Figure 111.

Figure 111. Framer, Single Network Processor, Co-Processor, and Fabric Interface Chip

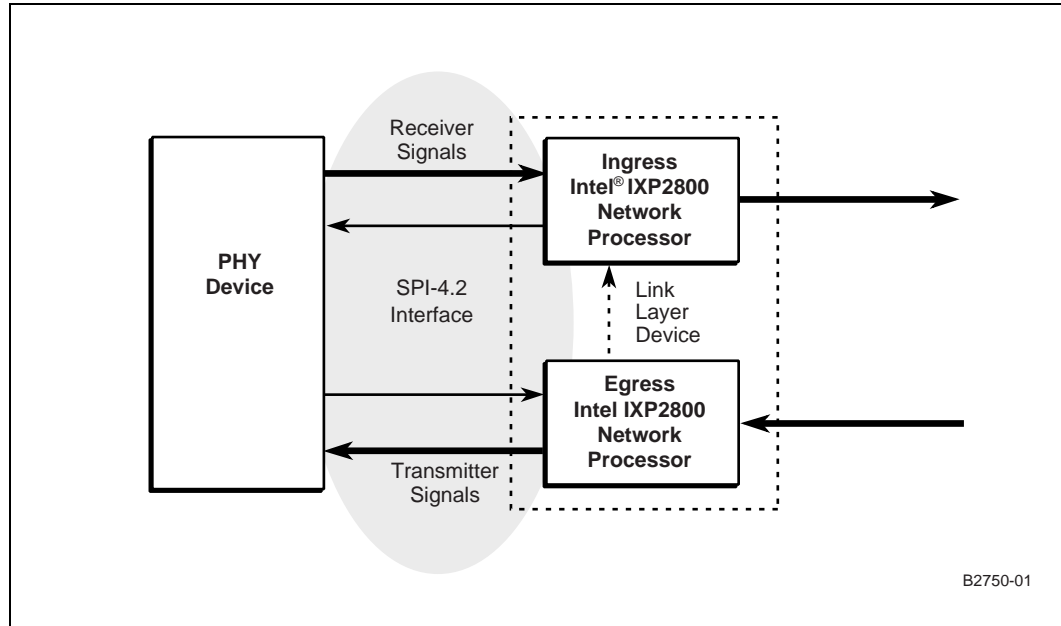




8.9.3 SPI-4.2 Support

Data is transferred across the SPI-4.2 interface in variously-sized bursts and encapsulated with a leading and trailing control word. The control words provide annotation of the data with port address (0-255) information, start-of-packet and end-of-packet markers, and an error detection code (DIP-4). Data must be transferred in 16-byte integer multiples, except for the final burst of a packet.

Figure 112. SPI-4.2 Interface Reference Model with Receiver and Transmitter Labels Corresponding to Link Layer Device Functions



The status interface transfers state as an array of state or calendar, two bits per port, for all of the supported ports. The status information provides for reporting one of three status states for each port (satisfied, hungry, and starving) corresponding to credit availability for the port. The mapping of calendar offset to port is flexible. Individual ports may be repeated multiple times for greater frequency of update.

8.9.3.1 SPI-4.2 Receiver

The network processor receiver stores received SPI-4.2 bursts into receiver buffers. The buffers may be configured as 128 buffers of 64 bytes, 64 buffers of 128 bytes, or 32 buffers of 256 bytes. Information from the control words, the length of the burst, and the TCP checksum of the data are stored in an additional eight bytes of control storage. The buffers support storage of bursts containing an amount of data that is less than or equal to the buffer size. A burst that is greater than the configured size of the buffers is stored in multiple buffers. Each buffer is made available to software as it becomes filled.

As the filling of each buffer completes, the buffer is dispatched to a thread of a Microengine that has been registered in a free list of threads, and the eight bytes of control information are forwarded to the register context of the thread. If no thread is currently available, the receiver waits for a new thread to become available as other buffers are also filled (and then also have “waiting queues”).



As threads complete processing of the data in a buffer, the buffer is returned to a free list. Subsequently, the thread also returns to a separate free list. The return of buffers and threads to the free lists may occur in a different order than the order of their removal.

All SPI-4.2 ports sharing the interface have equal access to the buffering resources. Flow control can transition to a non-starving state when 25%, 50%, 75%, or 87.5% of the buffers are consumed, as configured by `HWM_Control[RBUF_S_HWM]`. At this point, the remaining buffers are available and, additionally, 2K bytes of packed FIFO (corresponding to 128 SPI-4.2 credits) are available for incoming data storage. If receiver flow control is expected to be asserted and for a sufficiently large number of ports and values of `MaxBurst1` or `MaxBurst2`, it may be necessary for the PHY device to discard credits already granted if a state of Satisfied is reported by the network processor to the device, treating the Satisfied state more as an XOFF state. Otherwise, excessive credits may be outstanding for the storage available and receiver overruns may occur.

For more information about the SPI-4.2 receiver, see [Section 8.2.7](#).

8.9.3.2 SPI-4.2 Transmitter

The network processor transmitter transfers SPI-4.2 bursts from transmitter buffers. The buffers may be configured as 128 buffers of 64 bytes, 64 buffers of 128 bytes, or 32 buffers of 256 bytes. The control word information and other control information for the burst are stored in additional control storage. The buffers are always transmitted in a fixed order. Software can determine the index of the last buffer transmitted, and keep track of the last buffer committed to the transmitter. The transmitter buffers are used as a ring, with the “get index” updated by the transmitter and the “put index” updated due to committing a buffer element to transmission.

Each transmit buffer supports a limited gather capability to stitch together a protocol header and a payload. The buffer supports independent prefix (or prepended) data and payload data. The prefix data can begin at any offset from 0 to 7 and have a length of from 0 to 31 bytes. The payload begins at an offset of 0 to 7 bytes from the next octal-byte boundary following the prefix and can fill out the remainder of the buffer. For more complicated merging or shifting of data within a burst, the data should be passed through a Microengine to perform any arbitrary merging and/or shifting.

Buffers may be statically allocated to different ports in an inter-leaved fashion so that bandwidth availability is balanced for each of the ports. Transmit buffers may be flagged to be skipped if no data is available for a particular port.

The transmitter scheduler, implemented on a Microengine, is responsible for reacting to the status information provided by the PHY device. The status information can be read via registers. The status information is available in two formats: a single status per register and status for 16 ports in a single register. For more information, see [Section 8.3.4, “Transmit Flow Control Status” on page 270](#).

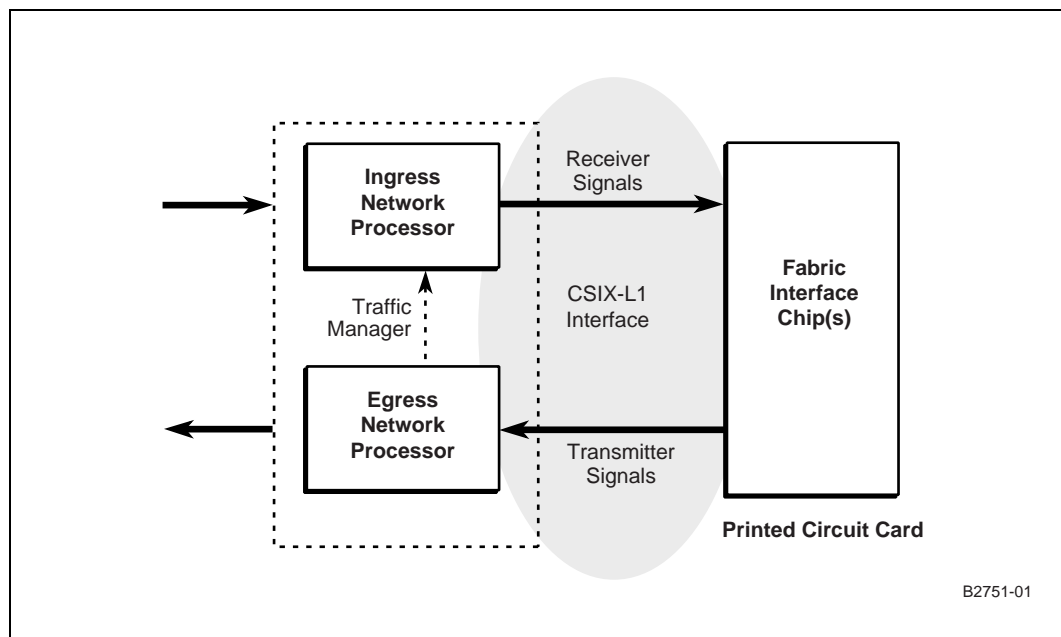


8.9.4 CSIX-L1 Protocol Support

8.9.4.1 CSIX-L1 Interface Reference Model: Traffic Manager and Fabric Interface Chip

The CSIX-L1 protocol operates between a Traffic Manager and a Fabric Interface Chip(s) across a full-duplex interface. It supports mechanisms to interface to a fabric that avoid congestion using virtual output queue (VOQ) flow control and enables a fabric that offers lossless, non-blocking transfer of data from ingress port to egress ports. Both data and control information pass over the receiver and transmitter interfaces.

Figure 113. CSIX-L1 Interface Reference Model with Receiver and Transmitter Labels Corresponding to Fabric Interface Chip Functions



The Traffic Manager on fabric ingress is responsible for segmentation of packet data and scheduling the transmission of data segments into the fabric. The fabric on ingress is responsible for influencing the scheduling of data transmission through link-level flow control and Virtual Output Queue (VOQ) flow control so that the fabric does not experience blocking or data loss due to congestion. The fabric on egress is responsible for scheduling the transfer of data to the Traffic Manager according to the flow control indications from the Traffic Manager.

The CSIX-L1 protocol supports addressing up to 4096 fabric ports and identifies up to 256 classes of unicast traffic. It optionally supports multicast and broadcast traffic, supporting identification of up to 256 queues of such traffic. Virtual output queue flow control is supported at the ingress to the fabric and the egress from the fabric.

The standard CSIX-L1 interface supports interface widths of 32, 64, 94, and 128 bits. A single clocked transfer of information across the interface is called a CWord. The CWord size is the width of the interface.



Information is passed across the interface in CFrames. CFrames are padded out to an integer multiple of CWords. CFrames consist of a 2-byte base header, an optional 4-byte extension header, a payload of 1 to 256 bytes, padding, and a 2-byte vertical parity. Transfers across the interface are protected by a horizontal parity. When there is no information to pass over the interface, an alternating sequence of Idle CFrames and Dead Cycles are passed across the interface.

There are 16 possible codes for CFrame types. Each CFrame type is either a data CFrame or a control CFrame. Data CFrame types include Unicast, Multicast Mask, Multicast ID, Multicast Binary Copy, and Broadcast. Control CFrames include Flow Control.

CSIX-L1 supports independent link-layer flow control for data CFrames and control CFrames by using “ready bits” (CRdy and DRdy) in the base header. The response time for link-level flow control is specified to be 32 interface clock ticks, but allows for additional time to complete transmission of any CFrame already in progress at the end of that interval.

8.9.4.2 Intel® IXP2800 Support of the CSIX-L1 Protocol

The adaptation of the CSIX-L1 protocol to the network processor physical interface has been accomplished in a straightforward manner.

8.9.4.2.1 Mapping to 16-Bit Wide DDR LVDS

The CSIX-L1 interface is built in units of 32 data bits. For each group of 32 data signals, there is a clock signal (RxClk, TxClk), a start-of-frame signal (RxSOF, TxSOF) and a horizontal-parity signal (RxPar, TxPar). If the CWord or interface width is greater than 32 bits, the assertion of the Start-of-Frame signal associated with each group of 32 data bits is used to synchronize the transfers across the independently clocked individual 32-bit interfaces.

The network processor supports 32-bit data transfers across two transfers or clock edges of the SPI-4.2 16-bit DDR LVDS data interface. The CSIX-L1 RxSOF and TxSOF signals are mapped to the SPI-4.2 TCTL and RCTL signals. For the transfer of CFrames, the start-of-frame signal is asserted on only the first edge of the 32-bit transfer. (Assertion of the start-of-frame signal for multiple contiguous clock edges denotes the start of a de-skew training sequence as described below.)

Receiver logic for the interface should align the start of 32-bit transfers to the assertion of the start-of-frame signal. The network processor always transmits the high order bits of a 32-bit transfer on the rising edge of the transmit clock, but a receiver may de-skew the signals and align the received data with the falling edge of the clock. The network processor receiver always aligns the received data according to the assertion of the start-of-frame signal.

The network processor supports CWord widths of 32, 64, 96, and 128 bits. It will pad out CFrames (including Idle CFrames) and Dead Cycles according to this CWord width. The physical interface remains just 16 data bits. The start-of-frame signal is only asserted for the high order 16 bits of the first 32-bit transfer; it is not asserted for each 32-bit transfer. Support for multiple CWord widths is intended to facilitate implementation of IXP2800-to-CSIX-L1 translator chips and to facilitate implementation of chips with native network processor interfaces, but with wider internal transfer widths.

The network processor supports a horizontal parity signal (RPAR, TPAR). The horizontal parity signal covers the 16 data bits that are transferred on each edge of the clock. It does not cover 32 bits as in CSIX-L1. Support for horizontal-parity requires an additional physical signal beyond that required for SPI-4.2. Checking of the horizontal parity can be optionally disabled on reception. If a fabric interface chip does not support TPAR, then the checking of RPAR should be disabled.

The network processor supports a variation of the standard CSIX-L1 vertical parity. Instead of a single vertical XOR for the calculation of the vertical parity, the network processor can be configured to calculate as DIP-16 code, as documented within the SPI-4.2 specification. If horizontal parity is not enabled for the interface, the use of the DIP-16 code is recommended to provide for better error coverage than that provided by a vertical parity.

8.9.4.2.2 **Support for Dual Chip, Full-Duplex Operation**

A dual-chip configuration of network processors consisting of an ingress and egress network processor, can present a full-duplex interface to a fabric interface chip, consistent with the expectations of the CSIX-L1 protocol. A flow control interface is supported between the ingress and egress chips to forward necessary flow control information from the egress network processor to the ingress network processor. Additional information can be transferred between the ingress and egress network processors through the PCI bus.

The flow control interface consists of a data transfer signal group, a serial signal for conveying the state of the CSIX-L1 “ready bits” (TXCSRB, RXCSRB), and a backpressure signal (TXCFC, RXCFC) to avoid overrunning the receiver in the ingress network processor. (The orientation of the signal names is consistent with the egress network processor, receiving CFrames from the fabric, and forwarding flow control information out through the transmit flow control pins.) The data transfer signal group consists of:

- Four data signals (TXCDAT[0..3], RXCDAT[0..3])
- A clock (TXCCLK, RXCCLK)
- A start-of-frame signal (TXCSOF, RXCSOF)
- A horizontal-parity signal (TXCPAR, RXCPAR)

The network processor receiver forwards Flow Control CFrames from the fabric in a cut-through fashion over the flow control interface. The flow control interface has one-fourth of the bandwidth of the network processor fabric data interface. The Crdy bit in the base header of the CSIX-L1 protocol (link-level flow control) prevents overflowing of the FIFO for transmitting out the flow control interface from the egress network processor. The fabric can implement a rate limit on the transmission of Flow Control CFrames to the egress network processor, consistent with the bandwidth available on the flow control interface. With a rate limit, the fabric can detect congestion of Flow Control CFrames earlier, instead of waiting for the assertion of cascaded backpressure signals.

The CRdy and DRdy bits of CFrames sent across the flow control interface are set to 0 on transmission and ignored upon reception at the ingress network processor. If no CFrames are available to send from the egress network processor to the ingress network processor, an alternating sequence of Idle CFrames and Dead Cycles is sent from the egress to the ingress network processor, consistent with the CSIX-L1 protocol.

The state of the CRdy and DRdy bits sent to the egress network processor by the fabric and the state of the CRdy and DRdy bits that should be sent to the fabric by the ingress network processor, reflecting the state of the egress network processor buffering, are sent through the TXCSRB signal and received through the RXCSRB signal. A new set of bits are conveyed every 10 clock edges or five clock cycles, of the interface. A de-assertion of a “ready bit” is forwarded immediately upon processing the “ready bit”. An assertion of a “ready bit” is forwarded only after all of the horizontal parities and the vertical parity of the CFrame are checked. A configuration of ingress and egress network processors is expected to respond to the de-assertion of a CRdy or DRdy bit within 32 clock cycles (RCLK), consistent with the formulation described for CSIX-L1.

The backpressure signal (TXCFC, RXCFC) is an asynchronous signal and is asserted by the ingress network processor to prevent overflow of the ingress network processor ingress flow control FIFO. If the egress network processor is so optionally configured, it will react to assertion of the backpressure signal for 32 clock cycles (64 edges) as a request for a de-skew training sequence to be transmitted on the flow control interface.

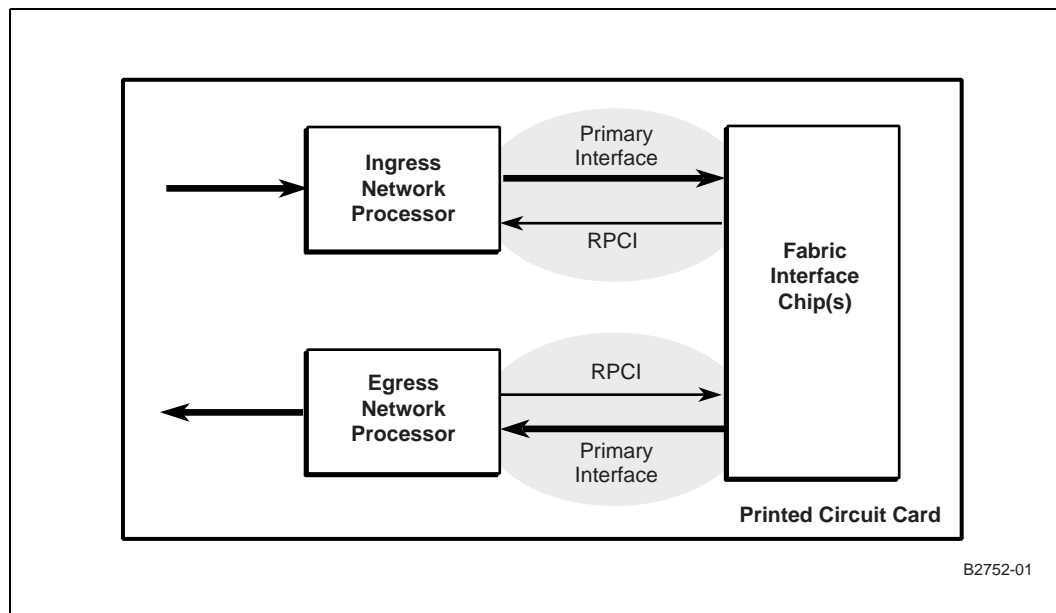
The flow control interface only supports a 32-bit CWord. Flow Control CFrames that are received by the egress network processor are stripped of any padding associated with large CWord widths and forwarded to the flow control interface.

The various options for parity calculation and checking supported on the data interface are supported on the flow control interface. Horizontal parity checking may be optionally disabled. The standard calculation of vertical parity may be replaced with a DIP-16 calculation.

8.9.4.2.3 Support for Simplex Operation

The network processor supports a mode of operation that supports the CSIX-L1 protocol, but offers an independent interface for the ingress and egress network processors. In this mode, the ingress and egress network processors each offer an independent full-duplex CSIX-L1 flavor of interface to the fabric, but the network processor-to-fabric interface on the egress network processor and the fabric-to-network processor interface of the ingress network processor are of reduced width, consisting of four (instead of 16) data signals. These narrow interfaces are referred to as Reverse Path Control Interfaces and use the same physical interface as the flow control interface in the dual-chip, full duplex configuration. They support the transfer of Flow Control CFrames and the CRdy and DRdy “ready” bits, but are not intended to support the transfer of data CFrames.

Figure 114. Reference Model for IXP2800 Support of the Simplex Configuration Using Independent Ingress and Egress Interfaces



The Reverse Path Control Interfaces (RPCI) support only the 32-bit CWord width of the dual chip, full duplex flow control interface. The variations of parity support provided by the data interface and the flow control interface are supported by the RPCI.

The transfer time of CFrames across the RPCI is four times that of the data interface. The latency of link-level flow control notifications depends on the frequency of sending new CFrame base headers. As such, the maximum size of CFrames supported on the RPCI should be limited to provide sufficient link-level flow control responsiveness.

The behavior of state machines for a full-duplex interface regarding interface initialization, link-level flow control, and requests to send a de-skew training sequence is supported by the data interface in combination with its reverse path control interface as if the two interfaces were equivalent to a full-duplex interface.

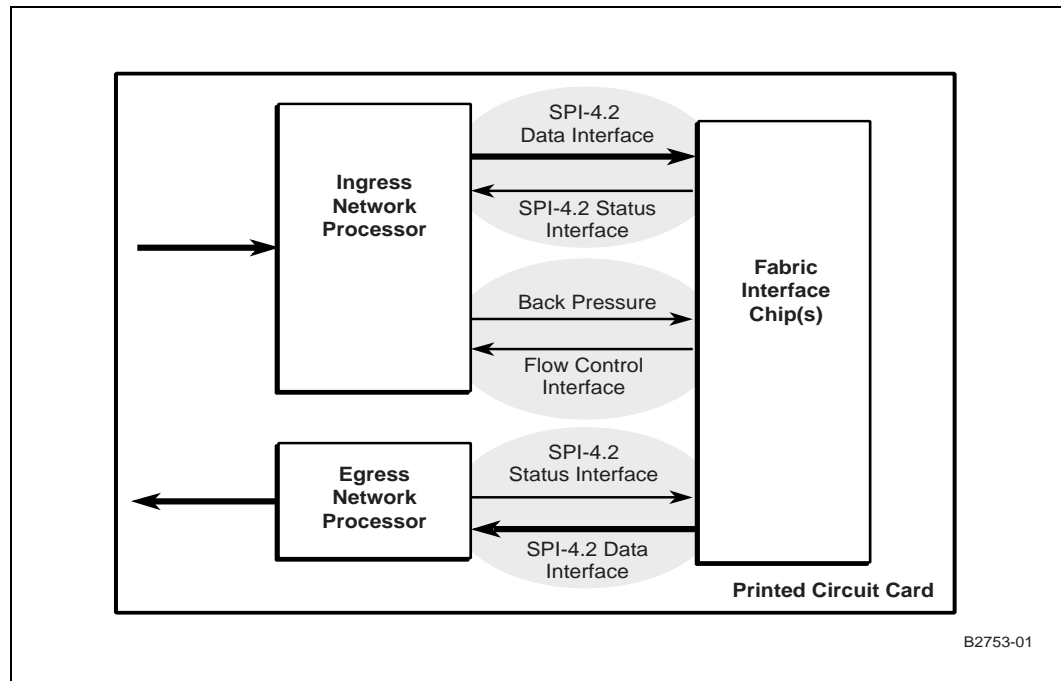
The simplex mode of interfacing to the ingress and egress network processor is an alternative to the dual chip full-duplex configuration. It provides earlier notification of Flow Control CFrame congestion within the ingress network processor and marginally less latency for delivery of Flow Control CFrames to the ingress network processor. It allows more of the bandwidth on the data interface to be used for the transfer of data CFrames as Flow Control CFrames are transferred on the RPCI.

The simplex configuration provides a straightforward mechanism for the egress network processor to send VOQ flow control to the fabric if the fabric supports such functionality. In the dual chip, full-duplex configuration, the egress network processor sends a request across the PCI to the ingress network processor, requesting that a Flow Control CFrame be sent to the fabric.

8.9.4.2.4 Support for Hybrid Simplex Operation

The SPI-4.2 interface may be used to transfer data to and from a fabric, although there is no standard protocol for such conveyance. The necessary addressing information for the fabric and egress network processor may be encoded within the address bits of the preceding control word or stored in the initial data words of the SPI-4.2 burst. The LVTTTL status interface may be used to provide link-level flow control for the data bursts. (The SPI-4.2 LVDS status interface cannot be used, because it shares the same pins with the fabric flow control interface.)

Figure 115. Reference Model for Hybrid Simplex Operation





The SPI-4.2 interface does not support a virtual output queue (VOQ) flow control mechanism. The Intel® IXP2800 Network Processor supports use of the CSIX-L1 protocol-based flow control interface (as used in the dual chip, full-duplex configuration) on the ingress network processor, while SPI-4.2 is operational on the data interface. This interface can provide VOQ flow control information from the fabric and allow the transmitter scheduler, implemented in a Microengine within the ingress network processor, to avoid sending data bursts to congested destinations.

The fabric should send alternating Idle CFrames and Dead Cycles when there are no Flow Control CFrames to transmit. The CRdy and DRdy “ready bits” should be set to 0 on transmission and are ignored on reception.

The fabric should respond to the RXCFC backpressure signal. In this mode of operation, the RXCSR signal that would normally receive the state of the CRdy and DRdy “ready bits” is not used. If dynamic de-skew is configured on the interface, and the backpressure signal is asserted for 32 clock cycles, the fabric sends a (de-skew) training sequence on the flow control interface. It may be acceptable in this configuration to operate the flow control interface at a sufficiently low clock rate that dynamic de-skew is not required.

Operation in the hybrid simplex mode for the ingress network processor is slightly more taxing on the transmit scheduler computation than the homogenous CSIX-L1 protocol configurations. The status reported for the data interface must be polled by the transmit scheduler. In this configuration, the response to link-level flow control is performed in software and is slower than in the homogenous CSIX-L1 protocol configurations where it is accomplished in hardware.

8.9.4.2.5 Support for Dynamic De-Skew Training

The SPI-4.2 interface incorporates a training sequence for dynamic de-skew of its signals relative to the source synchronous clock. This training sequence has been extended and incorporated into the CSIX-L1 protocol support of the Intel® IXP2800 Network Processor.

The training pattern for the 16-bit data interface consists of 20 words, 10 repetitions of 0x0fff followed by 10 repetitions of 0xf000. The CTL and PAR signals are asserted for the first 10 words and de-asserted for the second 10 words. The PROT signal (see below) is de-asserted for the first 10 words and asserted for the second 10 words. A training sequence consists of “alpha” repetitions of the training pattern. The idle control word that precedes a training sequence in SPI-4.2 is not used in conjunction with the CSIX-L1 protocol. See [Section 8.6.1](#) for more information.

A receiver should detect a training sequence in the context of the CSIX-L1 protocol implementation by the assertion of the start-of-frame signal for three adjacent clock edges and the correct value on the data signals for those three adjacent clock edges.

A receiver may request a training sequence to be sent by transmitting continuous Dead Cycles on the interface. Reception of two adjacent Dead Cycles triggers the transmission of a training sequence in the opposite direction. If an interface is sending Dead Cycles and a training sequence becomes pending, the interface must send the training sequence at a higher priority than the Dead Cycles. Otherwise, a deadlocked situation may arise.

In the simplex configuration, the request for training, and the response to it, occur between a primary interface and its associated reverse path control interface. In the dual chip, full-duplex configuration, requests for training and Dead Cycles are encoded across the flow control interface as either continuous Dead Cycles or continuous Idle CFrames, both of which violate the standard CSIX-L1 protocol.

The training pattern for the flow control data signals consists of 10 nibbles of 0xc followed by 10 nibbles of 0x3. The parity and serial “ready bits” signal is de-asserted for the first 10 nibbles and asserted for the second 10 nibbles. The start-of-frame signal is asserted for the first 10 nibbles and de-asserted for the second 10 nibbles. See [Section 8.6.2](#) for more information.

When a training sequence is received, the receiver should update the state of the received CRdy and DRdy “ready bits” to a de-asserted state until they are updated by a subsequent CFrame.

8.9.4.3 CSIX-L1 Protocol Receiver Support

The Intel® IXP2800 Network Processor receiver support for the CSIX-L1 protocol is similar to that for SPI-4.2. CFrames are stored in the receiver data buffers. The buffers are configured to be of a size of 64, 128, or 256 bytes. The contents of the CFrame base header and extension header are stored in separate storage with the reception status of the CFrame. Unlike SPI-4.2 data bursts, the entire CFrame must fit into a single buffer. The receiver does not progress to the next buffer to store subsequent parts of a single CFrame. (The buffer is required only to be sufficiently large to accommodate the payload, not the header, the padding, or the vertical parity.) Designated CFrame types, typically Flow Control CFrames, are forwarded in cut-through mode directly to the flow control egress FIFO and not stored in the receiver buffers.

The receiver resources are separately allocated to the processing of data and control CFrames. Separate free lists of buffers and Microengine threads for each category of CFrame type are maintained. The size of the buffers in each resource pool is separately configurable. The mapping of CFrame type to data or control category is completely configurable via the CSIX_Type_Map register. This register also allows for any types to be designated for cut-through forwarding to the flow control egress FIFO. Typically, only the Flow Control CFrame type is configured in this way.

The receiver buffers are partitioned into two pools via MSF_Rx_Control[RBUF_Partition], providing 75% of the buffer memory (6 Kbytes) for data CFrames and 25% of the buffer memory (2 Kbytes) for control CFrames. The number of buffers available per pool depends on the configured buffer size. For 64-byte buffers, there are 96 and 32 buffers, respectively. For 128-byte buffers, there are 48 and 16 buffers, respectively. For 256-byte buffers, there are 24 and 8 buffers, respectively.

As with SPI-4.2, link-level flow control for a buffer pool can be asserted by configuration when buffer consumption reaches 25%, 50%, 75%, or 87.5% within that pool. The receiver has an additional 1024 bytes of packed FIFO storage for each traffic category to accept additional CFrames after link-level flow control (CRdy or DRdy) is asserted. Link-level flow control for control CFrames (CRdy) is also asserted if the flow-control egress FIFO contents exceeds a threshold as configured by HWM_Control[FCEFIFO_HWM]. The threshold may be set to 16, 32, 64, or 128 32-bit words. The total capacity of the FIFO is 512 32-bit words.

Within the base header, the receiver hardware processes the CRdy bit, the DRdy bit, the Type field, and the Payload Length. *Only the Flow Control Frame CFrame is expected to lack the 32-bit extension header.* The receiver hardware validates the vertical parity of the CFrame and only writes it to the receiver buffer if the write operation also includes payload data. The hardware supports configuration options for processing all 16 CFrame types. In all other respects, processing of the CFrame contents is done entirely by software. Variations in the CSIX-L1 protocol are supported that only affect the software processing. These variations might include address swapping (egress port address swapping with ingress port address) and use of reserve bits to encode start and end of packets.

When the network processor is configured to forward Flow Control Frame CFrames to the flow control egress FIFO, software does not process those CFrames. Processor interrupts occur if there are reception errors, but the actual CFrames are not made available for further processing.

8.9.4.4 CSIX-L1 Protocol Transmitter Support

The Intel® IXP2800 Network Processor transmitter support for the CSIX-L1 protocol is similar to that for SPI-4.2. The transmitter fetches CFrames from transmitter buffers. An entire CFrame must fit within a single buffer. In the case of SPI-4.2, the array of transmitter buffers operates as a single ring. In the case of CSIX-L1 protocol support, the array of buffers operates as two rings, one for data CFrames and another for control CFrames. The partitioning of the transmitter buffers is configured via `MSF_Tx_Control[TBUF_Partition]`. The portion of the aggregate transmitter buffer storage (8 Kbytes) allocated to data CFrames is 75% (6 Kbytes), with the remainder (2 Kbytes) allocated to control CFrames. The size of the buffers within each partition is independently configurable to a size of 64, 128, or 256 bytes. The payload size of CFrames sent from the buffers may vary from 1 to the size of the buffer.

The CSIX-L1 protocol link-level flow control operates directly upon the hardware that processes the two (control and data) transmitter rings. The transmitter services the two rings in round-robin order when allowed by link-level flow control. The transmitter transmits Idle CFrames and Dead Cycles according to the CSIX-L1 protocol if there are no CFrames to transmit.

Virtual output queue flow control is accommodated by a transmit scheduler implemented on a Microengine. In all three network processor ingress configurations, Flow Control CFrames are loaded by hardware into the flow control ingress FIFO. Two state bits associated with this FIFO are distributed to all of the Microengines: (1) the FIFO is non-empty, and (2) the FIFO contains more than a threshold amount of CFrame 32-bit words (`HWM_Control[FCIFIFO_Int_HWM]`)

Any Microengine can perform transmitter scheduling by sensing the state associated with the flow control ingress FIFO, using the branch-on-state instruction. If the FIFO is not empty, the transmit scheduler processes some of the FIFO by performing a read of the FCIFIFO registers.

A single Microengine instruction can perform a block read of up to 16 32-bit words. The data for the read is likely to arrive after several subsequent scheduling decisions. The scheduler should incorporate the new information from the newly-read Flow Control CFrame(s) in its later scheduling decisions. If the FIFO state indicates that the threshold capacity has been exceeded, the scheduler should suspend further scheduling decisions until the FIFO is sufficiently processed, otherwise it risks making scheduling decisions with information that is stale.

The responsiveness of the network processor to VOQ flow control depends on the transmit pipeline length, from transmit scheduler to CFrames on the interface signals. For rates at or above 10 Gb/s, the pipeline length is likely to be 32 – 64 CFrames, assuming four pipeline stages (schedule, de-queue, data movement, and transmit) and 8 – 16 CFrames concurrently processed per stage.

In the simplex configuration, the egress network processor can send CFrames over the Reverse Path Control Interface. The CFrames are loaded into the flow control egress FIFO by performing writes of 32-bit words to the FCEFIFO registers. The base header, the extension header, the payload, the padding, and a dummy vertical parity must be written to the FIFO. The transmitter hardware calculates the actual vertical parity as the CFrame is transmitted.

Note: The transmitter hardware for the transmitter buffers and the flow control egress FIFO expect that *only* the Flow Control CFrame type does not have an extension header of 32 bits (all other types have this header). The hardware disregards the contents of the extension header or the payload.

The limited gather capability described for SPI-4.2 also is available for CFrames. A prefix header of up to 31 bytes and a disjoint payload is supported. The prefix header may start at an offset of 0 to 7 bytes. The payload may start at an offset of 0 to 7 bytes from the octal-byte boundary following the end of the prefix header. For more complicated merging or shifting of data within a CFrame, the data should be passed through a Microengine to perform any arbitrary merging and/or shifting.



8.9.4.5 Implementation of a Bridge Chip to CSIX-L1

The Intel® IXP2800 Network Processor support for the CSIX-L1 protocol in the dual chip, full-duplex configuration minimizes the difficulty in implementing a bridge chip to a standard CSIX-L1 interface. If dynamic de-skew training is not employed, the bridge chip can directly pass through the different CSIX-L1 protocol elements, CFrames, and Dead Cycles. The horizontal parity must be recalculated on each side of the bridge chip. If the standard CSIX-L1 interface implements a CWord width that is greater than 32 bits, it must implement a synchronization mechanism for aligning the received 32-bit portions of the CWord before passing the CWord to the network processor.

For transmitting the standard CSIX-L1 interface, the bridge chip must assert the start-of-frame signal for each 32-bit portion of the CWord, as the network processor only asserts it for the first 32-bit portion. If the bridge chip requires clock frequencies on the network processor interface and the standard CSIX-L1 interface to be appropriate, exact multiples of each other (2x for 32-bit CWord, 4x for 64-bit CWord, 6x for 96-bit CWord, and 8x for 128-bit CWord), then the bridge chip requires only minimal buffering and does not need to implement any flow control mechanisms.

A slightly more complicated bridge allows incorporating dynamic de-skew training and/or independent clock frequencies for the network processor and standard CSIX-L1 interfaces. The bridge chip must implement a control and data FIFO for each direction and the link-level flow control mechanisms specified in the protocol using CRdy and DRdy. The FIFOs must be large enough to accommodate the response latency of the link-level flow control mechanisms.

Idle CFrames and Dead Cycles are not directly passed through this more complicated bridge chip, but are discarded on reception and generated on transmission. The network processor interface of this bridge chip can support the dynamic de-skew training protocol extensions implemented on the network processor because it can send a training sequence to the network processor between CFrames without regard to CFrames arriving over the standard CSIX-L1 interface. (In the simpler bridge design, these CFrames must be forwarded immediately to the network processor.)



8.9.5 Dual Protocol (SPI and CSIX-L1) Support

In many system designs that are less bandwidth-intensive, a single network processor can forward and process data from the framer to the fabric and from the fabric to the framer. A bridge chip must pass data between the network processor and multiple physical devices. The network processor supports multiplexing SPI-4.2 and CSIX-L1 protocol elements over the same transmitter and receiver physical interfaces, differentiated by a protocol signal that is de-asserted for SPI-4.2 protocol elements and asserted for CSIX-L1 protocol elements.

In the dual protocol configuration, the CSIX-L1 configuration of the network processor corresponds to the dual chip, full duplex configuration. The flow control transmitter interface is looped back to the flow control receiver interface, either externally or internally. Only the LVTTTL status interface is available for the SPI-4.2 interface.

8.9.5.1 Dual Protocol Receiver Support

When the network processor receiver is configured for dual protocol support, the aggregate receiver buffer is partitioned in three ways: 50% for data CFrames (4 Kbytes), 37.5% for SPI-4.2 bursts (3 Kbytes) and 12.5% for control CFrames (1 Kbyte). The buffer sizes within each partition are independently configurable. Link-level flow control can be independently configured for assertion at thresholds of 25%, 50%, 75%, or 87.5%. For the traffic associated with each partition, an additional 680 bytes of packed FIFO storage is available to accommodate received traffic after assertion of link-level flow control.

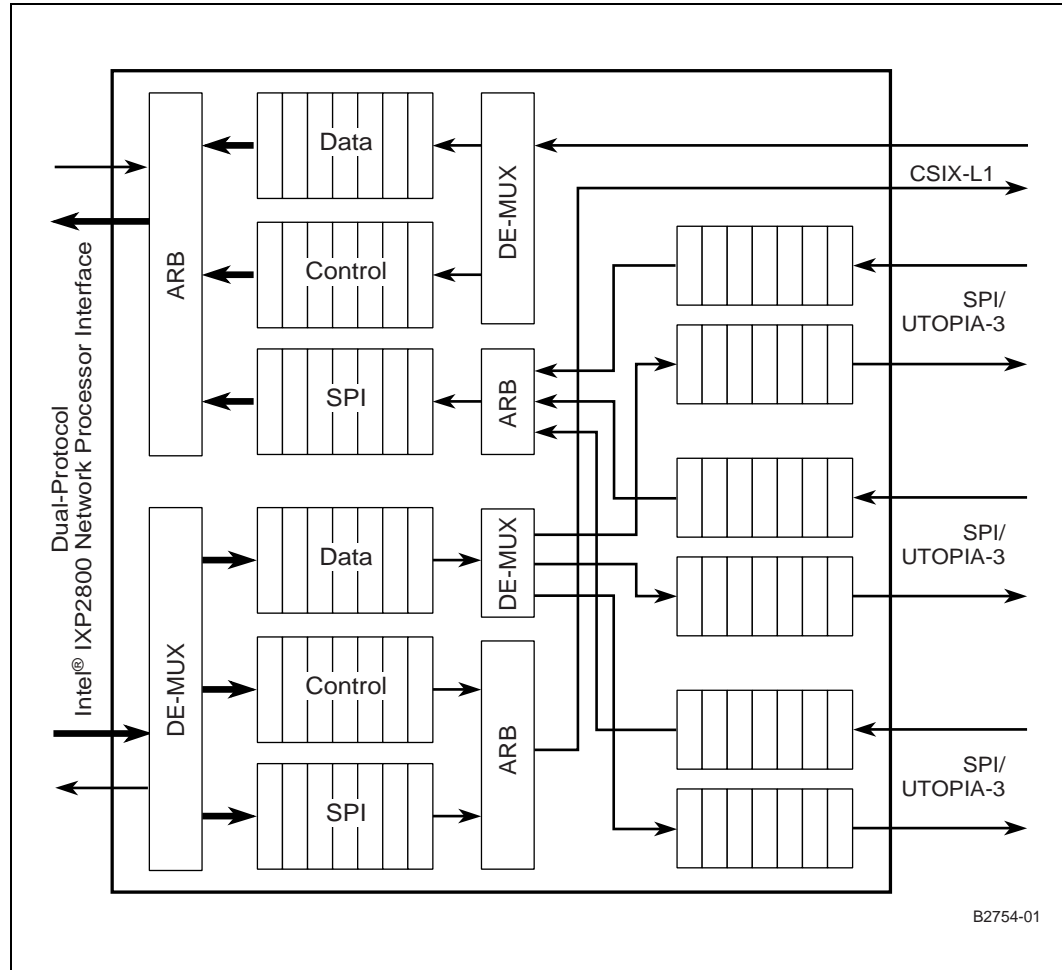
8.9.5.2 Dual Protocol Transmitter Support

When the network processor transmitter is configured for dual protocol support, the aggregate transmitter buffer is partitioned three ways, in the same proportions as the receiver. Each partition operates as a separate ring. The transmitter services each ring in round-robin order. If no CFrames are pending, an Idle CFrame is transmitted to update link-level flow control. If no SPI-4.2 bursts are pending, idle control words are not sent.

8.9.5.3 Implementation of a Bridge Chip to CSIX-L1 and SPI-4.2

A bridge chip can provide support for both standard CSIX-L1 and standard physical layer device interfaces such as SPI-3 or UTOPIA Level 3. The bridge chip must implement the functionality of the less trivial CSIX-L1 bridge chip described previously and additionally, implement bridge functionality between SPI-4.2 and the other physical device interfaces. The size of the FIFOs must be in accordance with the response times of the flow control mechanisms. Figure 116 is a block diagram of a dual protocol (SPI-4.2 and CSIX-L1) bridge chip.

Figure 116. Block Diagram of Dual Protocol (SPI-4.2 and CSIX-L1) Bridge Chip





8.9.6 Transmit State Machine

Table 114 describes the transmitter state machine by providing guidance in interfacing to the network processor. The state machine is described as three separate state machines for SPI-4.2, training, and CSIX-L1. When each machine is inactive, it tracks the states of the other two state machines.

8.9.6.1 SPI-4.2 Transmitter State Machine

The SPI-4.2 Transmit State Machine makes state transitions on each bus transfer of 16 bits, as described in Table 114.

Table 114. SPI-4.2 Transmitter State Machine Transitions on 16-Bit Bus Transfers

Current State	Next State	Conditions
Idle Control	Idle Control	No data pending and no training sequence pending, CSIX-L1 mode disabled.
	Payload Control	Data pending and no training sequence pending, CSIX-L1 mode disabled.
	Training	Training sequence pending, CSIX-L1 mode disabled.
	CSIX	CSIX-L1 mode enabled.
Payload Control	Data Burst	Always
Data Burst	Data Burst	Until end of burst as programmed by software.
	Payload Control	Data pending and no training sequence pending and CSIX-L1 mode not enabled.
	Idle Control	No data to send or training sequence pending or CSIX-L1 mode enabled.
Tracking Other State Machine States		
Training	Training	Training SM not entering CSIX-L1 or SPI state.
	CSIX	Training SM entering CSIX-L1 state.
	Payload Control	Training SM entering SPI state and data pending.
	Idle Control	Training SM entering SPI state and no data pending.
CSIX	CSIX	CSIX-L1 SM not entering Training or SPI state.
	Training	CSIX-L1 SM entering Training state.
	Payload Control	CSIX-L1 SM entering SPI state and data pending.
	Idle Control	CSIX-L1 SM entering SPI state and no data pending.



8.9.6.2 Training Transmitter State Machine

The Training State Machine makes state transitions on each bus transfer of 16 bits, as described in Table 115.

Table 115. Training Transmitter State Machine Transitions on 16-Bit Bus Transfers

Current State	Next State	Conditions
Training Control	Training Control	Until 10 control cycles.
	Training Data	After 10 control cycles.
Training Data	Training Data	Until 10 data cycles.
	Training Control	After 10 data cycles and repetitions of training sequence or new training sequence pending.
	CSIX	After 10 data cycles and no training sequence pending and CSIX-L1 mode enabled.
	SPI	After 10 data cycles and No training sequence pending and CSIX-L1 mode disabled.
Tracking Other State Machine States		
CSIX	CSIX	CSIX-L1 SM not entering SPI or Training state.
	SPI	CSIX-L1 SM entering SPI state.
	Training Control	CSIX-L1 SM entering Training state.
SPI	SPI	SPI SM not entering CSIX-L1 or Training state.
	CSIX	SPI SM entering CSIX-L1 state.
	Training Control	SPI SM entering Training state.

8.9.6.3 CSIX-L1 Transmitter State Machine

The CSIX-L1 Transmit State Machine makes state transitions on CWord boundaries. CWords can be configured to consist of 32, 64, 96, or 128 bits, corresponding to 2, 4, 6, or 8 bus transfers, as described in Table 116.

Table 116. CSIX-L1 Transmitter State Machine Transitions on CWord Boundaries (Sheet 1 of 2)

Current State	Next State	Conditions
SoF CWord	CFrame CWord	CFrame longer than a CWord.
	Dead Cycle	CFrame fits in a CWord.
CFrame CWord	CFrame CWord	CFrame remainder pending.
	SoF CWord	Un-flow-controlled CFrame pending, no training sequence pending, and SPI mode not enabled.
	Dead Cycle	No un-flow-controlled CFrame pending or training sequence pending or requesting training sequence or SPI mode enabled and data pending.
Dead Cycle	SoF CWord	Un-flow-controlled CFrame pending and no training sequence pending and no SPI data pending and not requesting training sequence.
	Idle CFrame	No un-flow-controlled CFrame pending and no training sequence pending and no SPI data pending and not requesting training sequence.



Table 116. CSIX-L1 Transmitter State Machine Transitions on CWord Boundaries (Sheet 2 of 2)

Current State	Next State	Conditions
	Dead Cycle	Requesting reception of training sequence and no training sequence pending.
	Training	Training sequence pending.
	SPI	Training sequence not pending and SPI data pending and not requesting training sequence.
Idle CFrame	Dead Cycle	Always.
Tracking Other State Machine States		
SPI	SPI	SPI SM not entering CSIX-L1 or Training state.
	SoF CWord	SPI SM entering CSIX-L1 state and un-flow-controlled CFrame pending.
	Idle CFrame	SPI SM entering CSIX-L1 state and un-flow-controlled CFrame not pending.
	Training	SPI SM entering Training state.
Training	Training	Training SM not entering CSIX-L1 or Training state.
	SoF CWord	Training SM entering CSIX-L1 state and un-flow-controlled CFrame pending.
	Idle CFrame	Training SM entering CSIX-L1 state and un-flow-controlled CFrame not pending.
	SPI	Training SM entering SPI state.

8.9.7 Dynamic De-Skew

The Intel® IXP2800 Network Processor supports optional dynamic de-skew for the signals of the 16-bit data interface and the signals of the 4-bit flow control interface or the signals of the 2-bit SPI-4.2 LVDS status interface. (The flow control interface and the LVDS status interface are alternate configurations of the same signal balls and pads. They share the same de-skew circuits.)

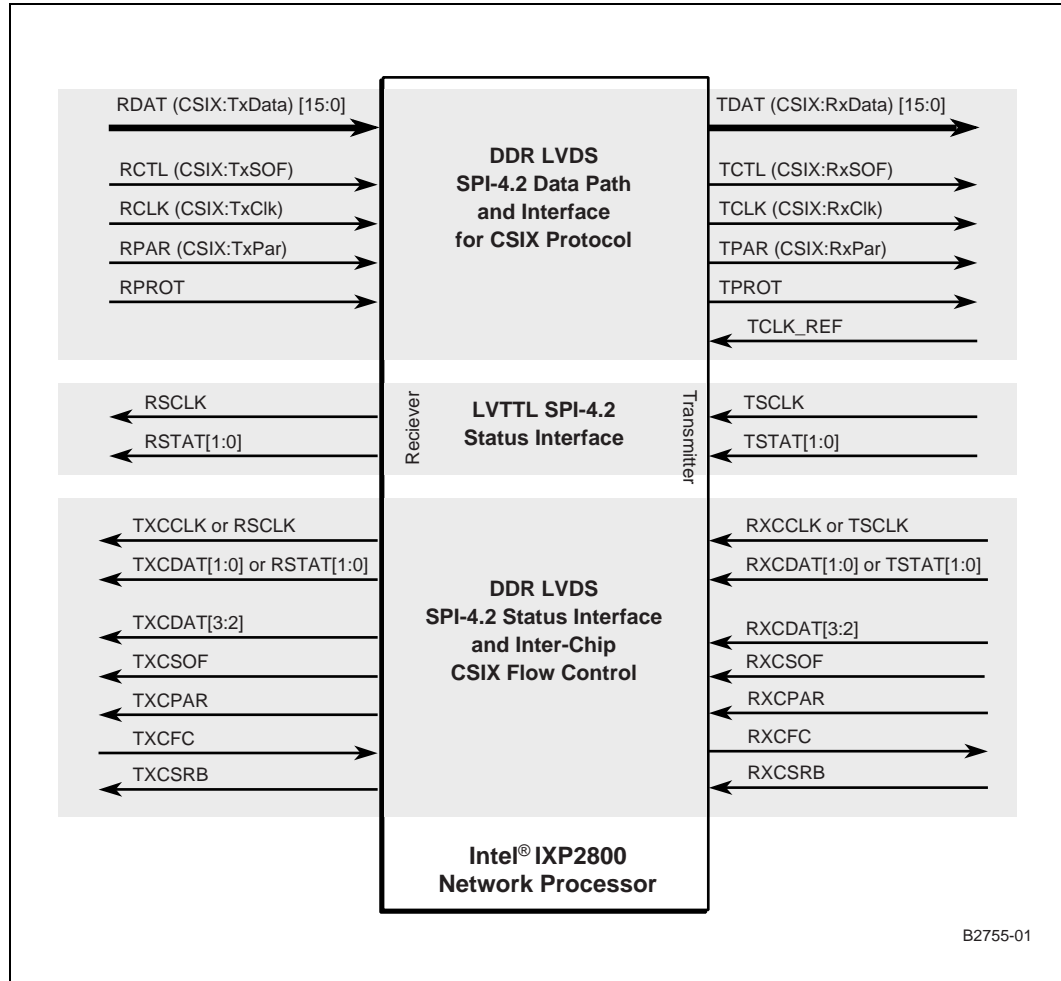
In both cases, eight evenly-spaced phases of the received clock are generated for each bit time. As the transition occurs during training a pattern, the best pair of clock phases is identified for sampling each received signal. An interpolated clock is generated from a pair of clock phases for each signal and that clock is used as a reference for sampling the data. This provides maximum quantization error in the sampling of the signals of 6.25%.



8.9.8 Summary of Receiver and Transmitter Signals

Figure 117 summarizes the Receiver and Transmitter Signals.

Figure 117. Summary of Receiver and Transmitter Signaling





This section contains information on the IXP2800 Network Processor PCI Unit.

9.1 Overview

The PCI Unit allows PCI target transactions to internal registers, SRAM, and DRAM. It also generates PCI initiator transactions from the DMA Engine, Intel XScale® core, and Microengines.

The PCI Unit main functional blocks are shown in [Figure 118](#) and include:

- PCI Core Logic
- PCI Bus Arbiter
- DRAM Interface Logic
- SRAM Interface Logic
- Mailbox and Message registers
- DMA Engine
- Intel XScale® core Direct Access to PCI

The main function of the PCI Unit is to transfer data between the PCI Bus and the internal devices, which are the Intel XScale® core, the internal registers, and memories.

These are the data transfer paths supported as shown in [Figure 119](#):

- PCI Slave read and write between PCI and internal buses
 - CSRs (PCI_CSR_BAR)
 - SRAM (PCI_SRAM_BAR)
 - DRAM (PCI_DRAM_BAR)
- Push/Pull Master (Intel XScale® core, Microengine, or PCI) accesses to internal registers within PCI unit
- DMA
 - Descriptor read from SRAM
 - Data transfers between PCI and DRAM
- Push/Pull Master (Intel XScale® core and Microengines) direct read and write to PCI Bus

Note: Detailed information about CSRs is contained in the *Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

Figure 118. PCI Functional Blocks

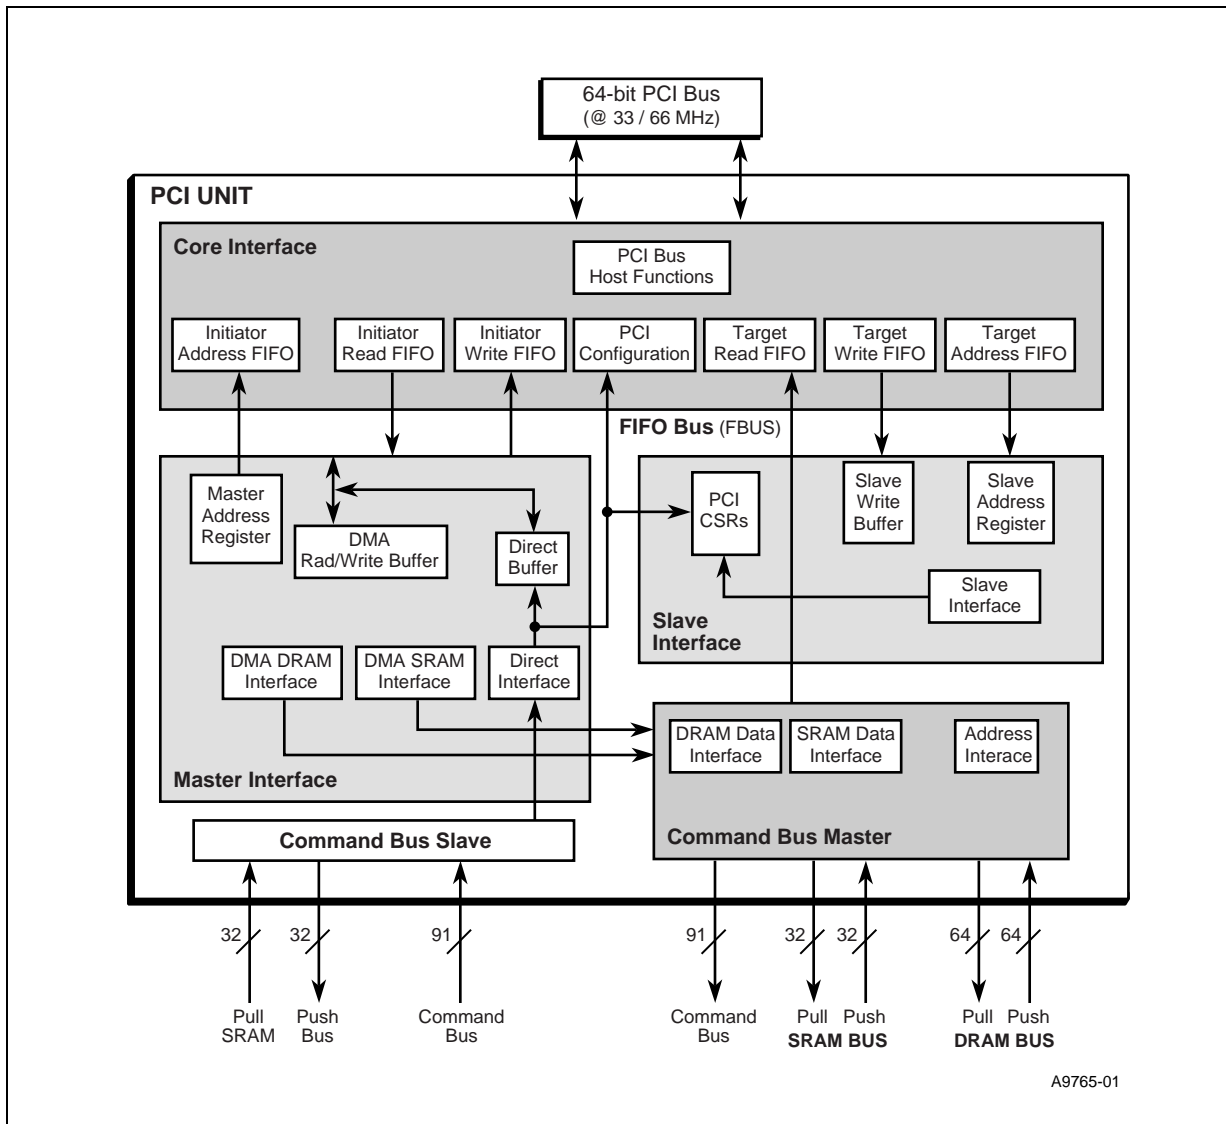





Figure 119. Data Access Paths



9.2 PCI Pin Protocol Interface Block

This block generates the PCI compliant protocol logic. It operates either as an initiator or a target device on the PCI Bus. As an initiator, all bus cycles are generated by the core. As a PCI target, the core responds to bus cycles that have been directed towards it.

On the PCI Bus, the interface supports interrupts, 64-bit data path, 32-bit addressing, and single configuration space. The local configuration registers are accessible from the PCI Bus or from the Intel XScale® core through an internal path.

The PCI block interfaces with the other sub-blocks with a FIFO bus called FBus. The FBus speed is the same as the internal Push/Pull bus speed. The FIFOs are implemented with clock synchronization logic between the PCI speed and the internal Push/Pull bus speed.

There are four data FIFOs and two address FIFOs in the core. The separate slave and master data FIFOs allows simultaneous operations and multiple outstanding PCI bus transfers. Table 117 lists the FIFO sizes. The target address FIFO latches up to four PCI read or write addresses.



If a read address is latched, the subsequent cycles will be retried and no address will be latched until the read completes. The initiator address FIFO can accumulate up to four addresses that can be PCI reads or writes.

These FIFOs are inside the PCI Core, which stores data received from the PCI Bus or data to be sent out to the PCI Bus. There are additional buffers implemented in other sub-blocks that buffers data to and from the internal push/pull buses.

Table 117. PCI Block FIFO Sizes

Location	Depth
Target Address	4
Target Write Data	8
Target Read Data	8
Initiator Address	4
Initiator Write Data	8
Initiator Read Data	8

Table 118 lists the maximum PCI Interface loading.

Table 118. Maximum Loading

Bus Interface	Maximum Number of Loads	Trace Length (inches)
PCI	Four loads at 66-MHz bus frequency Eight loads at 33-MHz bus frequency	5 to 7

9.2.1 PCI Commands

Table 119 lists the supported PCI commands and identifies them as either a target or initiator.

Table 119. PCI Commands (Sheet 1 of 2)

C_BE_L	Command	Support	
		Target	Initiator
0x0	Interrupt Acknowledge	Not Supported	Supported
0x1	Special Cycle	Not Supported	Supported
0x2	IO Read cycle	Not Supported	Supported
0x3	IO Write cycle	Not Supported	Supported
0x4	Reserved	–	–
0x5	Reserved	–	–
0x6	Memory Read	Supported	Supported
0x7	Memory Write	Supported	Supported
0x8	Reserved	–	–
0x9	Reserved	–	–
0xA	Configuration Read	Supported	Supported
0xB	Configuration Write	Supported	Supported



Table 119. PCI Commands (Sheet 2 of 2)

C_BE_L	Command	Support	
		Target	Initiator
0xC	Memory Read Multiple	Aliased as Memory Read except SRAM accesses where the number of Dwords to read is given by the cache line size.	Supported
0xD	Reserved	—	—
0xE	Memory read line	Aliased as Memory Read except SRAM accesses where the number of Dwords to read is given by the cache line size.	Supported
0xF	Memory Write and Invalidate	Aliased as Memory Write.	Not Supported

PCI functions not supported by the PCI Unit include:

- IO Space response as a target
- Cacheable memory
- VGA palette snooping
- PCI Lock Cycle
- Multi-function devices
- Dual Address cycle

9.2.2 IXP2800 Network Processor Initialization

When the IXP2800 Network Processor is a target, the internal CSR, DRAM, or SRAM address is generated when the PCI address matches the appropriate base address register. The window sizes to the SRAM and DRAM Base Address Registers (BARs) can be optionally set by PCI_SWIN and PCI_DWIN strap pins or mask registers depending on the state of the PROM_BOOT signal.

There are two initialization modes supported. They are determined by the PROM_BOOT signal sampled on the de-assertion edge of Chip Reset. If PROM_BOOT is asserted, then there is a boot prom in the system. The Intel XScale® core will boot from the prom and be able to program the BAR space mask registers. If PROM_BOOT is not asserted, the Intel XScale® core is held in reset and the BAR sizes are determined by strap pins.



9.2.2.1 Initialization by the Intel XScale® Core

The PCI unit is initialized to an inactive, disabled state until the Intel XScale® core has set the Initialize Complete bit in the Control register. This bit is set after the Intel XScale® core has initialized the various PCI base address and mask registers (which should occur within 1 ms of the end of PCI_RESET). The mask registers are used to initialize the PCI base address registers to values other than the default power-up values, which includes the base address visible to the PCI host and the prefetchable bit in the base registers (see [Table 120](#)).

Table 120. PCI BAR Programmable Sizes

Base Address Register	Address Space	Sizes
PCI_CSR_BAR	CSR	1 Mbyte
PCI_SRAM_BAR	SRAM	0 Bytes; 128, 256, or 512 Kbytes; 1, 2, 4, 8, 16, 32, 64, 128, or 256 Mbytes
PCI_DRAM_BAR	DRAM	0 Bytes; 1, 2, 4, 8, 16, 32, 64, 128, 256, or 512 Mbytes; 1 Gbyte

When the PCI unit is in the inactive state, it returns retry responses as the target of PCI configuration cycles if the PCI Unit is not configured as the PCI host. In the case of PCI Unit being configured as the PCI host, the PCI bus will be held in reset until the Intel XScale® core completes the PCI Bus configurations and clears the PCI Reset (as described in [Section 9.2.11](#)).

Note: During PCI bus enumeration initiated by the Intel XScale® core, reading a non-existent address (an address for which no target asserts DEVSEL) results in a Master Abort. The Master Abort then results in an Intel XScale® Core Data Abort Exception that must be handled by the enumeration software. When this occurs, the RMA bit in the PCI_CONTROL register and the RX_MA bit in the PCI_CMD_STAT register is set. The enumeration software must then clear these bits before continuing with the enumeration process.

9.2.2.2 Initialization by a PCI Host

In this mode, the PCI Unit is not hosting the PCI Bus regardless of the PCI_CFG[0] signal. The host processor is allowed to configure the internal CSRs while the Intel XScale® core is held in reset. The host processor configures the PCI address space, the memory controllers, and other interfaces. Also, the program code for the Intel XScale® core may be downloaded into local memory.

The host processor then clears the Intel XScale® core reset bit in the PCI Reset register. This de-asserts the internal reset signal to the Intel XScale® core and the core begins its initialization process. The PCI_SWIN and PCI_DWIN strap signals are used to select the window sizes to SRAM BAR and DRAM BAR (see [Table 121](#)).

Table 121. PCI BAR Sizes with PCI Host Initialization

Base Address Register	Address Space	Sizes
PCI_CSR_BAR	CSR	1 Mbyte
PCI_SRAM_BAR	SRAM	32, 64, 128, or 256 Mbytes
PCI_DRAM_BAR	DRAM	128, 256, or 512 Mbytes; 1 Gbyte



9.2.3 PCI Type 0 Configuration Cycles

A PCI access to a configuration register occurs when the following conditions are satisfied:

- PCI_IDSEL is asserted. (PCI_IDSEL only supports PCI_AD[23:16] bits).
- The PCI command is a configuration write or read.
- The PCI_AD [1:0] are 00.

A configuration register is selected by PCI_AD[7:2]. If the PCI master attempts to do a burst longer than one 32-bit Dword, the PCI unit signals a target disconnect. PCI unit does not issue PCI_ACK64 for configuration cycle.

9.2.3.1 Configuration Write

A *write* occurs if the PCI command is a Configuration Write. The PCI byte-enables determine which bytes are written. If a nonexistent configuration register is selected within the configuration register address range, the data is discarded and no error action is taken.

9.2.3.2 Configuration Read

A *read* occurs if the PCI command is a Configuration Read. The data from the configuration register selected by PCI_AD[7:2] is returned on PCI_AD[31:0]. If a nonexistent configuration register is selected within the configuration register address range, the data returned are zeros and no error action is taken.

9.2.4 PCI 64-Bit Bus Extension

The PCI Unit is in 64-bit mode when PCI_REQ64_L is sampled active on the de-assertion edge of PCI Reset. These are the general rules in assertions of PCI_REQ64_L and PCI_ACK64_L:

As a target:

1. PCI Unit asserts PCI_ACK64_L only in 64-bit mode.
2. PCI Unit asserts PCI_ACK64_L only to target cycles that matches the PCI_SRAM_BAR and PCI_DRAM_BAR and a 64-bit transaction is negotiated.
3. PCI Unit does not assert PCI_ACK64_L target cycles that matches the PCI_CSR_BAR even a 64-bit transaction is negotiated.

As an initiator:

1. PCI Unit asserts PCI_REQ64_L only in 64-bit mode.
2. PCI Unit asserts PCI_REQ64_L to negotiate a 64-bit transaction only if the address is double Dword aligned (PCI_AD[2] must be 0 during the address phase).
3. If the target responds to PCI_REQ64_L with PCI_ACK64_L de-asserted, PCI Unit will complete the transaction acting as a 32-bit master by not asserting PCI_REQ64_L on subsequent cycle.
4. If the target responds to PCI_REQ64_L with PCI_ACK64_L de-asserted and PCI_STOP_L asserted, PCI Unit will complete the transaction by not asserting PCI_REQ64_L on subsequent cycles.



9.2.5 PCI Target Cycles

The following PCI transactions are not supported by the PCI Unit as a target:

- IO read or write
- Type 1 configuration read or write
- Special cycle
- IACK cycle
- PCI Lock cycle
- Multi-function devices
- Dual Address cycle

9.2.5.1 PCI Accesses to CSR

A PCI access to a CSR occurs if the PCI address matches the CSR base address register (PCI_CSR_BAR). The PCI Bus will be disconnected after the first data-phase if the data is more than one data phase. For 64-bit CSR accesses, the PCI Unit will not assert PCI_ACK64_L on the PCI bus.

9.2.5.2 PCI Accesses to DRAM

A PCI access to DRAM occurs if the PCI address matches the DRAM base address register (PCI_DRAM_BAR).

9.2.5.3 PCI Accesses to SRAM

A PCI access to SRAM occurs if the PCI address matches the SRAM base address register (PCI_SRAM_BAR). The SRAM is organized as three distinct channel and the address is not contiguous. The PCI_SRAM_BAR programmed window size will be used as the total memory space. The upper two bits of the address will be used as channel number in addressing the particular channel and the remaining address bits will be used as the memory address.

9.2.5.4 Target Write Accesses from the PCI Bus

A PCI write occurs if the PCI address matches one of the base address registers and the PCI command is either a Memory Write or Memory Write and Invalidate. The core will store up to four write addresses into the target address FIFO along with the BAR IDs of the transaction. The write data will be stored into the target write FIFO. When either the address FIFO or data FIFO is full, a retry is forced on the PCI Bus in response to write accesses.

The FIFO data is forwarded to an internal slave buffer before being written into SRAM or DRAM. If the FIFO fills during the write, the address is crossing the 64-byte address boundary, or in the case of the command being a burst to the CSR space, the PCI unit signals target disconnect to the PCI master.



9.2.5.5 Target Read Accesses from the PCI Bus

A PCI read occurs if the PCI address matches one of the base address registers and the PCI command is either a Memory Read, Memory Read Line, or Memory Read Multiple.

The read is completed as a PCI delayed read. That is, on the first occurrence of the read, the PCI unit signals a retry to the PCI master. If there is no prior read pending, the PCI unit latches the address and command and places it into the target address FIFO. When the address reaches the head of the FIFO, the PCI unit reads the DRAM. Subsequent reads will also get retry responses until data is available.

When the read data is returned into the PCI Read FIFO, the PCI unit begins to decrement its discard timer. If the PCI bus master has not repeated the read by the time the timer reaches 0, the PCI unit discards the read data, invalidates the delayed read address and sets Discard Timer Expired (bit 16) in the Control register (PCI_CONTROL). If enabled, the PCI unit interrupts the Intel XScale® core. The discard timer counts 2^{15} (32,768) PCI clocks.

When the master repeats the read command, the PCI unit compares the address and checks that the command is a Memory Read, a Memory Read Line, or a Memory Read Multiple. If there is a match, the response is as follows:

- If the read data has not yet been read, the response is retry.
- If the read data has been read, assert `trdy_1` and deliver the data. If the master attempts to continue the burst past the amount of data read, the PCI unit signals a target disconnect.
- CSR reads are always 32-bit reads.
- If the discard timer has expired for a read, the subsequent read will be treated as a new read.

9.2.6 PCI Initiator Transactions

PCI master transactions are caused by either the Intel XScale® core loads and stores that fall into the various PCI address spaces, Microengine read and write commands, or by the DMA engine. The command register (PCI_COMMAND) bus master bit (BUS_MASTER) must be set for the PCI unit to perform any of the initiator transactions.

The PCI cycle is initiated when there is an entry in the PCI Core Interface initiator address FIFO. The core handshakes with the master interface with the FBus FIFO status signals. The PCI core supports both burst and non-burst master read transfers by the burst count inputs (FB_BstCnt[7:0]), driven by Master Interface to inform the core the burst size. For a Master write, FB_WBstonN indicates to the PCI core whether the transfers are burst or non-burst, on a 64-bit double Dword basis.

The PCI core supports read and write memory cycles as an initiator while taking care of all disconnect/retry situations on the PCI Bus.

9.2.6.1 PCI Request Operation

If an external arbiter is used (PCI_CFG_ARB[1] is not active), the `req_l[0]` and `gnt[0]` are connected to the `PCI_REQ_L` and `PCI_GNT_L` pins. Otherwise, they are connected to the internal arbiter.

The PCI unit asserts `req_l[0]` to act as a bus master on the PCI. If `gnt_l[0]` is asserted, the PCI unit can start a PCI transaction regardless of the state of `req_l[0]`. When the PCI unit requests the PCI bus, it performs a PCI transaction when `gnt_l[0]` is received. Once `req_l[0]` is asserted, the PCI unit



never de-asserts it prior to receiving `gnt_l[0]` or de-asserts it after receiving `gnt_l[0]` without doing a transaction. PCI Unit de-asserts `req_l[0]` for two cycles when it receives a retry or disconnect response from the target.

9.2.6.2 PCI Commands

The following PCI transactions are not generated by PCI Unit as an initiator:

- PCI Lock Cycle
- Dual Address cycle
- Memory Write and Invalidate

9.2.6.3 Initiator Write Transactions

The following general rules apply to the write command transactions:

- If the PCI unit receives either a target retry response or a target disconnect response before all of the write data has been delivered, it resumes the transaction at the first opportunity, using the address of the first undeliverable data.
- If the PCI unit receives a master abort, it discards all of the write data from that transaction and sets the status register (`PCI_STATUS`) received master abort bit, which, if enabled, interrupts the Intel XScale® core.
- If the PCI unit receives a target abort, it discards all of the remaining write data from that transaction, if any, and sets the status registers (`PCI_STATUS`) received target abort bit, which, if enabled, interrupts the Intel XScale® core.
- The PCI unit can dessert `frame_1` prior to delivering all data due to the master latency timer, If this occurs, it resumes the write at the first opportunity, using the address of the first undeliverable data.

9.2.6.4 Initiator Read Transactions

The following general rules apply to the read command transactions:

- If the PCI unit receives a target retry, it repeats the transaction at the first opportunity until the whole transaction is completed.
- If the PCI unit receives a master abort, it substitutes `0xFFFF FFFF` for the read data and sets the status register (`PCI_STATUS`) received master abort bit, which, if enabled, interrupts the Intel XScale® core.
- If the PCI unit receives a target abort, it sets the status registers (`PCI_STATUS`) received target abort bit, which, if enabled, interrupts the Intel XScale® core and does not try to get any more read data. PCI unit will substitute `0xFFFF FFFF` for the data which are not read and complete the cycle.

9.2.6.5 Initiator Latency Timer

When the PCI unit begins PCI transaction as an initiator, asserting `frame_1`, it begins to decrement its master latency timer. When the timer value reaches 0, the PCI unit checks the value of `gnt_l[0]`. If `gnt_l[0]` is de-asserted, the PCI unit de-asserts `frame_1` (if it is still asserted) at the earliest opportunity. This is normally the next data phase for all transactions.



9.2.6.6 Special Cycle

As an initiator, special cycles are broadcast to all PCI agents, so DEVSEL_L is not asserted and no error can be received.

9.2.7 PCI Fast Back-to-Back Cycles

The core supports fast back-to-back target cycles on the PCI Bus. The core does not generate initiator fast back-to-back cycles on the PCI Bus regardless of the value in the fast back-to-back enable bit of the Status and Command register in the PCI configuration space.

9.2.8 PCI Retry

As a slave, the PCI Unit generates retry on:

- A slave write when the Data write FIFO is full.
- When address FIFO is full
- Data read is handled as delay transactions. If the HOG_MODE bit is set in the PCI_CONTROL register, the bus will be held for 16 PCI clocks before asserting retry.

As an initiator, the core supports retry by maintaining an internal counter of the current address. On receiving a retry, the core de-asserts PciFrameN and then re-assert PciFrameN with the current address from the counter.

9.2.9 PCI Disconnect

As a slave, it disconnects for the following conditions:

- Bursted PCI configuration cycle.
- Bursted access to PCI_CSR_BAR.
- PCI reads past the amount of data in the read FIFO.
- PCI burst cycles that cross 1K PCI address boundary which includes PCI burst cycles that cross memory decodes from the core as a target to decodes that are outside the core (e.g., started inside a BAR and ends outside of that BAR).

As an initiator, the core supports retry and disconnect by maintaining an internal counter of the current address. On receiving a retry or disconnect, the core de-asserts PciFrameN and then re-assert PciFrameN with the current address + “current transfer byte size” from the counter.

9.2.10 PCI Built-In System Test

The IXP2800 Network Processor supports BIST when there is an external PCI host. The PCI host will set the STRT bit in the PCI_CACHE_LAT_HDR_BIST configuration register. An interrupt is generated to the Intel XScale® core if it is enabled by the Intel XScale® core Interrupt Enable register. The Intel XScale® core software can respond to the interrupt by running an application-specific test. Upon successful completion of the test, the Intel XScale® core will reset the STRT bit. If this bit is not reset two seconds after the PCI host sets the STRT bit, the host will indicate that the Network Processor failed the test.



9.2.11 PCI Central Functions

The CFG_RSTDIR pin is active high for enabling the PCI Unit central function.

The CFG_PCI_ARB(GPIO[2]) pin is the strap pin for the internal arbiter. When this strap pin is high during reset then the XPI Unit owns the arbitration.

The CFG_PCI_BOOT_HOST(GPIO[1]) pin is the strap pin for the PCI host. When PCI_BOOT_HOST is asserted during reset then PCI Unit will support as a PCI host.

Table 122. Legal Combinations of the Strap Pin Options

	CFG_PCI_BOOT_HOST (GPIO[1])	CFG_PCI_ARB (GPIO[2])	CFG_PCI_RSTDIR (Central function)	CFG_PROM_BOOT (GPIO[0])
OK	0	0	0	0
OK	0	0	0	1
OK	0	0	1	1
Not supported	0	1	0	x
OK	0	1	1	1
Not supported	1	0	0	x
OK	1	0	1	1
Not supported	1	1	0	x
OK	1	1	1	1

Note

- * CFG_PCI_RSTDIR = 1 then central function.
- * CFG_PCI_BOOT_HOST must be central function.
- * CFG_PCI_ARB must be central function.

9.2.11.1 PCI Interrupt Inputs

The PCI Unit supports two interrupt lines from the PCI Bus as host. One of the interrupt lines will be open-drain output and input. The other interrupt line will be selected as PCI interrupt input. Both the interrupt lines can be enabled in the Intel XScale® core Interrupt Enable register.

9.2.11.2 PCI Reset Output

If the IXP2800 Network Processor is central function (CFG_RSTDIR =1), PCI Unit will be asserting the PCI_RST_L after the system power-on. The Intel XScale® core has to write to the PCI External Reset bit in the IXP2800 Network Processor's Reset register to de-assert the PCI_RST_L. In this case, chip reset (CLK_NRESET) is driven by a signal other than PCI_RST_L.

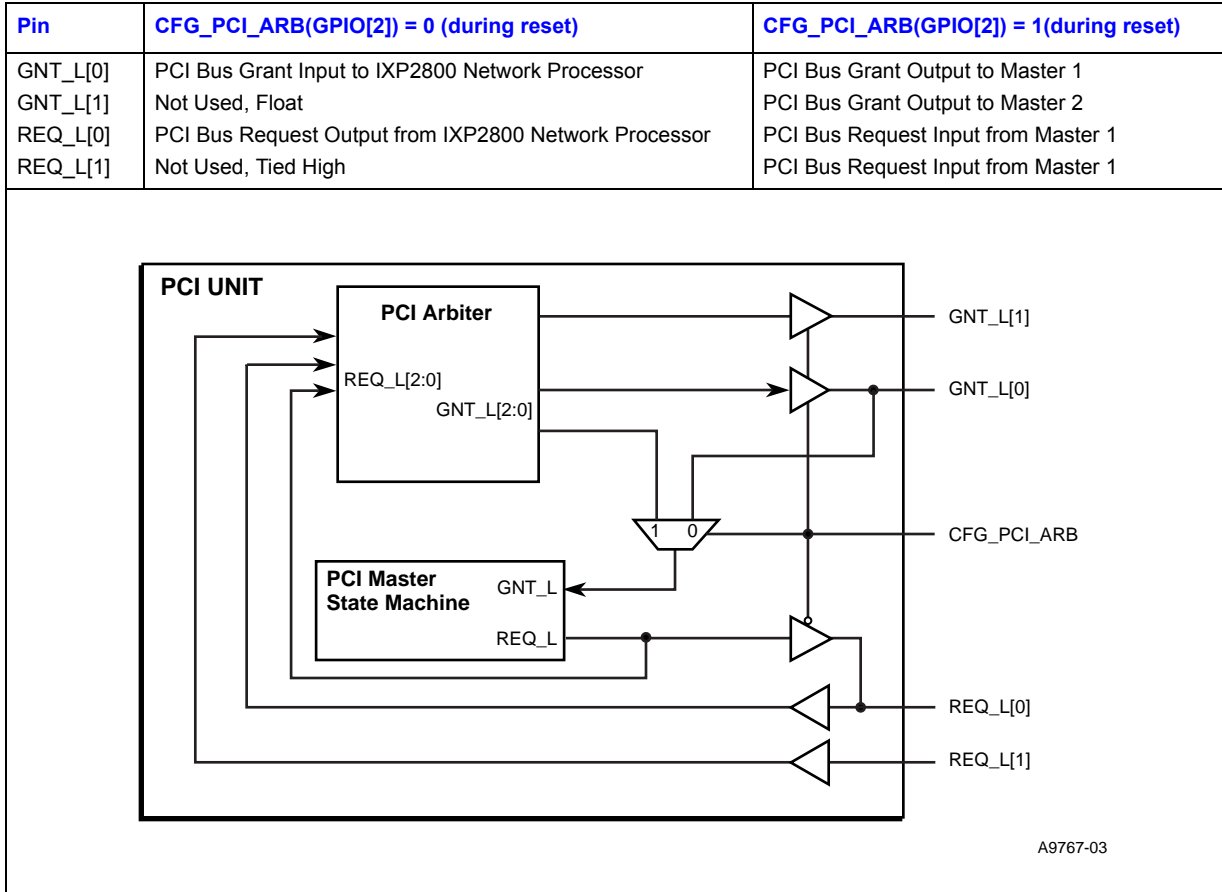
When the PCI Unit is not configured as the central function (CFG_RSTDIR =0), PCI_RST_L is used as a chip reset input.



9.2.11.3 PCI Internal Arbiter

The PCI unit contains a PCI bus arbiter that supports two external masters in addition to the PCI Unit's initiator interface. To enable the PCI arbiter, the CFG_PCI_ARB(GPIO[2]) strapping pin must be 1 during reset. As shown in Figure 120, the local bus request and grant pair become externally not visible. These signals will be made available to external debug pins for debug purpose.

Figure 120. PCI Arbiter Configuration Using CFG_PCI_ARB(GPIO[2])



The arbiter uses a simple round-robin priority algorithm. The arbiter asserts the grant signal corresponding to the next request in the round-robin during the current executing transaction on the PCI bus (this is also called hidden arbitration). If the arbiter detects that an initiator has failed to assert frame_1 after 16 cycles of both grant assertion and PCI bus idle condition, the arbiter de-asserts the grant. That master does not receive any more grants until it de-asserts its request for at least one PCI clock cycle. Bus parking is implemented in that the last bus grant will stay asserted if no request is pending.

To prevent bus contention, if the PCI bus is idle, the arbiter never asserts one grant signal in the same PCI cycle in which it de-asserts another. It de-asserts one grant, and then asserts the next grant after one full PCI clock cycle has elapsed to provide for bus driver turnaround.



9.3 Slave Interface Block

The slave interface logic supports internal slave devices interfacing to the target port of the FBus.

- CSR — register access cycles to local CSRs.
- DRAM — memory access cycles to the DRAM push/pull Bus.
- SRAM — memory access cycles to the SRAM push/pull Bus.

The slave port of the FBus is connected to a 64-byte write buffer to support bursts of up to 64 bytes to the memory interfaces. The slave read data are directly downloaded into the FBus read FIFO. See [Table 123](#).

Table 123. Slave Interface Buffer Sizes

Location	Slave Address	Slave Write	Slave Read
Buffer Depth	1	64 Byte	0
Usage	CSR, SRAM, DRAM	SRAM, DRAM	None

As a push/pull command bus master, the PCI Unit translates these accesses into different types of push/pull command. As the push/pull data bus target, the write data is sent through the pull data bus and the read data is received on the push data bus.

9.3.1 CSR Interface

The internal Control and Status registers data is directed to or from the Slave FIFO port of the PCI core FBus when the BAR id matches PCI_CSR_BAR (BAR0). The CSR accesses from the PCI Bus directed towards CSRs not in PCI Unit is translated into a push/pull CSR type command. PCI local CSRs are handled within the PCI Unit.

For writes, the data is sent when the pull bus is valid and the ID matches. The address is unloaded from the FBus target address FIFO as indication to the PCI core logic that the cycle is completed. The slave write buffer is not used for CSR access.

For reads, the data is loaded into the target receive FIFO as soon as the push bus is valid and the ID matches. The address is unloaded from the FBus address FIFO.

Note: Target reads to the Scratch unit must always be in multiples of 32-bit (PCI_CBE_L[3:0] =0x0) as the Scratch unit only supports 32-bit accesses.

One example of a PCI host access to internal registers is the initialization of internal registers and memory to enable the Intel XScale® core to boot off the DRAM in the absence of a boot up PROM.

The accesses to the CSRs inside the PCI Unit are completed internally without sending the transaction out to the push pull bus, just like the other internal register accesses.



9.3.2 SRAM Interface

The SRAM interface connects the FBus to the internal push/pull command bus and the SRAM push/pull data buses. Request to memory is sent on the command bus. Data request is received as valid push/pull ID sent by the SRAM push/pull data bus.


If the PCI_SRAM_BAR is used, the target state machine generates a request to the command bus for SRAM access. Once the grant is received, the address, then data is directed between the slave FIFOs of the PCI core and the SRAM push/pull bus.

9.3.2.1 SRAM Slave Writes

The slave write buffer is used to support memory burst accesses. The buffer is added to guarantee data transfer for each clock and burst size can be determined before memory request is issued. Data is assembled in the buffers before being sent to memory for SRAM write.

On the push/pull bus, AM access can start at any address and have length up to 16 Dwords as shown in Figure 121. For masked writes, only size 1 is supported to transfer up to four bytes.

Figure 121. Example of Target Write to SRAM of 68 Bytes



The slave interface also has to make sure there is enough data in the slave write buffer to complete the memory data transfer before making a memory request.



9.3.2.2 SRAM Slave Reads

For a slave read from SRAM, a 32-bit DWORD is fetched from the memory for memory read command, one cache line is fetched for memory read line command, and two cache lines are read for memory read multiple command. Cache line size is programmable in the `CACHE_LINE` field of the `PCI_CACHE_LAT_HDR_BIST` configuration register. If the computed read size is greater than 64 bytes, the PCI SRAM read will default to the maximum of 64 bytes. No pre-fetch is supported in that the PCI Unit will not read beyond the computed read size.

The PCI core resets the target read FIFO before issuing a memory read data request on FBus. The maximum size of SRAM data read is 64 bytes. The PCI core will disconnect at the 64-byte address boundary.

9.3.3 DRAM Interface

The memory is accessed using the push/pull mechanism. Request to memory is sent on the command bus. If the `PCI_DRAM_BAR` is used, the target state machine generates a request to the command bus for DRAM access with the address in the slave address FIFO. Once the push/pull request is received. The data is directed between the Slave FIFOs of the PCI core and DRAM push/pull bus.

9.3.3.1 DRAM Slave Writes

The slave write buffer is used to support memory burst accesses. The buffer is added to guarantee data transfer for each clock and burst size can be determined before memory request is issued. Data is assembled in the buffers before being sent to memory for memory write.

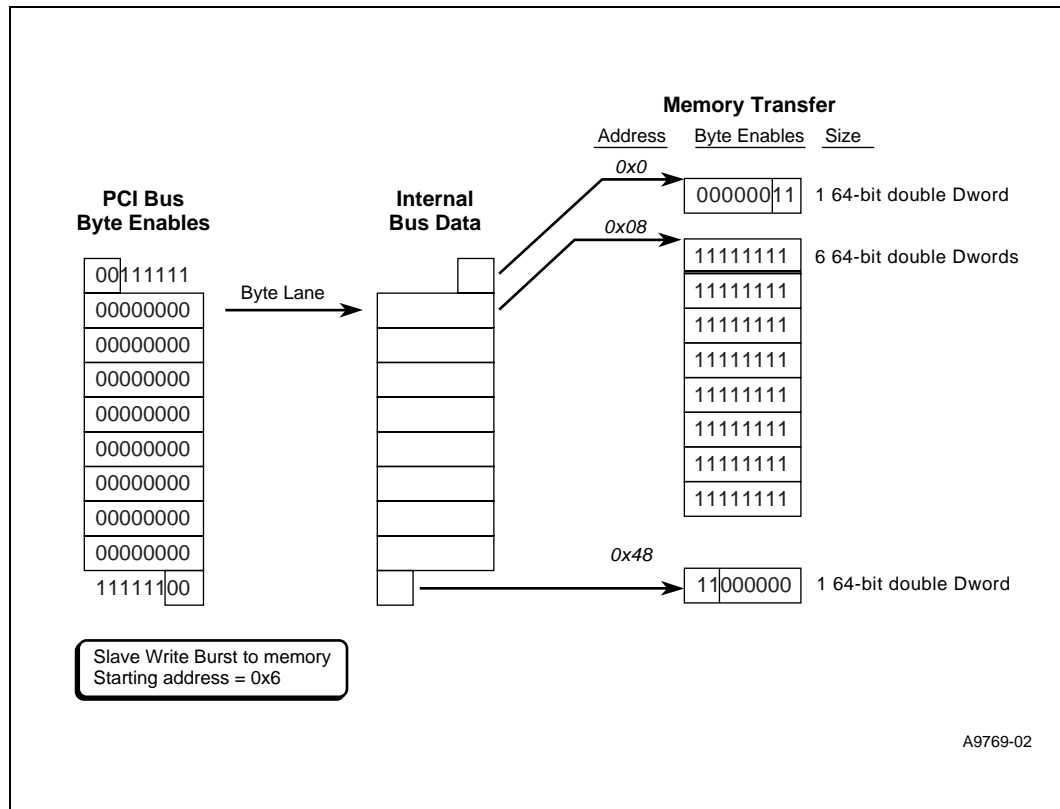
DRAM target write access is only required to be 8-byte address aligned and the address does not wrap around the 64-byte address boundary on a DRAM burst. Each 8-byte access that is a partial write to the memory, is treated as single write. Remaining writes of the 64-byte segment is written as one single burst. Transfers that cross a 64 -byte segment are split into separate transfers. [Figure 123](#) splits the 68-byte transfers into two partial 8-byte transfers to address 06 and address 48 and one 56-byte burst transfer in the first 64-byte segment from address 08 to 38 and one 8-byte transfer to address 40.

For write to DRAM on the push/pull bus, the burst must be broken down into address aligned smaller transfer sizes (see [Figure 122](#)).

The Target interface also must make sure there is enough data in the target write buffer to complete the memory data transfer before making a memory request.



Figure 122. Example of Target Write to DRAM of 68 Bytes

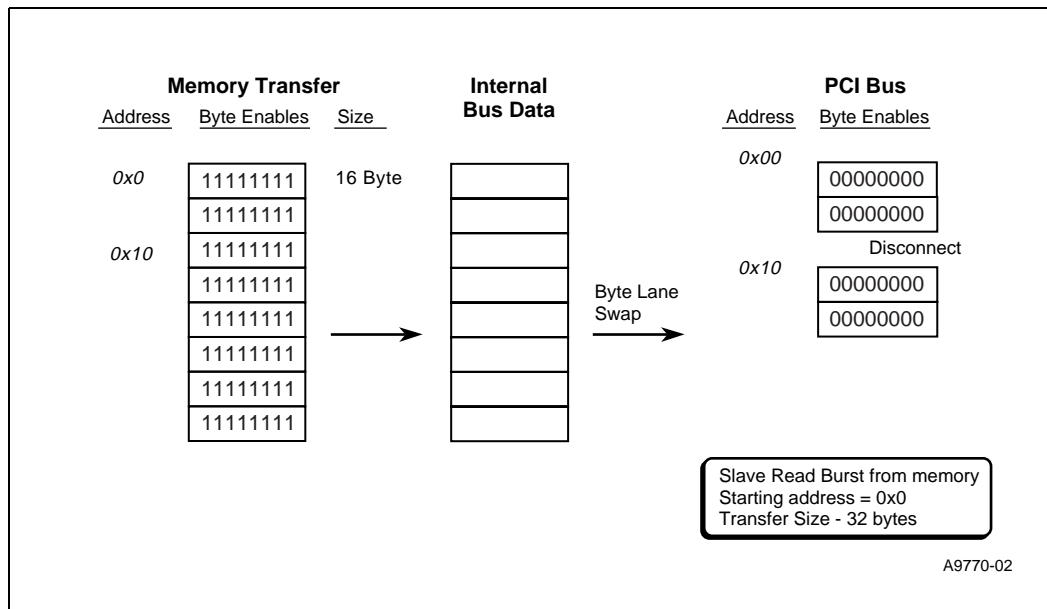


9.3.3.2 DRAM Slave Reads

For target reads from IXP2800 Network Processor memory, the entire 64-byte block is fetched from DRAM. For target reads from IXP2800/IXP2850 Network Processor memory, the block size is 16 bytes. Depending on the address for the target request, extra data is discarded at the beginning until the target address is reached. Also, extra data is discarded at the end of the transfer also when the burst ends in the middle of a data block. No pre-fetch is supported for DRAM access. See Figure 123.

Note: The IXP2800/IXP2850 always disconnects after transferring 16-bytes for DRAM target reads. The PCI core will also disconnect at a 64-byte address boundary.

Figure 123. Example of Target Read from DRAM Using 64-Byte Burst



The PCI core resets the read FIFO before issuing a memory read data request on FBus. The PCI core will disconnect at the 64-byte address boundary.

9.3.4 Mailbox and Doorbell Registers

Mailbox and Doorbell registers provide hardware support for communication between the Intel XScale® core and a device on the PCI Bus.

Four mailbox registers are provided so that messages can be passed between the Intel XScale® core and a PCI device. All four registers are 32 bits and can be read and written with byte resolution from both the Intel XScale® core and PCI. How the registers are used is application dependent and the messages are not used internally by the PCI Unit in any way. The mailbox registers are often used with the Doorbell interrupts.

Doorbell interrupts provide an efficient method of generating an interrupt as well as encoding the purpose of the interrupt. The PCI Unit supports an Intel XScale® core Doorbell register that is used by a PCI device to generate an Intel XScale® core FIQ and a separate PCI Doorbell register that is used by the Intel XScale® core to generate a PCI interrupt. A source generating the Doorbell interrupt can write a software defined bitmap to the register to indicate a specific purpose. This bitmap is translated into a single interrupt signal to the destination (either a PCI interrupt or a IXP2800 Network Processor interrupt). When an interrupt is received, the Doorbell registers can be read and the bit mask can be interpreted. If a larger bit mask is required than that is provided by the Doorbell register, the Mailbox registers can be used to pass up to four 32-bit blocks of data.



The doorbell interrupts are controlled through the registers shown in Table 124.

Table 124. Doorbell Interrupt Registers

Register Name	Description
Intel XScale® core Doorbell	Used to generate the Intel XScale® core Doorbell interrupts.
Intel XScale® core Doorbell Setup	Used to initialize the Intel XScale® core Doorbell register and for diagnostics.
PCI Doorbell	Used to generate the PCI Doorbell interrupts.
PCI Doorbell Setup	Used to initialize the PCI Doorbell register and for diagnostics.

The Intel XScale® core and PCI devices write to the corresponding DOORBELL register to generate up to 32 doorbell interrupts. Each bit in the DOORBELL register is implemented as an SR flip-flop. The Intel XScale® core writes a 1 to set the flip-flop and the PCI device writes a 1 to clear the flip-flop. Writing a 0 has no effect on the registers. The PCI interrupt signal is the output of an NOR functions of all the PCI DOORBELL register bits (outputs of the SR flip-flops). The Intel XScale® core interrupt signal is the output of an NAND function of all the Intel XScale® core DOORBELL register bits (outputs of the SR flip-flops).

To assert an interrupt (i.e., to “push a doorbell”):

- A write of 1 to the corresponding bit of the DOORBELL register generates an interrupt. This is the case for either PCI device or the Intel XScale® core, since writing 1 changes the doorbell bit to the proper asserted state (i.e., 0 for an Intel XScale® core interrupt and 1 for a PCI interrupt).

To dismiss an interrupt:

- A write of 1 to the corresponding bit of the DOORBELL register clears an interrupt. This is the case for either PCI device or the Intel XScale® core, since writing 1 changes the doorbell bit to the proper de-asserted state (i.e., 1 for an Intel XScale® core interrupt and 0 for a PCI interrupt).

Figure 124 and Figure 125 illustrate how a Doorbell interrupt is asserted and cleared by both the Intel XScale® core and a PCI device.

Figure 124. Generation of the Doorbell Interrupts to PCI

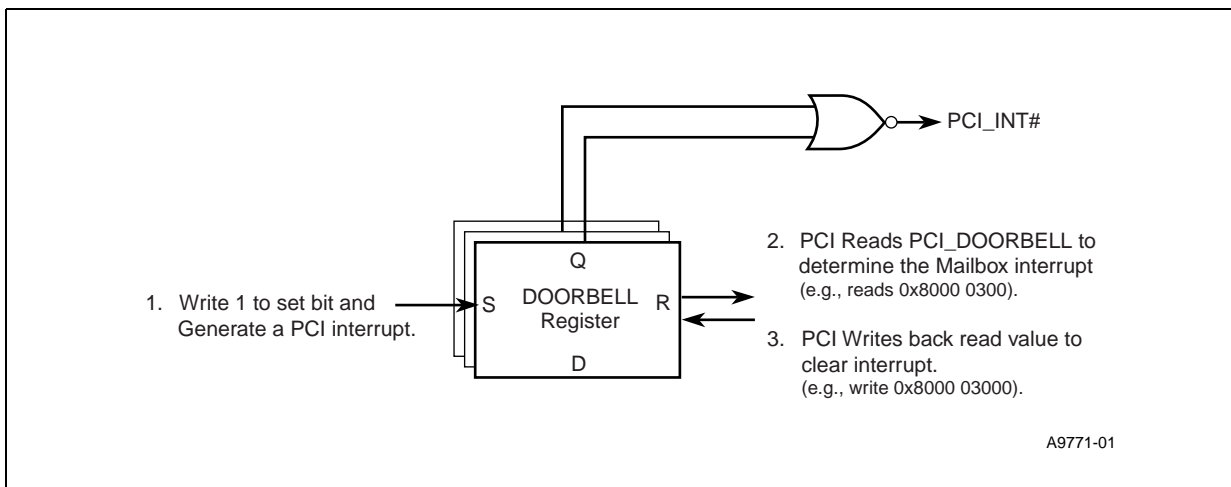
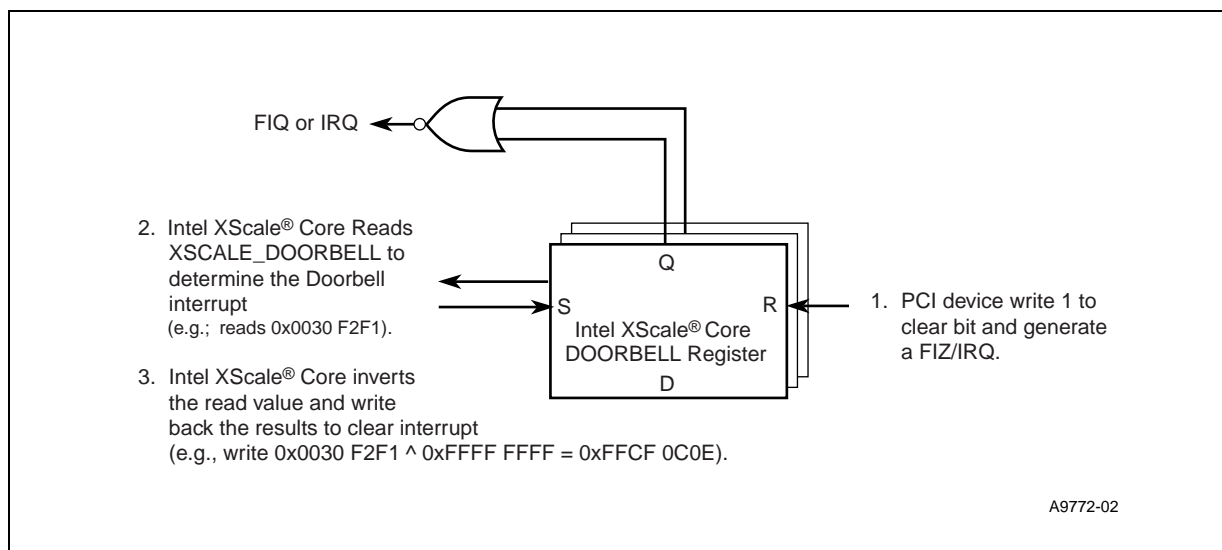


Figure 125. Generation of the Doorbell Interrupts to the Intel XScale® Core



The Doorbell Setup register allows the Intel XScale® core and a PCI device to perform two functions that are not possible using the Doorbell register. This register is used during setup and diagnostics and is not used during normal operations. First, it allows the Intel XScale® core and PCI device to clear an interrupt that it has generated to the other device. If the Intel XScale® core sets an interrupt to PCI device using the Doorbell register, the PCI device is the only one that can use the Doorbell register to clear the interrupt by writing one. With the Doorbell setup register, the Intel XScale® core can clear the interrupt by write 0 to it.

Second, it allows the Intel XScale® core and PCI device to generate a doorbell interrupt to itself. This can be used for diagnostic testing. Each bit in the Doorbell Setup register is mapped directly to the data input of the Doorbell register such that the data is directly written into the Doorbell register.

During system initialization, the doorbell registers must be initialized by clearing the interrupt bits in the Doorbell register using the Doorbell Setup register. This is done by writing zeros to the PCI Doorbell setup register and ones to the Intel XScale® core Doorbell setup register.



9.3.5 PCI Interrupt Pin

An external PCI interrupt can be generated in the following way:

- The Intel XScale® core initiates a Doorbell interrupt XSCALE_INT_ENABLE.
- One or more of the DMA channels have completed the DMA transfers.
- The PNI bit is cleared by the Intel XScale® core to generate a PCI interrupt
- An internal functional unit generates either an interrupt or an error directly to the PCI host.

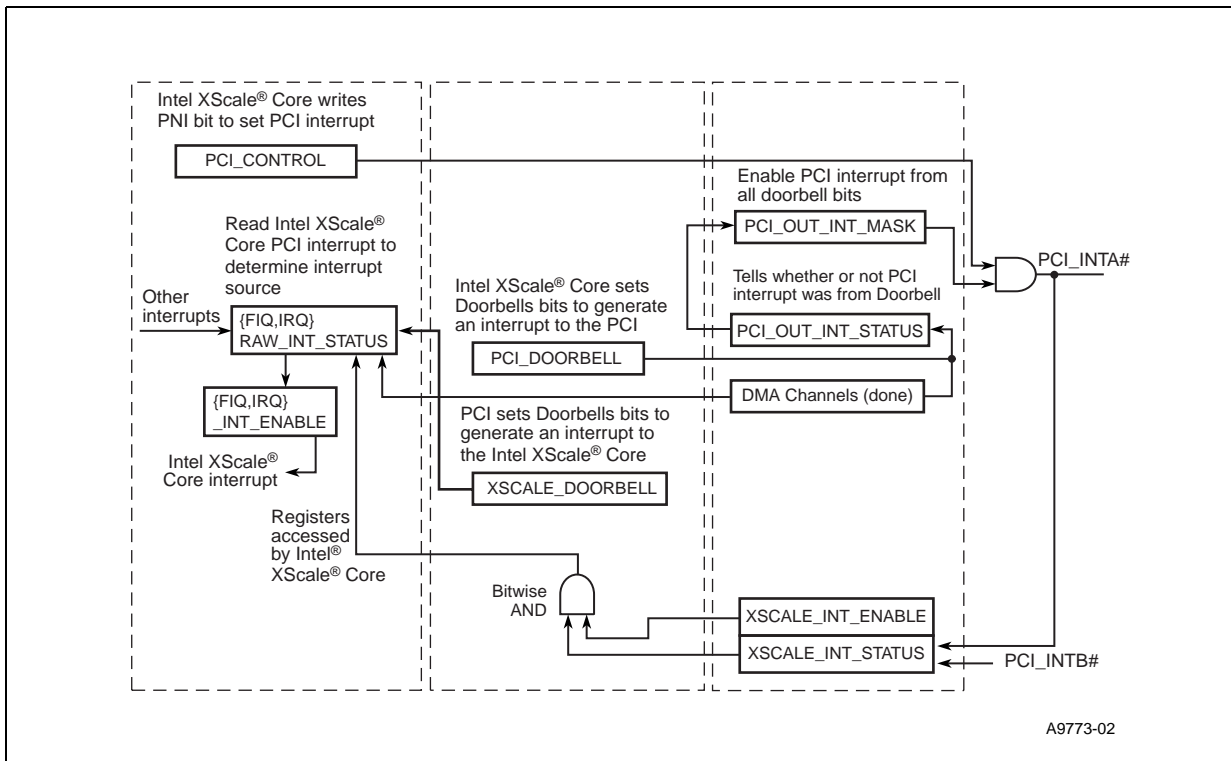
Table 125 describes how IRQ are generated for each silicon stepping.

Table 125. IRQ Interrupt Options by Stepping

Stepping	Description
A stepping	IRQ interrupts can be handled only by the Intel XScale® core.
B Stepping	IRQ interrupts can be handled by either the Intel XScale® core or a PCI host. Refer to the description of the PCI_OUT_INT_MASK and PCI_OUT_INT_STATUS registers in the <i>Intel® IXP2400 and IXP2800 Network Processor Programmer's Reference Manual</i> .

Figure 126 shows how PCI interrupts are managed via the PCI and the Intel XScale® core.

Figure 126. PCI Interrupts





9.4 Master Interface Block

The Master Interface consists of the DMA engine and the Push/pull target interface. Both can generate initiator PCI transactions.

9.4.1 DMA Interface

There are two DMA channels, each of which can move blocks of data from DRAM to the PCI or from the PCI to DRAM. The DMA channels read parameters from a list of descriptors in SRAM, perform the data movement to or from DRAM, and stop when the list is exhausted. The descriptors are loaded from predefined SRAM entries or may be set directly by CSR writes to DMA registers. There is no restriction on byte alignment of the source address or the destination address. For PCI to DRAM transfers, the PCI command is Memory Read, Memory Read line, or Memory Read Multiple. For DRAM to PCI transfers, the PCI command is Memory Write. Memory Write Invalidate is not supported.

DMA reads are unmasked reads (all byte enables asserted) from DRAM. After each transfer, the byte count is decremented by the number of bytes read, and the source address is incremental by one 64-bit double Dword. The whole data block is fetched from the DRAM. For a system using RDRAM (like the IXP2800 Network Processor), the block size is 16 bytes.

DMA reads are masked reads from the PCI and writes are masked for both the PCI and DRAM. When moving a block of data, the internal hardware adjusts the byte enables so that the data is aligned properly on block boundaries and that only the correct bytes are transferred if the initial and final data requires masking.

For DMA data, the DMA FIFO consists of two separate FBus initiator read FIFOs and two initiator write FIFOs, which are inside the PCI Core and three DMA buffers (corresponding to the DMA channels), which buffer data to and from the DRAM. Since there is no simultaneous DMA read and write outstanding, one shared 64-byte buffer is used for both read and write DRAM data

Up to two DMA channels are running at a time with three descriptors outstanding. The two DMA channels and the direct access channel to PCI Bus from Command Bus Master are contending to use the address, read and write FIFOs inside the Core.

Effectively, the active channels interleave bursts to or from the PCI Bus. Each channel is required to arbitrate for the PCI FIFOs after each PCI burst request.



9.4.1.1 Allocation of the DMA Channels

Static allocation are employed such that the DMA resources are controlled exclusively by a single device for each channel. The Intel XScale® core, a Microengine and the external PCI host can access the two DMA channels. The first two channels can function in one of the following modes, as determined by the DMA_INF_MODE register:

- The Intel XScale® core owns both DMA channel 1 and channel 2.
- The Microengines owns both DMA channel 1 and channel 2.
- PCI host owns both DMA channel 1 and channel 2.
- The Intel XScale® core owns both DMA channel 1 and channel 2.

The third channel can be allocated to either the Intel XScale® core, PCI host, or Microengines.

The DMA mode can be changed only by the Intel XScale® core under software control. The software should signal to suspend DMA transactions and wait until all DMA channels are free before changing the mode. Software should determine when all DMA channels are free either by polling XSCALE_INT_STATUS register bits DMA1 and DMA3 until both DMA channels are done.

9.4.1.2 Special Registers for Microengine Channels

Interrupts are generated at the end of DMA operation for the Intel XScale® core and PCI-initiated DMA. However, the Microengine does not provide the interrupt mechanism. The PCI Unit will instead use an “Auto-Push” mechanism to signal the particular Microengine on completion of DMA.

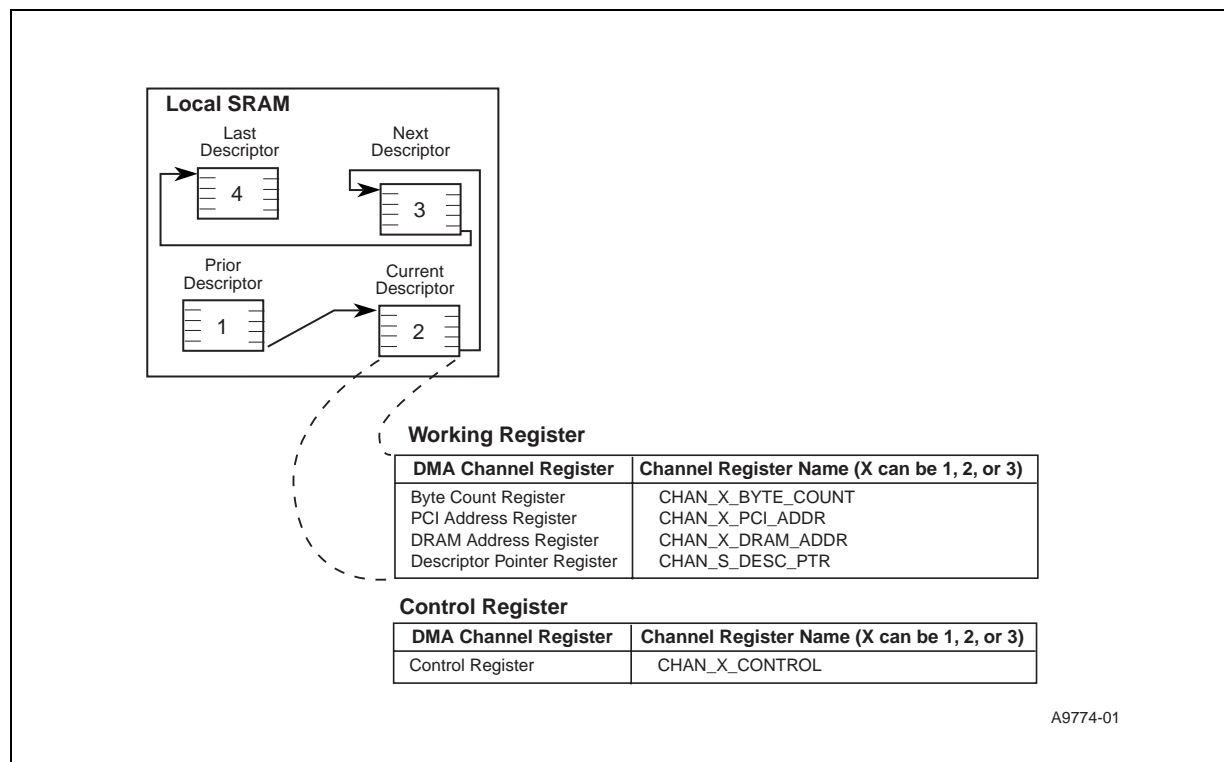
When the Microengine sets up the DMA channel, it would also write the CHAN_X_ME_PARAM with Microengine number, Context number, Register number, and Signal number. When the DMA channel completes, it writes some status information (Error or OK status) to the Microengine/Context/Register/Signal. PCI Unit will arbitrate for the SRAM Push bus. The Push ID is from the parameters in the register.

The ME_PUSH_STATUS reflects the DMA Done bit in each of the CHAN_X_CONTROL registers. The Auto-Push operation will proceed after the DMA is done for the particular DMA channel if the corresponding enable bit in the ME_PUSH_ENABLE is set.

9.4.1.3 DMA Descriptor

Each descriptor occupies four 32-bit Dwords and is aligned on a 16-byte boundary. The DMA channels read the descriptors from local SRAM into the four DMA working registers once the control register has been set to initiate the transaction. This control must be set explicitly. This starts the DMA transfer. The register names for the DMA channels are listed in [Figure 127](#).

Figure 127. DMA Descriptor Reads



After a descriptor is processed, the next descriptor is loaded in the working registers. This process repeats until the chain of descriptors is terminated (i.e., the End of Chain bit is set). See [Table 126](#).

Table 126. DMA Descriptor Format

Offset from Descriptor Pointer	Description
0x0	Byte Count
0x4	PCI Address
0x8	DRAM Address
0xC	Next Descriptor Address



9.4.1.4 DMA Channel Operation

Since a PCI device, Microengine, or the Intel XScale® core can access the internal CSRs and memory in a similar way, the DMA channel operation description that follows will apply to all channels. CHAN_1_, CHAN_2_, or CHAN_3_ can be placed before the name for the DMA registers.

The DMA channel owner can either set up the descriptors in SRAM or it can write the first descriptor directly to the DMA channel registers.

When descriptors and the descriptor list are in SRAM, the procedure is as follows:

1. The DMA channel owner writes the address of the first descriptor into the DMA Channel Descriptor Pointer register (DESC_PTR).
2. The DMA channel owner writes the DMA Channel Control register (CONTROL) with miscellaneous control information and also sets the channel enable bit (bit 0). The channel initial descriptor bit (bit 4) in the CONTROL register must also be cleared to indicate that the first descriptor is in SRAM.
3. Depending on the DMA channel number, the DMA channel reads the descriptor block into the corresponding DMA registers, BYTE_COUNT, PCI_ADDR, DRAM_ADDR, and DESC_PTR.
4. The DMA channel transfers the data until the byte count is exhausted, and then sets the channel transfer done (bit 2) in the CONTROL register.
5. If the end of chain bit (bit 31) in the BYTE_COUNT register is clear, the channel checks the Chain Pointer value. If the Chain Pointer value is not equal to 0, it reads the next descriptor and transfers the data (step 3 and 4 above). If the Chain Pointer value is equal to 0, it waits for the Descriptor Added bit of the Channel Control register to be set before reading the next descriptor and transfers the data (step 3 and 4 above). If bit 31 is set, the channel sets the channel chain done bit (bit 7) in the CONTROL register and then stops.
6. Proceed to the Channel End Operation. (See [Section 9.4.1.5](#).)

When single descriptors are written directly into the DMA channel registers, the procedure is as follows:

1. The DMA channel owner writes the descriptor values directly into the DMA channel registers. The end of chain bit (bit 31) in the BYTE_COUNT register must be set, and the value in the DESC_PTR register is not used.
2. The DMA channel owner writes the base address of the DMA transfer into the PCI_ADDR to specify the PCI starting address.
3. When the first descriptor is in the BYTE_COUNT register, the DRAM_ADDR register must be written with the address of the data to be moved.
4. The DMA channel owner writes the CONTROL register with miscellaneous control information, along with setting the channel enable bit (bit 0). The channel initial descriptor in register bit (bit 4) in the CONTROL register must also be set to indicate that the first descriptor is already in the channel descriptor registers.
5. The DMA channel transfers the data until the byte count is exhausted, and then sets the channel transfer done bit (bit 2) in the CONTROL register.
6. Since the end of the chain bit (bit 31) in the BYTE_COUNT register is set, the channel sets the channel chain done bit (bit 7) in the CONTROL register and then stops.
7. Proceed to the Channel End Operation. (See [Section 9.4.1.5](#).)



9.4.1.5 DMA Channel End Operation

1. Channel owned by PCI:
If not masked via the PCI Outbound Interrupt Mask register, the DMA channel interrupts the PCI host after the setting of the DMA done bit in the CHAN_X_CONTROL register, which is readable in the PCI Outbound Interrupt Status register.
2. Channel owned by the Intel XScale® core:
If enabled via the Intel XScale® core Interrupt Enable registers, the DMA channel interrupts the Intel XScale® core by setting the DMA channel done bit in the CHAN_X_CONTROL register, which is readable in the Intel XScale® core Interrupt Status register.
3. Channel owned by Microengine:
If enabled via the Microengine Auto-Push Enable registers, the DMA channel signals the Microengine after setting the DMA channel done bit in the CHAN_X_CONTROL register, which is readable in the Microengine Auto-Push Status register.

9.4.1.6 Adding Descriptor to an Unterminated Chain

It is possible to add a descriptor to a chain while a channel is running. To do so the chain should be left un-terminated, that is the last descriptor should have End of Chain clear, and the Chain Pointer value equal to 0. A new descriptor (descriptors) can be added to the chain by overwriting the Chain Pointer value of the un-terminated descriptor (in SRAM) with the Local Memory address of the (first) added descriptor (Note that the added descriptor must actually be valid in Local Memory prior to that). After updating the Chain Pointer field, the software must write a 1 to the Descriptor Added bit of the Channel Control register. This is necessary for the case where the channel was paused to reactivate the channel. However, software need not check the state of the channel before writing that bit; there is no side-effect of writing that bit in the case where the channel had not yet read the unlinked descriptor.

If the channel was paused or had read an unlinked Pointer, it will re-read the last descriptor processed (i.e., the one that originally had the 0 value for Chain Pointer) to get the address of the newly added descriptor.

A descriptor cannot be added to a descriptor that has End of Chain set.

9.4.1.7 DRAM to PCI Transfer

For a DRAM-to-PCI transfer, the DMA channel reads data from DRAM and places it into the DMA buffer for transfer to the FBus FIFO when the following conditions are met:

- There is at least free space for a read block in the buffer.
- The DRAM controller issues data valid on DRAM push data bus to the DMA engine.
- DMA transfer is not done.

Before data is stored into the DMA buffer, the DRAM starting address is evaluated. Extra data will be discarded in case the DRAM starting address does not start at aligned addresses. The lower address bits determine the byte enables for the first data double Dword. At the end of the DMA transfer, extra data will be discarded and byte enables are calculated for the last 64-bit double Dword. After the data is loaded into the buffer, the PCI starting address is evaluated and the buffer is shifted byte wise to align the starting DRAM data with the starting PCI starting address.



A 64-bit double Dword with byte enables is pushed into the FBus FIFO from the DMA buffers as soon as there is data available in the buffer and there is space in the FBus FIFO. The Core logic will transfer the exact number of bytes to the PCI Bus. The maximum burst size on the PCI bus varies according to the stepping and is described in [Table 127](#)

Table 127. PCI Maximum Burst Size

Stepping	Description
A Stepping	The maximum burst size is 64 bytes.
B Stepping	<p>The maximum burst size can be greater than 64 bytes for certain operations. The register PCI_IXP_PARAM configures the burst length for target write operations.</p> <p>The register CHAN_#_CONTROL configures the burst length for DMA read and write operations.</p> <p>The register PCI_CONTROL configures the atomic feature for target write operations of 64 bytes or fewer.</p> <p>Note: Bursts longer than 64 bytes are not supported for PCI target read operations.</p>

9.4.1.8 PCI to DRAM Transfer

The DMA channel issues a sequence of PCI read request commands through the FBus address FIFO to read the precise byte count from PCI.

The DMA engine will continue to load the DMA write buffer with FBus FIFO data as soon as data is available.

The DMA engine determines the largest size of memory request possible with the current DRAM address and remaining byte count. It also has to make sure there is enough data in the write buffer before sending the memory request.

9.4.2 Push/Pull Command Bus Target Interface

Through the command bus target interface, the command bus masters (PCI, Intel XScale® core, and Microengines) can access the PCI Unit internal registers including the local PCI configuration registers and the local PCI Unit CSRs. Also, the Microengine and the Intel XScale® core can issue transactions on the PCI bus. The requests are generated from the command master to the command bus arbiter. The arbiter selects a master and sends it a grant. That master then sends a command, which is passed through by the arbiter.

PCI Unit will issue the push and pull data responses to the SRAM push/pull data buses. When the read command is received, the PCI Unit will issue the push data request on the SRAM push data bus. When the write command is received, PCI Unit will issue the pull command on the SRAM pull data bus.

9.4.2.1 Command Bus Master Access to Local Configuration Registers

The configuration register within the PCI unit can be accessed by push/pull command bus access to configuration space through the FBus interface of the PCI core. When the IXP2800 Network Processor is a PCI host, these registers have to be accessed through this internal path and no PCI bus cycle will be generated.



9.4.2.2 Command Bus Master Access to Local Control and Status Registers

These are CSRs within the PCI Unit that are accessible from push/pull bus masters. The masters include the Intel XScale® core, Microengines. There is no PCI bus cycles generated. The CSRs within the PCI Unit can be accessed internally by external PCI devices.

9.4.2.3 Command Bus Master Direct Access to PCI Bus

The Intel XScale® core and Microengines are the only command bus masters that have direct access to the PCI bus as a PCI Bus initiator. The PCI Bus can be accessed by push/pull command bus access to PCI bus address space. The PCI Unit will share the internal SRAM push/pull data bus with SRAM for the data transfers.

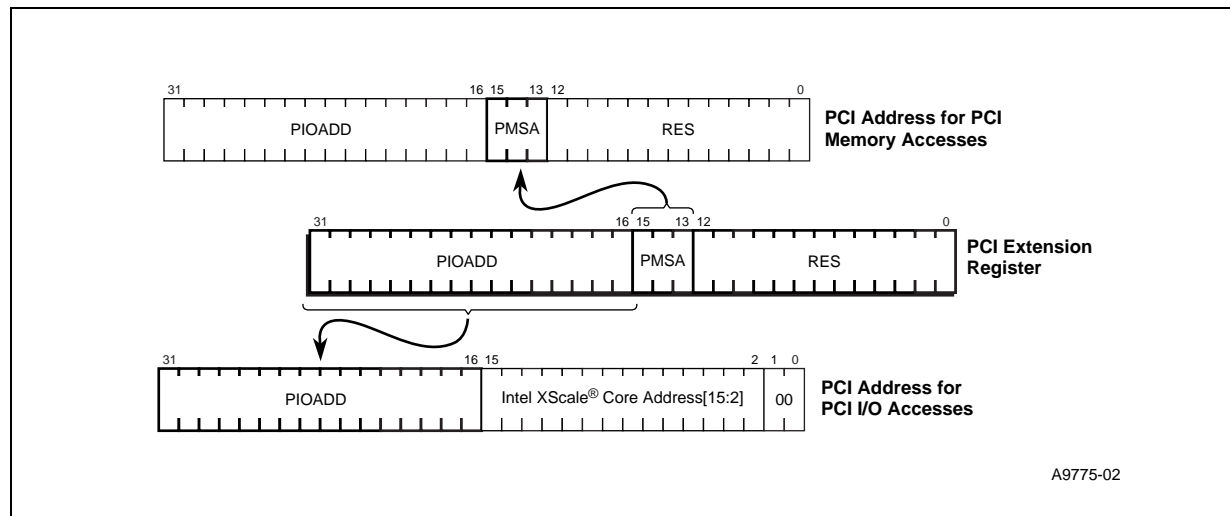
Data from the SRAM push/pull data bus is transferred through the master data port of the FBus interface of the PCI core. The PCI Core handles all of the PCI Bus protocol handshakes. The SRAM pull data received for a write command will be transferred to the Master write FIFO for PCI writes. For PCI reads, data is transferred from the read FIFO to the SRAM push data bus. A 32-byte Direct buffer is used to support up to 32 bytes of data responses to the direct access to PCI Bus.

The Command Bus Master access to the PCI bus will require internal arbitration to gain access to the data FIFOs inside the core, which are shared between the DMA engine and direct access to PCI.

9.4.2.3.1 PCI Address Generation for IO and MEM Cycles

When the push/pull command bus master is accessing the PCI Bus, the PCI address is generated based on the PCI address extension register (PCI_ADDR_EXT). Figure 128 shows how the address is generated from a Command Bus Master address.

Figure 128. PCI Address Generation for Command Bus Master to PCI






9.4.2.3.2 PCI Address Generation for Configuration Cycles

When a push/pull command bus master is accessing the PCI Bus to generate a configuration cycle, the PCI address is generated based on the a Command Bus Master address as shown in [Table 128](#) and [Figure 129](#):

Table 128. Command Bus Master Configuration Transactions

Cycle	Result
Type 1 Configuration Cycle	Command Bus address bits [31:24] are equal to 0xDA
Type 0 Configuration Cycle	Command Bus address bits [31:24] are equal to 0xDB.

Figure 129. PCI Address Generation for Command Bus Master to PCI Configuration Cycle



9.4.2.3.3 PCI Address Generation for Special and IACK Cycles

The PCI address is undefined for special and IACK PCI cycles.

9.4.2.3.4 PCI Enables

The PCI byte-enables are generated based on the Command Bus Master instruction, and the PCI unit does not change the states of the enables.

9.4.2.3.5 PCI Command

The PCI command is derived from the Command Bus Master address space map. The different spaces supported are listed in [Table 129](#):

Table 129. Command Bus Master Address Space Map to PCI

PCI Command	Intel XScale® Core Address Space
PCI Memory	0xE000 0000 – 0xFFFF FFFF
Local CSR	0xDF00 0000 – 0xDFFF FFFF
Local Configuration Register	0xDE00 0000 – 0xDEFF FFFF
PCI Special Cycle/PCI IACK Read	0xDC00 0000 – 0xDDFF FFFF
PCI Type 1 Configuration Cycle	0xDB00 0000 – 0xDBFF FFFF
PCI Type 0 Configuration Cycle	0xDA00 0000 – 0DAFF FFFF
PCI I/O	0xD800 0000 – 0xD8FF FFFF



9.5 PCI Unit Error Behavior

9.5.1 PCI Target Error Behavior

9.5.1.1 Target Access Has an Address Parity Error

1. If PCI_CMD_STAT[PERR_RESP] is not set, PCI Unit will ignore the parity error.
2. If PCI_CMD_STAT[PERR_RESP] is set:
 - a. PCI core will not claim the cycle regardless of internal device select signal.
 - b. PCI core will let the cycle terminate with master abort.
 - c. PCI core will not assert PCI_SERR_L.
 - d. Slave Interface sets PCI_CONTROL[TGT_ADR_ERR], which will interrupt the Intel XScale® core if enabled.

9.5.1.2 Initiator Asserts PCI_PERR_L in Response to One of Our Data Phases

1. Core does nothing.
2. Responsibility lies with the initiator to discard data, report this to the system, etc.

9.5.1.3 Discard Timer Expires on a Target Read

1. PCI unit discards the read data.
2. PCI Unit invalidates the delayed read address
3. PCI Unit sets Discard Timer Expired bit (DTX) in the PCI_CONTROL.
4. If enabled (XSCALE_INT_ENABLE [DTE]), the PCI unit interrupts the Intel XScale® core.

9.5.1.4 Target Access to the PCI_CSR_BAR Space Has Illegal Byte Enables

Note: The acceptable byte enables are:

1. PCI local CSRs - PCI_BE[3:0] = 0x0 or 0xF.
2. CSRs not in the PCI Unit - PCI_BE[3:0] = 0x0, 0xE, 0xD, 0xB, 0x7, 0xC, 0x3, or 0xF.

When byte-enables are detected, the hardware asserts the following error conditions:

1. Slave Interface will set PCI_CONTROL[TGT_CSR_BE].
2. Slave Interface will issue target abort for target read and drop the transaction for target write.



9.5.1.5 Target Write Access Receives Bad Parity PCI_PAR with the Data

1. If PCI_CMD_STAT[PERR_RESP] is not set, PCI Unit will ignore the parity error.
2. If PCI_CMD_STAT[PERR_RESP] is set:
 - a. core asserts PCI_PERR_L and sets PCI_CMD_STAT[PERR].
 - b. Slave Interface sets PCI_CONTROL[TGT_WR_PAR], which will interrupt the Intel XScale® core if enabled.
 - c. Data is discarded.

9.5.1.6 SRAM Responds with a Memory Error on One or More Data Phases on a Target Read

1. Slave Interface sets PCI_CONTROL[TGT_SRAM_ERR], which will interrupt the Intel XScale® core if enabled.
2. Assert PCI Target Abort at or before the data in question is driven on PCI.

9.5.1.7 DRAM Responds with a Memory Error on One or More Data Phases on a Target Read

1. Slave Interface sets PCI_CONTROL[TGT_DRAM_ERR], which will interrupt the Intel XScale® core if enabled.
2. Slave Interface asserts PCI Target Abort at or before the data in question is driven on PCI.

9.5.2 As a PCI Initiator During a DMA Transfer

9.5.2.1 DMA Read from DRAM (Memory-to-PCI Transaction) Gets a Memory Error

1. Set PCI_CONTROL[DMA_DRAM_ERR] which will interrupt the Intel XScale® core if enabled.
2. Master Interface terminates transaction before bad data is transferred (okay to terminate earlier).
3. Master Interface clears the Channel Enable bit in CHAN_X_CONTROL.
4. Master Interface sets DMA channel error bit in CHAN_X_CONTROL.
5. Master Interface does not reset the DMA CSRs; This leaves the descriptor pointer pointing to the DMA descriptor of the failed transfer.
6. Master Interface resets the state machines and DMA buffers.



9.5.2.2 DMA Read from SRAM (Descriptor Read) Gets a Memory Error

1. Set PCI_CONTROL[DMA_SRAM_ERR] which will interrupt the Intel XScale® core if enabled.
2. Master Interface clears the Channel Enable bit in CHAN_X_CONTROL.
3. Master Interface sets DMA channel error bit in CHAN_X_CONTROL.
4. Master Interface does not reset the DMA CSRs; This leaves the descriptor pointer pointing to the DMA descriptor of the failed transfer.
5. Master Interface resets the state machines and DMA buffers.

9.5.2.3 DMA from DRAM Transfer (Write to PCI) Receives PCI_PERR_L on PCI Bus

1. If PCI_CMD_STAT[PERR_RESP] is not set, PCI Unit will ignore the parity error.
2. If PCI_CMD_STAT[PERR_RESP] is set:
 - a. Master Interface sets PCI_CONTROL[DPE] which will interrupt the Intel XScale® core if enabled.
 - b. Master Interface clears the Channel Enable bit in CHAN_X_CONTROL.
 - c. Master Interface sets DMA channel error bit in CHAN_X_CONTROL.
 - d. Master Interface does not reset the DMA CSRs; This leaves the descriptor pointer pointing to the DMA descriptor of the failed transfer.
 - e. Master Interface resets the state machines and DMA buffers.
 - f. Core sets PCI_CMD_STAT[PERR] if properly enabled.

9.5.2.4 DMA To DRAM (Read from PCI) Has Bad Data Parity

1. If PCI_CMD_STAT[PERR_RESP] is not set, PCI Unit will ignore the parity error.
2. If PCI_CMD_STAT[PERR_RESP] is set:
 - a. Core asserts PCI_PERR_L on PCI if PCI_CMD_STAT[PERR_RESP] is set.
 - b. Master Interface sets PCI_CONTROL[DPED] which can interrupt the Intel XScale® core if enabled.
 - c. Master Interface clears the Channel Enable bit in CHAN_X_CONTROL.
 - d. Master Interface sets DMA channel error bit in CHAN_X_CONTROL.
 - e. Master Interface does not reset the DMA CSRs; This leaves the descriptor pointer pointing to the DMA descriptor of the failed transfer.
 - f. Master Interface resets the state machines and DMA buffers.



9.5.2.5 DMA Transfer Experiences a Master Abort (Time-Out) on PCI

Note: That is, nobody asserts DEVSEL during the DEVSEL window.

1. Master Interface sets PCI_CONTROL[RMA] which will interrupt the Intel XScale® core if enabled.
2. Master Interface clears the Channel Enable bit in CHAN_X_CONTROL.
3. Master Interface sets DMA channel error bit in CHAN_X_CONTROL.
4. Master Interface does not reset the DMA CSRs; This leaves the descriptor pointer pointing to the DMA descriptor of the failed transfer.
5. Master Interface resets the state machines and DMA buffers

9.5.2.6 DMA Transfer Receives a Target Abort Response During a Data Phase

1. Core terminates the transaction.
2. Master Interface sets PCI_CONTROL[RTA] which can interrupt the Intel XScale® core if enabled.
3. Master Interface clears the Channel Enable bit in CHAN_X_CONTROL.
4. Master Interface sets DMA channel error bit in CHAN_X_CONTROL.
5. Master Interface does not reset the DMA CSRs; This leaves the descriptor pointer pointing to the DMA descriptor of the failed transfer.
6. Master Interface resets the state machines and DMA buffers.

9.5.2.7 DMA Descriptor Has a 0x0 Word Count (Not an Error)

1. No data is transferred.
2. Descriptor is retired normally.

9.5.3 As a PCI Initiator During a Direct Access from the Intel XScale® Core or Microengine

9.5.3.1 Master Transfer Experiences a Master Abort (Time-Out) on PCI

1. Core aborts the transaction.
2. Master Interface sets PCI_CONTROL[RMA] which will interrupt the Intel XScale® core if enabled.

9.5.3.2 Master Transfer Receives a Target Abort Response During a Data Phase

1. Core aborts the transaction.
2. Master Interface sets PCI_CONTROL[RTA] which will interrupt the Intel XScale® core if enabled.



9.5.3.3 Master from the Intel XScale® Core or Microengine Transfer (Write to PCI) Receives PCI_PERR_L on PCI Bus

1. If PCI_CMD_STAT[PERR_RESP] is not set, PCI Unit will ignore the parity error.
2. If PCI_CMD_STAT[PERR_RESP] is set:
 - a. Core sets PCI_CMD_STAT[PERR].
 - b. Master Interface sets PCI_CONTROL[DPE] which will interrupt the Intel XScale® core if enabled.

9.5.3.4 Master Read from PCI (Read from PCI) Has Bad Data Parity

1. If PCI_CMD_STAT[PERR_RESP] is not set, PCI Unit will ignore the parity error.
2. If PCI_CMD_STAT[PERR_RESP] is set:
 - a. Core asserts PCI_PERR_L on PCI.
 - b. Master Interface sets PCI_CONTROL[DPED] which will interrupt the Intel XScale® core if enabled.
 - c. Data that has been read from PCI is sent to the Intel XScale® core or Microengine with a data error indication.

9.5.3.5 Master Transfer Receives PCI_SERR_L from the PCI Bus

Master Interface sets PCI_CONTROL[RSERR] which will interrupt the Intel XScale® core if enabled.

9.5.3.6 Intel XScale® Core Microengine Requests Direct Transfer when the PCI Bus is in Reset

Master Interface will complete the transfer and drop the write data and return all ones on the read data.

9.6 PCI Data Byte Lane Alignment

During any endian conversion, PCI does not need to do any longword swapping between two 32-bit longwords (LW1, LW0). But PCI may need to do byte swapping within the 32-bit longwords. Because of the different endian convention between PCI Bus and the memory, all data going between the PCI core FIFO and memory data bus passes through the byte lane reversal as shown in [Table 130](#) through [Table 137](#).

PCI allows byte-enable swapping only without the data swapping or allow data swapping only without byte enable swapping. When PCI handle the mis align data in above two cases, PCI will only care about valid data. So PCI will drive any data values for those mis-aligned invalid data portions.



Table 130. Byte Lane Alignment for 64-Bit PCI Data In (64 Bits PCI Little-Endian to Big-Endian with Swap)

PCI Data	IN[63:56]	IN[55:48]	IN[47:40]	IN[39:32]	IN[31:24]	IN[23:16]	IN[15:8]	IN[7:0]
SRAM Data	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 (32 bits) LW0 drive first			
DRAM Data	OUT[39:32]	OUT[47:40]	OUT[55:48]	OUT[63:56]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]

Table 131. Byte Lane Alignment for 64-Bit PCI Data In (64 Bits PCI Big-Endian to Big-Endian without Swap)

PCI Data	IN[39:32]	IN[47:40]	IN[55:48]	IN[63:56]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
SRAM Data	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 (32 bits) LW0 drive first			
DRAM Data	OUT[39:32]	OUT[47:40]	OUT[55:48]	OUT[63:56]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]

Table 132. Byte Lane Alignment for 32-Bit PCI Data In (32 Bits PCI Little-Endian to Big-Endian with Swap)

	PCI Add[2]=1				PCI Add[2]=0			
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
PCI Data	IN[31:24]	IN[23:16]	IN[15:8]	IN[7:0]	IN[31:24]	IN[23:16]	IN[15:8]	IN[7:0]
SRAM Data	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
DRAM Data	OUT[39:32]	OUT[47:40]	OUT[55:48]	OUT[63:56]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]

Table 133. Byte Lane Alignment for 32-Bit PCI Data In (32 Bits PCI Big-Endian to Big-Endian without Swap)

	PCI Add[2]=1				PCI Add[2]=0			
	Longword1 (32 bits) LW1 drive first				Longword0 ((32 bits) LW0 drive first			
PCI Data	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
SRAM Data	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
direct map pci to dram	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
DRAM Data	OUT[39:32]	OUT[47:40]	OUT[55:48]	OUT[63:56]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]



Table 134. Byte Lane Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Little Endian with Swap)

SRAM Data	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
DRAM Data	IN[39:32]	IN[47:40]	IN[55:48]	IN[63:56]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
PCI Side	OUT[63:56]	OUT[55:48]	OUT[47:40]	OUT[39:32]	OUT[31:24]	OUT[23:16]	OUT[15:8]	OUT[7:0]

Table 135. Byte Lane Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Big-Endian without Swap)

SRAM Data	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
DRAM Data	IN[39:32]	IN[47:40]	IN[55:48]	IN[63:56]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
direct map pci to dram	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
PCI Side	OUT[39:32]	OUT[47:40]	OUT[55:48]	OUT[63:56]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]

Table 136. Byte Lane Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Little Endian with Swap)

SRAM Data	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
DRAM Data	IN[39:32]	IN[47:40]	IN[55:48]	IN[63:56]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
PCI Data	OUT[31:24]	OUT[23:16]	OUT[15:8]	OUT[7:0]	OUT[31:24]	OUT[23:16]	OUT[15:8]	OUT[7:0]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
	PCI Add[2]=1				PCI Add[2]=0			

Table 137. Byte Lane Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Big-Endian without Swap)

SRAM Data	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
DRAM Data	IN[39:32]	IN[47:40]	IN[55:48]	IN[63:56]	IN[7:0]	IN[15:8]	IN[23:16]	IN[31:24]
PCI Data	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]	OUT[7:0]	OUT[15:8]	OUT[23:16]	OUT[31:24]
	Longword1 (32 bits) LW1 drive after LW0				Longword0 ((32 bits) LW0 drive first			
	PCI Add[2]=1				PCI Add[2]=0			



The BE_DEMI bit of the PCI_CONTROL register can be set to enable big-endian on the incoming data from the PCI Bus to both the SRAM and DRAM. The BE_DEMO bit of the PCI_CONTROL register can be set to enable big-endian on the outgoing data to the PCI Bus from both the SRAM and DRAM.

9.6.1 Endian for Byte Enable

During any endian conversion, PCI does not need to do any longword byte enable swapping between two 32-bit longwords (LW1, LW0). But PCI may need to do byte enable swapping within the 32-bit longword byte enable. Because of the different endian convention between PCI Bus and the memory, all data going between the PCI core FIFO and memory data bus passes through the byte lane reversal as shown in Table 138 through Table 145:

Table 138. Byte Enable Alignment for 64-Bit PCI Data In (64 Bits PCI Little-Endian to Big-Endian with Swap)

PCI Data	IN_BE[7]	IN_BE[6]	IN_BE[5]	IN_BE[4]	IN_BE[3]	IN_BE[2]	IN_BE[1]	IN_BE[0]
SRAM Data	OUT_BE[3]	OUT_BE[2]	OUT_BE[1]	OUT_BE[0]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
DRAM Data	OUT_BE[4]	OUT_BE[5]	OUT_BE[6]	OUT_BE[7]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]

Table 139. Byte Enable Alignment for 64-Bit PCI Data In (64 Bits PCI Big-Endian to Big-Endian without Swap)

PCI Data	IN_BE[4]	IN_BE[5]	IN_BE[6]	IN_BE[7]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
SRAM Data	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
DRAM Data	OUT_BE[4]	OUT_BE[5]	OUT_BE[6]	OUT_BE[7]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]

Table 140. Byte Enable Alignment for 32-Bit PCI Data In (32 bits PCI Little-Endian to Big-Endian with Swap)

	PCI Add[2]=1				PCI Add[2]=0			
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
PCI Data	IN_BE[3]	IN_BE[2]	IN_BE[1]	IN_BE[0]	IN_BE[3]	IN_BE[2]	IN_BE[1]	IN_BE[0]
SRAM Data	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
DRAM Data	OUT_BE[4]	OUT_BE[5]	OUT_BE[6]	OUT_BE[7]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]



Table 141. Byte Enable Alignment for 32-Bit PCI Data In (32 Bits PCI Big-Endian to Big-Endian without Swap)

	PCI Add[2]=1				PCI Add[2]=0			
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
PCI Data	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
SRAM Data	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
direct map pci to dram	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
DRAM Data	OUT_BE[4]	OUT_BE[5]	OUT_BE[6]	OUT_BE[7]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]

Table 142. Byte Enable Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Little Endian with Swap)

	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
SRAM Data	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
DRAM Data	IN_BE[4]	IN_BE[5]	IN_BE[6]	IN_BE[7]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
PCI Side	OUT_BE[7]	OUT_BE[6]	OUT_BE[5]	OUT_BE[4]	OUT_BE[3]	OUT_BE[2]	OUT_BE[1]	OUT_BE[0]

Table 143. Byte Enable Alignment for 64-Bit PCI Data Out (Big-Endian to 64 Bits PCI Big Endian without Swap)

SRAM Data	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
DRAM Data	IN_BE[4]	IN_BE[5]	IN_BE[6]	IN_BE[7]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
PCI Side	OUT_BE[4]	OUT_BE[5]	OUT_BE[6]	OUT_BE[7]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]

Table 144. Byte Enable Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Little Endian with Swap)

SRAM Data	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
DRAM Data	IN_BE[4]	IN_BE[5]	IN_BE[6]	IN_BE[7]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
PCI Data	OUT_BE[3]	OUT_BE[2]	OUT_BE[1]	OUT_BE[0]	OUT_BE[3]	OUT_BE[2]	OUT_BE[1]	OUT_BE[0]
	Longword1byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
	PCI Add[2]=1				PCI Add[2]=0			



Table 145. Byte Enable Alignment for 32-Bit PCI Data Out (Big-Endian to 32 Bits PCI Big Endian without Swap)

SRAM Data	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
	Longword1 byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
DRAM Data	IN_BE[4]	IN_BE[5]	IN_BE[6]	IN_BE[7]	IN_BE[0]	IN_BE[1]	IN_BE[2]	IN_BE[3]
PCI Data	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]	OUT_BE[0]	OUT_BE[1]	OUT_BE[2]	OUT_BE[3]
	Longword1 byte enable LW1 byte enable drive after LW0 byte enable				Longword0 byte enable LW0 byte enable drive first			
	PCI Add[2]=1				PCI Add[2]=0			

The BE_BEMI bit of the PCI_CONTROL register can be set to enable big-endian on the incoming byte enable from the PCI Bus to both the SRAM and DRAM. The BE_BEMO bit of the PCI_CONTROL register can be set to enable big-endian on the outgoing byte enable to the PCI Bus from both the SRAM and DRAM.

The B-stepping silicon provides a mechanism to enable byte swapping for PCI I/O operations as described in [Table 146](#).



Table 146. PCI I/O Cycles with Data Swap Enable

Stepping	Description			
A Stepping	A PCI IO cycle is treated like CSR where the data bytes are not swapped. It is sent in the same byte order whether the PCI bus is configured in Big-Endian or Little-Endian mode.			
B Stepping	When PCI_CONTROL[IEE] is 0, PCI data is sent in the same byte order whether the PCI bus is configured in Big-Endian or Little-Endian mode. When PCI_CONTROL[IEE] is 1, PCI IO data will follow the same memory space swapping rule. The address always follows the physical location, Example:			
	BEs not Swapped (1 byte access)		BEs Swapped (1 byte access)	
	ad[1:0]	BE3 BE2 BE1 BE0	ad[1:0]	BE3 BE2 BE1 BE0
	00	1 1 1 0	1 1	0 1 1 1
	01	1 1 0 1	1 0	1 0 1 1
	10	1 0 1 1	0 1	1 1 0 1
	11	0 1 1 1	0 0	1 1 1 0
	BEs not Swapped (2 byte access)		BEs Swapped (2 byte access)	
	ad[1:0]	BE3 BE2 BE1 BE0	ad[1:0]	BE3 BE2 BE1 BE0
	00	1 1 0 0	1 0	0 0 1 1
	01	1 0 0 1	0 1	1 0 0 1
	10	0 0 1 1	0 0	1 1 0 0
	BEs not Swapped (3 byte access)		BEs Swapped (3 byte access)	
	ad[1:0]	BE3 BE2 BE1 BE0	ad[1:0]	BE3 BE2 BE1 BE0
	00	1 0 0 0	0 1	0 0 0 1
	01	0 0 0 1	0 0	1 0 0 0
	BEs not Swapped (4 byte access)		BEs Swapped (4 byte access)	
	ad[1:0]	BE3 BE2 BE1 BE0	ad[1:0]	BE3 BE2 BE1 BE0
	00	0 0 0 0	0 0	0 0 0 0



Clocks and Reset

10

This section describes the IXP2800 Network Processor clocks and reset. Refer to the *Intel® IXP2800 Network Processor Hardware Initialization Reference Manual* for information about the initialization of all units of the IXP2800 Network Processor.

10.1 Clocks

The block diagram in [Figure 130](#) shows how the IXP2800 Network Processor implements an onboard clock generator to generate the internal clocks used by the various functional units in the device. It takes an external reference frequency and multiplies it to a higher frequency clock using a PLL. That clock is then divided down by a set of programmable dividers to provide clocks to SRAM and DRAM controllers.

The Intel XScale® core and Microengines get clocks using fixed divide ratios. The Media and Switch Fabric Interface clock is selected based on the strap pin (CFG_MSF_FREQ_SEL) so that when CFG_MSF_FREQ_SEL is high, an internally-generated clock using the programmable divider is used and when CFG_MSF_FREQ_SEL is low, an externally-received clock on the MSF interface is used.

The PCI controller uses external clocks. Each of the units also interfaces to internal buses, which run at ½ the Microengine frequency. [Figure 130](#) shows the overall clock generation and distribution and [Table 147](#) summarizes the clock usage.

Figure 130. Overall Clock Generation and Distribution

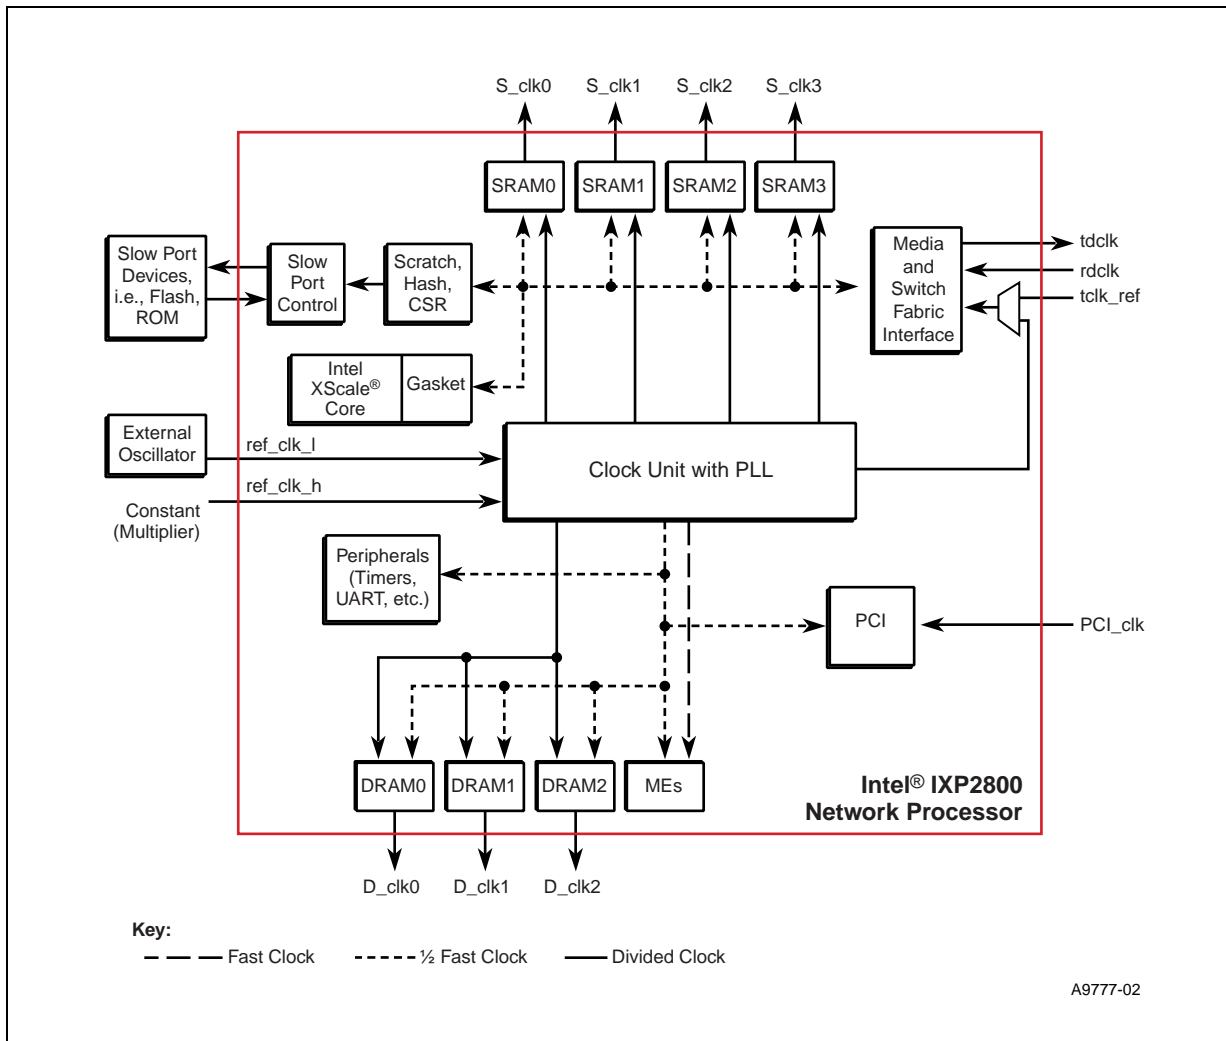


Table 147. Clock Usage Summary (Sheet 1 of 2)

Unit Name	Description	Comment
Microengine	Microengines internal.	
Internal Buses	Command/Push/Pull interface of DRAM, SRAM, Intel XScale® core, Peripheral, MSF, and PCI Units.	1/2 Microengine frequency.
Intel XScale® core	Intel XScale® core microprocessor, caches, microprocessor side of Gasket.	1/2 of Microengine frequency.
DRAM	DRAM pins and control logic (all of DRAM unit except Internal Bus interface).	Divide of Microengine frequency. All DRAM channels use the same frequency. Clocks are driven by the IXP2800 Network Processor to external DRAMs.



Table 147. Clock Usage Summary (Sheet 2 of 2)

Unit Name	Description	Comment
SRAM	SRAM pins and control logic (all of the SRAM unit except Internal Bus interface).	Divide of Microengine frequency. Each SRAM channel has its own frequency selection. Clocks are driven by the IXP2800 Network Processor to external SRAMs and/or Coprocessors.
Scratch, Hash, CSR	Scratch RAM, Hash Unit, CSR access block	1/2 of Microengine frequency. Note that Slowport has no clock. Timing for Slowport accesses is defined in Slowport registers.
MSF	Receive and Transmit pins and control logic.	The transmit clock for the Media and Switch interface can be derived in two different ways. <ul style="list-style-type: none"> From TCLK input signal (supplied by PHY device). Divided from internal clock. For details please refer to Chapter 8, "Media and Switch Fabric Interface" .
APB	APB logic	Divide of Microengine frequency.
PCI	PCI pins and control logic.	External reference. Either from Host system or on-board oscillator.

The fast frequency on the IXP2800 Network Processor is generated by an on-chip PLL that multiplies a reference frequency provided by an on-board LVDS oscillator (frequency 100 MHz) by a selectable multiplier. The multiplier is selected by using external strap pins SP_AD[5:0] and can be viewed by software via the STRAP_OPTIONS[CFG_PLL_MULT] CAP CSR register bits. The multiplier range is even multiples between 16 and 48, so the PLL can generate a 1.6 GHz to 4.8 GHz clock (with a 100-Mhz reference frequency).

The PLL output frequency is divided by 2 to get the Microengine clock and by 4 to get the Intel XScale® core and the internal Command/Push/Pull bus frequency. An additional division (after the divide by 2) is used to generate the clock frequencies for the other internal units. The divisors are programmable via the CLOCK_CONTROL CSR. APB divisor specified in the CLOCK_CONTROL CSR clock is scaled by 4 (i.e., a value of 2 in the CSR selects a divisor of 8).

[Table 148](#) shows the frequencies that are available based on a 100-Mhz oscillator and various values of PLL multipliers, for the supported divisor values of 3 to 15.



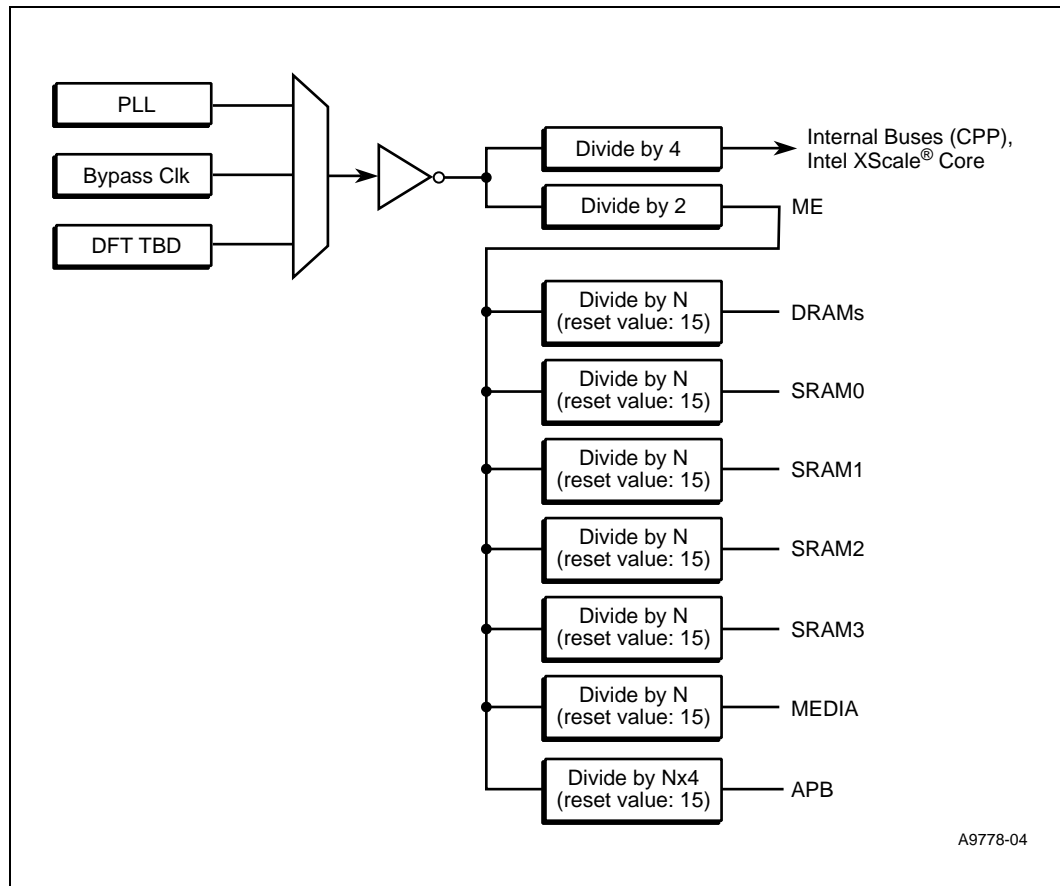
Table 148. Clock Rates Examples

Input Oscillator Frequency (MHz)		100							
PLL Output Frequency (MHz) [PLL Multiplier] ¹		2000 [20]	2200 [22]	2400 [24]	2600 [26]	2800 [28]	4000 [40]	4800 [48]	
Microengine Frequency ²		1000	1100	1200	1300	1400	2000	2400	
Intel XScale® core & Command/Push/Pull Bus Frequency ³		500	550	600	650	700	1000	1200	
Divide Ratio for other Units (except APB) ⁴	Divisor ⁵	2 ⁶	500	550	600	650	700	1000	1200
		3	333	367	400	433	467	666	800
		4	250	275	300	325	350	500	600
		5	200	220	240	260	280	400	480
		6	167	183	200	217	233	334	400
		7	143	157	171	186	200	286	342
		8	125	138	150	163	175	250	300
		9	111	122	133	144	156	222	266
		10	100	110	120	130	140	200	240
		11	91	100	109	118	127	182	218
		12	83	92	100	108	117	166	200
		13	77	85	92	100	107	154	184
		14	71	79	86	93	100	142	172
15	67	73	80	87	93	134	160		

1. This multiplier is selected via SP_AD[5:0] strap pins.
2. This frequency is the PLL output frequency divided by 2.
3. This frequency is the PLL output frequency divided by 4.
4. The ABP divisor specified in the CLOCK_CONTROL CAP CSR is scaled by an additional x4.
5. This divisor is selected via the CLOCK_CONTROL CAP CSR. The Base Frequency is the PLL output frequency divided by 2.
6. This divide ratio is only used by test logic. In the normal functional mode, this ratio is reserved for Push/Pull clocks only.

Figure 131 shows the clocks generation circuitry for the IXP2800 Network Processor. When the chip is powered up, bypass clock will be sent to all the units. After the PLL is locked, clock unit will switch all units from bypass clock to a fixed frequency clock which is generated by dividing PLL OUTPUT FREQUENCY by 16. Once Clock Control CSR is written, clock unit will replace fixed frequency clock with the defined clocks for different units.

Figure 131. IXP2800 Network Processor Clock Generation



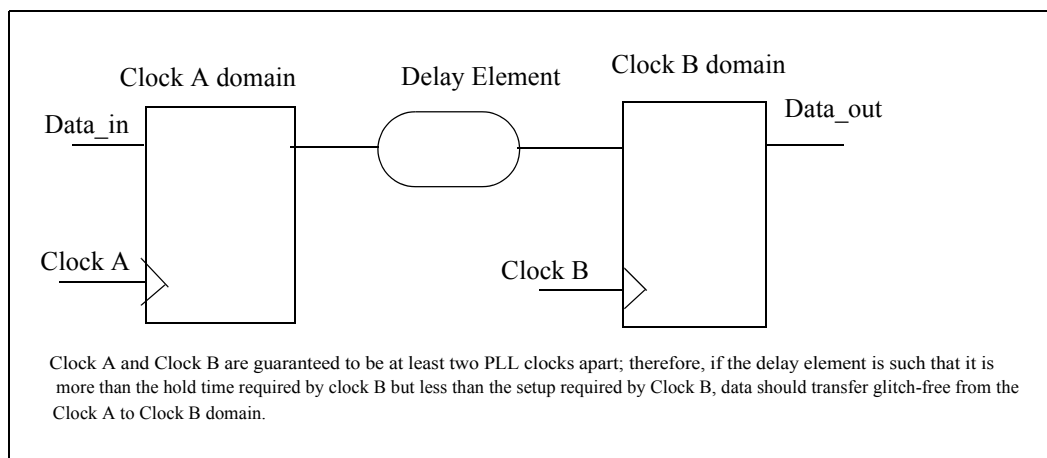
10.2 Synchronization Between Frequency Domains

Due to the internal design architecture of the IXP2800 Network Processor, it is guaranteed that one of the clock domains of an asynchronous transfer will be the Push/Pull domain (PLL/4). Additionally, all other clocks are derived by further dividing the Microengine clock (PLL/2n where n is 3 or more); refer to [Figure 132](#).

Note: The exception is the PCI unit where the PCI clock is fully asynchronous with the PP clock. Therefore in the PCI unit, data is synchronized using the usual 3-flop synchronization method.

Therefore, the clock A and clock B relationship will always be apart by at least two PLL clocks. To solve hold problem between clock A and clock B, a delay is added anytime data is transferred from clock A to clock B. The characteristic of this delay element is such that it is high enough to resolve any hold issue in fast environment but in the slow environment its delay is still less than two PLL clocks.

Figure 132. Synchronization Between Frequency Domains



10.3 Reset

The IXP2800 Network Processor can be reset four ways.

- Hardware Reset Using nRESET or PCI_RST_L.
- PCI-Initiated Reset.
- Watchdog Timer Initiated Reset.
- Software Initiated Reset.

10.3.1 Hardware Reset Using nRESET or PCI_RST_L

The IXP2800 Network Processor provides the nRESET pin so that it can be reset by an external device. Asserting this pin resets the internal functions and generates an external reset via the nRESET_OUT pin.

Upon power-up, nRESET (or PCI_RST_L) must remain asserted for 1ms after VDD is stable to properly reset the IXP2800 Network Processor and ensure that the external clocks are stable. While nRESET is asserted, the processor is held in reset. When nRESET is released, the Intel XScale® core begins executing from address 0x0. If PCI_RST_L is input to the chip, nRESET should be removed before or at the same time as PCI_RST_L.

All the strap options are latched with nRESET except for PCI strap option BOARD_IS_64 which is latched with PCI_RST_L only (by latching the status of REQ64_L at the trailing edge of PCI_RST_L).

If nRESET is asserted, while the Intel XScale® core is executing, the current instruction is terminated abnormally and the reset sequence is initiated.

The nRESET_OUT signal de-assertion depends upon settings of reset_out_strap and IXP_RESET_0[22] also called the EXTRST_EN bit. During power up, IXP_RESET_0[22] is reset to 0; therefore the value to be driven on nRESET_OUT is defined by reset_out_strap. When



“reset_out_strap” is sampled as 0 on the trailing edge of reset, nRESET_OUT is de-asserted based on the value of IXP_RESET_0[15] which is written by software. If “reset_out_strap” is sampled as 1 on the trailing edge of reset, nRESET_OUT is de-asserted after PLL locks.

During normal function mode, if software wants to pull nRESET_OUT high, it should set IXP_RESET_0[22] = 1 and then set IXP_RESET_0[15] = 1. To pull nRESET_OUT low, software should set the IXP_RESET_0[15] bit back to 0.

Figure 133. Reset Out Behavior

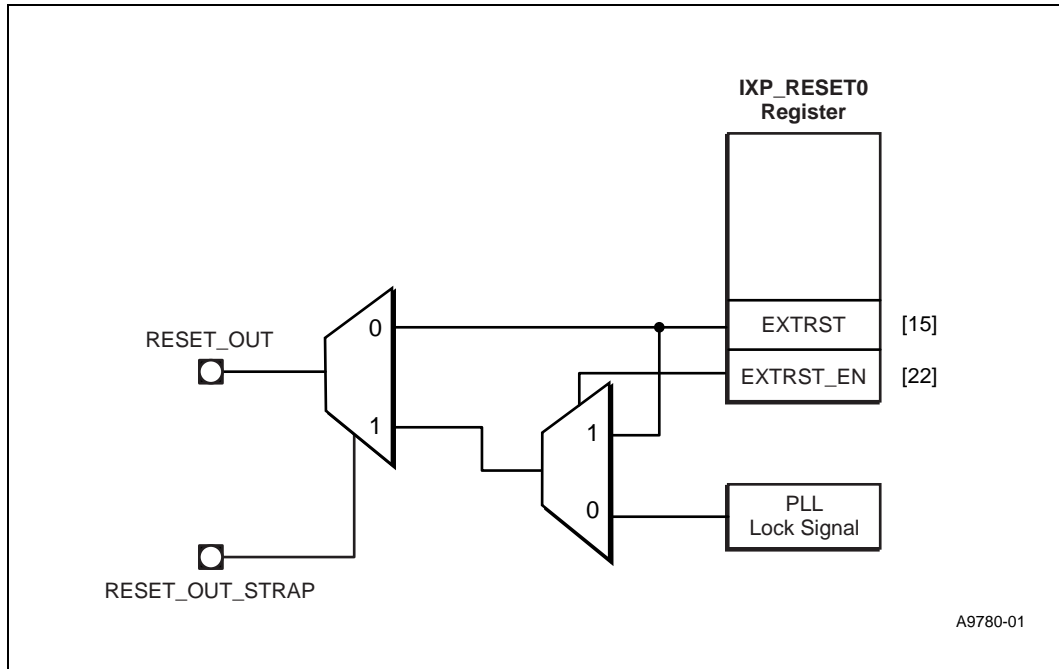
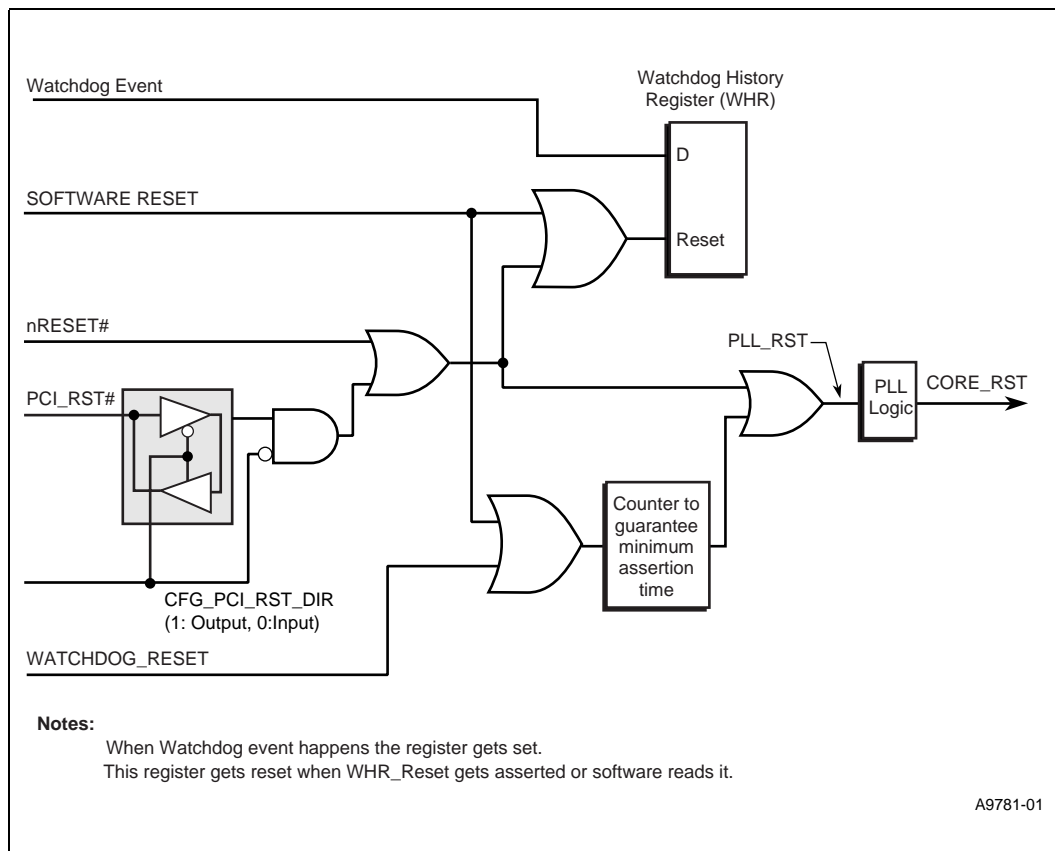


Figure 134. Reset Generation



10.3.2 PCI-Initiated Reset

CFG_RST_DIR is not asserted and PCI_RST_L is asserted.

When the CFG_RST_DIR strap pin is not asserted (sampled 0), PCI_RST_L is input to the IXP2800 Network Processor and is used to reset all the internal functions. Its behavior is the same as a hardware reset using nRESET pin.

10.3.3 Watchdog Timer-Initiated Reset

The IXP2800 Network Processor provides a watchdog timer that can cause a reset if the Watchdog timer expires and the Watchdog enable bit WDE in the Timer Watchdog Enable register is also set. The Intel XScale® core should be programmed to reset the watch dog timer periodically to ensure that it does not expire. If a watchdog timer expires, it is assumed that the Intel XScale® core has ceased executing instructions properly. When the timer expires, the Watchdog History register bit[0] is set which can be read by the software later on.

The following sections define IXP2800 Network Processor behavior for the watchdog event.



10.3.3.1 Slave Network Processor (Non-Central Function)

- If the Watchdog timer reset enable bit set to 1, Watchdog reset will trigger the soft reset
- If the Watchdog timer reset enable bit set to 0, Watchdog reset will trigger the PCI interrupt to external PCI host (if interrupt is enabled by PCI Outbound Interrupt Mask Register[3]). External PCI host can check the IXP2800 error status and log the error then reset the Slave IXP2800 Network Processor only or reset all the PCI devices (assert the PCI_RST_L).
- If the Watchdog history bit is already set when a new watchdog event happens, the Watchdog timer reset enable bit is disregarded and a soft reset is generated.

10.3.3.2 Master Network Processor (PCI Host, Central Function)

- If the Watchdog timer reset enable bit is set to 1, Watchdog reset will trigger the soft reset and set the watchdog history bit.
- If the Watchdog timer reset enable bit is set to 0, check the watchdog history bit. If is already set, generate soft reset. If the watchdog history bit is not set already, watchdog reset will just set the watchdog history bit and no further action is taken.

10.3.3.3 Master Network Processor (Central Function)

- If the Watchdog timer reset enable bit is set to 0, Watchdog reset will trigger the PCI interrupt to external PCI host (if interrupt is enabled by PCI Outbound Interrupt Mask Register[3]).
- If the Watchdog history bit is already set when a new watchdog event happens, the Watchdog timer reset enable bit is disregarded, and a soft reset is generated.
- If the Watchdog timer reset enable bit is set to 1, Watchdog reset will trigger the soft reset.

10.3.4 Software-Initiated Reset

The Intel XScale® core or external PCI bus master can reset specific functions in the IXP2800 Network Processor by writing to the IXP_RESET0 and IXP_RESET1 registers. All the individual microengines and specific units can be reset individually in this fashion.

Software reset initiated by the Reset All bit in the IXP_RESET0 register behaves almost the same as hardware resets in the sense that PLL and rest of the core gets reset. The only difference between soft reset and hard reset is that a 512-cycle counter is added at the output of the RESET_ALL bit going to the PLL unit for chip reset generation. The PCI unit in the meantime detects the bus idle condition and generates a local reset. This local reset is removed once chip reset is generated and chip reset then takes over the reset function of PCI unit.

Both hardware and software resets (software reset after 512 cycles delay) combined generate PLL_RST for the PLL logic. During the assertion of PLL_RST, PLL block remains in the bypass mode and passes the incoming clock directly to the core logic. At this time everyone inside the core gets the same basic clock. The Clock Control register is reset to 0x0FFF_FFFF using the same signal.

Once the PLL_RST signal goes away, the PLL starts generating divide_by_2 clock for the Microengines, divide_by_4 clock for the Intel XScale® core and divide_by_16 clock for the rest of the chip (not using divide_by_4 clock) after inserting 16 – 32 idle clocks. Once the clock control CSR is written by software, the PLL block detects it by finding a change in value of this register.



Once in operation, if the watchdog timer expires with watchdog timer enable bit WDE from Timer Watchdog Enable register set, a reset pulse from the watchdog timer logic goes to PLL unit after passing through a counter to guarantee minimum assertion time, which in turn resets the IXP_RESETE_n registers that cause the entire chip to be reset.

Figure 134 explains the reset generation for the PLL logic and for the rest of the core. CORE_RST is used inside the IXP2800 to reset everything; PLL_RST can be disabled.

10.3.5 Reset Removal Operation Based on CFG_PROM_BOOT

Reset removal based on the CFG_PROM_BOOT strap option (BOOT_PROM) can be divided into two parts:

1. When CFG_PROM_BOOT is 1 (BOOT_PROM is present).
2. When CFG_PROM_BOOT is 0 (BOOT_PROM is not present).

10.3.5.1 When CFG_PROM_BOOT is 1 (BOOT_PROM is Present)

After CORE_RST is de-asserted, reset from the Intel XScale® core, SHaC, and CMDARB is removed. Once the Intel XScale® core reset is removed, the Intel XScale® core starts initializing the chip. The Intel XScale® core writes the 'clock control CSR' to define the operating frequencies of different units. The Intel XScale® core writes IXP_RESETE0[21] to allow the PCI logic to start accepting transactions on the PCI bus as part of initialization process.

10.3.5.2 When CFG_PROM_BOOT is 0 (BOOT_PROM is Not Present)

After CORE_RST is de-asserted, IXP_RESETE0[21] is set, allowing the PCI unit to start accepting transactions on the PCI bus. In this mode, the Intel XScale® core is kept in reset. Reset from DRAM logic is removed by the PCI host by writing 0 to specific bits in the IXP_RESETE0 register.

10.3.6 Strap Pins

The IXP2800 Strap pins for reset and initialization operation are described in Table 149.



Table 149. IXP2800 Network Processor Strap Pins

Signal	Name	Description
CFG_RST_DIR	RST_DIR	PCI_RST direction pin: (Also called PCI_HOST) Need to be a dedicated pin. 1—IXP2800 Network Processor is the host supporting central function. PCI_RST_L is output. 0—IXP2800 Network Processor is not central function. PCI_RST_L is input. This pin is stored at XSC[31] (XScale_Control register) at the trailing edge of reset.
CFG_PROM_BOOT	GPIO[0]	PCI PROM BOOT Pin: 1—IXP2800 Network Processor will boot from PROM: Whether Intel XScale® core will configure the system or not will be defined by CFG_PCI_BOOT_HOST strap option. 0—IXP2800 Network Processor will not boot from PROM. So after host has downloaded image of boot code into DRAM, Intel XScale® core will boot from DRAM address 0. This pin is stored at XSC[29] (XScale_Control register) at the trailing edge of reset.
CFG_PCI_BOOT_HOST	GPIO[1]	PCI BOOT HOST Pin: 1—IXP2800 Network Processor will configure the PCI system. 0—IXP2800 Network Processor will not configure the PCI system. This pin is stored at XSC[28] (XScale_Control register) at the trailing edge of reset.
CFG_PCI_ARB	GPIO[2]	PCI Arbiter Pin: 1—IXP2800 Network Processor is the arbiter on the PCI bus. 0—IXP2800 Network Processor is not the arbiter on the PCI bus.
PLL_MULT[5:0]	SP_AD[5:0]	PLL Multiplier Valid values are 010000-110000 for a multiplier range of 16 – 48. Other values will result in undefined behavior by PLL.
RESET_OUT_STRAP	SP_AD[7]	When 1: nRESET_OUT is removed after PLL locks. When 0: nRESET_OUT is removed by software using bit IXP_RESET0[17].
CFG_PCI_SWIN[1:0]	GPIO[6:5]	SRAM Bar Window: 11—SRAM BAR size of 256 Mbytes 10—SRAM BAR size of 128 Mbytes 01—SRAM BAR size of 64 Mbytes 00—SRAM BAR size of 32 Mbytes
CFG_PCI_DWIN[1:0]	GPIO[4:3]	DRAM BAR Window: 11—DRAM BAR size of 1024 Mbytes 10—DRAM BAR size of 512 Mbytes 01—DRAM BAR size of 256 Mbytes 00—DRAM BAR size of 128 Mbytes
CFG_MS_FREQ_SEL	SP_AD[6]	Select source of MSF Tx Clock: 0—TCLK_Ref input pin 1—Internally generated clock



Table 150 lists the supported Strap combinations of CFG_PROM_BOOT, CFG_RST_DIR, and CFG_PCI_BOOT_HIST.

Table 150. Supported Strap Combinations

CFG_PROM_BOOT, CFG_RST_DIR, CFG_PCI_BOOT_HOST	Result
000	Allowed
001	Allowed
010	Not allowed
011	Not allowed
100	Allowed
101	Allowed
110	Allowed
111	Allowed

One more restriction in the PCI unit is that, if the IXP2800 Network Processor is a PCI_HOST or PCI_ARBITER, it should also be PCI_CENTRAL_FUNCTION.

10.3.7 Powerup Reset Sequence

When the system is powered up, bypass clock is sent to all the units as the chip begins to power up. It will merely be used to allow a gradual power up and to begin clocking state elements to remove possible circuit contention. When PLL gets locked after nRESET is de-asserted, it will start generating divide_by_16 clocks for all the units. Reset from the IXP_RESET register is also removed at the same time. When software updates the clock count register, clocks are again stopped for 32 cycles and then start again.

The reset sequence described above is the same in the case when reset happens through the PCI_RST_L signal and CFG_RST_DIR is asserted.

Once in operation, if watchdog timer expires with watchdog timer enable bit (bit [0] in the Timer Watchdog Enable register ON, a reset pulse from the watchdog timer logic resets the IXP_RESET n registers and in turn causes the entire network processor to be reset.

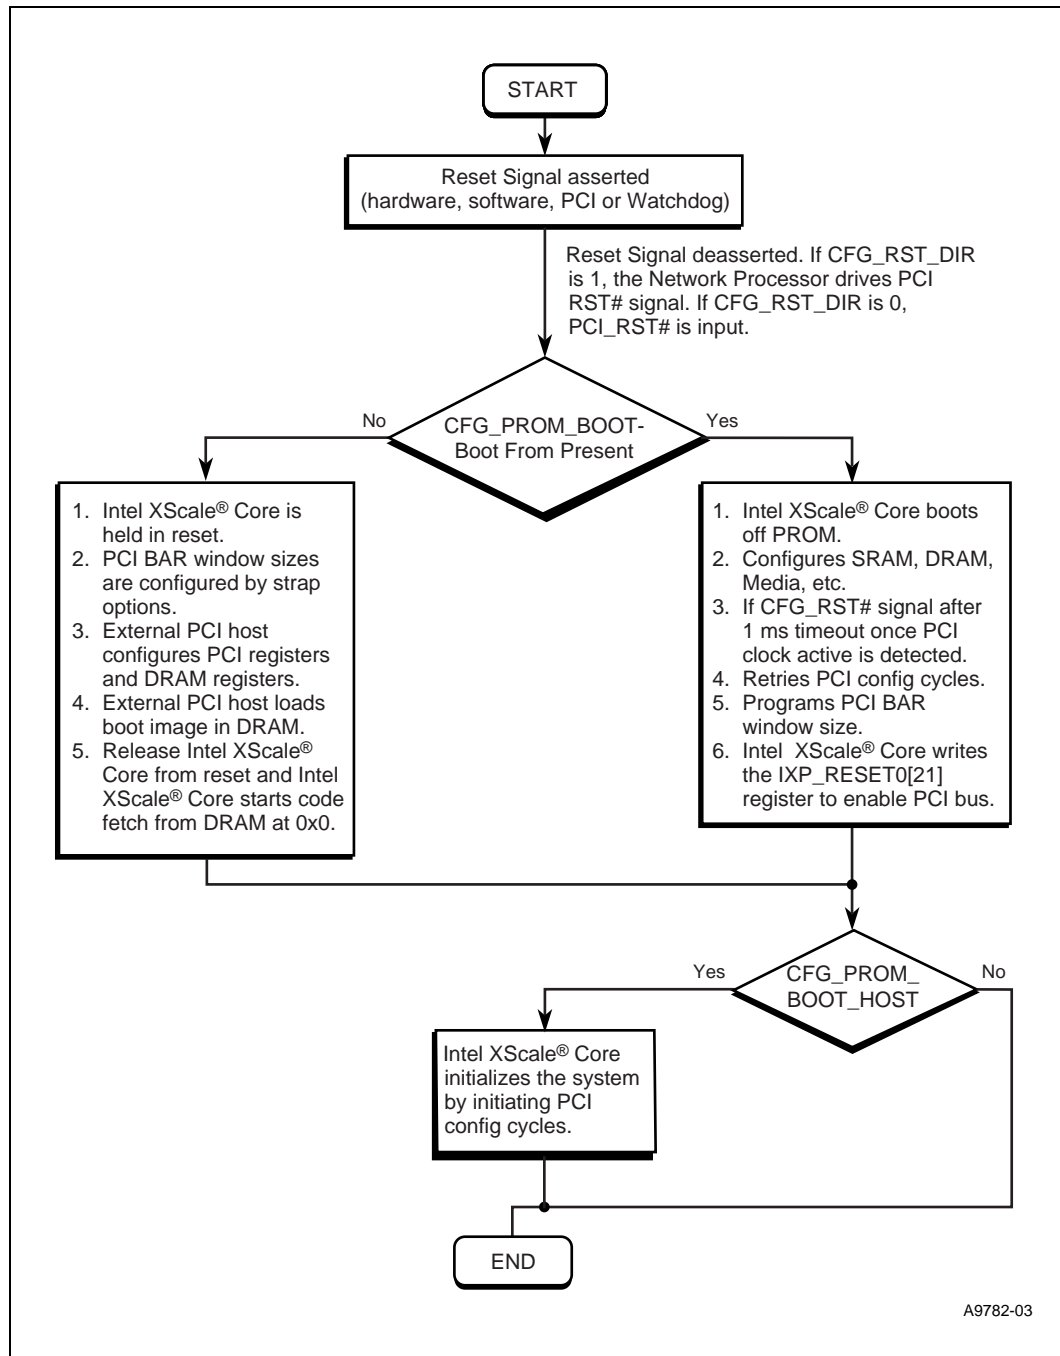
10.4 Boot Mode

The IXP2800 can boot in following two modes:

- Flash ROM
- PCI Host Download

Figure 135 shows the IXP2800 Network Processor Boot process.

Figure 135. Boot Process



A9782-03



10.4.1 Flash ROM

At power up, if FLASH_ROM is present, strap pin CFG_PROM_BOOT should be sampled 1 (should be pulled up). Therefore after reset being removed by the PLL logic from the IXP_RESET0 register, the Intel XScale® core reset is automatically removed. Flash Alias Disable (bit [8] of Misc Control register) information is used by the Intel XScale® core gasket to decide where to forward address 0 from the Intel XScale® core when the Intel XScale® core wakes up and starts accessing the code from address 0. In this mode, since “flash alias disable” bit is reset to 0, the Intel XScale® core gasket will convert access to address 0 to PROM access from address 0 using the CAP command. Based on the code residing inside PROM, the Intel XScale® core starts removing reset from SRAM, PCI, DRAM, Microengines etc. by writing 0 in their corresponding bit location of IXP_RESET n register and then initializing their configuration registers.

Boot code in PROM can change flash alias disable bit to 1 anytime to map DRAM at address 0 and therefore block further accesses to PROM at address 0. This change should be done before putting any data in DRAM at address 0.

The Intel XScale® core also sets different BARs inside PCI unit to define memory requirements for different windows.

The Intel XScale® core behavior as a host is controlled by CFG_PCI_BOOT_HOST strap option. If CFG_PCI_BOOT_HOST is sampled asserted in the de-asserting edge of reset, the Intel XScale® core will behave as boot host and configure the PCI system.

10.4.2 PCI Host Download

At power up, if FLASH_ROM is not present, strap pin CFG_PROM_BOOT should be sampled 0 (should be pulled down). In this mode CFG_RST_DIR pin should be 0 at power up signaling PCI_RST_L pin is an input that behaves as global chip reset.

1. Even after reset is removed by the PLL logic from IXP_RESET0 register (after PCI_RST_L reset is de-asserted), the Intel XScale® core reset is not removed.
2. PCI Reset through IXP_RESET0 [16] is removed automatically after being set and reset being removed.
3. IXP_RESET0[21] is set after PCI_RST_L has been removed and PLL_LOCK is sampled asserted.
4. Once IXP_RESET0[21] is set, PCI unit starts responding to transactions.
5. PCI Host first configures CSR, SRAM and DRAM base address registers after reading size requirements for these BARs. The size for CSR, SRAM and DRAM is defined by the use of Strap pins. Pre-fetchability for the window is defined by bit [3] of the respective BAR registers; therefore when host reads these registers, bit [3] is returned as 0 for CSR, SRAM and DRAM defining CSRs and also if SRAM and DRAM are to be non-prefetchable. Type Bits [2:0] are always Read-Only and return the value of 0x0 when read for CSR, SRAM and DRAM BAR registers.
6. PCI Host also programs Clock Control CSR, for PLL unit to generate proper clocks for SRAM, DRAM and other units.

Once these base address registers have been programmed, PCI host programs DRAM channels by initializing SDRAM_CSR, SDRAM_MEMCTL0, SDRAM_MEMCTL1 and SDRAM_MEMINIT registers. Once these registers have been programmed, PCI host writes the BOOT Code in DRAM starting at DRAM address 0. PCI Host can also program other registers if required. Once the boot



code is written in DRAM, PCI host writes 1 at bit [8] of Misc_Control register called Flash Alias Disable (Reset value 0). The Alias Disable bit can be wired to the Intel XScale® core gasket directly so that gasket knows how to transform address 0 from the Intel XScale® core. After writing 1 at Flash Alias Disable bit, host removes reset from the Intel XScale® core by writing 0 in bit [0] of IXP_RESET0 register. The Intel XScale® core starts booting from address 0, which is now directed by the gasket to DRAM.

10.5 Initialization

Refer to the *Intel® IXP2800 Network Processor Hardware Initialization Reference Manual* for information about the initialization of all units of the IXP2800 Network Processor.





Performance Monitor Unit

11

11.1 Introduction

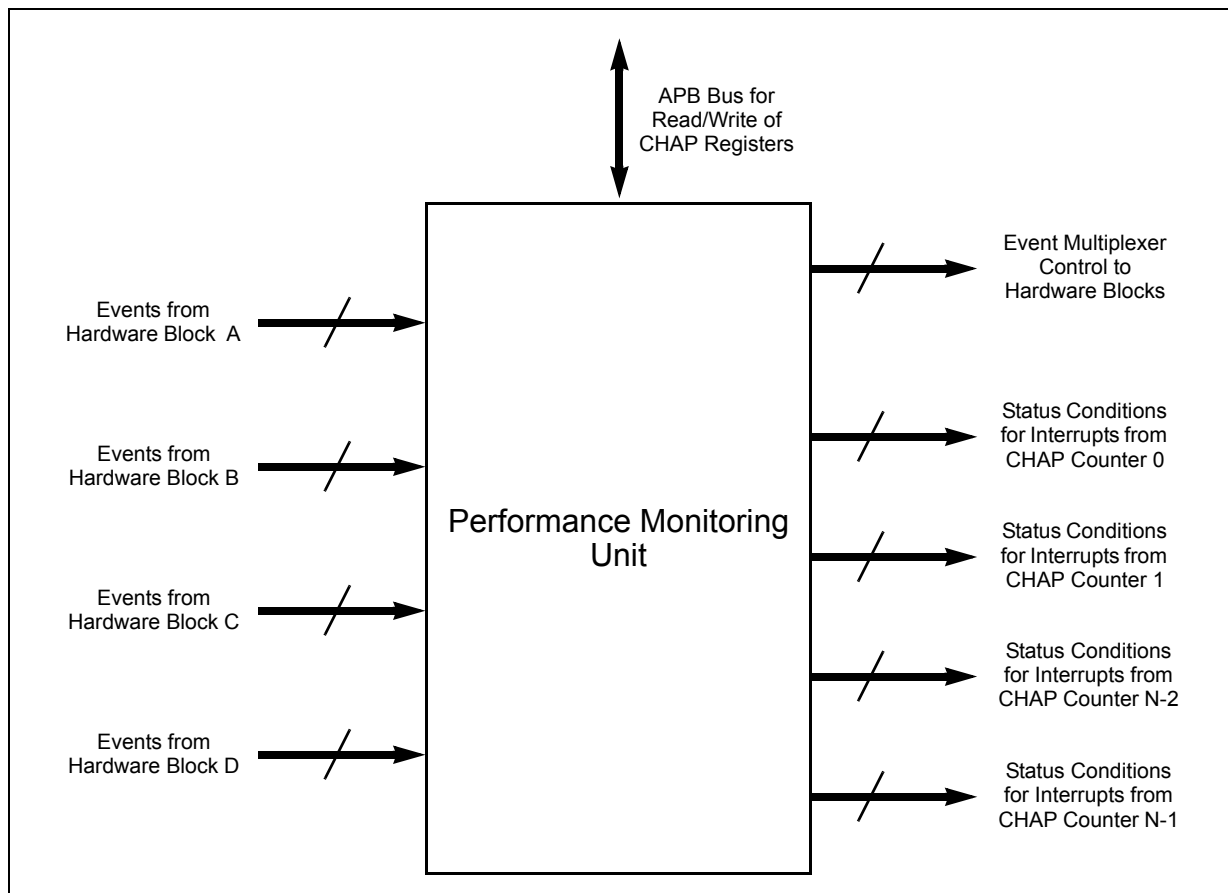
The Performance Monitor Unit (PMU) is a hardware block consisting of counters and comparators that can be programmed and controlled by using a set of configured registers to monitor and to fine tune performance of different hardware units in the IXP2800 Network Processor. The total number of such counters needed is determined based on the different events and functions that must be monitored concurrently. Observation of such events on the chip is used for statistical analysis, uncovering bottlenecks, and to tune the software to fit the hardware resources.

11.1.1 Motivation for Performance Monitors

For a given set of functionality, a measure of performance is very important in making decisions on feature sets to be supported, and to tune the embedded software on the chip. An accurate estimate of latency and speed in hardware blocks enables firmware and software designers to understand the limitations of the chip and to make prudent judgments about its software architecture. The current generation does not provide any performance monitor hooks.

Since IXP2800 Network Processors are targeted for high performance segments (OC-48 and above), the need for tuning the software to get the most out of the hardware resources becomes extremely critical. The performance monitors provide valuable insight into the chip by providing real-time data on latency and utilization of various resources. See [Figure 136](#) for the Performance Monitor Interface Block Diagram.

Figure 136. Performance Monitor Interface Block Diagram



11.1.2 Motivation for Choosing CHAP Counters

The Chipset Hardware Architecture Performance (CHAP) counters enable statistics gathering of internal hardware events in real-time. This implementation provides users with direct event counting and timing for performance monitoring purposes, and provides enough visibility into the internal architecture to perform utilization studies and workload characterization.

This implementation can also be used for chipset validation, higher-performing future chipsets, and applications tuned to the current chipset. The goal is that this will benefit both internal and external hardware and software development. The primary motivation for selecting the CHAP architecture for use in the IXP2800 Network Processor product family is that it has been designed and validated in several Intel desktop chipsets and the framework also provides a software suite that may be reused with little modification.



11.1.3 Functional Overview of CHAP Counters

At the heart of the CHAP counter's functionality are counters, each with associated registers. Each counter has a corresponding command, event, status, and data register. The smallest implementation has two counters, but if justified for a particular product, this architecture can support many more counters. The primary consideration is available silicon area. The memory-mapped space currently defined can accommodate registers for 256 counters. It can be configured for more, but that is beyond what is currently practical.

Signals that represent events from throughout the chip are routed to the CHAP unit. Software can select events that are recorded during a measurement session. The number of counters in an implementation defines the number of events that can be recorded simultaneously. Software and hardware events can control the starting, stopping, and sampling of the counters. This can be done in a time-based (polling) or an event-based fashion. Each counter can be incremented or decremented by different events. In addition to simple counting of events, the unit can provide data for histograms, queue analysis, and conditional event counting (for example, the number of times that event A happens before the first event B takes place).

When a counter is sampled, the current value of the counter is latched into the corresponding data register. The command, event, status, and data registers are accessible via standard Advanced Peripheral Bus (APB) memory-mapped registers, to facilitate high-speed sampling.

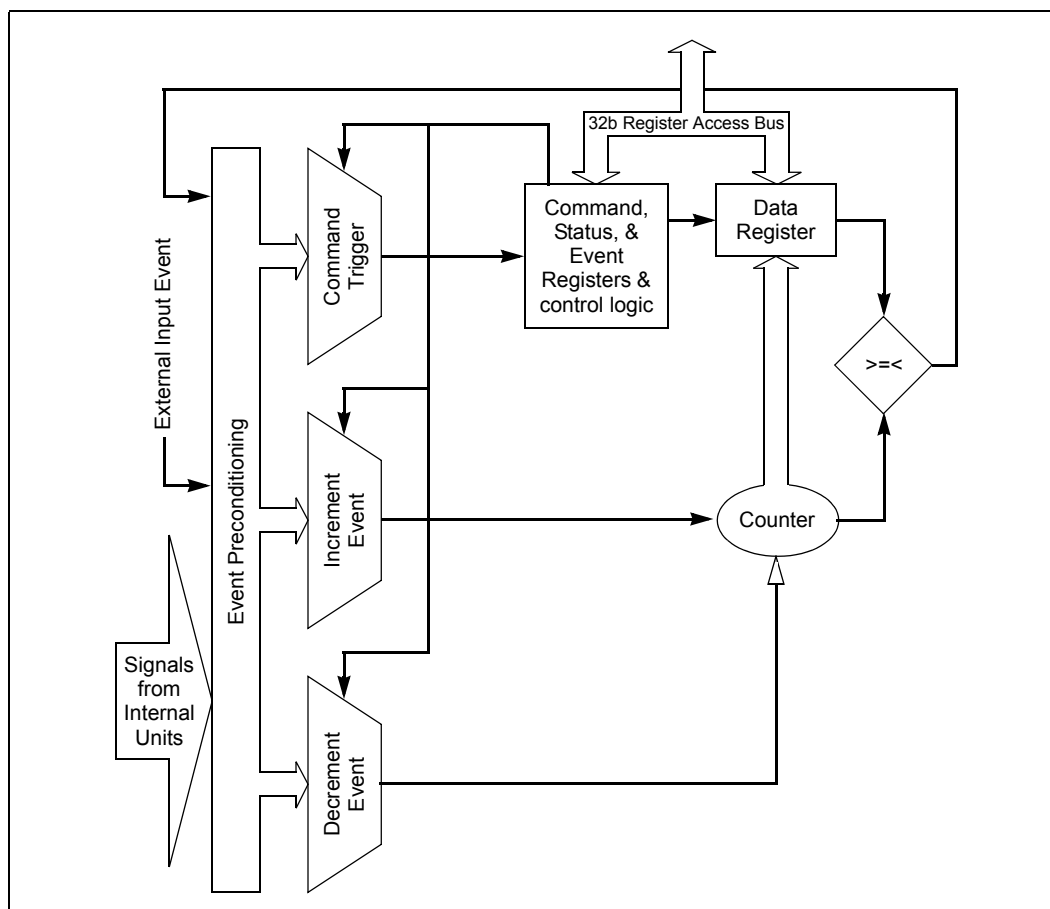
Two optional external pins allow for external visibility and control of the counters. The output pin signals that one of the following conditions generated an interrupt from any one of the counters:

- A programmable threshold condition was true.
- A command was triggered to begin.
- A counter overflow or underflow occurred.

The input pin allows an external source to control when a CHAP command is executed.

Figure 137 represents a single counter block. The multiplexers, registers, and all other logic are repeated for each counter that is present. There is a threshold event from each counter block that feeds into each multiplexer.

Figure 137. Block Diagram of a Single CHAP Counter



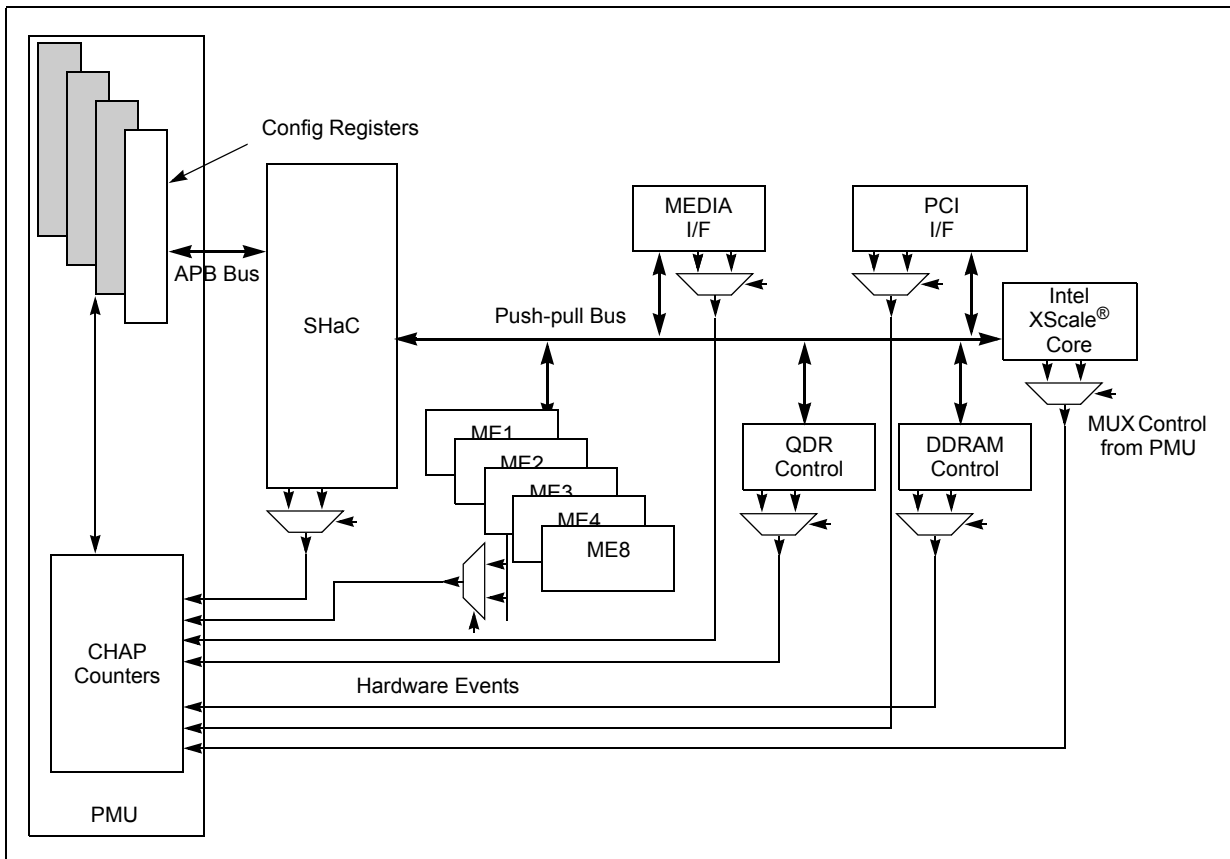
11.1.4 Basic Operation of the Performance Monitor Unit

At power-up, the Intel XScale® core invokes the performance monitoring software code. The PMU software has the application code to generate different types of data, such as histograms and graphs. It also has a device driver to configure and read data from the PMU in the IXP2800 Network Processor. This software programs the configuration registers in the PMU block to perform a certain set of monitoring and data collection. PMU CHAP counters execute the commands programmed by the Intel XScale® core and they collect various types of data such as latency and counts. Upon collection, it triggers an interrupt to the Intel XScale® core to indicate the completion of monitoring.

The Intel XScale® core either periodically monitors the PMU registers or waits for an interrupt to collect the observed data. The Intel XScale® core uses the APB to communicate with the PMU configuration registers.

Figure 138 represents a block diagram of the IXP2800 Network Processor and Performance Monitor Unit's (PMU) in relation to other hardware blocks in the chip.

Figure 138. Basic Block Diagram of IXP2800 Network Processor with PMU



11.1.5 Definition of CHAP Terminology

Duration Count	The counter is incremented for each clock for which the event signal is asserted as logic high.
MMR	Memory Mapped register.
OA	Observation Architecture. The predecessor to CHAP counters that facilitates the counting of hardware events.
Occurrence Count	The counter is incremented each time a rising edge of the event signal is detected.
Preconditioning	Altering a design block signal that represents an event such that it can be counted by the CHAP unit. The most common preconditioning is likely to be a 'one-shot' to count occurrences.
RO (register)	Read Only. If a register is read-only, writes to this register have no effect.
R/W (register)	Read/Write. A register with this attribute can be read and written.
WO (register)	Write Once. Once written, a register with this attribute becomes Read Only. This register can only be cleared by a Reset.
WC (register)	Write Clear. A register bit with this attribute can be read and written. However, a write of 1 clears (sets to 0) the corresponding bit and a write of 0 has no effect.

11.1.6 Definition of Clock Domains

The following abbreviations are used in the events table under clock domain.

P_CLK	The Command Push/Pull Clock also known as the Chassis clock. This clock is derived from the Microengine (ME) Clock. It is one-half of the Microengine clock.
T_CLK	Microengine Clock.
MTS_CLK	MSF Flow Control Status LVTTTL Clock TS_CLK.
MRX_CLK	MSF Flow Control Receive LVDS Clock RX_CLK.
MR_CLK	MSF Receive Data Clock R_CLK.
MT_CLK	MSF Transmit Data Clock T_CLK.
MTX_CLK	MSF Flow Control Transmit LVDS Clock TX_CLK.
D_CLK	DRAM Clock.
S_CLK	SRAM Clock.
APB_CLK	Advance Peripheral Bus Clock.

11.2 Interface and CSR Description

CAP is a standard logic block provided as part of the Network Processor that provides a method of interfacing to the ARM APB. This bus supports standard APB peripherals such as PMU, UART, Timers, and GPIO as well as CSRs that do not need to be accessed by the Microengines.

As shown in Figure 139, CAP uses three bus interfaces to support these modes. CAP supports a target ID of 0101, which Microengine assemblers should identify as a CSR instruction.

Figure 139. CAP Interface to the APB

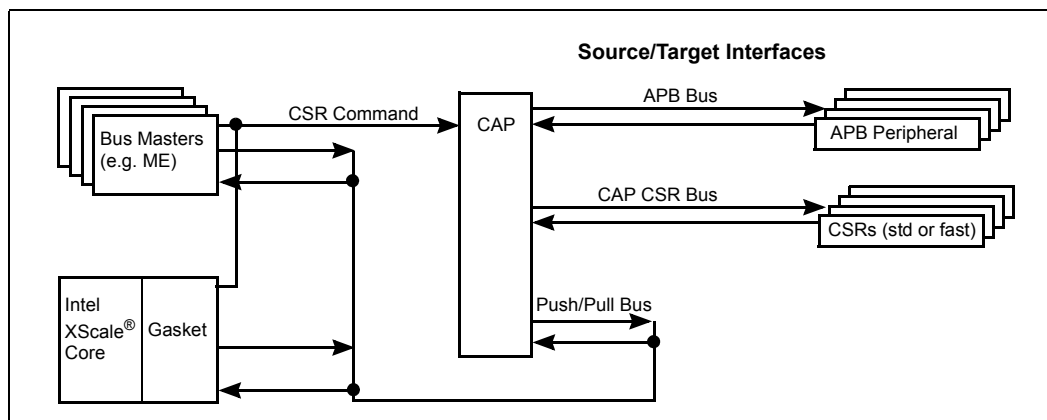


Table 151 shows the Intel XScale® core and Microengine instructions used to access devices on these buses and it shows which buses are used during the operation. For example, to read an APB peripheral such as a UART CSR, a Microengine would execute a `csr[read]` instruction and the Intel XScale® core would execute a `Load (ld)` instruction. Data is then moved between the CSR and the Intel XScale® core/Microengine by first reading the CSR via the APB and then writing the result to the Intel XScale® core/Microengine via the Push Bus.



Table 151. APB Usage

Accessing	Read Operation	Write Operation
APB Peripheral	Access Method: <ul style="list-style-type: none"> • Microengine: csr[read] • Intel XScale® core: Id 	Access Method: <ul style="list-style-type: none"> • Microengine: csr[write] • Intel XScale® core: st
	Bus Usages: <ul style="list-style-type: none"> • Read source: APB • Write destination: Push bus 	Bus Usages: <ul style="list-style-type: none"> • Read source: Pull Bus • Write destination: APB

11.2.1 APB Peripheral

The APB is part of the AMD* controller Bus Architecture (AMBA) hierarchy of buses that is optimized for minimal power consumption and reduced design complexity. The PMU needs to operate as an APB peripheral, interfacing with rest of the chip via the APB. The PMU needs to have an APB interface unit, which can perform a APB reads and writes to enable data transfer to and from the PMU registers.

11.2.2 CAP Description

11.2.2.1 Selecting the Access Mode

The CAP selects the appropriate access mode based on the COMMAND and ADDRESS fields from the Command Bus.

11.2.2.2 PMU CSR

Please refer to *Intel IXP2400 and IXP2800 Network Processor Programmer's Reference Manual*.

11.2.2.3 CAP Writes

For an APB write, CAP arbitrates for the S_Pull_Bus, pulls the write data from the source identified in PP_ID (either a Microengine transfer register or Intel XScale® core write buffer), and puts it into the CAP Pull Data FIFO. It then drives the address and writes data onto the appropriate bus. CAP CSRs locally decode the address to match their own. CAP generates a separate APB — devices select signal for each CAP device (up to 15 devices). If the write is to an APB CSR, the Control Logic maintains valid signaling until the APB_RDY_H signal is returned. (The APB RDY signal is an extension to the APB specification specifically added for the Network Processor).

CAP supports write operations with burst counts greater than 1. CAP looks at the length field on the command bus and breaks each count into a separate APB write cycle, incrementing the CSR number for each bus access.

11.2.2.4 CAP Reads

For an APB read, CAP drives the address, write, select, and enable signals, and waits for the acknowledge signal (APB_RDY_H) from the APB device. For a CAP CSR read, CAP drives the address, which controls a tree of multiplexers to select the appropriate CSR. CAP then waits for the



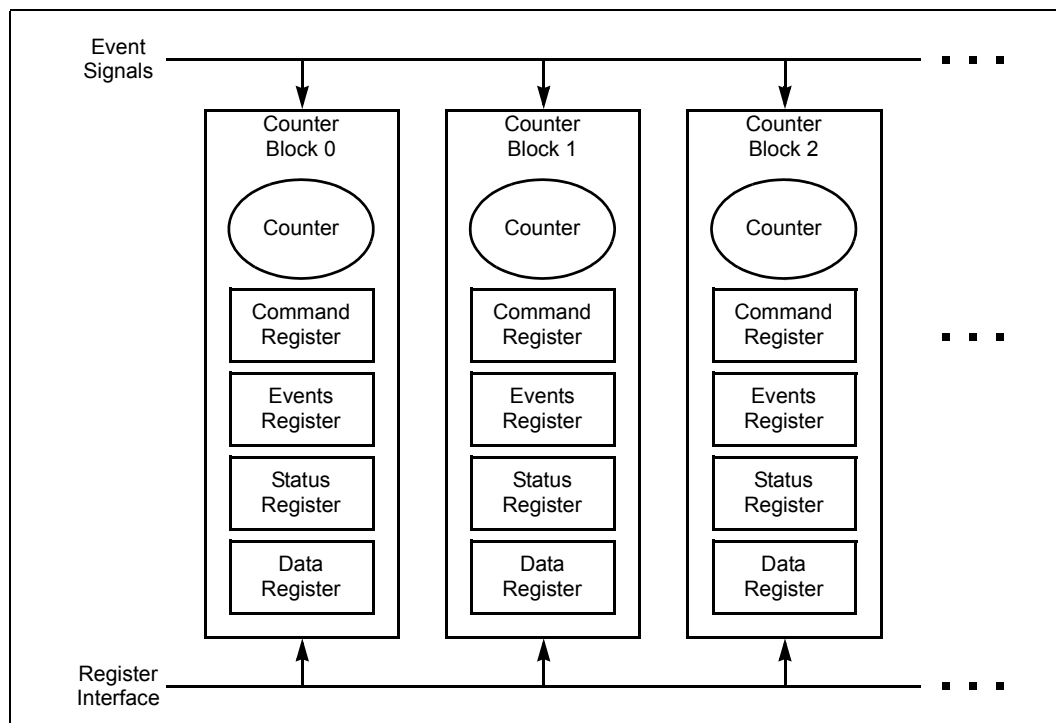
acknowledge signal (CAP_CSR_RD_RDY). When the data is returned, CAP puts the read data into the Push Data FIFO, arbitrates for the S_Push_Bus, and then the Push/Pull Arbiter pushes the data to the destination identified in PP_ID.

11.2.3 Configuration Registers

Because the CHAP unit resides on the APB, the offset associated with each of these registers is relative to the Memory Base Address that the configuration software sets in the PMUADR register.

Each counter has one command, one event, one status, and one data register associated with it. Each counter is “packaged” with these four registers in a “counter block”. Each implementation selects the number of counters it will implement, and therefore how many counter blocks (or slices) it will have. These registers are numbered 0 through N - 1 where N represents the number of counters - 1. See [Figure 140](#).

Figure 140. Conceptual Diagram of Counter Array



11.3 Performance Measurements

There are several measurements that can be made on each of the hardware blocks. These measurements together would enable improvements in hardware and software implementation and architectural issues. [Table 152](#) describes the different blocks and their associated performance measurement events.



Table 152. Hardware Blocks and Their Performance Measurement Events (Sheet 1 of 2)

Hardware Block	Performance Measurement Event	Description
Intel XScale® Core		
	DRAM Read Head of Queue Latency Histogram	The Intel XScale® core generates a read or write command to the DRAM primarily to either push or pull data of the DDRAM. These commands are scheduled to the DRAM through the push-pull arbiter through a command FIFO in the gasket. The DRAM-read head of queue enables the PMU to monitor when the read and write commands posted by the Intel XScale® core in the gasket gets fetched and delivered to DDRAM.
	SRAM Read Head of Queue Latency Histogram	The Intel XScale® core generates a read or write command to the SRAM primarily to either push or pull data of the SRAM. These commands are scheduled to the SRAM through the push-pull arbiter through a command FIFO in the gasket. The SRAM-read head of queue enables the PMU to monitor when the read and write commands posted by the Intel XScale® core in the gasket gets fetched and delivered to SRAM.
	Interrupts	Number of interrupts seen. Histogram of time between interrupts.
Microengines		
	Command FIFO Number of Commands	These statistics give the number of the commands issued by the Microengine in a particular period of time. It also can count each different thread.
	Control Store Measures	Count time between two microstore locations (locations can be set by instrumentation software). Histogram time between two microstore locations (locations can be set by instrumentation software)
	Execution Unit Status	Histogram of stall time. Histogram of aborted time. Histogram of swapped out time. Histogram of idle time.
	Command FIFO Head of Queue Wait Time Histogram (Latency)	This is to measure the latency of a command, which is at the head of the queue and is waiting to be sent out to the destination over the chassis.
SRAM		
	SRAM Commands	A count of SRAM commands received. These are maskable by command type such as Put and Get.
	SRAM Bytes, Cycles Busy	This measurement describes the number of bytes transferred and the SRAM busy time.
	Queue Depth Histogram	This measurement analyzes the different queues such as ordered, priority, push queue, pull queue, read lock fail, and HW queues, and provides information about utilization.
DRAM		
	DRAM Commands	This measurement lists the total commands issued to the DRAM, and they can be counted based on command type and error type.
	DRAM Bytes, Cycles Busy	This measurement indicates the DRAM busy time and bytes transferred.
	Maskable by Read/Write, Microengine, PCI, or the Intel XScale® Core	This measurement indicates the different accesses that are initiated to the DRAM. These measurements could be for all the accesses to the memory or can be masked using a specific source such as PCI, the Intel XScale® core, or Microengine. This can further be measured based on read or write cycles.



Table 152. Hardware Blocks and Their Performance Measurement Events (Sheet 2 of 2)

Hardware Block	Performance Measurement Event	Description
Chassis/Push-Pull		
	Command Bus Utilization	These statistics give the number of the command requests issued by the different Masters in a particular period of time. This measurement also indicates how long it takes to issue the grant from the request being issued by the different Masters.
	Push and Pull Bus Utilization	This measurement keeps track of the number of accesses issued and how long it takes to send the data to its destination.
Hash		
	Number of Accesses by Command Type	This measurement indicates the number of hash accesses issued; this count is maskable, based on command type.
	Latency of Histogram	This monitors the latency through each of the HASH queues.
Scratch		
	Number of Accesses by Command Type	This measurement indicates the number of Scratch accesses issued and this count is maskable, based on command type.
	Number of Bytes Transfer	This measurement indicates total number of bytes transferred to or from Scratch.
	Latency of Histogram	This measurement indicates the latency of performing read or write from the Scratch. Latency in command executions may also be measured.
PCI		
	Master Accesses	These statistics give the number of Master accesses that were generated by the PCI blocks. This measurement can be counted based on individual command type.
	Slave Accesses	These statistics give the number of Slave accesses that were generated by the PCI blocks. This measurement can be counted based on individual command type.
	Master/Slave Read Byte Count	This statistics give the total number of bytes of data that were generated by the PCI Master/Slave reads access. This measurement can be counted based on individual command type.
	Master/Slave Write Byte Count	These statistics give the total number of bytes of data that were generated by the PCI Master/Slave write accesses. This measurement can be counted based on individual command type.
	Burst Size Histogram	These statistics give a histogram of the number of various burst sizes.
Media Interface		
	TBUF Occupancy Histogram	This measurement shows the occupancy rate at different depths of the FIFO. This can help in better utilization of TBUF.
	RBUF Occupancy Histogram	This measurement shows the occupancy rate at different depths of the FIFO. This can help in better utilization of RBUF.
	Packet/Cell/Frame Count on a Per-Port Basis	This measurement gives the count of number of packets or cells or frames transferred in Transmitting mode. This measurement gives the count of number of packets or cells or frames transferred in the receiving mode. This may be measured using a per-port basis.
	Inter-arrival Time for Packets on a Per-Port Basis	This measurement can provide information on gaps between packets, thereby indicating effective line rate.
	Burst Size Histogram	This measurement gives the various burst sizes of packets being transmitted and received.



11.4 Events Monitored in Hardware

Tables in this section describe the events that can be measured, including the name of the event and the Event Selection Code (ESC). Refer to [Section 11.4](#) for tables showing event selection codes.

The acronyms in the event names typically represent unit names. The guidelines for which events a particular component must implement are provided in the following sections.

11.4.1 Queue Statistics Events

11.4.1.1 Queue Latency

Latency of Queue is an indicator of the control logic performance in terms of effective execution of the commands in the Control/Command queue or, the control logic's ability to effectively transfer data from the Data Queue.

This kind of monitoring needs observation of specific events such as:

- Enqueue into the Queue
This event indicates when an entry was made to the queue.
- Dequeue into the Queue
This event indicates when an entry was removed from the queue. The time period between when a particular entry was made into the queue and when the entry was removed from the queue indicates the latency of the queue for that entry.
- Queue Full Event
This event indicates when the queue has no room for additional entries.
- Queue Empty Event
This event indicates when the queue has no entries.

Queue Full and Queue Empty events can be used to determine Queue Utilization and bandwidth available in the queue to determine how to handle more traffic.

11.4.1.2 Queue Utilization

Utilization of Queue is determined by observing the percentage of time each queue is operating at a particular threshold level. Based on Queue size, multiple threshold values can be predetermined and monitored. The result of these observations can be used to provide histograms for Queue utilization. This kind of observation helps us better utilize the available resources in the queue.

11.4.2 Count Events

11.4.2.1 Hardware Block Execution Count

On each of the hardware blocks, events of importance such as number of commands executed, number of bytes transferred, total amount of clock-blocks that are free, and the total amount of time all of the contexts in the Microengine were idle can be counted as statistics, for managing the available resources.



11.4.3 Design Block Select Definitions

Once an event is defined, its definition must remain consistent between products. If the definition changes, it should have a new event selection code. This document contains the master list of all ESCs in all CHAP-enabled products. Not all of the ESCs in this document are listed in numerical order. The recommendation is to group similar events within the following ESC ranges. See [Table 153](#).

Table 153. PMU Design Unit Selection (Sheet 1 of 2)

Target Device	Target ID	PMU Design Group Block #	Description
Null	xxx xxx	0000	Null (False) Event
PMU_Counter	xxx xxx	0001 (PMU)	CHAP Counters Internal Threshold Events Event bit 0 CHAP Counter 0 Event bit 1 CHAP Counter 1 Event bit 2 CHAP Counter 2 Event bit 3 CHAP Counter 3 Event bit 4 CHAP Counter 4 Event bit 5 CHAP Counter 5
SRAM Group			
SRAM_DP1	001 001	0010 (SRAM Group) one and only one will be selected from same group	SRAM channel 0 SRAM channel 1 SRAM channel 2 SRAM channel 3 SRAM d-push SRAM d-pull
SRAM_DP0	001 010		
SRAM_CH3	001 011		
SRAM_CH2	001 100		
SRAM_CH1	001 101		
SRAM_CH0	001 110		
DRAM Group			
DRAM_CR1	010 000	0011 (DRAM) one and only one will be selected from same group	DRAM channel 0 DRAM channel 1 DRAM channel 2 DRAM d-push DRAM d-pull
DRAM_CR0	010 001		
DRAM_DPLA	010 010		
DRAM_DPSPA	010 011		
DRAM_CH2	010 100		
DRAM_CH1	010 101		
DRAM_CH0	010 110		
XPI	000 001	0100 (XPI)	XPI
SHaC	000 010	0101	
MSF	000 011	0110	Media



Table 153. PMU Design Unit Selection (Sheet 2 of 2)

Target Device	Target ID	PMU Design Group Block #	Description
Intel XScale® core	000 100	0111	Intel XScale® core
PCI	000 101	1000	PCI
ME Cluster 0 Group			
ME07	100 111	1001 (MEC0) one and only one will be selected from same group	ME Channel 0
ME06	100 110		ME00
ME05	100 101		ME01
ME04	100 100		ME02
ME03	100 011		ME03
ME02	100 010		ME04
ME01	100 001		ME05
ME00	100 000		ME06
ME Cluster 1 Group			
ME17	110 111	1010 (MEC1) one and only one will be selected from same group	ME Channel 1
ME16	110 110		ME10
ME15	110 101		ME11
ME14	110 100		ME12
ME13	110 011		ME13
ME12	110 010		ME14
ME11	110 001		ME15
ME10	110 000		ME16
		1011-1111	Reserved

11.4.4 Null Event

Not an actual event. When used as an increment or decrement event, no action takes place. When used as a Command Trigger, it causes the command to be triggered immediately after the command register is written to by the software. Also called False Event. Not reserved.



11.4.5 Threshold Events

These are the outputs of the threshold comparators. When the value in a data register is compared to its corresponding counter value and the condition is true, a threshold event is generated. This results in:

- A pulse on the signal lines that are routed to the event's input port (one signal line from each comparator).
- One piece of functionality this enables is to allow for CHAP commands to be completed only when a Threshold Event occurs. In other words, a Threshold Event can be used as a Command Trigger to control the execution of any CHAP command (start, stop, sample, etc.). See [Table 154](#).

Table 154. Chap Counter Threshold Events (Design Block # 0001)

Multiplexer #	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
000	Counter 0 Threshold	P_CLK	single	separate	Threshold Condition True on Event Counter 0
001	Counter 1 Threshold	P_CLK	single	separate	Threshold Condition True on Event Counter 1
010	Counter 2 Threshold	P_CLK	single	separate	Threshold Condition True on Event Counter 2
011	Counter 3 Threshold	P_CLK	single	separate	Threshold Condition True on Event Counter 3
100	Counter 4 Threshold	P_CLK	single	separate	Threshold Condition True on Event Counter 4
101	Counter 5 Threshold	P_CLK	single	separate	Threshold Condition True on Event Counter 5



11.4.6 External Input Events

11.4.6.1 XPI Events Target ID(000001) / Design Block #(0100)

Table 155. XPI PMU Event List (Sheet 1 of 4)

Event Number	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
0	XPI_RD_P	APB_CLK	single	separate	It includes all the read accesses, PMU, timer, GPIO, UART, and Slowport.
1	XPI_WR_P	APB_CLK	single	separate	It includes all the write accesses, PMU, timer, GPIO, UART, and Slowport.
2	PMU_RD_P	APB_CLK	single	separate	It executes the read access to the PMU unit.
3	PMU_WR_P	APB_CLK	single	separate	It executes the write access to the PMU unit.
4	UART_RD_P	APB_CLK	single	separate	It executes the read access to the UART unit.
5	UART_WR_P	APB_CLK	single	separate	It executes the write access to the UART unit.
6	GPIO_RD_P	APB_CLK	single	separate	It executes the read access to the GPIO unit.
7	GPIO_WR_P	APB_CLK	single	separate	It executes the write access to the GPIO unit.
8	TIMER_RD_P	APB_CLK	single	separate	It executes the read access to the Timer unit.
9	TIMER_WR_P	APB_CLK	single	separate	It executes the write access to the Timer unit.
10	SPDEV_RD_P	APB_CLK	single	separate	It executes the read access to the Slowport Device.
11	SPDEV_WR_P	APB_CLK	single	separate	It executes the write access to the Slowport Device.
12	SPCSR_RD_P	APB_CLK	single	separate	It executes the read access to the Slowport CSR.
13	SPCSR_WR_P	APB_CLK	single	separate	It executes the write access to the Slowport CSR.
14	TM0_UF_P	APB_CLK	single	separate	It shows the occurrence of timer 1 counter underflow.
15	TM1_UF_P	APB_CLK	single	separate	It shows the occurrence of timer 2 counter underflow.
16	TM2_UF_P	APB_CLK	single	separate	It shows the occurrence of timer 3 counter underflow.
17	TM3_UF_P	APB_CLK	single	separate	It shows the occurrence of timer 4 counter underflow.
18	IDLE0_0_P	APB_CLK	single	separate	It displays the idle state of the state machine 0 for the mode 0 of Slowport.
19	START0_1_P	APB_CLK	single	separate	It enters the start state of the state machine 0 for the mode 0 of Slowport.
20	ADDR10_3_P	APB_CLK	single	separate	It enters the first address state, AD[9:2], of the state machine 0 for the mode 0 of Slowport.
21	ADDR20_2_P	APB_CLK	single	separate	It enters the second address state, AD[17:10], of the state machine 0 for the mode 0 of Slowport.
22	ADDR30_6_P	APB_CLK	single	separate	It enters the third address state, AD[24:18], of the state machine 0 for the mode 0 of Slowport.
23	SETUP0_4_P	APB_CLK	single	separate	It enters data setup state of the state machine 0 for the mode 0 of Slowport.
24	PULW0_5_P	APB_CLK	single	separate	It enters data duration state of the state machine 0 for the mode 0 of Slowport.
25	HOLD0_D_P	APB_CLK	single	separate	It enters data hold state of the state machine 0 for the mode 0 of Slowport.



Table 155. XPI PMU Event List (Sheet 2 of 4)

26	TURNA0_C_P	APB_CLK	single	separate	It enters the termination state of the state machine 0 for the mode 0 of Slowport.
27	IDLE1_0_P	APB_CLK	single	separate	It displays the idle state of the state machine 1 for the mode 1 of Slowport.
28	START1_1_P	APB_CLK	single	separate	It enters the start state of the state machine 1 for the mode 1 of Slowport.
29	ADDR11_3_P	APB_CLK	single	separate	It enters the first address state, AD[7:0], of the state machine 1 for the mode 1 of Slowport.
30	ADDR21_2_P	APB_CLK	single	separate	It enters the second address state, AD[15:8], of the state machine 1 for the mode 1 of Slowport.
31	ADDR31_6_P	APB_CLK	single	separate	It enters the second address state, AD[23:16], of the state machine 1 for the mode 1 of Slowport.
32	ADDR41_7_P	APB_CLK	single	separate	It enters the second address state, AD[24], of the state machine 1 for the mode 1 of Slowport.
33	WRDATA1_5_P	APB_CLK	single	separate	It unpacks the data from the APB onto the Slowport bus for the state machine 1 for the mode 1 of Slowport.
34	PULW1_4_P	APB_CLK	single	separate	It enters the pulse width of the data transaction cycle for the state machine 1 for the mode 1 of Slowport.
35	CHPSEL1_C_P	APB_CLK	single	separate	It enters the chip select assertion pulse width when the state machine 1 is active for the mode 1 of Slowport.
36	OUTEN1_E_P	APB_CLK	single	separate	It enters the cycle when the OE is asserted during running on the state machine 1 for the mode 1 of Slowport.
37	PKDATA1_F_P	APB_CLK	single	separate	It enters the read data packing state when the state machine 1 is active for the mode 1 of Slowport.
38	LADATA1_D_P	APB_CLK	single	separate	It enters the data capturing cycle when the state machine 1 is active for the mode 1 of Slowport.
39	READY1_9_P	APB_CLK	single	separate	It enters the acknowledge state to terminate the read cycle when the state machine 1 is active for the mode 1 of Slowport.
40	TURNA1_8_P	APB_CLK	single	separate	It enters the turnaround state of the transaction when the state machine 1 is active for the mode 1 of Slowport.
41	IDLE2_0_P	APB_CLK	single	separate	It displays the idle state of the state machine 2 for the mode 2 of Slowport.
42	START2_1_P	APB_CLK	single	separate	It enters the start state of the state machine 2 for the mode 2 of Slowport.
43	ADDR12_3_P	APB_CLK	single	separate	It enters the first address state, AD[7:0], of the state machine 2 for the mode 2 of Slowport.
44	ADDR22_2_P	APB_CLK	single	separate	It enters the second address state, AD[15:8], of the state machine 2 for the mode 2 of Slowport.
45	ADDR32_6_P	APB_CLK	single	separate	It enters the second address state, AD[23:16], of the state machine 2 for the mode 2 of Slowport.
46	ADDR42_7_P	APB_CLK	single	separate	It enters the second address state, AD[24], of the state machine 2 for the mode 2 of Slowport.
47	WRDATA2_5_P	APB_CLK	single	separate	It unpacks the data from the APB onto the Slowport bus for the state machine 2 for the mode 2 of Slowport.



Table 155. XPI PMU Event List (Sheet 3 of 4)

48	SETUP2_4_P	APB_CLK	single	separate	It enters the pulse width of the data transaction cycle for the state machine 2 for the mode 2 of Slowport.
49	PULW2_C_P	APB_CLK	single	separate	It enters the pulse width of the data transaction cycle for the state machine 2 for the mode 2 of Slowport.
50	HOLD2_E_P	APB_CLK	single	separate	It enters the data hold period for the state machine 2 for the mode 2 of Slowport.
51	OUTEN2_F_P	APB_CLK	single	separate	It starts to assert the OE when the state machine 2 is active for the mode 2 of Slowport.
52	PKDATA2_D_P	APB_CLK	single	separate	It enters the read data packing state during the active state machine 2 for the mode 2 of Slowport.
53	LADATA2_9_P	APB_CLK	single	separate	It enters the data capturing cycle during the active state machine 2 for the mode 2 of Slowport.
54	READY2_B_P	APB_CLK	single	separate	It enters the acknowledge state to terminate the read cycle when the state machine 2 is active for the mode 2 of Slowport.
55	TURNA2_8_P	APB_CLK	single	separate	It enters the turnaround state of the transaction when the state machine 2 is active for the mode 2 of Slowport.
56	IDLE3_0_P	APB_CLK	single	separate	It displays the idle state of the state machine 3 for the mode 3 of Slowport.
57	START3_1_P	APB_CLK	single	separate	It enters the start state of the state machine 3 for the mode 3 of Slowport.
58	ADDR13_3_P	APB_CLK	single	separate	It enters the first address state, AD[7:0], of the state machine 3 for the mode 3 of Slowport.
59	ADDR23_2_P	APB_CLK	single	separate	It enters the second address state, AD[15:8], of the state machine 3 for the mode 3 of Slowport.
60	ADDR33_6_P	APB_CLK	single	separate	It enters the second address state, AD[23:16], of the state machine 3 for the mode 3 of Slowport.
61	ADDR43_7_P	APB_CLK	single	separate	It enters the second address state, AD[24], of the state machine 3 for the mode 3 of Slowport.
62	WRDATA3_5_P	APB_CLK	single	separate	It unpacks the data from the APB onto the Slowport bus for the state machine 3 for the mode 3 of Slowport.
63	SETUP3_4_P	APB_CLK	single	separate	It enters the pulse width of the data transaction cycle for the state machine 3 for the mode 3 of Slowport.
64	PULW3_C_P	APB_CLK	single	separate	It enters the pulse width of the data transaction cycle for the state machine 3 for the mode 3 of Slowport.
65	HOLD3_E_P	APB_CLK	single	separate	It enters the data hold period for the state machine 3 for the mode 3 of Slowport.
66	OUTEN3_F_P	APB_CLK	single	separate	It starts to assert the OE when the state machine 3 is active for the mode 3 of Slowport.
67	PKDATA3_D_P	APB_CLK	single	separate	It enters the read data packing state during the active state machine 3 for the mode 3 of Slowport.
68	LADATA3_B_P	APB_CLK	single	separate	It enters the data capturing cycle during the active state machine 3 for the mode 3 of Slowport.
69	READY3_9_P	APB_CLK	single	separate	It enters the acknowledge state to terminate the read cycle when the state machine 3 is active for the mode 3 of Slowport.



Table 155. XPI PMU Event List (Sheet 4 of 4)

70	TURNA3_8_P	APB_CLK	single	separate	It enters the turnaround state of the transaction when the state machine 3 is active for the mode 3 of Slowport.
71	IDLE4_0_P	APB_CLK	single	separate	It displays the idle state of the state machine 4 for the mode 4 of Slowport.
72	START4_1_P	APB_CLK	single	separate	It enters the start state of the state machine 4 for the mode 4 of Slowport.
73	ADDR14_3_P	APB_CLK	single	separate	It enters the first address state, AD[7:0], of the state machine 4 for the mode 4 of Slowport.
74	ADDR24_2_P	APB_CLK	single	separate	It enters the second address state, AD[15:8], of the state machine 4 for the mode 4 of Slowport.
75	ADDR34_6_P	APB_CLK	single	separate	It enters the second address state, AD[23:16], of the state machine 4 for the mode 4 of Slowport.
76	ADDR44_7_P	APB_CLK	single	separate	It enters the second address state, AD[24], of the state machine 4 for the mode 4 of Slowport.
77	WRDATA4_5_P	APB_CLK	single	separate	It unpacks the data from the APB onto the Slowport bus for the state machine 4 for the mode 4 of Slowport.
78	SETUP4_4_P	APB_CLK	single	separate	It enters the pulse width of the data transaction cycle for the state machine 4 for the mode 4 of Slowport.
79	PULW4_C_P	APB_CLK	single	separate	It enters the pulse width of the data transaction cycle for the state machine 4 for the mode 4 of Slowport.
80	HOLD4_E_P	APB_CLK	single	separate	It enters the data hold period for the state machine 4 for the mode 4 of Slowport.
81	OUTEN4_F_P	APB_CLK	single	separate	It starts to assert the OE when the state machine 4 is active for the mode 4 of Slowport.
82	PKDATA4_D_P	APB_CLK	single	separate	It enters the read data packing state during the active state machine 4 for the mode 4 of Slowport.
83	LADATA4_B_P	APB_CLK	single	separate	It enters the data capturing cycle during the active state machine 4 for the mode 4 of Slowport.
84	READY4_9_P	APB_CLK	single	separate	It enters the acknowledge state to terminate the read cycle when the state machine 4 is active for the mode 4 of Slowport.
85	TURN4_8_P	APB_CLK	single	separate	It enters the turnaround state of the transaction when the state machine 4 is active for the mode 4 of Slowport.



11.4.6.2 SHaC Events Target ID(000010) / Design Block #(0101)

Table 156. SHaC PMU Event List (Sheet 1 of 4)

Event Number	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
0	Scratch Cmd_Inlet_Fifo Not_Empty	P_CLK	single	separate	Scratch Command Inlet FIFO Not Empty
1	Scratch Cmd_Inlet_Fifo Full	P_CLK	single	separate	Scratch Command Inlet FIFO Full
2	Scratch Cmd_Inlet_Fifo Enqueue	P_CLK	single	separate	Scratch Command Inlet FIFO Enqueue
3	Scratch Cmd_Inlet_Fifo Dequeue	P_CLK	single	separate	Scratch Command Inlet FIFO Dequeue
4	Scratch Cmd_Pipe Not_Empty	P_CLK	single	separate	Scratch Command Pipe Not Empty
5	Scratch Cmd_Pipe Full	P_CLK	single	separate	Scratch Command Pipe Full
6	Scratch Cmd_Pipe Enqueue	P_CLK	single	separate	Scratch Command Pipe Enqueue
7	Scratch Cmd_Pipe Dequeue	P_CLK	single	separate	Scratch Command Pipe Dequeue
8	Scratch Pull_Data_Fifo 0 Full	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 0 Full
9	Scratch Pull_Data_Fifo 1 Full	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 1 Full
10	Hash Pull_Data_Fifo 0 Full	P_CLK	single	separate	Hash Pull Data FIFO Cluster 0 Full
11	Hash Pull_Data_Fifo 1 Full	P_CLK	single	separate	Hash Pull Data FIFO Cluster 1 Full
12	Scratch Pull_Data_Fifo 0 Not_Empty	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 0 Not Empty
13	Scratch Pull_Data_Fifo 0 Enqueue	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 0 Enqueue
14	Scratch Pull_Data_Fifo 0 Dequeue	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 0 Dequeue
15	Scratch Pull_Data_Fifo 1 Not_Empty	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 1 Not Empty
16	Scratch Pull_Data_Fifo 1 Enqueue	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 1 Enqueue
17	Scratch Pull_Data_Fifo 1 Dequeue	P_CLK	single	separate	Scratch Pull Data FIFO Cluster 1 Dequeue
18	Scratch State Machine Idle	P_CLK	single	separate	Scratch State Machine Idle
19	Scratch RAM Write	P_CLK	single	separate	Scratch RAM Write
20	Scratch RAM Read	P_CLK	single	separate	Scratch RAM Read
21	Scratch Ring_0 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_0_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_0_STATUS indicates full.



Table 156. SHaC PMU Event List (Sheet 2 of 4)

22	Scratch Ring_1 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_1_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_1_STATUS indicates full.
23	Scratch Ring_2 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1 RING_2_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0 RING_2_STATUS indicates full.
24	Scratch Ring_3 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1 RING_3_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0 RING_3_STATUS indicates full.
25	Scratch Ring_4 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1 RING_4_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_4_STATUS indicates full.
26	Scratch Ring_5 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_5_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_5_STATUS indicates full.
27	Scratch Ring_6 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_6_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_6_STATUS indicates full.
28	Scratch Ring_7 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_7_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_7_STATUS indicates full.
29	Scratch Ring_8 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_8_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_8_STATUS indicates full.
30	Scratch Ring_9 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_9_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_9_STATUS indicates full.
31	Scratch Ring_10 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_10_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_10_STATUS indicates full.
32	Scratch Ring_11 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_11_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_11_STATUS indicates full.
33	Scratch Ring_12 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_12_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_12_STATUS indicates full.
34	Scratch Ring_13 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_13_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_13_STATUS indicates full.



Table 156. SHaC PMU Event List (Sheet 3 of 4)

35	Scratch Ring_14 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_14_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_14_STATUS indicates full.
36	Scratch Ring_15 Status	P_CLK	single	separate	If SCRATCH_RING_BASE_x[26] = 1, RING_15_STATUS indicates empty. If SCRATCH_RING_BASE_x[26] = 0, RING_15_STATUS indicates full.
37	CAP CSR Write	P_CLK	single	separate	CAP CSR Write
38	CAP CSR Fast Write	P_CLK	single	separate	CAP CSR Fast Write
39	CAP CSR Read	P_CLK	single	separate	CAP CSR Read
40	DEQUEUE APB data	P_CLK	single	separate	Dequeue APB Data
41	apb_push_cmd_wph	P_CLK	single	separate	APB Push Command
42	APB_PUSH_DATA_REQ_RPH	P_CLK	single	separate	APB Push Data Request
43	APB pull1 FIFO dequeue	P_CLK	single	separate	APB Pull Cluster 1 FIFO Dequeue
44	apb_deq_pull1_data_wph	P_CLK	single	separate	APB Pull Cluster 1 Data Dequeue
45	data valid in apb pull1 FIFO	P_CLK	single	separate	APB Pull Cluster 1 Data FIFO Valid
46	APB pull0 FIFO dequeue	P_CLK	single	separate	APB Pull Cluster 0 FIFO Dequeue
47	SCR_APB_TAKE_PULL0_DATA_WPH	P_CLK	single	separate	APB Pull Cluster 0 Data
48	data valid in apb pull0 FIFO	P_CLK	single	separate	APB Pull Cluster 0 Data FIFO Valid
49	CAP APB read	P_CLK	single	separate	CAP APB Read
50	CAP APB write	P_CLK	single	separate	CAP APB Write
51	APB cmd dequeue	P_CLK	single	separate	APB Command Dequeue
52	APB CMD FIFO enqueue	P_CLK	single	separate	APB Command FIFO Enqueue
53	APB CMD FIFO FULL	P_CLK	single	separate	APB Command FIFO Full
54	APB CMD valid	P_CLK	single	separate	APB Command Valid
55	Hash Pull_Data_Fifo 0 Not_Empty	P_CLK	single	separate	Hash Pull Data FIFO Cluster 0 Not Empty
56	Hash Pull_Data_Fifo 0 Enqueue	P_CLK	single	separate	Hash Pull Data FIFO Cluster 0 Enqueue
57	Hash Pull_Data_Fifo 0 Dequeue	P_CLK	single	separate	Hash Pull Data FIFO Cluster 0 Dequeue
58	Hash Pull_Data_Fifo 1 Not_Empty	P_CLK	single	separate	Hash Pull Data FIFO Cluster 1 Not Empty
59	Hash Pull_Data_Fifo 1 Enqueue	P_CLK	single	separate	Hash Pull Data FIFO Cluster 1 Enqueue
60	Hash Pull_Data_Fifo 1 Dequeue	P_CLK	single	separate	Hash Pull Data FIFO Cluster 1 Dequeue
61	Hash Active	P_CLK	single	separate	Hash Active
62	Hash Cmd_Pipe Not_Empty	P_CLK	single	separate	Hash Command Pipe Not Empty



Table 156. SHaC PMU Event List (Sheet 4 of 4)

63	Hash Cmd_Pipe Full	P_CLK	single	separate	Hash Command Pipe Full
64	Hash Push_Data_Pipe Not_Empty	P_CLK	single	separate	Hash Push Data Pipe Not Empty
65	Hash Push_Data_Pipe Full	P_CLK	single	separate	Hash Push Data Pipe Full

11.4.6.3 IXP2800 Network Processor MSF Events Target ID(00011) / Design Block #(0110)

Table 157. IXP2800 Network Processor MSF PMU Event List (Sheet 1 of 6)

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
0	inlet command FIFO enqueue	P_CLK	pulse	separate	
1	inlet command FIFO dequeue	P_CLK	pulse	separate	
2	inlet command FIFO full	P_CLK	level	separate	
3	inlet command FIFO not empty	P_CLK	level	separate	
4	read command FIFO enqueue	P_CLK	pulse	separate	
5	read command FIFO dequeue	P_CLK	pulse	separate	
6	read command FIFO full	P_CLK	level	separate	
7	read command FIFO not empty	P_CLK	level	separate	
8	write command FIFO enqueue	P_CLK	pulse	separate	
9	write command FIFO dequeue	P_CLK	pulse	separate	
10	write command FIFO full	P_CLK	level	separate	
11	write command FIFO not empty	P_CLK	level	separate	
12	S_PULL data FIFO 0 enqueue	P_CLK	pulse	separate	
13	S_PULL data FIFO 0 dequeue	P_CLK	pulse	separate	
14	S_PULL data FIFO 0 full	P_CLK	level	separate	
15	S_PULL data FIFO 0 not empty	P_CLK	level	separate	
16	Received Data Training	P_CLK	level	separate	
17	Received Calendar Training	P_CLK	level	separate	
18	Received Flow Control Training	P_CLK	level	separate	



Table 157. IXP2800 Network Processor MSF PMU Event List (Sheet 2 of 6)

19	reserved				
20	S_PULL data FIFO 1 enqueue	P_CLK	pulse	separate	
21	S_PULL data FIFO 1 dequeue	P_CLK	pulse	separate	
22	S_PULL data FIFO 1 full	P_CLK	level	separate	
23	S_PULL data FIFO 1 not empty	P_CLK	level	separate	
24	Tbuffer Partition 0 full	P_CLK	level	separate	Indicates partition 0 of the tbuffer is full
25	Tbuffer Partition 1 full	P_CLK	level	separate	Indicates partition 1 of the tbuffer is full
26	Tbuffer Partition 2 full	P_CLK	level	separate	Indicates partition 2 of the tbuffer is full
27	reserved				
28	Rx_Thread_Freelist 0 enqueue	P_CLK	pulse	separate	
29	Rx_Thread_Freelist 0 dequeue	P_CLK	pulse	separate	
30	Rx_Thread_Freelist 0 full	P_CLK	level	separate	
31	Rx_Thread_Freelist 0 not empty	P_CLK	level	separate	
32	Rx_Thread_Freelist 1 enqueue	P_CLK	pulse	separate	
33	Rx_Thread_Freelist 1 dequeue	P_CLK	pulse	separate	
34	Rx_Thread_Freelist 1 full	P_CLK	level	separate	
35	Rx_Thread_Freelist 1 not empty	P_CLK	level	separate	
36	Rx_Thread_Freelist 2 enqueue	P_CLK	pulse	separate	
37	Rx_Thread_Freelist 2 dequeue	P_CLK	pulse	separate	
38	Rx_Thread_Freelist 2 empty	P_CLK	level	separate	
39	Rx_Thread_Freelist 2 not full	P_CLK	level	separate	
40	reserved				
41	reserved		I		
42	reserved				
43	Detect No Calendar	MTS_CLK	level	separate	Indicates that a framing pattern has been received on the TSTAT inputs for greater than 32 clock cycles; the valid signal from the MTS_CLK domain is synchronized; as such, it yields an approximate value.
44	Detect FC_IDLE	MRX_CLK	level	separate	Indicates that an idle cycle has been received on the RXCDAT inputs for greater than 2 clock cycles; the valid signal from the MTS_CLK domain is synchronized; as such, it yields an approximate value.



Table 157. IXP2800 Network Processor MSF PMU Event List (Sheet 3 of 6)

45	Detect FC_DEAD	MRX_CLK	level	separate	Indicates that a dead cycle has been received on the RXCDAT inputs for greater than 2 clock cycles; the valid signal from the MTS_CLK domain is synchronized; as such, it yields an approximate value.
46	Detect C_IDLE	MR_CLK	level	separate	Indicates that an idle cycle has been received on the RDAT inputs for greater than 2 clock cycles; the valid signal from the MTS_CLK domain is synchronized; as such, it yields an approximate value.
47	Detect C_DEAD	MR_CLK	level	separate	Indicates that a dead cycle has been received on the RDAT inputs for greater than 2 clock cycles; the valid signal from the MTS_CLK domain is synchronized; as such, it yields an approximate value.
48	Detect CFC sustained	MTX_CLK	level	separate	Indicates that the CFC input flag has been asserted for greater than 32 clock cycles; the valid signal from the MTX_CLK domain is synchronized; as such, it yields an approximate value.
49	Rbuffer Partition 0 empty	P_CLK	level	separate	Indicates that partition 0 of the rbuffer is empty.
50	Rbuffer Partition 1 empty	P_CLK	level	separate	Indicates that partition 1 of the rbuffer is empty.
51	Rbuffer Partition 2 empty	P_CLK	level	separate	Indicates that partition 2 of the rbuffer is empty.
52	Full Element List enqueue	P_CLK	pulse	separate	
53	Full Element List dequeue	P_CLK	pulse	separate	
54	Full Element List full	P_CLK	level	separate	
55	Full Element List not empty	P_CLK	level	separate	
56	Rbuffer Partition 0 full	P_CLK	level	separate	Indicates that partition 0 of the rbuffer is full.
57	Rbuffer Partition 1 full	P_CLK	level	separate	Indicates that partition 1 of the rbuffer is full.
58	Rbuffer Partition 2 full	P_CLK	level	separate	Indicates that partition 2 of the rbuffer is full.
59	reserved				
60	Rx_Valid[0] is set	P_CLK		separate	
61	Rx_Valid[8] is set	P_CLK		separate	
62	Rx_Valid[16] is set	P_CLK		separate	
63	Rx_Valid[24] is set	P_CLK		separate	
64	Rx_Valid[32] is set	P_CLK		separate	
65	Rx_Valid[48] is set	P_CLK		separate	
66	Rx_Valid[64] is set	P_CLK		separate	
67	Rx_Valid[96] is set	P_CLK		separate	
68	Data CFrame received	P_CLK	pulse	separate	Indicates that the CSIX DATA state machine after the Receive input FIFO has received a CSIX DATA CFRAME.
69	Control CFrame received	P_CLK	pulse	separate	Indicates that the CSIX CONTROL state machine after the Receive input FIFO has received a CSIX CONTROL CFRAME.



Table 157. IXP2800 Network Processor MSF PMU Event List (Sheet 4 of 6)

70	SPI-4 Packet received	P_CLK	pulse	separate	Indicates that the SPI-4 state machine after the Receive input FIFO has received an SPI-4 packet.
71	reserved				
72	Data CFrame transmitted	P_CLK	level	separate	Indicates that the transmit buffer state machine is writing a CSIX Data CFRAME into the transmit FIFO; One P_CLK cycle indicates a 32-bit write into the transmit FIFO.
73	Control CFrame transmitted	P_CLK	level	separate	Indicates the transmit buffer state machine is writing a CSIX Control CFRAME into the transmit FIFO; One P_CLK cycle indicates a 32-bit write into the transmit FIFO.
74	SPI-4 CFrame transmitted	P_CLK	level	separate	Indicates the transmit buffer state machine is writing an SPI-4 Packet into the transmit FIFO; One P_CLK cycle indicates a 32-bit write into the transmit FIFO.
75	reserved				
76	Tx_Valid[0] is set	P_CLK	level	separate	
77	Tx_Valid[8] is set	P_CLK	level	separate	
78	Tx_Valid[16] is set	P_CLK	level	separate	
79	Tx_Valid[24] is set	P_CLK	level	separate	
80	Tx_Valid[32] is set	P_CLK	level	separate	
81	Tx_Valid[48] is set	P_CLK	level	separate	
82	Tx_Valid[64] is set	P_CLK	level	separate	
83	Tx_Valid[96] is set	P_CLK	level	separate	
84	Tbuffer Partition 0 empty	P_CLK	level	separate	Indicates partition 0 of the tbuffer is empty.
85	Tbuffer Partition 1 empty	P_CLK	level	separate	Indicates partition 1 of the tbuffer is empty.
86	Tbuffer Partition 2 empty	P_CLK	level	separate	Indicates partition 2 of the tbuffer is empty.
87	reserved				
88	D_PUSH_DATA write to TBUF	P_CLK	pulse	separate	Each write is in units of quadwords (8 bytes).
89	S_PULL_DATA_0 write to TBUF	P_CLK	pulse	separate	Each write is in units of longwords (4 bytes).
90	S_PULL_DATA_1 write to TBUF	P_CLK	pulse	separate	Each write is in units of longwords (4 bytes).
91	CSR write	P_CLK	pulse	separate	
92	D_PULL_DATA read from RBUF	P_CLK	pulse	separate	Each read is in units of quadwords (8 bytes).
93	S_PUSH_DATA read from RBUF	P_CLK	pulse	separate	Each read is in units of longwords (4 bytes).
94	CSR read	P_CLK	pulse	separate	
95	CSR fast write	P_CLK	pulse	separate	
96	RX Autopush Asserts for Null and Non-Null Autopushes	P_CLK	pulse	separate	



Table 157. IXP2800 Network Processor MSF PMU Event List (Sheet 5 of 6)

97	Rx null autopush	P_CLK	pulse	separate	
98	Tx skip	P_CLK	pulse	separate	An mpacket was dropped due to the Tx_Skip bit being set in the Transmit Control Word.
99	SF_CRDY	P_CLK	level	separate	Only valid in CSIX receive mode and indicates how much of the time the switch fabric is able to receive control CFrames.
100	SF_DRDY	P_CLK	level	separate	Only valid in CSIX receive mode and indicates how much of the time the switch fabric is able to receive data CFrames.
101	TM_CRDY	P_CLK	level	separate	Only valid in CSIX receive mode; indicates how much of the time the egress processor is able to receive control CFrames.
102	TM_DRDY	P_CLK	level	separate	Only valid in CSIX receive mode; indicates how much of the time the egress processor is able to receive data CFrames.
103	FCIFIFO enqueue	P_CLK	pulse	separate	
104	FCIFIFO dequeue	P_CLK	pulse	separate	
105	FCIFIFO error	P_CLK	pulse	separate	Indicates that a bad CFrame was received on the CBus (horizontal or vertical parity error, premature RxSOF); only valid in CSIX transmit mode.
106	FCIFIFO synchronizing FIFO error	P_CLK	pulse	separate	Indicates that the CBus ingress logic encountered a FCIFIFO full condition while enqueueing a CFrame into FCIFIFO.
107	Vertical parity error	P_CLK	pulse	separate	Only valid in CSIX receive mode.
108	Horizontal parity	P_CLK	pulse	separate	Only valid in CSIX receive mode.
109	Dip 4 Parity Error	P_CLK	pulse	separate	Only valid in SPI-4 receive mode.
110	Dip 2 Parity Error	P_CLK	pulse	separate	Only valid in SPI-4 receive mode.
111	reserved				
112	CSIX DATA receive active	MR_CLK	level	separate	Indicates a valid CSIX DATA CFRAME received on the RX_DATA bus and may be used to measure bus utilization; the active signal from the MR_CLK domain is synchronized; as such, it yields an approximate value.
113	CSIX CONTROL receive active	MR_CLK	level	separate	Indicates a valid CSIX CONTROL CFRAME received on the RX_DATA bus and may be used to measure bus utilization; the active signal from the MR_CLK domain is synchronized; as such, it yields an approximate value.
114	SPI-4 receive active	MR_CLK	level	separate	Indicates a valid SPI-4 Packet received on the RX_DATA bus and may be used to measure bus utilization; the active signal from the MR_CLK domain is synchronized; as such, it yields an approximate value.



Table 157. IXP2800 Network Processor MSF PMU Event List (Sheet 6 of 6)

115	FCE receive active	MR_CLK	level	separate	Indicates a valid Flow Control Packet received on the RX_DATA bus and may be used to measure bus utilization; the active signal from the MR_CLK domain is synchronized; as such, it yields an approximate value.
116	CSIX DATA transmit active	MT_CLK	level	separate	Indicates valid transmit data on the TX_DATA bus and may be used to measure bus utilization; the valid signal from the MT_CLK domain is synchronized; as such, it yields an approximate value.
117	CSIX CONTROL transmit active	MT_CLK	level	separate	Indicates valid transmit data on the TX_DATA bus and may be used to measure bus utilization; the valid signal from the MT_CLK domain is synchronized; as such, it yields an approximate value.
118	SPI-4 transmit active	MT_CLK	level	separate	Indicates valid transmit data on the TX_DATA bus and may be used to measure bus utilization; the valid signal from the MT_CLK domain is synchronized; as such, it yields an approximate value.
119	FCE transmit active	MTX_CLK	level	separate	Indicates valid transmit data on the TXC_DATA bus and may be used to measure bus utilization; the valid signal from the MTX_CLK domain is synchronized; as such, it yields an approximate value.
120	FCI receive active	MRX_CLK	level	separate	Indicates a valid Flow Control Packet received on the RXC_DATA bus and may be used to measure bus utilization; the active signal from the MRX_CLK domain is synchronized; as such, it yields an approximate value.
121	Receive FIFO error	MR_CLK	pulse	separate	The receive FIFO has experienced an underflow or overflow. A pulse from the MR_CLK clock domain is converted to a pulse in the P_CLK clock domain.
122	reserved				
123	reserved				
124	reserved				
125	reserved				
126	reserved				
127	reserved				



11.4.6.4 Intel XScale® Core Events Target ID(000100) / Design Block #(0111)

Table 158. Intel XScale® Core Gasket PMU Event List (Sheet 1 of 4)

Event Number	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
0	XG_CFIFO_WR_EVEN_XS	P_CLK	single	separate	XG command FIFO even enqueue
1	reserved				
2	XG_DFIFO_WR_EVEN_XS	P_CLK	single	separate	XG DRAM data FIFO even enqueue
3	reserved				
4	XG_SFIFO_WR_EVEN_XS	P_CLK	single	separate	XG SRAM data FIFO even enqueue
5	reserved				
6	XG_LCFIFO_WR_EVEN_XS	P_CLK	single	separate	XG lcsr command FIFO even enqueue
7	XG_LCFIFO_WR_ODD_XS	P_CLK	single	separate	XG lcsr command FIFO odd enqueue
8	XG_LDFIFO_WR_EVEN_XS	P_CLK	single	separate	XG lcsr data FIFO even enqueue
9	XG_LDFIFO_WR_ODD_XS	P_CLK	single	separate	XG lcsr data FIFO odd enqueue
10	XG_LCSR_RD_EVEN_XS	P_CLK	single	separate	XG lcsr return data FIFO even dequeue
11	XG_LCSR_RD_ODD_XS	P_CLK	single	separate	XG lcsr return data FIFO odd dequeue
12	XG_LCSR_RD_OR_XS	P_CLK	single	separate	XG lcsr return data FIFO even_or_odd dequeue
13	XG_PUFF0_RD_EVEN_XS	P_CLK	single	separate	XG push fifo0 even dequeue
14	XG_PUFF0_RD_ODD_XS	P_CLK	single	separate	XG push fifo0 odd dequeue
15	XG_PUFF0_RD_OR_XS	P_CLK	single	separate	XG push fifo0 even_or_odd dequeue
16	XG_PUFF1_RD_EVEN_XS	P_CLK	single	separate	XG push fifo1 even dequeue
17	XG_PUFF1_RD_ODD_XS	P_CLK	single	separate	XG push fifo1 odd dequeue
18	XG_PUFF1_RD_OR_XS	P_CLK	single	separate	XG push fifo1 even_or_odd dequeue
19	XG_PUFF2_RD_EVEN_XS	P_CLK	single	separate	XG push fifo2 even dequeue
20	XG_PUFF2_RD_ODD_XS	P_CLK	single	separate	XG push fifo2 odd dequeue
21	XG_PUFF2_RD_OR_XS	P_CLK	single	separate	XG push fifo2 even_or_odd dequeue
22	XG_PUFF3_RD_EVEN_XS	P_CLK	single	separate	XG push fifo3 even dequeue
23	XG_PUFF3_RD_ODD_XS	P_CLK	single	separate	XG push fifo3 odd dequeue
24	XG_PUFF3_RD_OR_XS	P_CLK	single	separate	XG push fifo3 even_or_odd dequeue
25	XG_PUFF4_RD_EVEN_XS	P_CLK	single	separate	XG push fifo4 even dequeue
26	XG_PUFF4_RD_ODD_XS	P_CLK	single	separate	XG push fifo4 odd dequeue
27	XG_PUFF4_RD_OR_XS	P_CLK	single	separate	XG push fifo4 even_or_odd dequeue
28	XG_SYNC_ST_XS	P_CLK	single	separate	XG in sync. state
29	reserved				
30	reserved				
31	reserved				



Table 158. Intel XScale® Core Gasket PMU Event List (Sheet 2 of 4)

32	reserved				
33	reserved				
34	XG_CFIFO_EMPTYN_CPP	P_CLK	single	separate	XG command FIFO empty flag
35	XG_DFIFO_EMPTYN_CPP	P_CLK	single	separate	XG DRAM data FIFO empty flag
36	XG_SFIFO_EMPTYN_CPP	P_CLK	single	separate	XG SRAM data FIFO empty flag
37	XG_LCFIFO_EMPTYN_CPP	P_CLK	single	separate	XG lcsr command FIFO empty flag
38	XG_LDFIFO_EMPTYN_CPP	P_CLK	single	separate	XG lcsr data FIFO empty flag
39	reserved				
40	XG_OFIFO_EMPTYN_CPP	P_CLK	single	separate	XG cpp command FIFO empty flag
41	XG_OFIFO_FULLN_CPP	P_CLK	single	separate	XG cpp command FIFO full flag
42	XG_DP_EMPTYN_CPP	P_CLK	single	separate	XG DRAM pull data FIFO empty flag
43	XG_SP_EMPTYN_CPP	P_CLK	single	separate	XG SRAM pull data FIFO empty flag
44	XG_HASH_48_CPP	P_CLK	single	separate	hash_48 command on cpp bus
45	XG_HASH_64_CPP	P_CLK	single	separate	hash_64 command on cpp bus
46	XG_HASH_128_CPP	P_CLK	single	separate	hash_128 command on cpp bus
47	XG_LCSR_FIQ_CPP	P_CLK	single	separate	XG FIQ generated by interrupt CSR
48	XG_LCSR_IRQ_CPP	P_CLK	single	separate	XG IRQ generated by interrupt CSR
49	XG_CFIFO_RD_CPP	P_CLK	single	separate	XG command FIFO dequeue
50	XG_DFIFO_RD_CPP	P_CLK	single	separate	XG DRAM data FIFO dequeue
51	XG_SFIFO_RD_CPP	P_CLK	single	separate	XG SRAM data FIFO dequeue
52	XG_LCFIFO_RD_CPP	P_CLK	single	separate	XG lcsr command FIFO dequeue
53	XG_LDFIFO_RD_CPP	P_CLK	single	separate	XG lcsr data FIFO dequeue
54	XG_LCSR_WR_CPP	P_CLK	single	separate	XG lcsr return data FIFO enqueue
55	XG_OFIFO_RD_CPP	P_CLK	single	separate	XG cpp command FIFO dequeue
56	XG_OFIFO_WR_CPP	P_CLK	single	separate	XG cpp command FIFO enqueue
57	XG_DPDATA_WR_CPP	P_CLK	single	separate	XG DRAM pull data FIFO enqueue
58	XG_DPDATA_RD_CPP	P_CLK	single	separate	XG DRAM pull data FIFO dequeue
59	XG_SPDATA_WR_CPP	P_CLK	single	separate	XG SRAM pull data FIFO enqueue
60	XG_SPDATA_RD_CPP	P_CLK	single	separate	XG SRAM pull data FIFO dequeue
61	XG_PUFF0_WR_CPP	P_CLK	single	separate	XG push fifo0 enqueue
62	XG_PUFF1_WR_CPP	P_CLK	single	separate	XG push fifo1 enqueue
63	XG_PUFF2_WR_CPP	P_CLK	single	separate	XG push fifo2 enqueue
64	XG_PUFF3_WR_CPP	P_CLK	single	separate	XG push fifo3 enqueue
65	XG_PUFF4_WR_CPP	P_CLK	single	separate	XG push fifo4 enqueue
66	XG_SRAM_RD_CPP	P_CLK	single	separate	XG SRAM read command on cpp bus
67	XG_SRAM_RD_1_CPP	P_CLK	single	separate	XG SRAM read length=1 on cpp bus
68	XG_SRAM_RD_8_CPP	P_CLK	single	separate	XG SRAM read length=8 on cpp bus
69	XG_SRAM_WR_CPP	P_CLK	single	separate	XG SRAM write command on cpp bus
70	XG_SRAM_WR_1_CPP	P_CLK	single	separate	XG SRAM write length=1 on cpp bus



Table 158. Intel XScale® Core Gasket PMU Event List (Sheet 3 of 4)

71	XG_SRAM_WR_2_CPP	P_CLK	single	separate	XG SRAM write length=2 on cpp bus
72	XG_SRAM_WR_3_CPP	P_CLK	single	separate	XG SRAM write length=3 on cpp bus
73	XG_SRAM_WR_4_CPP	P_CLK	single	separate	XG SRAM write length=4 on cpp bus
74	XG_SRAM_CSR_RD_CPP	P_CLK	single	separate	XG SRAM csr read command on cpp bus
75	XG_SRAM_CSR_WR_CPP	P_CLK	single	separate	XG SRAM csr write command on cpp bus
76	XG_SRAM_ATOM_CPP	P_CLK	single	separate	XG SRAM atomic command on cpp bus
77	XG_SRAM_GET_CPP	P_CLK	single	separate	XG SRAM get command on cpp bus
78	XG_SRAM_PUT_CPP	P_CLK	single	separate	XG SRAM put command on cpp bus
79	XG_SRAM_ENQ_CPP	P_CLK	single	separate	XG SRAM enq command on cpp bus
80	XG_SRAM_DEQ_CPP	P_CLK	single	separate	XG SRAM deq command on cpp bus
81	XG_S0_ACC_CPP	P_CLK	single	separate	XG SRAM channel0 access on cpp bus
82	XG_S1_ACC_CPP	P_CLK	single	separate	XG SRAM channel1 access on cpp bus
83	XG_S2_ACC_CPP	P_CLK	single	separate	XG SRAM channel2 access on cpp bus
84	XG_S3_ACC_CPP	P_CLK	single	separate	XG SRAM channel3 access on cpp bus
85	XG_SCR_RD_CPP	P_CLK	single	separate	XG scratch read command on cpp bus
86	XG_SCR_RD_1_CPP	P_CLK	single	separate	XG scratch read length=1 on cpp bus
87	XG_SCR_RD_8_CPP	P_CLK	single	separate	XG scratch read length=8 on cpp bus
88	XG_SCR_WR_CPP	P_CLK	single	separate	XG scratch write command on cpp bus
89	XG_SCR_WR_1_CPP	P_CLK	single	separate	XG scratch write length=1 on cpp bus
90	XG_SCR_WR_2_CPP	P_CLK	single	separate	XG scratch write length=2 on cpp bus
91	XG_SCR_WR_3_CPP	P_CLK	single	separate	XG scratch write length=3 on cpp bus
92	XG_SCR_WR_4_CPP	P_CLK	single	separate	XG scratch write length=4 on cpp bus
93	XG_SCR_ATOM_CPP	P_CLK	single	separate	XG scratch atomic command on cpp bus
94	XG_SCR_GET_CPP	P_CLK	single	separate	XG scratch get command on cpp bus
95	XG_SCR_PUT_CPP	P_CLK	single	separate	XG scratch put command on cpp bus
96	XG_DRAM_RD_CPP	P_CLK	single	separate	XG DRAM read command on cpp bus
97	XG_DRAM_RD_1_CPP	P_CLK	single	separate	XG DRAM read length=1 on cpp bus
98	XG_DRAM_RD_4_CPP	P_CLK	single	separate	XG DRAM read length=4 on cpp bus
99	XG_DRAM_WR_CPP	P_CLK	single	separate	XG DRAM write on cpp bus
100	XG_DRAM_WR_1_CPP	P_CLK	single	separate	XG DRAM write length=1 on cpp bus
101	XG_DRAM_WR_2_CPP	P_CLK	single	separate	XG DRAM write length=2 on cpp bus
102	XG_DRAM_CSR_RD_CPP	P_CLK	single	separate	XG DRAM csr read command on cpp bus
103	XG_DRAM_CSR_WR_CPP	P_CLK	single	separate	XG DRAM csr write command on cpp bus
104	XG_MSF_RD_CPP	P_CLK	single	separate	XG msf read command on cpp bus
105	XG_MSF_RD_1_CPP	P_CLK	single	separate	XG msf read length=1 on cpp bus
106	reserved				
107	XG_MSF_WR_CPP	P_CLK	single	separate	XG msf write command on cpp bus
108	XG_MSF_WR_1_CPP	P_CLK	single	separate	XG msf write length=1 on cpp bus
109	XG_MSF_WR_2_CPP	P_CLK	single	separate	XG msf write length=2 on cpp bus



Table 158. Intel XScale® Core Gasket PMU Event List (Sheet 4 of 4)

110	XG_MSF_WR_3_CPP	P_CLK	single	separate	XG msf write length=3 on cpp bus
111	XG_MSF_WR_4_CPP	P_CLK	single	separate	XG msf write length=4 on cpp bus
112	XG_PCI_RD_CPP	P_CLK	single	separate	XG pci read command on cpp bus
113	XG_PCI_RD_1_CPP	P_CLK	single	separate	XG pci read length=1 on cpp bus
114	XG_PCI_RD_8_CPP	P_CLK	single	separate	XG pci read length=8 on cpp bus
115	XG_PCI_WR_CPP	P_CLK	single	separate	XG pci write command on cpp bus
116	XG_PCI_WR_1_CPP	P_CLK	single	separate	XG pci write length=1 on cpp bus
117	XG_PCI_WR_2_CPP	P_CLK	single	separate	XG pci write length=2 on cpp bus
118	XG_PCI_WR_3_CPP	P_CLK	single	separate	XG pci write length=3 on cpp bus
119	XG_PCI_WR_4_CPP	P_CLK	single	separate	XG pci write length=4 on cpp bus
120	XG_CAP_RD_CPP	P_CLK	single	separate	XG cap read command on cpp bus
121	XG_CAP_RD_1_CPP	P_CLK	single	separate	XG cap read length=1 on cpp bus
122	XG_CAP_RD_8_CPP	P_CLK	single	separate	XG cap read length=8 on cpp bus
123	XG_CAP_WR_CPP	P_CLK	single	separate	XG cap write command on cpp bus
124	XG_CAP_WR_1_CPP	P_CLK	single	separate	XG cap write length=1 on cpp bus
125	reserved				
126	reserved				
127	reserved				

11.4.6.5 PCI Events Target ID(000101) / Design Block #(1000)

Table 159. PCI PMU Event List (Sheet 1 of 5)

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
0	PCI_TGT_AFIFO_FULL	PCI_CLK	single	separate	PCI Target Address FIFO Full
1	PCI_TGT_AFIFO_NEMPTY	P_CLK	single	separate	PCI Target Address FIFO Not Empty
2	PCI_TGT_AFIFO_WR	PCI_CLK	single	separate	PCI Target Address FIFO Write
3	PCI_TGT_AFIFO_RD	P_CLK	single	separate	PCI Target Address FIFO Read
4	PCI_TGT_RFIFO_FULL	P_CLK	single	separate	PCI Target Read FIFO Full
5	PCI_TGT_RFIFO_NEMPTY	PCI_CLK	single	separate	PCI Target Read FIFO Not Empty
6	PCI_TGT_RFIFO_WR	P_CLK	single	separate	PCI Target Read FIFO Write
7	PCI_TGT_RFIFO_RD	PCI_CLK	single	separate	PCI Target Read FIFO Read
8	PCI_TGT_WFIFO_FULL	PCI_CLK	single	separate	PCI Target Write FIFO Full
9	PCI_TGT_WFIFO_NEMPTY	P_CLK	single	separate	PCI Target Write FIFO Not Empty
10	PCI_TGT_WFIFO_WR	PCI_CLK	single	separate	PCI Target Write FIFO Write
11	PCI_TGT_WFIFO_RD	P_CLK	single	separate	PCI Target Write FIFO Read
12	PCI_TGT_WBUF_FULL	P_CLK	single	separate	PCI Target Write Buffer Full



Table 159. PCI PMU Event List (Sheet 2 of 5)

13	PCI_TGT_WBUF_NEMPTY	P_CLK	single	separate	PCI Target Write Buffer Not Empty
14	PCI_TGT_WBUF_WR	P_CLK	single	separate	PCI Target Write Buffer Write
15	PCI_TGT_WBUF_RD	P_CLK	single	separate	PCI Target Write Buffer Read
16	PCI_MST_AFIFO_FULL	P_CLK	single	separate	PCI Master Address FIFO Full
17	PCI_MST_AFIFO_NEMPTY	PCI_CLK	single	separate	PCI Master Address FIFO Not Empty
18	PCI_MST_AFIFO_WR	P_CLK	single	separate	PCI Master Address FIFO Write
19	PCI_MST_AFIFO_RD	PCI_CLK	single	separate	PCI Master Address FIFO Read
20	PCI_MST_RFIFO_FULL	PCI_CLK	single	separate	PCI Master Read FIFO Full
21	PCI_MST_RFIFO_NEMPTY	P_CLK	single	separate	PCI Master Read FIFO Not Empty
22	PCI_MST_RFIFO_WR	PCI_CLK	single	separate	PCI Master Read FIFO Write
23	PCI_MST_RFIFO_RD	P_CLK	single	separate	PCI Master Read FIFO Read
24	PCI_MST_WFIFO_FULL	P_CLK	single	separate	PCI Master Write FIFO Full
25	PCI_MST_WFIFO_NEMPTY	PCI_CLK	single	separate	PCI Master Write FIFO Not Empty
26	PCI_MST_WFIFO_WR	P_CLK	single	separate	PCI Master Write FIFO Write
27	PCI_MST_WFIFO_RD	PCI_CLK	single	separate	PCI Master Write FIFO Read
28	PCI_DMA1_BUF_FULL	P_CLK	single	separate	PCI_DMA_Channel 1
29	PCI_DMA1_BUF_NEMPTY	P_CLK	single	separate	PCI DMA Channel 1 Buffer Not Empty
30	PCI_DMA1_BUF_WR	P_CLK	single	separate	PCI DMA Channel 1 Buffer Write
31	PCI_DMA1_BUF_RD	P_CLK	single	separate	PCI DMA Channel 1 Buffer Read
32	reserved				
33	reserved				
34	reserved				
35	reserved				
36	PCI_DMA3_BUF_FULL	P_CLK	single	separate	PCI_DMA_Channel 3
37	PCI_DMA3_BUF_NEMPTY	P_CLK	single	separate	
38	PCI_DMA3_BUF_WR	P_CLK	single	separate	
39	PCI_DMA3_BUF_RD	P_CLK	single	separate	
40	PCI_TCMD_FIFO_FULL	P_CLK	single	separate	PCI TARGET Command Fifo
41	PCI_TCMD_FIFO_NEMPTY	P_CLK	single	separate	
42	PCI_TCMD_FIFO_WR	P_CLK	single	separate	
43	PCI_TCMD_FIFO_RD	P_CLK	single	separate	
44	PCI_TDATA_FIFO_FULL	P_CLK	single	separate	PCI Push/Pull Data Fifo
45	PCI_TDATA_FIFO_NEMPTY	P_CLK	single	separate	
46	PCI_TDATA_FIFO_WR	P_CLK	single	separate	
47	PCI_TDATA_FIFO_RD	P_CLK	single	separate	
48	PCI_CSR_WRITE	P_CLK	single	separate	PCI Write to PCI_CSR_BAR
49	PCI_CSR_READ	P_CLK	single	separate	
50	PCI_DRAM_WRITE	P_CLK	single	separate	PCI Write to PCI_DRAM_BAR
51	PCI_DRAM_READ	P_CLK	single	separate	



Table 159. PCI PMU Event List (Sheet 3 of 5)

52	PCI_DRAM_BURST_WRITE	P_CLK	single	separate	PCI Burst Write to PCI_CSR_BAR
53	PCI_DRAM_BURST_READ	P_CLK	single	separate	PCI Burst Read to PCI_CSR_BAR
54	PCI_SRAM_WRITE	P_CLK	single	separate	PCI Write to PCI_SRAM_BAR
55	PCI_SRAM_READ	P_CLK	single	separate	
56	PCI_SRAM_BURST_WRITE	P_CLK	single	separate	PCI Burst Write to PCI_SRAM_BAR
57	PCI_SRAM_BURST_READ	P_CLK	single	separate	
58	PCI_CSR_CMD	P_CLK	single	separate	PCI CSR Command Generated
59	PCI_CSR_PUSH	P_CLK	single	separate	PCI CSR Push Command
60	PCI_CSR_PULL	P_CLK	single	separate	PCI CSR Pull Command
61	PCI_SRAM_CMD	P_CLK	single	separate	PCI SRAM Command
62	PCI_SRAM_PUSH	P_CLK	single	separate	PCI SRAM Push Command
63	PCI_SRAM_PULL	P_CLK	single	separate	PCI SRAM Pull Command
64	PCI_DRAM_CMD	P_CLK	single	separate	PCI DRAM Command
65	PCI_DRAM_PUSH	P_CLK	single	separate	
66	PCI_DRAM_PULL	P_CLK	single	separate	
67	PCI_CSR_2PCI_WR	P_CLK	single	separate	PCI Target Write to PCI local CSR
68	PCI_CSR_2PCI_RD	P_CLK	single	separate	
69	PCI_CSR_2CFG_WR	PCI_CLK	single	separate	PCI Target Write to PCI local Config CSR
70	PCI_CSR_2CFG_RD	PCI_CLK	single	separate	
71	PCI_CSR_2SRAM_WR	P_CLK	single	separate	PCI Target Write to SRAM CSR
72	PCI_CSR_2SRAM_RD	P_CLK	single	separate	
73	PCI_CSR_2DRAM_WR	P_CLK	single	separate	PCI Target Write to DRAM CSR
74	PCI_CSR_2DRAM_RD	P_CLK	single	separate	
75	PCI_CSR_2CAP_WR	P_CLK	single	separate	PCI Target Write to CAPCSR
76	PCI_CSR_2CAP_RD	P_CLK	single	separate	
77	PCI_CSR_2MSF_WR	P_CLK	single	separate	PCI Target Write to MSFCSR
78	PCI_CSR_2MSF_RD	P_CLK	single	separate	
79	PCI_CSR_2SCRAPE_WR	P_CLK	single	separate	PCI Target Write to Scrape CSR
80	PCI_CSR_2SCRAPE_RD	P_CLK	single	separate	
81	PCI_CSR_2SCRATCH_RING_WR	P_CLK	single	separate	PCI Target Write to Scratch Ring CSR
82	PCI_CSR_2SCRATCH_RING_RD	P_CLK	single	separate	
83	PCI_CSR_2SRAM_RING_WR	P_CLK	single	separate	PCI Target Write to SRAM Ring CSR
84	PCI_CSR_2SRAM_RING_RD	P_CLK	single	separate	
85	PCI_XS_LCFG_RD	P_CLK	single	separate	PCI Intel XScale® Core Read Local Config CSR
86	PCI_XS_LCFG_WR	P_CLK	single	separate	
87	PCI_XS_CSR_RD	P_CLK	single	separate	PCI Intel XScale® Core Read Local CSR
88	PCI_XS_CSR_WR	P_CLK	single	separate	



Table 159. PCI PMU Event List (Sheet 4 of 5)

89	PCI_XS_CFG_RD	P_CLK	single	separate	PCI Intel XScale® Core Read PCI Bus Config Space
90	PCI_XS_CFG_WR	P_CLK	single	separate	
91	PCI_XS_MEM_RD	P_CLK	single	separate	PCI Intel XScale® Core Read PCI Bus Memory Space
92	PCI_XS_MEM_WR	P_CLK	single	separate	
93	PCI_XS_BURST_RD	P_CLK	single	separate	PCI Intel XScale® Core Burst Read PCI Bus Memory Space
94	PCI_XS_BURST_WR	P_CLK	single	separate	
95	PCI_XS_IO_RD	P_CLK	single	separate	PCI Intel XScale® Core Read PCI Bus I/O Space
96	PCI_XS_IO_WR	P_CLK	single	separate	
97	PCI_XS_SPEC	P_CLK	single	separate	PCI Intel XScale® Core Read PCI Bus as Special
98	PCI_XS_IACK	P_CLK	single	separate	PCI Intel XScale® Core Read PCI Bus as IACK
99	PCI_ME_CSR_RD	P_CLK	single	separate	PCI ME Read Local CSR
100	PCI_ME_CSR_WR	P_CLK	single	separate	
101	PCI_ME_MEM_RD	P_CLK	single	separate	PCI ME Read PCI Bus Memory Space
102	PCI_ME_MEM_WR	P_CLK	single	separate	
103	PCI_ME_BURST_RD	P_CLK	single	separate	PCI ME Burst Read PCI Bus Memory Space
104	PCI_ME_BURST_WR	P_CLK	single	separate	
105	PCI_MST_CFG_RD	P_CLK	single	separate	PCI Initiator Read PCI Bus Config Space
106	PCI_MST_CFG_WR	P_CLK	single	separate	
107	PCI_MST_MEM_RD	P_CLK	single	separate	PCI Initiator Read PCI Bus Memory Space
108	PCI_MST_MEM_WR	P_CLK	single	separate	
109	PCI_MST_BURST_RD	P_CLK	single	separate	PCI Initiator Burst Read PCI Bus Memory Space
110	PCI_MST_BURST_WR	P_CLK	single	separate	
111	PCI_MST_IO_READ	P_CLK	single	separate	PCI Initiator Read PCI Bus I/O Space
112	PCI_MST_IO_WRITE	P_CLK	single	separate	
113	PCI_MST_SPEC	P_CLK	single	separate	PCI Initiator Read PCI Bus As a Special Cycle
114	PCI_MST_IACK	P_CLK	single	separate	PCI Initiator Read PCI Bus As IACK Cycle
115	PCI_MST_READ_LINE	P_CLK	single	separate	PCI Initiator Read Line Command to PCI
116	PCI_MST_READ_MULT	P_CLK	single	separate	PCI Initiator Read Line Multiple Command to PCI
117	PCI_ARB_REQ[2]	PCI_CLK	single	separate	Internal Arbiter PCI Bus Request 2



Table 159. PCI PMU Event List (Sheet 5 of 5)

118	PCI_ARB_GNT[2]	PCI_CLK	single	separate	Internal Arbiter PCI Bus Grant 2
119	PCI_ARB_REQ[1]	PCI_CLK	single	separate	
120	PCI_ARB_GNT[1]	PCI_CLK	single	separate	
121	PCI_ARB_REQ[0]	PCI_CLK	single	separate	
122	PCI_ARB_GNT[0]	PCI_CLK	single	separate	
123	PCI_TGT_STATE[4]	P_CLK	single	separate	PCI Target State Machine State Bit 4
124	PCI_TGT_STATE[3]	P_CLK	single	separate	
125	PCI_TGT_STATE[2]	P_CLK	single	separate	
126	PCI_TGT_STATE[1]	P_CLK	single	separate	
127	PCI_TGT_STATE[0]	P_CLK	single	separate	

11.4.6.6 ME00 Events Target ID(100000) / Design Block #(1001)

Table 160. ME00 PMU Event List (Sheet 1 of 2)

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
0	ME_FIFO_ENQ_EVEN	T_CLK	single	separate	Even version of Command FIFO Enqueue (pair with event #8)
1	ME_IDLE_EVEN	T_CLK	single	separate	Even version of No Thread running in Microengine (pair with event #9)
2	ME_EXECUTING_EVEN	T_CLK	single	separate	Even version of Valid Instruction (pair with event #10)
3	ME_STALL_EVEN	T_CLK	single	separate	Even version of Microengine stall caused by FIFO Full (pair with event #11)
4	ME_CTX_SWAPPING_EVEN	T_CLK	single	separate	Even version of Occurrence of context swap (pair with event #12)
5	ME_INST_ABORT_EVEN	T_CLK	single	separate	Even version of Instruction aborted due to branch taken (pair with event #13)
6	ME_FIFO_ENQ_ODD	T_CLK	single	separate	Odd version of Command FIFO Enqueue (pair with event #0)
7	ME_IDLE_ODD	T_CLK	single	separate	Odd version of No Thread running in Microengine (pair with event #3)
8	ME_EXECUTING_ODD	T_CLK	single	separate	Odd version of Valid Instruction (pair with event #4)
9	ME_STALL_ODD	T_CLK	single	separate	Odd version of Microengine stall caused by FIFO Full (pair with event #5)
10	ME_CTX_SWAPPING_ODD	T_CLK	single	separate	Odd version of Occurrence of context swap (pair with event #6)
11	ME_INST_ABORT_ODD	T_CLK	single	separate	Odd version of Instruction aborted due to branch (pair with event #7)



Table 160. ME00 PMU Event List (Sheet 2 of 2)

12	ME_FIFO_DEQ	P_CLK	single	separate	Command FIFO Dequeue
13	ME_FIFO_NOT_EMPTY	P_CLK	single	separate	Command FIFO <i>not</i> empty
<p>Note:</p> <ol style="list-style-type: none"> All the Microengine have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.7 ME01 Events Target ID(100001) / Design Block #(1001)

Table 161. ME01 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.8 ME02 Events Target ID(100010) / Design Block #(1001)

Table 162. ME02 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.9 ME03 Events Target ID(100011) / Design Block #(1001)

Table 163. ME03 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.10 ME04 Events Target ID(100100) / Design Block #(1001)

Table 164. ME04 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.11 ME05 Events Target ID(100101) / Design Block #(1001)

Table 165. ME05 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.12 ME06 Events Target ID(100110) / Design Block #(1001)

Table 166. ME06 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.13 ME07 Events Target ID(100111) / Design Block #(1001)

Table 167. ME07 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.14 ME10 Events Target ID(110000) / Design Block #(1010)

Table 168. ME10 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.15 ME11 Events Target ID(110001) / Design Block #(1010)

Table 169. ME11 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.16 ME12 Events Target ID(110010) / Design Block #(1010)

Table 170. ME12 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.17 ME13 Events Target ID(110011) / Design Block #(1010)

Table 171. ME13 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.18 ME14 Events Target ID(110100) / Design Block #(1010)

Table 172. ME14 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.19 ME15 Events Target ID(110101) / Design Block #(1010)

Table 173. ME15 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.20 ME16 Events Target ID(100110) / Design Block #(1010)

Table 174. ME16 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					

11.4.6.21 ME17 Events Target ID(110111) / Design Block #(1010)

Table 175. ME17 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the Microengines have the same event list. CC_Enable bit[2:0] is PMU_CTX_Monitor in Microengine CSR, This field holds the number of context to be monitored. The event count only reflects the events that occur when this context is executing. CC_Enable[2:0] = 000, select context number 0, CC_Enable[2:0] = 001, select context number 1, CC_Enable[2:0] = 111, select context number 7. 1.4 GHz events are sampled by the PMU at a 700 MHz rate. For this reason, all 1.4 GHz events have both an even and an odd event. To determine the total number of 1.4 GHz events, the occurrences of the even events and odd events should be added together. For IXP2800 Network Processor Rev B, CC_Enable[3] must be set to 1 on all 16 Microengines for proper PMU functionality. 					



11.4.6.22 SRAM DP1 Events Target ID(001001) / Design Block #(0010)

Table 176. SRAM DP1 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> 1. SRAM DP1/DP0 push/pull arbiter has same event lists. 2. S_CLK = SRAM clock domain 3. P_CLK = PP clock domain <p>signals that begin with sps_ correspond to S_Push Arb signals that begin with spl_ correspond to S_Pull Arb</p> <p>signals that contain _pc_ (after the unit designation) correspond to the PCI target interface signals that contain _m_ (after the unit designation) correspond to the MSF target interface signals that contain _sh_ (after the unit designation) correspond to the SHaC target interface signals that contain _s0_ (after the unit designation) correspond to the SRAM0 target interface signals that contain _s1_ (after the unit designation) correspond to the SRAM1 target interface signals that contain _s2_ (after the unit designation) correspond to the SRAM2 target interface signals that contain _s3_ (after the unit designation) correspond to the SRAM3 target interface</p>					

11.4.6.23 SRAM DP0 Events Target ID(001010) / Design Block #(0010)

Table 177. SRAM DP0 PMU Event List (Sheet 1 of 3)

Event Number	Event Name	Clock Domain	Single pulse/Long pulse	Burst	Description
0	sps_pc_cmd_valid_rph	P_CLK	Long	separate	PCI Push Command Queue FIFO Valid
1	sps_pc_enq_wph	P_CLK	single	separate	PCI Push Command Queue FIFO Enqueue
2	sps_pc_deq_wph	P_CLK	single	separate	PCI Push Command Queue FIFO Dequeue
3	sps_pc_push_q_full_wph	P_CLK	Long	separate	PCI Push Command Queue FIFO Full
4	sps_m_cmd_valid_rph	P_CLK	Long	separate	MSF Push Command Queue FIFO Valid
5	sps_m_enq_wph	P_CLK	single	separate	MSF Push Command Queue FIFO Enqueue
6	sps_m_deq_wph	P_CLK	single	separate	MSF Push Command Queue FIFO Dequeue
7	sps_m_push_q_full_wph	P_CLK	Long	separate	MSF Push Command Queue FIFO Full
8	sps_sh_cmd_valid_rph	P_CLK	Long	separate	SHaC Push Command Queue FIFO Valid
9	sps_sh_enq_wph	P_CLK	single	separate	SHaC Push Command Queue FIFO Enqueue
10	sps_sh_deq_wph	P_CLK	single	separate	SHaC Push Command Queue FIFO Dequeue
11	sps_sh_push_q_full_wph	P_CLK	Long	separate	SHaC Push Command Queue FIFO Full
12	sps_s0_cmd_valid_rph	P_CLK	Long	separate	SRAM0 Push Command Queue FIFO Valid



Table 177. SRAM DP0 PMU Event List (Sheet 2 of 3)

13	sps_s0_enq_wph	P_CLK	single	separate	SRAM0 Push Command Queue FIFO Enqueue
14	sps_s0_deq_wph	P_CLK	single	separate	SRAM0 Push Command Queue FIFO Dequeue
15	sps_s0_push_q_full_wph	P_CLK	Long	separate	SRAM0 Push Command Queue FIFO Full
16	sps_s1_cmd_valid_rph	P_CLK	Long	separate	SRAM1 Push Command Queue FIFO Valid
17	sps_s1_enq_wph	P_CLK	single	separate	SRAM1 Push Command Queue FIFO Enqueue
18	sps_s1_deq_wph	P_CLK	single	separate	SRAM1 Push Command Queue FIFO Dequeue
19	sps_s1_push_q_full_wph	P_CLK	Long	separate	SRAM1 Push Command Queue FIFO Full
20	sps_s2_cmd_valid_rph	P_CLK	Long	separate	SRAM2 Push Command Queue FIFO Valid
21	sps_s2_enq_wph	P_CLK	single	separate	SRAM2 Push Command Queue FIFO Enqueue
22	sps_s2_deq_wph	P_CLK	single	separate	SRAM2 Push Command Queue FIFO Dequeue
23	sps_s2_push_q_full_wph	P_CLK	Long	separate	SRAM2 Push Command Queue FIFO Full
24	sps_s3_cmd_valid_rph	P_CLK	Long	separate	SRAM3 Push Command Queue FIFO Valid
25	sps_s3_enq_wph	P_CLK	single	separate	SRAM3 Push Command Queue FIFO Enqueue
26	sps_s3_deq_wph	P_CLK	single	separate	SRAM3 Push Command Queue FIFO Dequeue
27	sps_s3_push_q_full_wph	P_CLK	Long	separate	SRAM3 Push Command Queue FIFO Full
28	spl_pc_cmd_valid_rph	P_CLK	Long	separate	PCI Pull Command Queue FIFO Valid
29	spl_pc_enq_cmd_wph	P_CLK	single	separate	PCI Pull Command Queue FIFO Enqueue
30	spl_pc_deq_wph	P_CLK	single	separate	PCI Pull Command Queue FIFO Dequeue
31	spl_pc_cmd_que_full_wph	P_CLK	Long	separate	PCI Pull Command Queue FIFO Full
32	spl_m_cmd_valid_rph	P_CLK	Long	separate	MSF Pull Command Queue FIFO Valid
33	spl_m_enq_cmd_wph	P_CLK	single	separate	MSF Pull Command Queue FIFO Enqueue
34	spl_m_deq_wph	P_CLK	single	separate	MSF Pull Command Queue FIFO Dequeue
35	spl_m_cmd_que_full_wph	P_CLK	Long	separate	MSF Pull Command Queue FIFO Full
36	spl_sh_cmd_valid_rph	P_CLK	Long	separate	SHaC Pull Command Queue FIFO Valid
37	spl_sh_enq_cmd_wph	P_CLK	single	separate	SHaC Pull Command Queue FIFO Enqueue
38	spl_sh_deq_wph	P_CLK	single	separate	SHaC Pull Command Queue FIFO Dequeue
39	spl_sh_cmd_que_full_wph	P_CLK	Long	separate	SHaC Pull Command Queue FIFO Full
40	spl_s0_cmd_valid_rph	P_CLK	Long	separate	SRAM0 Pull Command Queue FIFO Valid
41	spl_s0_enq_cmd_wph	P_CLK	single	separate	SRAM0 Pull Command Queue FIFO Enqueue
42	spl_s0_deq_wph	P_CLK	single	separate	SRAM0 Pull Command Queue FIFO Dequeue
43	spl_s0_cmd_que_full_wph	P_CLK	Long	separate	SRAM0 Pull Command Queue FIFO Full
44	spl_s1_cmd_valid_rph	P_CLK	Long	separate	SRAM1 Pull Command Queue FIFO Valid



Table 177. SRAM DP0 PMU Event List (Sheet 3 of 3)

45	spl_s1_enq_cmd_wph	P_CLK	single	separate	SRAM1 Pull Command Queue FIFO Enqueue
46	spl_s1_deq_wph	P_CLK	single	separate	SRAM1 Pull Command Queue FIFO Dequeue
47	spl_s1_cmd_que_full_wph	P_CLK	Long	separate	SRAM1 Pull Command Queue FIFO Full
48	spl_s2_cmd_valid_rph	P_CLK	Long	separate	SRAM2 Pull Command Queue FIFO Valid
49	spl_s2_enq_cmd_wph	P_CLK	single	separate	SRAM2 Pull Command Queue FIFO Enqueue
50	spl_s2_deq_wph	P_CLK	single	separate	SRAM2 Pull Command Queue FIFO Dequeue
51	spl_s2_cmd_que_full_wph	P_CLK	Long	separate	SRAM2 Pull Command Queue FIFO Full
52	spl_s3_cmd_valid_rph	P_CLK	Long	separate	SRAM3 Pull Command Queue FIFO Valid
53	spl_s3_enq_cmd_wph	P_CLK	single	separate	SRAM3 Pull Command Queue FIFO Enqueue
54	spl_s3_deq_wph	P_CLK	single	separate	SRAM3 Pull Command Queue FIFO Dequeue
55	spl_s3_cmd_que_full_wph	P_CLK	Long	separate	SRAM3 Pull Command Queue FIFO Full

11.4.6.24 SRAM CH3 Events Target ID(001011) / Design Block #(0010)

Table 178. SRAM CH3 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the SRAM Channel has same event lists. S_CLK = SRAM clock domain P_CLK = PP clock domain <p>signals that begin with sps_ correspond to S_Push Arb signals that begin with spl_ correspond to S_Pull Arb</p> <p>signals that contain _pc_ (after the unit designation) correspond to the PCI target interface signals that contain _m_ (after the unit designation) correspond to the MSF target interface signals that contain _sh_ (after the unit designation) correspond to the SHaC target interface signals that contain _s0_ (after the unit designation) correspond to the SRAM0 target interface signals that contain _s1_ (after the unit designation) correspond to the SRAM1 target interface signals that contain _s2_ (after the unit designation) correspond to the SRAM2 target interface signals that contain _s3_ (after the unit designation) correspond to the SRAM3 target interface</p>					



11.4.6.25 SRAM CH2 Events Target ID(001100) / Design Block #(0010)

Table 179. SRAM CH3 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the SRAM Channel has same event lists. S_CLK = SRAM clock domain P_CLK = PP clock domain <p>signals that begin with sps_ correspond to S_Push Arb signals that begin with spl_ correspond to S-Pull Arb</p> <p>signals that contain _pc_ (after the unit designation) correspond to the PCI target interface signals that contain _m_ (after the unit designation) correspond to the MSF target interface signals that contain _sh_ (after the unit designation) correspond to the SHaC target interface signals that contain _s0_ (after the unit designation) correspond to the SRAM0 target interface signals that contain _s1_ (after the unit designation) correspond to the SRAM1 target interface signals that contain _s2_ (after the unit designation) correspond to the SRAM2 target interface signals that contain _s3_ (after the unit designation) correspond to the SRAM3 target interface</p>					

11.4.6.26 SRAM CH1 Events Target ID(001101) / Design Block #(0010)

Table 180. SRAM CH3 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
<p>Note:</p> <ol style="list-style-type: none"> All the SRAM Channel has same event lists. S_CLK = SRAM clock domain P_CLK = PP clock domain <p>signals that begin with sps_ correspond to S-Push Arb signals that begin with spl_ correspond to S-Pull Arb</p> <p>signals that contain _pc_ (after the unit designation) correspond to the PCI target interface signals that contain _m_ (after the unit designation) correspond to the MSF target interface signals that contain _sh_ (after the unit designation) correspond to the SHaC target interface signals that contain _s0_ (after the unit designation) correspond to the SRAM0 target interface signals that contain _s1_ (after the unit designation) correspond to the SRAM1 target interface signals that contain _s2_ (after the unit designation) correspond to the SRAM2 target interface signals that contain _s3_ (after the unit designation) correspond to the SRAM3 target interface</p>					



11.4.6.27 SRAM CH0 Events Target ID(001110) / Design Block #(0010)

Table 181. SRAM CH0 PMU Event List (Sheet 1 of 2)

Event Number	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
0	QDR I/O Read	S_CLK	single	separate	QDR I/O Read
1	QDR I/O Write	S_CLK	single	separate	QDR I/O Write
2	Read Cmd Dispatched	P_CLK	single	separate	Read Cmd Dispatched
3	Write Cmd Dispatched	P_CLK	single	separate	Write Cmd Dispatched
4	Swap Cmd Dispatched	P_CLK	single	separate	Swap Cmd Dispatched
5	Set Dispatched	P_CLK	single	separate	Set Dispatched
6	Clear Cmd Dispatched	P_CLK	single	separate	Clear Cmd Dispatched
7	Add Cmd Dispatched	P_CLK	single	separate	Add Cmd Dispatched
8	Sub Cmd Dispatched	P_CLK	single	separate	Sub Cmd Dispatched
9	Incr Cmd Dispatched	P_CLK	single	separate	Incr Cmd Dispatched
10	Decr Cmd Dispatched	P_CLK	single	separate	Decr Cmd Dispatched
11	Ring Cmd Dispatched	P_CLK	single	separate	Ring Cmd Dispatched
12	Jour Cmd Dispatched	P_CLK	single	separate	Jour Cmd Dispatched
13	Deq Cmd Dispatched	P_CLK	single	separate	Deq Cmd Dispatched
14	Enq Cmd Dispatched	P_CLK	single	separate	Enq Cmd Dispatched
15	CSR Cmd Dispatched	P_CLK	single	separate	CSR Cmd Dispatched
16	WQDesc Cmd Dispatched	P_CLK	single	separate	WQDesc Cmd Dispatched
17	RQDesc Cmd Dispatched	P_CLK	single	separate	RQDesc Cmd Dispatched
18	FIFO Dequeue – CmdA0 Inlet Q	P_CLK	single	separate	FIFO Dequeue – CmdA0 Inlet Q
19	FIFO Enqueue – CmdA0 Inlet Q	P_CLK	single	separate	FIFO Enqueue – CmdA0 Inlet Q
20	FIFO Valid – CmdA0 Inlet Q	P_CLK	long	separate	FIFO Valid – CmdA0 Inlet Q
21	FIFO Full – CmdA1 Inlet Q	P_CLK	long	separate	FIFO Full – CmdA1 Inlet Q
22	FIFO Dequeue – CmdA1 Inlet Q	P_CLK	single	separate	FIFO Dequeue – CmdA1 Inlet Q
23	FIFO Enqueue – CmdA1 Inlet Q	P_CLK	single	separate	FIFO Enqueue – CmdA1 Inlet Q
24	FIFO Valid – CmdA1 Inlet Q	P_CLK	long	separate	FIFO Valid – CmdA1 Inlet Q
25	FIFO Full – CmdA1 Inlet Q	P_CLK	long	separate	FIFO Full – CmdA1 Inlet Q
26	FIFO Dequeue – Wr Cmd Q	S_CLK	single	separate	FIFO Dequeue – Wr Cmd Q
27	FIFO Enqueue – Wr Cmd Q	P_CLK	single	separate	FIFO Enqueue – Wr Cmd Q
28	FIFO Valid – Wr Cmd Q	S_CLK	long	separate	FIFO Valid – Wr Cmd Q
29	FIFO Full – Wr Cmd Q	P_CLK	long	separate	FIFO Full – Wr Cmd Q
30	FIFO Dequeue – Queue Cmd Q	S_CLK	single	separate	FIFO Dequeue – Queue Cmd Q
31	FIFO Enqueue – Queue Cmd Q	P_CLK	single	separate	FIFO Enqueue – Queue Cmd Q
32	FIFO Valid – Queue Cmd Q	S_CLK	long	separate	FIFO Valid – Queue Cmd Q



Table 181. SRAM CH0 PMU Event List (Sheet 2 of 2)

33	FIFO Full – Queue Cmd Q	P_CLK	long	separate	FIFO Full – Queue Cmd Q
34	FIFO Dequeue – Rd Cmd Q	S_CLK	single	separate	FIFO Dequeue – Rd Cmd Q
35	FIFO Enqueue – Rd Cmd Q	P_CLK	single	separate	FIFO Enqueue – Rd Cmd Q
36	FIFO Valid – Rd Cmd Q	S_CLK	long	separate	FIFO Valid – Rd Cmd Q
37	FIFO Full – Rd Cmd Q	P_CLK	long	separate	FIFO Full – Rd Cmd Q
38	FIFO Dequeue – Oref Cmd Q	S_CLK	single	separate	FIFO Dequeue – Oref Cmd Q
39	FIFO Enqueue – Oref Cmd Q	P_CLK	single	separate	FIFO Enqueue – Oref Cmd Q
40	FIFO Valid – Oref Cmd Q	S_CLK	long	separate	FIFO Valid – Oref Cmd Q
41	FIFO Full – Oref Cmd Q	P_CLK	long	separate	FIFO Full – Oref Cmd Q
42	FIFO Dequeue – SP0 Pull Data Q	S_CLK	single	separate	FIFO Dequeue – SP0 Pull Data Q
43	FIFO Enqueue – SP0 Pull Data Q	P_CLK	single	separate	FIFO Enqueue – SP0 Pull Data Q
44	FIFO Valid – SP0 Pull Data Q	S_CLK	long	separate	FIFO Valid – SP0 Pull Data Q
45	FIFO Full – SP0 Pull Data Q	P_CLK	long	separate	FIFO Full – SP0 Pull Data Q
46	FIFO Dequeue – SP1 Pull Data Q	S_CLK	single	separate	FIFO Dequeue – SP1 Pull Data Q
47	FIFO Enqueue – SP1 Pull Data Q	P_CLK	single	separate	FIFO Enqueue – SP1 Pull Data Q
48	FIFO Valid – SP1 Pull Data Q	S_CLK	long	separate	FIFO Valid – SP1 Pull Data Q
49	FIFO Full – SP1 Pull Data Q	P_CLK	long	separate	FIFO Full – SP1 Pull Data Q
50	FIFO Dequeue – Push ID/Data Q	P_CLK	single	separate	FIFO Dequeue – Push ID/Data Q
51	FIFO Enqueue – Push ID/Data Q	S_CLK	single	separate	FIFO Enqueue – Push ID/Data Q
52	FIFO Valid – Push ID/Data Q	P_CLK	long	separate	FIFO Valid – Push ID/Data Q
53	FIFO Full – Push ID/Data Q	S_CLK	long	separate	FIFO Full – Push ID/Data Q

11.4.6.28 DRAM DPLA Events Target ID(010010) / Design Block #(0011)

Table 182. IXP2800 Network Processor Dram DPLA PMU Event List (Sheet 1 of 2)

Event Number	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
0	d0_enq_id_wph	P_CLK	single	separate	Enqueue d0 cmd
1	d0_deq_id_wph	P_CLK	single	separate	Dequeue d0 cmd
2	dram_req_rph[0]	P_CLK	single	separate	d0 has a valid req
3	next_d0_full_wph	P_CLK	single	separate	d0 FIFO hit the full threshold
4	d1_enq_id_wph	P_CLK	single	separate	Enqueue d1 cmd
5	d1_deq_id_wph	P_CLK	single	separate	Dequeue d1 cmd
6	dram_req_rph[1]	P_CLK	single	separate	d1 has a valid req
7	next_d1_full_wph	P_CLK	single	separate	d1 FIFO hit the full threshold
8	d2_enq_id_wph	P_CLK	single	separate	Enqueue d2 cmd



Table 182. IXP2800 Network Processor Dram DPLA PMU Event List (Sheet 2 of 2)

9	d2_deq_id_wph	P_CLK	single	separate	Dequeue d2 cmd
10	dram_req_rph[2]	P_CLK	single	separate	d2 has a valid req
11	next_d2_full_wph	P_CLK	single	separate	d2 FIFO hit the full threshold
12	cr0_enq_id_wph	P_CLK	single	separate	Enqueue cr0 cmd
13	cr0_deq_id_wph	P_CLK	single	separate	Dequeue cr0 cmd
14	dram_req_rph[3]	P_CLK	single	separate	cr0 has a valid req
15	next_cr0_full_wph	P_CLK	single	separate	cr0 FIFO hit the full threshold
16	cr1_enq_id_wph	P_CLK	single	separate	Enqueue cr1 cmd
17	cr1_deq_id_wph	P_CLK	single	separate	Dequeue cr1 cmd
18	dram_req_rph[4]	P_CLK	single	separate	cr1 has a valid req
19	next_cr1_full_wph	P_CLK	single	separate	cr1 FIFO hit the full threshold

11.4.6.29 DRAM DPSA Events Target ID(010011) / Design Block #(0011)

Table 183. IXP2800 Network Processor Dram DPSA PMU Event List (Sheet 1 of 2)

Event Number	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
0	d0_enq_id_wph	P_CLK	single	separate	Enqueue d0 cmd/data
1	d0_deq_id_wph	P_CLK	single	separate	Dequeue d0 cmd/data
2	dram_req_rph[0]	P_CLK	single	separate	d0 has a valid req(not empty)
3	next_d0_full_wph	P_CLK	single	separate	d0 FIFO hit the full threshold
4	d1_enq_id_wph	P_CLK	single	separate	Enqueue d1 cmd/data
5	d1_deq_id_wph	P_CLK	single	separate	Dequeue d1 cmd/data
6	dram_req_rph[1]	P_CLK	single	separate	d1 has a valid req
7	next_d1_full_wph	P_CLK	single	separate	d1 FIFO hit the full threshold
8	d2_enq_id_wph	P_CLK	single	separate	Enqueue d2 cmd/data
9	d2_deq_id_wph	P_CLK	single	separate	Dequeue d2 cmd/data
10	dram_req_rph[2]	P_CLK	single	separate	d2 has a valid req
11	next_d2_full_wph	P_CLK	single	separate	d2 FIFO hit the full threshold
12	cr0_enq_id_wph	P_CLK	single	separate	Enqueue cr0 cmd/data
13	cr0_deq_id_wph	P_CLK	single	separate	Dequeue cr0 cmd/data
14	dram_req_rph[3]	P_CLK	single	separate	cr0 has a valid req
15	next_cr0_full_wph	P_CLK	single	separate	cr0 FIFO hit the full threshold
16	cr1_enq_id_wph	P_CLK	single	separate	Enqueue cr1 cmd/data



Table 183. IXP2800 Network Processor Dram DPSA PMU Event List (Sheet 2 of 2)

17	cr1_deq_id_wph	P_CLK	single	separate	Dequeue cr1 cmd/data
18	dram_req_rph[4]	P_CLK	single	separate	cr1 has a valid req
19	next_cr1_full_wph	P_CLK	single	separate	cr1 FIFO hit the full threshold

11.4.6.30 IXP2800 Network Processor DRAM CH2 Events Target ID(010100) / Design Block #(0011)

Table 184. IXP2800 Network Processor Dram CH2 PMU Event List (Sheet 1 of 5)

Event Number	Event Name	Clock Domain	Single pulse/ Long pulse	Burst	Description
0	DATA_ERROR_SYNC_RPH	P_CLK	single	separate	Indicates that a data error has been detected on the RDRAM read data from the RMC. The signal asserts for both correctable and uncorrectable errors.
1	GET_PULL_DATA_SYNC_RPH	P_CLK	single	separate	Asserts when the RMC is accepting RDRAM write data from the d_app_unit block.
2	TAKE_PUSH_DATA_SYNC_RPH	P_CLK	single	separate	Asserts when the RMC is driving RDRAM read data to the d_push_bus_if block.
3	START_SYNC_RPH	P_CLK	single	separate	Input to RMC, asserted to request memory transaction and deasserted when the RMC is ready to accept a command (i.e., when RMC asserts GETC_SYNC_RPH).
4	GETC_SYNC_RPH	P_CLK	single	separate	Output from RMC, indicates the RMC is ready to accept a command for the RDRAM channel.
5	reserved				
6	reserved				
7	dps_push_ctrl_fifo_full_rph	P_CLK	single	separate	Active when the push_control FIFO is nearly full, i.e., > 6 entries 1:0 2'H0
8	push_ctrl_fifo_enq_rph	P_CLK	single	separate	Active when enqueueing push control and status data to the FIFOs in d_push_bus_if.
9	DPS_ENQ_PUSH_DATA_RPH	P_CLK	single	separate	Active when enqueueing data from the RMC into the push data FIFO.
10	valid_push_data_rp	P_CLK	single	separate	Is (data_valid AND !create_databackup), where data_valid indicates data available in the d_push_bus_if data FIFO; create_databackup asserts when the push arbiter FIFO (in dp_unit block) gets nearly full. When it asserts, it prevents dequeuing from the d_push_bus_if data FIFO.
11	push_ctrl_fifo_empty_rph	P_CLK	single	separate	Active when the push control FIFO is empty.
12	deq_push_data_wph	P_CLK	single	separate	Asserts to dequeue from the data FIFO in the d_push_bus_if block.
13	deq_csr_data_wph	P_CLK	single	separate	Pulses active when reading from a CSR instead of DRAM.



Table 184. IXP2800 Network Processor Dram CH2 PMU Event List (Sheet 2 of 5)

14	deq_push_ctrl_wph	P_CLK	single	separate	Active when dequeuing from the push control FIFO; occurs on the last cycle of a burst or on the only cycle of a single transfer.
15	d_push_ctrl_fsm/ single_xfer_wph	P_CLK	single	separate	Active if the data is about to be transferred from d_push data FIFO to the dp_unit FIFO is length 0, i.e., a single 8-byte transfer.
16	d_push_ctrl_fsm/ data_128_bit_aligned	P_CLK	single	separate	Active if the data is about to be transferred from d_push data FIFO to the dp_unit FIFO is quad-word (128-bit) aligned.
17	perf_data_fifo_full	P_CLK	single	separate	Asserts when the data FIFO in the d_push_bus_if block has > 4 entries (Data from the RDRAM is enqueued into this FIFO; dequeued data is written to the push bus arbiter FIFO in dp_unit block.)
18	reserved				
19	reserved				
20	DPL_RMW_BANK3_READ_ DATA_AVAIL_RPH	P_CLK	single	separate	Indicates that the read data for a read-modify-write operation is available in the d_pull_bus_if block. This signal is deasserted when the data and command
21	DPL_RMW_BANK2_READ_ DATA_AVAIL_RPH	P_CLK	single	separate	Indicates that the read data for a read-modify-write operation is available in the d_pull_bus_if block. This signal is deasserted when the data and command
22	DPL_RMW_BANK1_READ_ DATA_AVAIL_RPH	P_CLK	single	separate	Indicates that the read data for a read-modify-write operation is available in the d_pull_bus_if block. This signal is deasserted when the data and command
23	DPL_RMW_BANK0_READ_ DATA_AVAIL_RPH	P_CLK	single	separate	Indicates that the read data for a read-modify-write operation is available in the d_pull_bus_if block. This signal is deasserted when the data and command
24	addr_128bit_aligned	P_CLK	single	separate	Indicates that bit 3 of the DRAM command's address at the head of the pull control FIFO (i.e., about to be dequeued) is low. This command is for the pull data which is about to be enqueued into a pull data bank FIFO.
25	b3_empty_rph	P_CLK	single	separate	Indicates that the pull data's bank FIFO is empty.
26	b2_empty_rph	P_CLK	single	separate	Indicates that the pull data's bank FIFO is empty.
27	b1_empty_rph	P_CLK	single	separate	Indicates that the pull data's bank FIFO is empty.
28	b0_empty_rph	P_CLK	single	separate	indicates that the pull data's bank FIFO is empty
29	b3_full_rph	P_CLK	single	separate	Indicates that the pull data's bank FIFO has > 0xf entries in it, i.e., is full
30	b2_full_rph	P_CLK	single	separate	Indicates that the pull data's bank FIFO has > 0xf entries in it, i.e., is full.
31	b1_full_rph	P_CLK	single	separate	Indicates that the pull data's bank FIFO has > 0xf entries in it, i.e., is full.
32	b0_full_rph	P_CLK	single	separate	Indicates that the pull data's bank FIFO has > 0xf entries in it, i.e., is full.



Table 184. IXP2800 Network Processor Dram CH2 PMU Event List (Sheet 3 of 5)

33	DAP_DEQ_B3_DATA_RPH	P_CLK	single	separate	Indicates pull data and command are being dequeued from the data and command bank FIFOs to the RMC (the command and data FIFOs used in tandem for pulls to supply the address and data respectively).
34	DAP_DEQ_B2_DATA_RPH	P_CLK	single	separate	
35	DAP_DEQ_B1_DATA_RPH	P_CLK	single	separate	
36	DAP_DEQ_B0_DATA_RPH	P_CLK	single	separate	Indicates pull data and command are being dequeued from the data and command bank FIFOs to the RMC (the command and data FIFOs used in tandem for pulls to supply the address and data respectively).
37	csr_wr_data_avail	P_CLK	single	separate	Indicates that CSR write data is ready to be latched into the CSR.
38	bank3_enq_wph	P_CLK	single	separate	Indicates pull data is being enqueued to a bank FIFO in the pull bus interface block.
39	bank2_enq_wph	P_CLK	single	separate	Indicates pull data is being enqueued to a bank FIFO in the pull bus interface block.
40	bank1_enq_wph	P_CLK	single	separate	Indicates pull data is being enqueued to a bank FIFO in the pull bus interface block.
41	bank0_enq_wph	P_CLK	single	separate	Indicates pull data is being enqueued to a bank FIFO in the pull bus interface block.
42	reserved				
43	DCB_BANK3_CMD_AVAL_RDH	D_CLK	single	separate	Indicates that this bank FIFO has a command available.
44	DCB_BANK2_CMD_AVAL_RDH	D_CLK	single	separate	Indicates that this bank FIFO has a command available.
45	DCB_BANK1_CMD_AVAL_RDH	D_CLK	single	separate	Indicates that this bank FIFO has a command available.
46	DCB_BANK0_CMD_AVAL_RDH	D_CLK	single	separate	Indicates that this bank FIFO has a command available.
47	DAP_CSR_READ_CMD_TAKEN_WDH	D_CLK	single	separate	Indicates dequeuing of a CSR read command and clears the CSR read request signal coming out of d_command_bus_if.
48	DAP_BANK3_CMD_DEQ_WDH	D_CLK	single	separate	Active to dequeue a DRAM command from bank N's FIFO, generated by d_app block.
49	DAP_BANK2_CMD_DEQ_WDH	D_CLK	single	separate	Active to dequeue a DRAM command from bank N's FIFO, generated by d_app block.
50	DAP_BANK1_CMD_DEQ_WDH	D_CLK	single	separate	Active to dequeue a DRAM command from bank N's FIFO, generated by d_app block.
51	DAP_BANK0_CMD_DEQ_WDH	D_CLK	single	separate	Active to dequeue a DRAM command from bank N's FIFO, generated by d_app block.
52	split_cmd_wph	P_CLK	single	separate	Active if the command will cross a 128-byte boundary and thus be split across channels.
53	reserved				
54	reserved				
55	reserved				
56	reserved				



Table 184. IXP2800 Network Processor Dram CH2 PMU Event List (Sheet 4 of 5)

57	reserved				
58	reserved				
59	deq_split_cmd_fifo_wph	P_CLK	single	separate	Active when dequeuing from the split inlet FIFO.
60	deq_inlet_fifo1_wph	P_CLK	single	separate	Active when dequeuing from the inlet FIFO.
61	deq_inlet_fifo_wph	P_CLK	single	separate	Active when dequeuing from either the inlet or split-inlet FIFO.
62	DCB_PULL_CTRL_AVAL_WPH	P_CLK	single	separate	Indicates the pull control FIFO has ≥ 1 entry.
63	inlet_cmd_aval_rph	P_CLK	single	separate	Indicates a command is available in the non-split inlet FIFO.
64	split_fifo_not_empty	P_CLK	single	separate	Indicates a command is available in the "split inlet" FIFO (split refers to a command being split across channels).
65	bank3_pull_ok_wph	P_CLK	single	separate	Indicates that this bank's data FIFO has enough room to accommodate the size of the next pull command in the inlet FIFO.
66	bank2_pull_ok_wph	P_CLK	single	separate	Indicates that this bank's data FIFO has enough room to accommodate the size of the next pull command in the inlet FIFO.
67	bank1_pull_ok_wph	P_CLK	single	separate	Indicates that this bank's data FIFO has enough room to accommodate the size of the next pull command in the inlet FIFO.
68	bank0_pull_ok_wph	P_CLK	single	separate	Indicates that this bank's data FIFO has enough room to accommodate the size of the next pull command in the inlet FIFO.
69	csr_q_full_wph	P_CLK	single	separate	Indicates that a CSR access is in process.
70	DXDP_CMD_Q_FULL_RPH	P_CLK	single	separate	Indicates the command inlet FIFO contains > 8 entries.
71	pull_ctrl_fifo_full	P_CLK	single	separate	Indicates that there are > 6 outstanding pull requests.
72	bank3_cmd_q_full_rph	P_CLK	single	separate	Indicates the bank command FIFO contains > 6 entries.
73	bank2_cmd_q_full_rph	P_CLK	single	separate	Indicates the bank command FIFO contains > 6 entries.
74	bank1_cmd_q_full_rph	P_CLK	single	separate	Indicates the bank command FIFO contains > 6 entries.
75	bank0_cmd_q_full_rph	P_CLK	single	separate	Indicates the bank command FIFO contains > 6 entries.
76	valid_write_req_wph	P_CLK	single	separate	Indicates a DRAM write is being passed from the inlet FIFO to a bank FIFO. The DRAM write may be: DRAM RBUF read, DRAM write, or CSR write.
77	csr_q_full_en_wph	P_CLK	single	separate	Pulses at both the start of a CSR read/write and at the completion of a CSR read/write.
78	push_rmw_wr_cmd_wph	P_CLK	single	separate	Indicates the command being passed from the inlet FIFO to a bank FIFO is a read-modify-write.
79	bank3_enq_wph	P_CLK	single	separate	Indicates this channel is enqueueing a DRAM command for bank3.



Table 184. IXP2800 Network Processor Dram CH2 PMU Event List (Sheet 5 of 5)

80	bank2_enq_wph	P_CLK	single	separate	Indicates this channel is enqueueing a DRAM command for bank2.
81	bank1_enq_wph	P_CLK	single	separate	Indicates this channel is enqueueing a DRAM command for bank1.
82	bank0_enq_wph	P_CLK	single	separate	Indicates this channel is enqueueing a DRAM command for bank0.
83	push2split_cmd_fifo_wph	P_CLK	single	separate	Indicates this channel is enqueueing a DRAM command which is split between two channels.
84	push2inlet_fifo_wph	P_CLK	single	separate	Indicates this channel is enqueueing a DRAM command which fits entirely in this channel.
85	valid_dram_cmd_wph	P_CLK	single	separate	Indicates the command bus' target ID is DRAM.
86-127	reserved				

11.4.6.31 IXP2800 Network Processor DRAM CH1 Events Target ID(010101) / Design Block #(0011)

Table 185. IXP2800 Network Processor Dram CH1 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
Note: 1. All the Dram Channels have same event lists. Please reference Channel 2 Event lists.					

11.4.6.32 IXP2800 Network Processor DRAM CH0 Events Target ID(010110) / Design Block #(0011)

Table 186. IXP2800 Network Processor Dram CH0 PMU Event List

Event Number	Event Name	Clock Domain	Pulse/Level	Burst	Description
Note: 1. All the Dram Channels have same event lists. Please reference Channel 2 Event lists.					

