#### Intel<sup>®</sup> 80333 I/O Processor

**Developer's Manual** 

March 2005

Document Number: 305432001US

Information in this document is provided in connection with Intel<sup>®</sup> products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. 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 products are not intended for use in medical, life saving, or life sustaining 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.

Intel<sup>®</sup> internal code names are subject to change.

THIS SPECIFICATION, THE Intel<sup>®</sup> 80333 I/O Processor IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

#### Copyright © Intel Corporation, 2005

AlertVIEW, i960, AnyPoint, AppChoice, BoardWatch, BunnyPeople, CablePort, Celeron, Chips, Commerce Cart, CT Connect, CT Media, Dialogic, DM3, EtherExpress, ETOX, FlashFile, GatherRound, i386, i486, iCat, iCOMP, Insight960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX4, IntelSX2, Intel ChatPad, Intel Create&Share, Intel Dot.Station, Intel GigaBlade, Intel InBusiness, Intel Inside, Intel Inside logo, Intel NetBurst, Intel Nete IstrataFlash, Intel Packet Concert, Intel SingleDriver, Intel SprateGitep, Intel StrataFlash, Intel reamStation, Intel WebOutfitter, Intel Xeon, Intel Xscale, Itanium, JobAnalyst, LANDesk, LanRover, MCS, MMX, MMX logo, NetPort, NetportExpress, Optimizer logo, OverDrive, Paragon, PC Dads, PC Parents, Pentium, Pentium II Xeon, Pentium III Xeon, Performance at Your Command, ProShare, RemoteExpress, Screamline, Shiva, SmartDie, Solutions960, Sound Mark, StorageExpress, The Computer Inside, The Journey Inside, This Way In, TokenExpress, Trillium, Vivonic, and VTune are trademarks or 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.



1	Introd	duction.			.41		
	1.1	About 7	This Docu	ment	. 41		
		1.1.1	How To F	Read This Document	.41		
		1.1.2	Other Re	levant Documents	.41		
	1.2	About t	he Intel <sup>®</sup> 8	30333 I/O Processor	.42		
	1.3	Features					
		1.3.1	Intel XSc	ale <sup>®</sup> Core	.44		
		1.3.2	PCI Expr	ess-to-PCI Bridge Units	.44		
		1.3.3	Address	Translation Units	.45		
		1.3.4	Memory	Controller	.45		
		1.3.5	Application	on Accelerator Unit	. 45		
		1.3.6	Periphera	al Bus Interface	.45		
		1.3.7	DMA Cor	ntroller	.45		
		1.3.8	I <sup>2</sup> C Bus I	nterface Unit	.46		
		1.3.9	Messagir	ng Unit	.46		
		1.3.10	Internal E	Bus	.46		
		1.3.11	UART Ur	nit	.46		
		1.3.12	Interrupt	Controller Unit	.46		
		1.3.13	GPIO		.46		
		1.3.14	SMBus L	Jnit	.46		
	1.4	Terminology and Conventions4					
		1.4.1	Represer	nting Numbers	.47		
		1.4.2	Fields		.47		
		1.4.3	Specifyin	g Bit and Signal Values	.47		
		1.4.4	Signal Na	ame Conventions	.48		
		1.4.5	Ierminol	ogy	.48		
2	PCI E	xpress-	to-PCI Br	idges	.49		
	2.1	Overvie	ew		.49		
		2.1.1	PCI Expr	ess Interface Features	.49		
		2.1.2	PCI-X Int	erface Features	.50		
		2.1.3	Power M	anagement	.50		
		2.1.4	80333 in	Server Systems	.50		
	2.2	Signal I	Descriptio	n	.51		
		2.2.1	PCI Expr	ess Interface	.52		
		2.2.2	PCI Bus	Interface (Two Instances)	.53		
		2.2.3	PCI Bus	Interface 64-Bit Extension (Two Interfaces)	.55		
		2.2.4	PCI Bus	Interface Clocks and, Reset and Power Management (Two Interfaces)	. 56		
		2.2.5	Reset St	raps	. 57		
	2.3	Program	mming Mo	del and Addressing	. 58		
		2.3.1	Program	ming Model	. 58		
		2.3.2	Addressi	ng	. 59		
			2.3.2.1	Addressable 80333 Spaces	.59		
			2.3.2.2	Chipset Configuration Model	.60		
			2.3.2.3	Secondary PCI Devices	.61		
			2.3.2.4	Configuration Space Access	.61		
				2.3.2.4.1 PUI Express Configuration Access	.61		



			2.3.2.4.2 PCI Configuration Access	63
		2.3.2.5	I/O Space Access Mechanism	64
		2.3.2.6	Memory Space Access Mechanism	66
		2.3.2.7	VGA Addressing	69
2.4	Transa	action Ord	ering	70
	2.4.1	80333 T	ransaction Ordering	70
		2411	Inhound Transaction Ordering	70
		2.4.1.1	Outbound Transaction Ordering	70
25		2.4.1.2	PCL-X Bridge Interfaces	72
2.0	251		torfood	72
	2.3.1			12
		2.5.1.1		12
		2.5.1.2	I ransactions Supported	73
				73
		0540	2.5.1.2.2 PCI-X Mode	
		2.5.1.3	Read Transactions	
			2.5.1.3.1 Prefetchable	
		0 5 4 4		
		2.5.1.4	Configuration Transactions	
		2.5.1.5		
		2.5.1.6	I ransaction Termination	
			2.5.1.6.1 PCI Mode Transaction Termination	
		0 5 4 7	2.5.1.6.2 PCI-X Mode Transaction Termination	
		2.5.1.7	Arbitration	80
		2.5.1.8	PCI-X Protocol Specifics	82
			2.5.1.8.1 Attributes	82
			2.5.1.8.2 4 GB and 4 K Page Crossover	82
			2.5.1.8.3 Wait States	
			2.5.1.8.4 Split Transactions	82
	2.5.2	LOCK C	ycles	83
		2.5.2.1	Error Logging and Escalation	84
			2.5.2.1.1 Bridge and APIC Error Escalation Rules	
			on PCI Express	84
2.6	Power	Managem	nent	86
	2.6.1	PCI Exp	ress Link Power Management	86
		2.6.1.1	Hardware Controlled Active State Power Management	86
		2.6.1.2	OS-Driven PCI-PM Compatible Power Management	86
	262	PCI Bus	Power Management	87
	263	80333 D	evice Power Management	
	2.0.0	Power M	Ionagoment Event Signaling	
07			ad Error Llondling	
2.1	RADE	eatures a	nu Error Hanuling	90
	2.7.1	RASFea	atures	90
		2.7.1.1	PCI Express Error Handling	90
		2.7.1.2	PCI Express Hot-Plug	90
		2.7.1.3	Internal RAM Error Protection	90
		2.7.1.4	PCI Error Protection	90
		2.7.1.5	Advanced Error Reporting	90
	2.7.2	Error Ha	Indling	91
		2.7.2.1	Completion Required Transaction Termination Translation	
			Between Interfaces	91
			2.7.2.1.1 Outbound Transactions – PCI Express to PCI/X	91
			2.7.2.1.2 Inbound Transactions – PCI/X-to-PCI Express	92
			2.7.2.1.3 Error Logging and Escalation	93
			2.7.2.1.4 Summary of Error Reporting	94
		2.7.2.2	Data Poisoning	96

2.8	BIOS/I	Device Driv	ver Recommendations	97
	2.8.1	Standard	Hot-Plug Controller Initialization Requirements	97
	2.8.2	80333-Si	pecific Bridge Register Initialization	98
	283	Peer-to-F	Peer Memory Read Support	98
	284		Bit Usage Guidelines	98
	2.0.4	Porforma	unco Considerations	00
	2.0.5	2951		30
2.0	Confin	2.0.3.1	rietere	90
2.9	Coning		gisiers	99
	2.9.1			99
	2.9.2	Accesses	s to Reserved Registers	99
	2.9.3	80333-S	pecificBridge Register Initializatio	n100
	2.9.4	80333 - 3	Specific Bridge Register Initialization	101
	2.9.5	Configura	ation Registers	102
		2.9.5.1	Identifiers - ID (Offset 00)	105
		2.9.5.2	Command - PCICMD (Offset 04)	106
		2.9.5.3	Primary Device Status - PSTS (Offset 06)	108
		2.9.5.4	Revision ID - REVID (Offset 08)	110
		2.9.5.5	Class Code - CC (Offset 09)	110
		2.9.5.6	Cache Line Size - CLS (Offset 0C)	111
		2.9.5.7	Primary Master Latency Timer - PMLT (Offset 0D)	111
		2.9.5.8	Header Type - HEADTYP (Offset 0E)	111
		2.9.5.9	SHPC 64-Bit Base Address Register - SHPC_BAR (Offset 10)	112
		2.9.5.10	Bus Numbers - BNUM (Offset 18)	112
		2.9.5.11	Secondary Master Latency Timer - SMLT (Offset 1B)	113
		2.9.5.12	I/O Base and Limit - IOBL (Offset 1C)	113
		2.9.5.13	Secondary Status - SSTS (Offset 1E)	114
		2.9.5.14	Memory Base and Limit - MBL (Offset 20)	115
		2.9.5.15	Prefetchable Memory Base and Limit - PMBL (Unset 24)	
		2.9.5.10	Prefetchable Memory Limit Upper 32 Bits - PMBU32 (Offect 20)	
		2.9.5.17	VO Pass and Limit Upper 16 Pite UOPLU16 (Offset 20)	
		2.9.5.10	Capabilities List Pointer - $CCAP$ (Offset 3/)	
		2.9.5.19	Interrupt Information - INTR (Offset 3C)	
		2.9.5.20	Bridge Control - BCTRL (Offset 3E)	118
		29522	Bridge Configuration Register - BCNF (Offset 40)	120
		29523	Multi-Transaction Timer - MTT (Offset 42)	122
		29524	PCI Clock Control - PCI KC (Offset 43)	122
		2.9.5.25	PCI Express Capability Identifier - EXP_CAPID (Offset 44)	
		2.9.5.26	PCI Express Next Item Pointer - EXP_NXTP (Offset 45)	
		2.9.5.27	PCI Express Capability - EXP_CAP (Offset 46)	123
		2.9.5.28	PCI Express Device Capabilities Register - EXP DCAP (Offset 48).	123
		2.9.5.29	PCI Express Device Control Register - EXP_DCTL (Offset 4C)	124
		2.9.5.30	PCI Express Device Status Register - EXP_DSTS (Offset 4E)	125
		2.9.5.31	PCI Express Link Capabilities Register - EXP_LCAP (Offset 50)	125
		2.9.5.32	PCI Express Link Control Register - EXP_LCTL (Offset 54)	126
		2.9.5.33	PCI Express Link Status Register - EXP_LSTS (Offset 56)	127
		2.9.5.34	MSI Capability Identifier - MSI_CAPID (Offset 5C)	127
		2.9.5.35	Next Item Pointer - MSI_NXTP (Offset 5D)	127
		2.9.5.36	MSI Message Control - MSI_MC (Offset 5E)	128
		2.9.5.37	MSI Message Address - MSI_MA (Offset 60)	128
		2.9.5.38	MSI Message Data - MSI_MD (Offset 68)	128
		2.9.5.39	Power Management Capabilities Identifier - PM_CAPID (Offset 6C)	128
		2.9.5.40	Power Management Next Item Pointer - PM_NXTP (Offset 6D)	129
		2.9.5.41	Power Management Capabilities - PM_CAP (Offset 6E)	129



2.9.5.42	Power Management Control/Status Register -
	PM_CSR (Offset 70)130
2.9.5.43	Power Management Bridge Support Extensions -
	PM BSE (Offset 72)
2.9.5.44	Power Management Data Field - PM_DATA (Offset 73)
2.9.5.45	SHPC Capability Identifier - SHPC CAPID (Offset 78)
2.9.5.46	SHPC Next Item Pointer - SHPC NXTP (Offset 79)
29547	SHPC DWORD Select Register - SHPC DWSEL (Offset 7A) 132
29548	SHPC Status - SHPC, STS (Offset 7B)
29549	SHPC DWORD - SHPC DWORD - (Offset 7C) $133$
2 9 5 50	PCI-X Capabilities Identifier - PX CAPID (Offset D8) 133
2.0.5.50	PCI_X Next Item Pointer - PX_NIXTP (Offset D0)
2.3.5.51	PCLY Secondary Status - PY SSTS (Offset DA)
2.3.5.52	DCLV Bridge Status DV BSTS (Offset DC)
2.9.5.55	PCI-X Druge Status - PX_DSTS (Offset DC)
2.9.5.54	PCI-X Opsilearn Split Transaction Control _ PX_05TC (Offset E0) 155
2.9.5.55	PCI-X Downstream Spill Hansaction Control - PX_DSTC (Oliset E4) 130 Bridge Initialization Degister BINIT (Offect EC)
2.9.5.50	Diluge Initialization Register - Diniti (Unset FC)
2.9.5.57	PCI Express Advanced Error Capability Identifier -
~ ~ ~	EXPAERR_CAPID (Offset 100)139
2.9.5.58	PCI Express Uncorrectable Error Status -
	ERRUNC_STS (Offset 104)140
2.9.5.59	PCI Express Uncorrectable Error Mask -
	ERRUNC_MSK (Offset 108)141
2.9.5.60	PCI Express Uncorrectable Error Severity -
	ERRUNC_SEV (Offset 10C)142
2.9.5.61	PCI Express Correctable Error Status -
	ERRCOR_STS (Offset 110)143
2.9.5.62	PCI Express Correctable Error Mask -
	ERRCOR_MSK (Offset 114)143
2.9.5.63	Advanced Error Control and Capability Register -
	ADVERR_CTL (Offset 118)
2.9.5.64	PCI Express Transaction Header Log -
	HDR LOG (Offset 11C-12B)
2.9.5.65	Uncorrectable PCI/X Status Register -
	PCI-XERRUNC STS (Offset 12C)
2.9.5.66	Uncorrectable PCI/X Error Mask Register -
	CIXERRUNC MSK (Offset 130)
29567	Uncorrectable PCI/X Error Severity Register -
2.0.0.07	PCI-XERRUNC SEV (Offset 134) 149
29568	Lincorrectable PCI/X Error Pointer -
2.5.5.00	PCI-XERRUNC PTR (Offset 138)
20560	Lincorrectable PCI/X Error Transaction Header Log
2.9.5.09	DCLYHDD LOG (Offect 13C-14B)
20570	POI-ANDR_LOG (Oliset 130-14b)
2.9.5.70	DILUTECIADIE FCI/A Data ETIOI LOU -
00574	PCI-XD_LOG (OIISet 140-153)
2.9.5.71	Other PCI/X Error Logs and Control -
00570	PCI-XERRLOGUTL(UIISet 154)
2.9.5.72	Internal Arbiter Control Register -
~ ~ ~	ARB_CNTRL (Offset 16A)
2.9.5.73	Strap Status Register - SSR (Offset 170)
2.9.5.74	Power Budgeting Enhanced Capability Header -
	PWRBGT_CAPID (Offset 300)
2.9.5.75	Power Budgeting Data Select Register -
	PWRBGT_DSEL (Offset 304)156
2.9.5.76	Power Budgeting Data Register -
	PWRBGT_DATA (Offset 308)156

			2.9.5.77	Power Budgeting Capability Register -	4
			<u> </u>	PWRBGI_CAP (Offset 30C)	.157
			2.9.5.78	Power Budgeting Information Registers 23:0 -	4
				PWRBG1_INFO[23:0] (Offset 314-370)	.157
3	Addre	ess Trar	nslation U	Init	.159
	3.1	Overvie	ewwe		.159
	3.2	ATU A	ddress Tra	anslation	.162
		3.2.1	Inbound	Transactions	.164
			3211	Inbound Address Translation	165
			3.2.1.2	Inbound Write Transaction	.168
			3.2.1.3	Inbound Read Transaction	.170
			3.2.1.4	Inbound Configuration Cycle Translation	.173
			3.2.1.5	Discard Timers	.174
		3.2.2	Outboun	d Transactions- Single Address Cycle (SAC)	
			Internal E	Bus Transactions	.175
			3.2.2.1	Outbound Address Translation - Single Address Cycle (SAC)	
				Internal Bus Transactions	.175
			3.2.2.2	Outbound Address Translation Windows- Single Address Cycle (SAG	C)
				Internal Bus Transactions	.176
			3.2.2.3	Direct Addressing Window - Single Address Cycle	
				(SAC) Transactions	.181
			3.2.2.4	Outbound DMA Transactions	.182
		3.2.3	Outboun	d Write Transaction	.183
		3.2.4	Outboun	d Read Transaction	.185
		3.2.5	Outboun	d Configuration Cycle Translation	.186
	3.3	Messa	ging Unit		.187
	3.4	Expans	sion ROM	Translation Unit	.188
	3.5	ATU Q	ueue Arch	itecture	.189
		3.5.1	Inbound	Queues	.189
			3.5.1.1	Inbound Write Queue Structure	.189
			3.5.1.2	Inbound Read Queue Structure	.190
			3.5.1.3	Inbound Delayed Write Queue	.191
			3.5.1.4	Inbound Transaction Queues Command Translation Summary	.191
		3.5.2	Outboun	d Queues	.192
		3.5.3	Transact	ion Ordering	.193
			3.5.3.1	Transaction Ordering Summary	.196
	3.6	Private	Device C	ontrol	.198
		3.6.1	Private T	ype 0 Commands on the Secondary Interface	.198
		3.6.2	Private N	lemory Space	.199
	3.7	ATU Ei	rror Condit	tions	.200
		3.7.1	Address	and Attribute Parity Errors on the PCI Interface	.201
		3.7.2	Data Par	ity Errors on the PCI Interface	.202
			3.7.2.1	Outbound Read Request Data Parity Errors	.203
				3.7.2.1.1 Immediate Data Transfer	.203
				3.7.2.1.2 Split Response Termination	.204
			3.7.2.2	Outbound Write Request Data Parity Errors	.205
				3.7.2.2.1 Outbound Writes that are not MSI	007
				(Message Signaled Interrupts)	.205
			2722	3.1.2.2.2 INST OUTDOUND WITTES	.206
			3.1.Z.J	Inbound Configuration Write Completion Measure Date Darity France	.207
			3.1.2.4	moound Configuration write Completion Message Data Pality Enois	201



		3.7.2.5	Inbound Read Request Data Parity Errors	207
			3.7.2.5.1 Immediate Data Transfer	207
			3.7.2.5.2 Split Response Termination	207
		3.7.2.6	Inbound Write Request Data Parity Errors	207
		3.7.2.7	Outbound Read Completion Data Parity Errors	208
		3.7.2.8	Outbound Split Write Data Parity Error Message	209
		3.7.2.9	Inbound Configuration Write Request	210
			3.7.2.9.1 Conventional PCI Mode	210
			3.7.2.9.2 PCI-X Mode	211
		3.7.2.10	Split Completion Messages	212
	3.7.3	Master A	borts on the PCI Interface	213
		3.7.3.1	Master Aborts for Outbound Read or Write Request	213
		3.7.3.2	Inbound Read Completion or Inbound	
			Configuration Write Completion Message	214
		3.7.3.3	Master-Aborts Signaled by the ATU as a Target	214
			3.7.3.3.1 Address Parity Errors	214
			3.7.3.3.2 Internal Bus Master-Abort	214
	3.7.4	Target A	borts on the PCI Interface	215
		3.7.4.1	Target Aborts for Outbound Read Request	
			or Outbound Write Request	215
		3.7.4.2	Inbound Read Completion or Inbound	-
			Configuration Write Completion Message	216
		3.7.4.3	Target-Aborts Signaled by the ATU as a Target	216
			3.7.4.3.1 Internal Bus Master Abort	216
			3.7.4.3.2 Internal Bus Target Abort	216
			3.7.4.3.3 Inbound EROM Memory Write	216
	3.7.5	Corrupte	d or Unexpected Split Completions	217
		3.7.5.1	Completer Address	217
		3.7.5.2	Completer Attributes	217
	3.7.6	SERR# /	Assertion and Detection	218
	377	Internal I	Bus Error Conditions	219
	0	3771	Master Abort on the Internal Rus	219
		0.7.7.1	3 7 7 1 1 Inhound Write Request	219
			3.7.7.1.2 Inbound Read Request	220
		3772	Target Abort on the Internal Rus	221
		0.7.7.2	37721 Conventional Mode	221
			37722 PCI-X Mode	221
	378		or Summary	222
38	Mossar		ad Interrunte	226
2.0	Vital Dr			220
3.9			id	
	3.9.1	Configur	Ing Vital Product Data Operation	
	3.9.2	Accessir	ng Vital Product Data	228
		3.9.2.1	Reading Vital Product Data	228
		3.9.2.2	Writing Vital Product Data	229
3.10	Registe	er Definitio	ons	230
	3.10.1	ATU Ver	ndor ID Register - ATUVID	238
	3.10.2	ATU Dev	/ice ID Register - ATUDID	239
	3.10.3	ATU Cor	nmand Register - ATUCMD	240
	3.10.4	ATU Sta	tus Register - ATUSR	241
	3.10.5	ATU Rev	vision ID Register - ATURID	
	3 10 6	ATU Cla	ss Code Register - ATUCCR	244
	3 10 7		shalina Siza Radistar - ATLICI SR	244 215
	3.10.7	AIUUat	Melline Size Reyister - ATUGLOR	

3.10.8 ATU Latency Timer Register - ATULT	.246
3.10.9 ATU Header Type Register - ATUHTR	.247
3.10.10 ATU BIST Register - ATUBISTR	.248
3.10.11 Inbound ATU Base Address Register 0 - IABAR0	.249
3.10.12 Inbound ATU Upper Base Address Register 0 - IAUBAR0	.250
3.10.13 Inbound ATU Base Address Register 1 - IABAR1	251
3 10 14 Inbound ATU Upper Base Address Register 1 - IAUBAR1	252
3 10 15 Inbound ATH Base Address Register 2 - IABAR2	253
3 10 16 Inbound ATLU Inper Base Address Register 2 - IAUBAR2	254
3 10 17 ATLI Subsystem Vendor ID Register - ASVIR	255
3 10 18 ATLI Subsystem ID Pagistar - ASIP	255
2 10 10 Expansion DOM Pass Address Projector EPRAP	250
2.10.19 Expansion ROM base Address Register - ERDAR	201
3.10.20 ATO Capabilities Pointer Register - ATO_Cap_Ptr	.200
3.10.21 Determining block Sizes for base Address Registers	.209
3.10.22 ATU Interrupt Line Register - ATUILR	.261
3.10.23 ATU Interrupt Pin Register - ATUPR	262
3.10.24 ATU Minimum Grant Register - ATUMGNT	.263
3.10.25 ATU Maximum Latency Register - ATUMLAT	.264
3.10.26 Inbound ATU Limit Register 0 - IALR0	.265
3.10.27 Inbound ATU Translate Value Register 0 - IATVR0	.266
3.10.28 Expansion ROM Limit Register - ERLR	.267
3.10.29 Expansion ROM Translate Value Register - ERTVR	.268
3.10.30 Inbound ATU Limit Register 1 - IALR1	.269
3.10.31 Inbound ATU Limit Register 2 - IALR2	.270
3.10.32 Inbound ATU Translate Value Register 2 - IATVR2	.271
3.10.33 Outbound I/O Window Translate Value Register - OIOWTVR	.272
3.10.34 Outbound Memory Window Translate Value	
Register 0 - OMWTVR0	.273
3.10.35 Outbound Upper 32-bit Memory Window	
Translate Value Register 0 - OUMWTVR0	.274
3.10.36 Outbound Memory Window Translate Value	
Register 1 - OMWTVR1	.275
3.10.37 Outbound Upper 32-bit Memory Window	
Translate Value Register 1 - OUMWTVR1	.276
3.10.38 Outbound Upper 32-bit Direct Window	
Translate Value Register - OUDWTVR	.277
3.10.39 PCI Express-to-PCI Bridge Secondary	
A-Segment Bus Number Register - PEBSABNR	.278
3.10.40 PCI Express-to-PCI Bridge Secondary	
B-Segment Bus Number Register - PEBSBBNR	.279
3.10.41 PCI Express-to-PCI Bridge Primary Bus	
Number Register - PEBPBNR	.280
3.10.42 PCI Express-to-PCI Bridge Device	
Number Register - PEBDNUM	.281
3.10.43 ATU Configuration Register - ATUCR	.282
3.10.44 PCI Configuration and Status Register - PCSR	.283
3.10.45 ATU Interrupt Status Register - ATUISR	.287
3.10.46 ATU Interrupt Mask Register - ATUIMR	.289
3.10.47 Inbound ATU Base Address Register 3 - IABAR3	292
3.10.48 Inbound ATU Upper Base Address Register 3 - IAUBAR3	293
3.10.49 Inbound ATU Limit Register 3 - IALR3	.294



	3.10.5	0 Inbound ATU Translate Value Register 3 - IATVR3	295
	3.10.5	1 Outbound Configuration Cycle Address Register - OCCAR	296
	3.10.5	2 Outbound Configuration Cycle Data Register - OCCDR	
	3.10.5	3 VPD Capability Identifier Register - VPD_CAPID	
	3.10.5	4 VPD Next Item Pointer Register - VPD NXTP	
	3.10.5	5 VPD Address Register - VPD AR	
	3.10.5	6 VPD Data Register - VPD DR	
	3.10.5	7 Power Management Capability Identifier	
		Register - PM CAPID	
	3.10.5	8 Power Management Next Item Pointer Register - PM_NXTP	
	3.10.5	9 Power Management Capabilities Register - PM_CAP	
	3.10.6	0 Power Management Control/Status Register - PM CSR	
	3.10.6	1 MSI Capability Registers	
	3.10.6	2 PCI-X Capability Identifier Register - PX_CAPID	
	3.10.6	3 PCI-X Next Item Pointer Register - PX_NXTP	
	3.10.6	4 PCI-X Command Register - PX_CMD	
	3.10.6	5 PCI-X Status Register - PX_SR	
	3 10 6	6 PCI Interrupt Routing Select Register - PIRSR	312
	3 10 6	7 PCI-A Drive Strength Control Register - PADSCR	313
	3 10 6	8 PCI-A Drive Strength Value Register - PADSVR	314
	3 10 6	9 PCI-B Drive Strength Control Register - PBDSCR	315
	3 10 7	0 PCI-B Drive Strength Value Register - PBDSVR	316
	0.10.7		
Mes	saging L	Jnit	317
4.1	Overvi	iew	
4.2	Theory	y of Operation	
	4.2.1	Transaction Ordering	
4.3	Messa	age Registers	
	4.3.1	Outbound Messages	
	4.3.2	Inbound Messages	
4.4	Doorb	ell Registers	
	4.4.1	Outbound Doorbells	
	4.4.2	Inbound Doorbells	
4.5	Circula	ar Queues	
-	4.5.1	Inbound Free Queue	
	4.5.2	Inbound Post Queue	
	4.5.3	Outbound Post Queue	
	454	Outbound Free Queue	331
4.6	Index	Registers	
47	Messa	aging Unit Error Conditions	332
4 8	Messa	age-Signaled Interrupts	333
1.0	4 8 1	Level-Triggered Versus Edge-Triggered Interrupts	
49	Regist	ter Definitions	
1.0	491	Inhound Message Register - IMRx	336
	4.0.7	Outbound Message Register - OMRy	337
	402	Inhound Doorhell Register - IDR	
	4.3.3 4 0 1	Inbound Interrunt Status Register - IISP	
	7.3.4 105	Inbound Interrupt Otatus Negister - IION	2/0
	4.9.0	Authound Doorball Pagister - ADP	
	4.9.0	Outbound Interrunt Status Pagistar OISP	
	4.9.7	Outound Interrupt Status Register - OISR	

4

		4.9.8	Outbound	Interrupt Mask Register - OIMR	343
		4.9.9	MU Confi	guration Register - MUCR	344
		4.9.10	Queue Ba	ase Address Register - QBAR	345
		4.9.11	Inbound I	Free Head Pointer Register - IFHPR	346
		4.9.12	Inbound I	Free Tail Pointer Register - IFTPR	347
		4.9.13	Inbound I	Post Head Pointer Register - IPHPR	348
		4.9.14	Inbound I	Post Tail Pointer Register - IPTPR	349
		4.9.15	Outbound	Free Head Pointer Register - OFHPR	350
		4.9.16	Outbound	Free Tail Pointer Register - OFTPR	351
		4.9.17	Outbound	Post Head Pointer Register - OPHPR	352
		4.9.18	Outbound	Post Tail Pointer Register - OPTPR	353
		4.9.19	Index Ad	dress Register - IAR	354
		4.9.20	MSI Capa	ability Identifier Register - MSI CAPID	355
		4.9.21	MSI Next	Item Pointer Register - MSI NXTP	356
		4.9.22	MSI Mes	sage Control Register - MSI MCR	357
		4.9.23	MSI Mes	sage Address Register - MSI MAR	358
		4.9.24	MSI Mes	sage Upper Address Register - MSI_MUAR	359
		4.9.25	MSI Mes	sage Data Register- MSI MDR	360
	4.10	Power/I	Default Sta	atus	360
-	1.4.13		0 D		004
5	Intel )	(Scale <sup>°</sup>	Core Bus	interface Unit	361
	5.1	Overvie	WW		361
	5.2	Theory	of Operati	on	362
		5.2.1	Functiona	al Blocks	362
			5.2.1.1	Core Processor Port Address Decode	363
			5.2.1.2	Transaction Queues	363
	5.3	Addres	sing		364
		5.3.1	Bus Widt	٦	364
	5.4	Internal	Bus Com	mands	365
	5.5	Feature	S		365
		5.5.1	Multi-Trai	nsaction Timer	365
		5.5.2	ATU Acce	esses	365
	5.6	Interrup	ots and Err	or Conditions	366
		5.6.1	Internal E	us Errors	366
			5.6.1.1	Internal Bus Master-Abort	366
			5.6.1.2	Internal Bus Target-Abort	366
			5.6.1.3	PCI Master-Abort	367
			5.6.1.4	PUI Target-Abort	367
		560	0.0.1.0		307
		5.0.Z		J EIIOIS	300
			5.6.2.1	Involid Addross Region Error	200
	57	Pocot o	0.0.2.2	invalid Address Region Endi	260
	5.7	574	Posot		260
		570		~	260
	БŶ	D.1.Z		UII	270
	0.0			ið	370
		0.0.1 5.0.1		Addross Dogistor DEAD	3/1 272
		0.0.Z		AUUIESS REYISTEI - DEAR	313 171
		J.O.J	BIO CONT	IUI REYISIEI - DIUCR	374



6.1     Overview.     375       6.2     Theory of Operation.     377       6.3     DMA Transfer     379       6.3.1     Chain Descriptors     380       6.3.2     Chaining DMA Descriptors     381       6.3.3     Initiating DMA Transfers     382       6.3.4     Scatter Gather DMA Transfers     382       6.3.4     Scatter Gather DMA Transfers     383       6.3.5     Synchronizing a Program to Chained Transfers     386       6.4.3     Local Memory to Chained Transfers     386       6.4.1     PCI-to-Local Memory Transfers     386       6.4.2     Local Memory to Local Memory DMA     387       6.4.4     Exclusive Access     387       6.5     Data Queues     388       6.6.1     6.4.1     Holinged Data Transfers     388       6.6.2     6.4.2     Gather Unaligned Data Transfers     388       6.6.2     6.4.3     Local Memory to PCI Transfers     389       6.7     CRC Generation     390     6.7.1     CRC Mode Configuration and Operation     390       6.7.1     CRC Mode Configuration and Operation     390     391<	6	DMA Controller Unit							
6.2     Theory of Operation		6.1	Overview						
6.3     DMA Transfer		6.2	Theory of Operation						
6.3.1     Chaining DMA Descriptors		6.3	DMA Transfer						
6.3.2     Chaining DMA Descriptors			6.3.1 Chain Descriptors	380					
6.3.4       Scatter Gather DMA Transfers       382         6.3.4       Scatter Gather DMA Transfers       383         6.3.5       Synchronizing a Program to Chained Transfers       384         6.3.6       Appending to The End of a Chain       385         6.4       Data Transfers       386         6.4.1       PCI-to-Local Memory to PCI Transfers       386         6.4.2       Local Memory to Docal Memory DMA       387         6.4.4       Exclusive Access       387         6.5       Data Alignment       388         6.6.1       64-bit Unaligned Data Transfers       388         6.6.2       64/32-bit Unaligned Data Transfers       388         6.6.2       64/32-bit Unaligned Data Transfers       389         6.7.1       CRC Generation       390         6.7.2       CRC 32C Algorithm       390         6.8       Channel Programming Examples       393         6.10.1       Software Start DMA Controller Initialization       393         6.10.2       Software Start DMA Controller Initialization       393         6.12.1       Poter Acdress Register x - CRX       396         6.12.1       Poter I Erors			6.3.2 Chaining DMA Descriptors						
6.3.4       Scatter Gather DMA Transfers       383         6.3.5       Synchronizing a Program to Chained Transfers       384         6.3.6       Appending to The End of a Chain.       385         6.4       Data Transfers.       386         6.4.1       PCI-to-Local Memory to PCI Transfers.       386         6.4.1       PCI-to-Local Memory to Local Memory DMA.       387         6.4.2       Local Memory to Local Memory DMA.       387         6.4.4       Exclusive Access       387         6.5       Data Queues.       387         6.6       Data Queues.       388         6.6.2       64/32-bit Unaligned Data Transfers.       388         6.6.2       CRC Generation       390         6.7       CRC Generation       390         6.7       CRC Mode Configuration and Operation       390         6.9       Programming Model State Diagram       392         6.10       DMA Channel Programming Examples <td></td> <td>6.3.3 Initiating DMA Transfers</td> <td> 382</td>			6.3.3 Initiating DMA Transfers	382					
6.3.5       Synchronizing a Program to Chained Transfers.			6.3.4 Scatter Gather DMA Transfers	383					
6.3.6     Appending to The End of a Chain.     385       6.4     Data Transfers.     386       6.4.1     PCI-to-Local Memory Transfers.     386       6.4.2     Local Memory to PCI Transfers.     386       6.4.3     Local Memory to Local Memory DMA.     387       6.5     Data Queues.     387       6.5     Data Queues.     387       6.6     6.4.4     Exclusive Access.     387       6.5     Data Alignment.     388       6.6.1     64-bit Unaligned Data Transfers.     388       6.6.2     64/32-bit Unaligned Data Transfers.     389       6.7     CRC Generation.     390       6.7.1     CRC Mode Configuration and Operation.     390       6.8     Channel Priority.     391       6.9     Programming Model State Diagram.     392       6.10     DMA Channel Programming Examples.     393       6.10.1     Software Suspend Channel.     393       6.10.2     Software Suspend Channel.     394       6.11     Interrupts.     395       6.12.2     Internal Bus Errors.     396       6.12.4     PCI Error			6.3.5 Synchronizing a Program to Chained Transfers						
6.4     Data Transfers.     386       6.4.1     PCI-to-Local Memory Transfers.     386       6.4.2     Local Memory to PCI Transfers.     386       6.4.3     Local Memory to Local Memory DMA.     387       6.4.4     Exclusive Access     387       6.5     Data Queues.     387       6.6     Data Alignment.     388       6.6.1     64.432-bit Unaligned Data Transfers.     388       6.6.2     64/32-bit Unaligned Data Transfers.     389       6.7     CRC Generation     390       6.7.1     CRC Mode Configuration and Operation.     390       6.7.2     CRC-32C Algorithm.     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples.     393       6.10.1     Software Start DMA Transfer     393       6.10.2     Software Suspend Channel.     394       6.11     Interrupts.     396       6.12     Error Conditions.     396       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx.     399       6.14.2     Channel Stat			6.3.6 Appending to The End of a Chain	385					
6.4.1     PCI-to-Local Memory Transfers.     386       6.4.2     Local Memory to PCI Transfers.     386       6.4.3     Local Memory to Local Memory DMA.     387       6.4.4     Exclusive Access.     387       6.5     Data Alignment.     388       6.6.1     64-bit Unaligned Data Transfers.     388       6.6.1     64-bit Unaligned Data Transfers.     388       6.6.2     64/32-bit Unaligned Data Transfers.     388       6.6.2     64/32-bit Unaligned Data Transfers.     388       6.6.2     64/32-bit Unaligned Data Transfers.     388       6.6.1     CRC Generation.     390       6.7.1     CRC Mode Configuration and Operation.     390       6.7.2     CRC-32C Algorithm.     391       6.9     Programming Model State Diagram.     392       6.10     Software Start DMA Transfer.     393       6.10.2     Software Suspend Channel.     394       6.10.2     Software Suspend Channel.     394       6.11     Programming Examples     395       6.12     Error Conditions.     396       6.12.1     PCI Errors.     396       <		6.4	Data Transfers						
6.4.2     Local Memory to Local Memory DMA			6.4.1 PCI-to-Local Memory Transfers	386					
6.4.3     Local Memory to Local Memory DMA			6.4.2 Local Memory to PCI Transfers	386					
6.4.4     Exclusive Access     387       6.5     Data Queues     387       6.6     Data Alignment     388       6.6.1     64-bit Unaligned Data Transfers     388       6.6.2     64/32-bit Unaligned Data Transfers     389       6.7     CRC Generation     390       6.7.1     CRC Mode Configuration and Operation     390       6.7.2     CRC-32C Algorithm     390       6.8     Channel Priority     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Suspend Channel     394       6.11     Interrupts     395       6.12     Fror Conditions     396       6.12.2     Interrupts     396       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Channel Status Register x - CCRx     398       6.14.4     Next Descriptor Address Register x - NDARs     400       6.14.5     PCI Address Register x			6.4.3 Local Memory to Local Memory DMA						
6.5     Data Queues			6.4.4 Exclusive Access						
6.6     Data Alignment     388       6.6.1     64-bit Unaligned Data Transfers     389       6.6.2     64/32-bit Unaligned Data Transfers     389       6.7     CRC Generation     390       6.7.1     CRC Mode Configuration and Operation     390       6.7.2     CRC-32C Algorithm     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.1     Software Start DMA Transfer     393       6.10.2     Software Suspend Channel     394       6.11     Interrupts     395       6.12     Error Conditions     396       6.12.2     Interrupts     396       6.13     Power-up/Default Status     398       6.14     Channel Control Register x - CCRx     399       6.14.2     Channel Status Register x - CCRx     399       6.14.3     Descriptor Address Register x - NDARs     400       6.14.4     Next Descriptor Address Register x - NDARs     401       6.14.5     PCI Address Register x - NDARs     402 <t< td=""><td></td><td>6.5</td><td>Data Queues</td><td></td></t<>		6.5	Data Queues						
6.6.1     64-bit Unaligned Data Transfers.     388       6.6.2     64/32-bit Unaligned Data Transfers.     389       6.7     CRC Generation     390       6.7.1     CRC Mode Configuration and Operation     390       6.7.2     CRC-32C Algorithm     390       6.8     Channel Priority     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Start DMA Transfer     393       6.10.3     Software Suspend Channel     394       6.11     Interrupts     395       6.12     Error Conditions     396       6.12.1     PCI Errors     396       6.12.1     PCI Errors     396       6.12.1     PCI Errors     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Channel Control Register x - CCRx     399       6.14.3     Descriptor Address Register x - DARx     400       6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5		6.6	Data Alignment						
6.6.2     64/32-bit Unaligned Data Transfers.     389       6.7     CRC Generation     390       6.7.1     CRC Mode Configuration and Operation     390       6.7.2     CRC-32C Algorithm     390       6.8     Channel Priority     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Suspend Channel     394       6.11     Interrupts.     395       6.12     Error Conditions.     396       6.12.1     PCI Errors.     396       6.12.2     Internupt.     397       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Channel Status Register x - CRx     400       6.14.3     Descriptor Address Register x - NDARs     401       6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - PADRx     403       6.14.6     PCI Upper Address Register x - DCRx     406			6.6.1 64-bit Unaligned Data Transfers	388					
6.7     CRC Generation     390       6.7.1     CRC Mode Configuration and Operation     390       6.7.2     CRC-32C Algorithm     390       6.8     Channel Priority     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Start DMA Transfer     393       6.10.3     Software Suspend Channel     394       6.11     Interrupts     396       6.12.2     Error Conditions     396       6.12.2     Internal Bus Errors     397       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Channel Status Register x - DARx     400       6.14.3     Descriptor Address Register x - DARx     400       6.14.4     Next pescriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - DARx     403       6.14.6     PCI Upper Address Register x - DARx     404       6.14.7     Local Address Register x - DARx     404			6.6.2 64/32-bit Unaligned Data Transfers						
6.7.1     CRC Mode Configuration and Operation     390       6.7.2     CRC-32C Algorithm     390       6.8     Channel Priority     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Suspend Channel     394       6.11     Interrupts     395       6.10.3     Software Suspend Channel     394       6.11     Interrupts     395       6.12     Error Conditions     396       6.12.1     PCI Errors     396       6.12.2     Internul Bus Errors     397       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     398       6.14.1     Channel Status Register x - CRx     400       6.14.2     Channel Status Register x - DARx     400       6.14.3     Descriptor Address Register x - NDARs     402       6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - PADRx     403       6.		6.7	CRC Generation						
6.7.2     CRC-32C Algorithm     390       6.8     Channel Priority     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Start DMA Transfer     393       6.10.3     Software Suspend Channel     394       6.11     Interrupts     395       6.12.1     PCI Error Conditions     396       6.12.2     Internal Bus Errors     397       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Internel Status Register x - CCRx     399       6.14.3     Descriptor Address Register x - DARx     400       6.14.3     Descriptor Address Register x - NDARs     402       6.14.4     Next Descriptor Address Register x - NDARs     403       6.14.5     PCI Address Register x - PADRx     404       6.14.6     PCI Upper Address Register x - DCRx     404       6.14.7     Local Address Register x - DCRx     405       6.14.8     Byte Count Register x - DCRx			6.7.1 CRC Mode Configuration and Operation						
6.8     Channel Priority     391       6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Start DMA Transfer     393       6.10.3     Software Suspend Channel     394       6.11     Interrupts     395       6.12.2     Software Suspend Channel     396       6.12.1     PCI Errors     396       6.12.2     Internal Bus Errors     396       6.12.1     PCI Errors     396       6.12.2     Internal Bus Errors     397       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Channel Status Register x - CCRx     400       6.14.3     Descriptor Address Register x - NDARs     402       6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - NDARs     403       6.14.6     PCI Upper Address Register x - NDARs     404       6.14.7     Local Address Register x - DCRx     404			6.7.2 CRC-32C Algorithm	390					
6.9     Programming Model State Diagram     392       6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Start DMA Transfer     393       6.10.3     Software Suspend Channel     394       6.11     Interrupts     395       6.12     Error Conditions     396       6.12.1     PCI Errors     396       6.12.1     PCI Errors     396       6.12.1     Internal Bus Errors     397       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     398       6.14.1     Channel Status Register x - CCRx     399       6.14.2     Channel Status Register x - CCRx     399       6.14.2     Channel Status Register x - DARx     400       6.14.3     Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - PADRx     403       6.14.6     PCI Upper Address Register x - PADRx     404       6.14.7     Local Address Register x - DCRx     404       6.14.8     Byte Count Register x - DCRx     406		6.8	Channel Priority						
6.10     DMA Channel Programming Examples     393       6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Start DMA Transfer     393       6.10.3     Software Suspend Channel     394       6.11     Interrupts     395       6.12     Error Conditions     396       6.12.1     PCI Errors     396       6.12.2     Internupts     397       6.13     Power-up/Default Status     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Channel Control Register x - CCRx     399       6.14.3     Descriptor Address Register x - DARx     400       6.14.3     Descriptor Address Register x - NDARs     402       6.14.4     Next Descriptor Address Register x - PADRx     403       6.14.5     PCI Address Register x - PADRx     404       6.14.6     PCI Upper Address Register x - PADRx     404       6.14.7     Local Address Register x - DCRx     404       6.14.8     Byte Count Register x - DCRx     406       6.14.9     Descriptor Control Register x - DCRx     406       6.14.9     Descriptor Control Regist		6.9	Programming Model State Diagram						
6.10.1     Software DMA Controller Initialization     393       6.10.2     Software Start DMA Transfer.     393       6.10.3     Software Suspend Channel     394       6.11     Interrupts.     395       6.12     Error Conditions.     396       6.12.1     PCI Errors.     396       6.12.2     Internal Bus Errors.     397       6.13     Power-up/Default Status.     398       6.14.1     Channel Control Register x - CCRx.     399       6.14.2     Channel Status Register x - CSRx     400       6.14.3     Descriptor Address Register x - NDARs     401       6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - PADRx     403       6.14.6     PCI Upper Address Register x - PADRx     404       6.14.7     Local Address Register x - DCRx     405       6.14.8     Byte Count Register x - DCRx     406       6.14.9     Descriptor Control Register x - DCRx     406       6.14.9     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview.     409 <td></td> <td rowspan="3">6.10</td> <td colspan="5">DMA Channel Programming Examples</td>		6.10	DMA Channel Programming Examples						
6.10.2     Software Start DMA Transfer			6.10.1 Software DMA Controller Initialization						
6.10.3     Software Suspend Channel			6.10.2 Software Start DMA Transfer	393					
6.11     Interrupts			6.10.3 Software Suspend Channel	394					
6.12     Error Conditions		6.11	Interrupts	395					
6.12.1     PCI Errors     396       6.12.2     Internal Bus Errors     397       6.13     Power-up/Default Status     398       6.14     Register Definitions     398       6.14.1     Channel Control Register x - CCRx     399       6.14.2     Channel Control Register x - CCRx     399       6.14.2     Channel Status Register x - CCRx     400       6.14.3     Descriptor Address Register x - DARx     401       6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - PADRx     403       6.14.6     PCI Upper Address Register x - PUADRx     404       6.14.7     Local Address Register x - BCRx     405       6.14.8     Byte Count Register x - BCRx     406       6.14.9     Descriptor Control Register x - DCRx     407       6.14.9.1     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview     409       7.2     Theory of Operation     410       7.3     Hardware-Assist XOR Unit     412       7.3.1     Data Transfer     412 <td></td> <td>6.12</td> <td>Error Conditions</td> <td> 396</td>		6.12	Error Conditions	396					
6.12.2     Internal Bus Errors.     397       6.13     Power-up/Default Status.     398       6.14     Register Definitions.     398       6.14     Channel Control Register x - CCRx.     399       6.14.1     Channel Control Register x - CCRx.     399       6.14.2     Channel Status Register x - CCRx.     400       6.14.3     Descriptor Address Register x - DARx     401       6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - PADRx     403       6.14.6     PCI Upper Address Register x - PUADRx     404       6.14.7     Local Address Register x - PADRx     403       6.14.8     Byte Count Register x - BCRx     406       6.14.9     Descriptor Control Register x - DCRx     407       6.14.9.1     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview.     409       7.2     Theory of Operation.     410       7.3     Hardware-Assist XOR Unit.     412       7.3.1     Data Transfer     412			6.12.1 PCI Errors	396					
6.13Power-up/Default Status3986.14Register Definitions3986.14.1Channel Control Register x - CCRx3996.14.2Channel Status Register x - CSRx4006.14.3Descriptor Address Register x - DARx4016.14.4Next Descriptor Address Register x - NDARs4026.14.5PCI Address Register x - PADRx4036.14.6PCI Upper Address Register x - PUADRx4046.14.7Local Address Register x - PUADRx4056.14.8Byte Count Register x - BCRx4066.14.9Descriptor Control Register x - DCRx4076.14.9.1PCI Transactions Support4087Application Accelerator Unit4097.1Overview4097.2Theory of Operation4107.3Hardware-Assist XOR Unit4127.3.1Data Transfer412			6.12.2 Internal Bus Errors	397					
6.14     Register Definitions		6.13	Power-up/Default Status						
6.14.1Channel Control Register x - CCRx3996.14.2Channel Status Register x - CSRx4006.14.3Descriptor Address Register x - DARx4016.14.4Next Descriptor Address Register x - NDARs4026.14.5PCI Address Register x - PADRx4036.14.6PCI Upper Address Register x - PUADRx4046.14.7Local Address Register x - PUADRx4056.14.8Byte Count Register x - BCRx4066.14.9Descriptor Control Register x - DCRx4076.14.9.1PCI Transactions Support4087Application Accelerator Unit4097.1Overview4097.2Theory of Operation4107.3Hardware-Assist XOR Unit4127.3.1Data Transfer412		6.14	Register Definitions						
6.14.2Channel Status Register x - CSRx4006.14.3Descriptor Address Register x - DARx4016.14.4Next Descriptor Address Register x - NDARs4026.14.5PCI Address Register x - PADRx4036.14.6PCI Upper Address Register x - PUADRx4046.14.7Local Address Register x - PUADRx4056.14.8Byte Count Register x - BCRx4066.14.9Descriptor Control Register x - DCRx4076.14.9.1PCI Transactions Support4087Application Accelerator Unit4097.1Overview4097.2Theory of Operation4107.3Hardware-Assist XOR Unit4127.3.1Data Transfer412			6.14.1 Channel Control Register x - CCRx	399					
6.14.3 Descriptor Address Register x - DARx4016.14.4 Next Descriptor Address Register x - NDARs4026.14.5 PCI Address Register x - PADRx4036.14.6 PCI Upper Address Register x - PUADRx4046.14.7 Local Address Register x - LADRx4056.14.8 Byte Count Register x - BCRx4066.14.9 Descriptor Control Register x - DCRx4076.14.9.1 PCI Transactions Support4087 Application Accelerator Unit4097.1 Overview4097.2 Theory of Operation4107.3 Hardware-Assist XOR Unit4127.3.1 Data Transfer412			6.14.2 Channel Status Register x - CSRx	400					
6.14.4     Next Descriptor Address Register x - NDARs     402       6.14.5     PCI Address Register x - PADRx     403       6.14.6     PCI Upper Address Register x - PUADRx     404       6.14.7     Local Address Register x - LADRx     405       6.14.8     Byte Count Register x - BCRx     406       6.14.9     Descriptor Control Register x - DCRx     407       6.14.9.1     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview     409       7.2     Theory of Operation     410       7.3     Hardware-Assist XOR Unit     412       7.3.1     Data Transfer     412			6.14.3 Descriptor Address Register x - DARx	401					
6.14.5     PCI Address Register x - PADRx     403       6.14.6     PCI Upper Address Register x - PUADRx     404       6.14.7     Local Address Register x - LADRx     405       6.14.8     Byte Count Register x - BCRx     406       6.14.9     Descriptor Control Register x - DCRx     407       6.14.9.1     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview     409       7.2     Theory of Operation     410       7.3     Hardware-Assist XOR Unit     412       7.3.1     Data Transfer     412			6.14.4 Next Descriptor Address Register x - NDARs	402					
6.14.6     PCI Upper Address Register x - PUADRx     404       6.14.7     Local Address Register x - LADRx     405       6.14.8     Byte Count Register x - BCRx     406       6.14.9     Descriptor Control Register x - DCRx     407       6.14.9.1     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview     409       7.2     Theory of Operation     410       7.3     Hardware-Assist XOR Unit     412       7.3.1     Data Transfer     412			6.14.5 PCI Address Register x - PADRx	403					
6.14.7     Local Address Register x - LADRx     405       6.14.8     Byte Count Register x - BCRx     406       6.14.9     Descriptor Control Register x - DCRx     407       6.14.9.1     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview     409       7.2     Theory of Operation     410       7.3     Hardware-Assist XOR Unit     412       7.3.1     Data Transfer     412			6.14.6 PCI Upper Address Register x - PUADRx	404					
6.14.8     Byte Count Register x - BCRx     406       6.14.9     Descriptor Control Register x - DCRx     407       6.14.9.1     PCI Transactions Support     408       7     Application Accelerator Unit     409       7.1     Overview     409       7.2     Theory of Operation     410       7.3     Hardware-Assist XOR Unit     412       7.3.1     Data Transfer     412			6.14.7 Local Address Register x - LADRx	405					
6.14.9     Descriptor Control Register x - DCRx			6.14.8 Byte Count Register x - BCRx	406					
6.14.9.1 PCI Transactions Support			6.14.9 Descriptor Control Register x - DCRx	407					
7Application Accelerator Unit4097.1Overview4097.2Theory of Operation4107.3Hardware-Assist XOR Unit4127.3.1Data Transfer412			6.14.9.1 PCI Transactions Support	408					
7.1     Overview	7	Appli	ication Accelerator Unit	409					
7.2Theory of Operation		7.1	Overview	409					
7.3 Hardware-Assist XOR Unit		7.2	Theory of Operation						
7.3.1 Data Transfer		7.3	Hardware-Assist XOR Unit	412					
			7.3.1 Data Transfer	412					

March 2005

	7.3.2	Chain D	escriptors	
		7.3.2.1	Principle / Four-Source Descriptor Format	
		7.3.2.2	Eight-Source Descriptor Format	
		7.3.2.3	Sixteen-Source Descriptor Format	
		7.3.2.4	Thirty-two-Source Descriptor Format	
		7.3.2.5	Dual-XOR-Transfer Descriptor Format	
		7.3.2.6	P+Q Three-Source Descriptor Format	
		7.3.2.1	P+Q SIX-Source Descriptor Format	
		7320	P+Q Twelve-Source Descriptor Format	
	733	Descript	or Summany	/21
	734	Descript	or Chaining	
71		Descriptor Dr		
7.4		Soottor (	Cothor Transforc	
	7.4.1	Sunchro	pizing a Program to Chained Operation	
	7.4.2	Appondi	ng to The End of a Chain	
75	7.4.3 A A On	Appendi		
7.5	7 5 1		accina	
	7.5.1		essing	
	7.5.2		Proveration with P(O PAID & Mode)	
	7.5.5			
	7.5.4		sult Buffer Check	
	7.5.5	Zero Re	sult Buffer Check	
	7.5.0	Momory	Block Fill Operation	
76	7.5.7 Drogra	mming M	adal Stata Diagram	
7.0	Applier	tion Acco	Norator Drightly	
7.0	Applica			
7.0		GA bit Ur	palianad Data Transfors	
70	7.0.1 Drogra	04-Dit Oi mming th	Application Accelerator	
7.9	7 0 1	Applicati	ion Accolorator Initialization	
	7.9.1	Succond	ding and Recuming the Application Accelerator	
	7.9.2	Appondi	and and Resulting the Application Accelerator	
	7.9.3	Appendi	ng Descriptor for Dual XOP Operations	
	7.9.4	Appendi	ng Descriptor for Memory Block Fill Operations	
	7.9.5	Appendi	ng Descriptor for Zero Pocult Puffor Check	
7 10	7.9.0	Appendi		
7.10	Error C	onditions		
7.11	Dowor.	un/Dofaul	It Status	
7.12	Pogist	or Dofinitio		
1.15	7 12 1		otor Control Pagistor - ACP	
	7 12 2	Accelera	ator Status Pagister - ASR	
	7 12 2	Accelera	ator Descriptor Address	
	1.15.5	Register		468
	7 13 /	Accelera	ator Next Descriptor Address	
	7.10.4	Register	· - ANDAR	469
	7.13.5	Data / Se	ource Address Register1 - D/SAR1/POSAR1	
	7.13.6	Source A	Address Registers 232 - SAR232	
	7.13.7	P+Q RA	ID-6 Source Address Registers 2.16 -	
		PQSAR	2.16	
	7.13.8	P+Q RA	ID-6 Galois Field Multiplier Registers 15	
		- GFMR	15474	



		7.13.9	Destinati	on Address Register - DAR	476
		7.13.10	0 Accelera	tor Byte Count Register - ABCR	477
		7.13.1	1 Accelera	tor Descriptor Control Register - ADCR	478
		7.13.12	2 Extended	d Descriptor Control Register 0 - EDCR0	482
		7 13 1	3 Extended	d Descriptor Control Register 1 - EDCR1	484
		7 13 1	4 Extended	d Descriptor Control Register 2 - EDCP2	486
		7.13.14			
8	Mem	ory Con	troller		489
	8.1	Overvi	ew		489
	8.2	Glossa	ary		490
	8.3	Theory	of Operat	tion	491
		8.3.1	Function	al Blocks	491
			8311	Transaction Ports	492
			0.0111	8.3.1.1.1 Core Processor Port	492
				8 3 1 1 2 Internal Bus Port	492
			8312	Address Decode Blocks	493
			0101112	8 3 1 2 1 DDR SDRAM Memory Space	493
				8.3.1.2.2 Memory-Manned Register Space	493
				8 3 1 2 3 Core Processor Port Address Decode	493
				8.3.1.2.4 Internal Rus Port Address Decode	493
			8.3.1.3	Memory Transaction Queues.	494
				8.3.1.3.1 Core Processor Memory Transaction Queue (CMTQ)	494
				8.3.1.3.2 Internal Bus Memory Transaction Queue (IBMTQ)	494
			8.3.1.4	Configuration Registers	
			8.3.1.5	Refresh Counter	
			8.3.1.6	Memory Controller Arbiter (MARB)	
			8.3.1.7	DDR SDRAM Control Block	496
				8.3.1.7.1 Page Control Block	496
				8.3.1.7.2 DDR SDRAM State Machine and Pipeline Queues	496
				8.3.1.7.3 Error Correction Logic	496
		8.3.2	MCU Art	pitration and Configuration	497
			8.3.2.1	MCU Port Priority	497
			8.3.2.2	MCU Port Transaction Count	498
			8.3.2.3	Core Processor Port Preemption	498
			8.3.2.4	Core Port Transaction Ordering	
			8.3.2.5	IB Port Ordering	
			8.3.2.6	MCU Port Coherency	
		8.3.3	DDR SD	RAM Memory Support	500
			8331	DDR SDRAM Interface	500
			8332	DDR SDRAM Addressing	504
			8333	DDR SDRAM Configuration	506
			8334	32-bit Data Bus Width	510
			8.3.3.5	Page Hit/Miss Determination	
			8336	On DIMM Termination	512
			8.3.3.7	DDR SDRAM Commands	
			8.3.3.8	DDR SDRAM Initialization	
			8.3.3.9	DDR SDRAM Mode Programming	
			8.3.3.10	Intel <sup>®</sup> 80333 I/O ProcessorDDR SDRAM Read Cvcle	
			8.3.3.11	DDR SDRAM Write Cycle	
			8.3.3.12	DDR SDRAM Refresh Cycle	
		8.3.4	Error Co	rrection and Detection	
		0.011	8341	ECC Generation	526
			8342	ECC Generation for Partial Writes	
			0.0.1.2		

		8.3.4.3	ECC Checking	530
		8.3.4.4	Scrubbing	534
			8.3.4.4.1 ECC Example Using the H-Matrix	534
		8.3.4.5	ECC Disabled	535
		8.3.4.6	ECC lesting	535
	8.3.5	Overlapp	bing Memory Regions	536
	8.3.6	DDR SD	RAM Clocking	536
8.4	Power	Failure Mo	ode	537
	8.4.1	Theory o	f Operation	537
	8.4.2	Power Fa	ailure Sequence	538
		8.4.2.1	Power Failure Impact on the System	538
		8.4.2.2	System Assumptions	538
	8.4.3	Memory	Controller Response to PWRGD	539
		8.4.3.1	External Logic Required for Power Failure	541
			8.4.3.1.1 Assertion of PWRGD During Power Failure	541
			8.4.3.1.2 Distinguishing Between a Power Up and	
			a Power Failure <b>PWRGD</b> Assertion	541
		8.4.3.2	P_RST# Usage Versus I_RST#	541
8.5	Interrup	ots/Error C	Conditions	542
	8.5.1	Single-Bi	it Error Detection	543
	8.5.2	Multi-bit	Error Detection	544
8.6	Reset (	Conditions	5	544
8.7	Registe	er Definitio	ons	545
	8.7.1	SDRAM	Initialization Register - SDIR	546
	8.7.2	SDRAM	Control Register 0 - SDCR0	547
	873	SDRAM	Control Register 1 - SDCR1	549
	874	SDRAM	Base Register - SDBR	551
	875	SDRAM	Boundary Register 0 - SBR0	552
	876	SDRAM	Boundary Register 1 - SBR1	553
	0.7.0 8.7.7		32-bit Pogion Sizo Pogistor - S32SP	554
	0.7.7		SZ-DIL REGION SIZE REGISLEL - SSZSR	
	0.7.0		niloi Register - ECCR	555
	0.7.9		J REGISIEIS - ELOGU, ELOGI	550
	8.7.10		Dress Registers - ECARU, ECART	557
	8.7.11	ECC Tes		558
	8.7.12	Memory	Controller Interrupt Status Register - MCISR	559
	8.7.13	MCU Po	rt Transaction Count Register - MPTCR	560
	8.7.14	MCU Pre	eemption Control Register - MPCR	561
	8.7.15	Frequence	cy Register - RFR	562
	8.7.16	DCAL Co	ontrol and Status Register - DCALCSR	563
	8.7.17	DCAL Ac	ddress Register - DCALADDR	565
	8.7.18	DCAL Da	ata Registers 17:0 - DCALDATA[17:0]	566
		8.7.18.1	Opcode: EMRS OCD Adjust/Drive Commands	567
		8.7.18.2	Opcode: Receive Enable Calibration	568
		8.7.18.3	Opcode: DQS Calibration	570
	8.7.19	Receive	Enable Delay Register - RCVDLY	571
	8.7.20	Slave Lo	w Mix 0 - SLVLMIX0	572
	8.7.21	Slave Lo	w Mix 1 - SLVLMIX1	573
	8.7.22	Slave Hig	gh Mix 0 - SLVHMIX0	574
	8.7.23	Slave Hig	gh Mix 1 - SLVHMIX1	575
	8.7.24	Slave Le	ngth - SLVLEN	576
	8.7.25	Master M	/ix - MASTMIX	577



		8.7.26 Master Length - MASTLEN	578
		8.7.27 DDR Drive Strength Status Register - DDRDSSR	579
		8.7.28 DDR Drive Strength Control Register - DDRDSCR	580
		8.7.29 DDR Miscellaneous Pad Control Register - DDRMPCR	581
9	Perip	oheral Bus Interface Unit	
	9.1	Overview	584
	9.1	Perinheral Rus Signals	585
	5.2	9.2.1 Address/Data Signal Definitions	585
		9.2.2 Control/Status Signal Definitions	585
		9.2.3 Bus Width	586
		9.2.4 Detailed Signal Descriptions	587
		925 Flash Memory Support	588
		9251 Flash Read Cycle	589
		9.2.5.2 Flash Write Cycle	
	9.3	Intel XScale <sup>®</sup> Core PCI Memory Boot Support	
	9.4	Register Definitions	
	••••	9.4.1 PBI Control Register - PBCR	
		9.4.2 Determining Block Sizes for Memory Windows	
		9.4.3 PBI Base Address Register 0 - PBBAR0	
		9.4.4 PBI Limit Register 0 - PBLR0	
		9.4.5 PBI Base Address Register 1 - PBBAR1	
		9.4.6 PBI Limit Register 1 - PBLR1	
		9.4.7 PBI Memory-less Boot Registers 2:0 - PMBR[2:0]	
		9.4.8 PBI Drive Strength Control Register - PBDSCR	600
10	I <sup>2</sup> C B	Bus Interface Units	
10	I <sup>2</sup> C B	Bus Interface Units	
10	<b>I<sup>2</sup>C B</b> 10.1	Bus Interface Units Overview	601 601
10	<b>I<sup>2</sup>C B</b> 10.1 10.2	Bus Interface Units Overview I <sup>2</sup> C Interface	601 601 601
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units. Overview I <sup>2</sup> C Interface Theory of Operation 10.3.1. Operational Blocks	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units. Overview I <sup>2</sup> C Interface Theory of Operation 10.3.1 Operational Blocks	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.3 1	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units       Overview       I <sup>2</sup> C Interface       Theory of Operation       10.3.1 Operational Blocks       10.3.2 I <sup>2</sup> C Bus Interface Modes       10.3.3 Start and Stop Bus States       10.3.4 START Condition       10.3.5 No START or STOP Condition	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.3.1 START Condition.       10.3.3.2 No START or STOP Condition .       10.3.3 STOP Condition	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.3.1 START Condition.       10.3.3.2 No START or STOP Condition .       10.3.3 STOP Condition.	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States .       10.3.1 START Condition.       10.3.2 No START or STOP Condition .       10.3.3 STOP Condition.       10.3.4 Serial Clock Line (SCL) Generation	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.3.1 START Condition.       10.3.3.2 No START or STOP Condition .       10.3.3.3 STOP Condition.       10.4.1 Serial Clock Line (SCL) Generation .       10.4.2 Data and Addressing Management .	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.3.1 START Condition.       10.3.3.2 No START or STOP Condition .       10.3.3.3 STOP Condition.       10.4.1 Serial Clock Line (SCL) Generation .       10.4.2 Data and Addressing Management .       10.4.2.1 Addressing a Slave Device .	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units       Overview       I <sup>2</sup> C Interface       Theory of Operation       10.3.1 Operational Blocks       10.3.2 I <sup>2</sup> C Bus Interface Modes       10.3.3 Start and Stop Bus States       10.3.3.1 START Condition       10.3.3.2 No START or STOP Condition       10.3.3.3 STOP Condition       10.4.1 Serial Clock Line (SCL) Generation       10.4.2 Data and Addressing Management       10.4.3 I <sup>2</sup> C Acknowledge       10.4.4 Arbitration       10.4.4.1 SCL Arbitration       10.4.4.2 SDA Arbitration	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units       Overview       I <sup>2</sup> C Interface       Theory of Operation       10.3.1 Operational Blocks       10.3.2 I <sup>2</sup> C Bus Interface Modes       10.3.3 Start and Stop Bus States       10.3.3.1 START Condition       10.3.3.2 No START or STOP Condition       10.3.3.3 STOP Condition       10.4.1 Serial Clock Line (SCL) Generation       10.4.2 Data and Addressing Management       10.4.3 I <sup>2</sup> C Acknowledge       10.4.4 Arbitration       10.4.4.1 SCL Arbitration       10.4.4.2 SDA Arbitration       10.4.5 Master Operations       10.4.6 Slave Operations	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.4 Dynamic State States       10.3.5 No START condition.       10.3.6 Dynamic State States       10.3.7 No START or STOP Condition       10.3.8 STOP Condition.       10.4.1 Serial Clock Line (SCL) Generation       10.4.2 Data and Addressing Management.       10.4.2.1 Addressing a Slave Device.       10.4.3 I <sup>2</sup> C Acknowledge.       10.4.4 Arbitration       10.4.4 SDA Arbitration       10.4.5 Master Operations.       10.4.6 Slave Operations.       10.4.7 General Call Address.       Slave Mode Programming Examples	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.3.1 START Condition.       10.3.3.2 No START or STOP Condition       10.3.3 STOP Condition.       10.4.1 Serial Clock Line (SCL) Generation       10.4.2 Data and Addressing Management.       10.4.3 I <sup>2</sup> C Acknowledge       10.4.4 Arbitration       10.4.5 Master Operations.       10.4.6 Slave Operations.       10.4.7 General Call Address.       Slave Mode Programming Examples       10.5.1 Initialize Unit	
10	<b>I<sup>2</sup>C B</b> 10.1 10.2 10.3	Bus Interface Units.       Overview.       I <sup>2</sup> C Interface.       Theory of Operation.       10.3.1 Operational Blocks.       10.3.2 I <sup>2</sup> C Bus Interface Modes.       10.3.3 Start and Stop Bus States       10.3.1 START Condition.       10.3.2 No START or STOP Condition       10.3.3 STOP Condition       10.4.1 Serial Clock Line (SCL) Generation       10.4.2 Data and Addressing Management       10.4.3 I <sup>2</sup> C Acknowledge.       10.4.4 rbitration       10.4.5 Master Operations       10.4.6 Slave Operations       10.4.7 General Call Address       Slave Mode Programming Examples       10.5.1 Initialize Unit	

	10.6	Master Programming Examples	.622
		10.6.2 Write 1 Byte as a Master	622
		10.6.3 Read 1 Byte as a Master	622
		10.6.4 Write 2 Bytes and Repeated Start Read 1 Byte as a Master	.623
		10.6.5 Read 2 Bytes as a Master - Send STOP Using the Abort	.624
	10.7	Glitch Suppression Logic	.625
	10.8	Reset Conditions	.625
	10.9	Register Definitions	.625
		10.9.1 I <sup>2</sup> C Control Register x - ICRx	.626
		10.9.2 I <sup>2</sup> C Status Register x - ISRx	.628
		10.9.3 I <sup>2</sup> C Slave Address Register x - ISARx	.630
		10.9.4 I <sup>2</sup> C Data Buffer Register x - IDBRx	.631
		10.9.5 I <sup>2</sup> C Bus Monitor Register x - IBMRx	.632
11	SMBı	us Interface Unit	.633
	11.1	Overview	.633
	11.2	SMBus Interface	.633
	11.3	System Management Bus Interface	.634
		11.3.1 SMBus Controller	.635
		11.3.1.1 SMBus Commands	.635
		11.3.1.2 Initialization Sequence	.636
		11.3.2 SMBus Signaling	.637
		11.3.2.1 Overview	.637
		11.3.2.2 Waveforms	.637
		11.3.2.2.1 Stall Pliase	638
		11.3.2.2.3ACK/NACK	.638
		11.3.2.2.4Wait States	.638
		11.3.3 Data Transfer Examples	.639
		11.3.3.1 Configuration and Memory Reads	.640
		11.3.3.2 Configuration and Memory Writes	.642
		11.3.4 Error Handling	.644
		11.3.5 SMBus Interface Reset	.644
		11.3.6 Configuration Access Arbitration	.644
	11.4	Register Definitions	.645
		11.4.1 Register Summary	.645
		11.4.2 SMBus Controller Command - SM_CMD	.646
		11.4.3 SMBus Controller Byte Count Register- SM_BC	.646
		11.4.4 SMDus Controller Address Register 3 - SM_ADDR3	.047
		11.4.5 SMBus Controller Address Register 1 - SM_ADDR2	6/7
		11.4.7 SMBus Controller Address Register 0 - SM_ADDR0	648
		11.4.8 SMBus Controller Data Registers - SM_ADDR0	648
		11 4.9 SMBus Controller Status - SM_STS	649
		11.4.10 SMBus Enable Register - SMBER	.650
12	UART	 ۲۶	.651
	12.1	Overview	651
	12.1	12 1 1 Compatibility with 16550 and 16750	652
	12.2	Signal Descriptions	.653
		о г <sup></sup>	

	12.3	Theory of Operation	654
		12.3.1 FIFO Interrupt Mode Operation	655
		12.3.1.1 Receiver Interrupt	655
		12.3.1.2 Transmit Interrupt	655
		12.3.2 Removing Trailing Bytes In Interrupt Mode	656
		12.3.2.1 Character Timeout Interrupt	656
		12.3.3 FIFO Polled Mode Operation	656
		12.3.3.1 Receive Data Service	656
		12.3.3.2 Transmit Data Service	656
		12.3.4 Autoflow Control	657
		12.3.4.1 RTS Autoflow	657
		12.3.4.2 CTS Autoflow	657
		12.3.5 Auto-Baud-Rate Detection	658
		12.3.6 Manual Baud Rate Selection	659
	12.4	Register Descriptions	660
		12.4.1 UART x Receive Buffer Register	662
		12.4.2 UART x Transmit Holding Register	662
		12.4.3 UART x Interrupt Enable Register	663
		12.4.4 UART x Interrupt Identification Register	664
		12.4.5 UART x FIFO Control Register	666
		12.4.6 UART x Line Control Register	668
		12.4.7 UART x Modem Control Register	670
		12.4.8 UART x Line Status Register	672
		12.4.9 UART x Modem Status Register	675
		12.4.10 UART x Scratchpad Register	676
		12.4.11 Divisor Latch Registers.	677
		12.4.12 UART x FIFO Occupancy Register	678
		12.4.13 UART x Auto-Baud Control Register	679
		12.4.14 UART x Auto-Baud Count Register	680
12	A rhit	rotion Unit	601
13	Arbit		001
	13.1	Arbitration Overview	681
	13.2	Internal Bus Arbiter Overview	682
		13.2.1 Theory of Operation	683
		13.2.1.1 Priority Mechanism	683
		13.2.1.2 Priority Example with Three Bus Initiators	684
		13.2.1.3 Priority Example with Six Bus Initiators	685
		13.2.1.4 Arbitration Signalling Protocol	686
		13.2.1.5 Internal Bus Arbitration Parking	687
		13.2.2 Intel XScale® Core Arbitration	688
	40.0	13.2.2.1 Multi-Transaction Timers	688
	13.3	Reset Conditions	689
	13.4	Register Definitions	689
		13.4.1 Internal Arbitration Control Register - IACR	690
		13.4.2 Multi-Transaction Timer Register 1 - MTTR1	691
		13.4.3 Multi-Transaction Timer Register 2 - MTTR2	692
14	Stand	dard Hot-Plug Controller	693
-	1/1		602
	14.1	1/11 DCI Standard Hat Dive Fastures	602
	140	14.1.1 FOI Stallualu HULFIUY FEatules	601
	14.2		094

	14.2.1	Slot Control Signals	694
		14.2.1.1 Output Control	694
		14.2.1.2 Input Control	694
	14.2.2	Parallel Mode 1-slot No-Glue Operation	695
	14.2.3	One-Slot-No-Glue Parallel Mode Operation	696
		14.2.3.1 Driving Bus To Ground When PCI Card is Disconnected	696
		14.2.3.2 Aborting Outbound PCI Cycles When Card is Disconnected	697
	14.2.4	Switch Debounce	698
	14.2.5	Enable/Disable Sequencing	698
	14.2.6	Secondary Bus Initialization and SHPC Role	698
		14.2.6.1 In-box Architecture	698
		14.2.6.2 Remote-I/O-Box Architecture	698
	14.2.7	M66EN Pin Handling	698
14.3	Interrup	ots	699
	14.3.1	MSI and Pin Interrupts	699
	14.3.2	ACPI Support	699
14.4	Error H	andling	700
14.5	Assum	ptions and Intel <sup>®</sup> 80333 I/O Processor Requirements	700
	14.5.1	MRL Opening During the Sequence	700
	14.5.2	Power Fault	700
14.6	SHPC	Working Register Set	701
	14.6.1	Access Disposition	701
	14.6.2	Register Description	702
		14.6.2.1 Standard Hot-Plug Controller Base Offset Register -	
		SHPC_BASEOFF (Offset 00)	702
		14.6.2.2 Slots Available Register 1 - SLOTS_AVAIL1 (Offset 04)	702
		14.6.2.3 Slots Available Register 2 - SLOTS_AVAIL2 (Offset 08)	703
		14.6.2.4 Slot Configuration Register - SLOT_CONFIG (Offset 0C)	703
		14.6.2.5 Secondary Bus Configuration Register - SBUS_CFG (Offset 10)	704
		14.0.2.0 SHPC MSI CONTRI REGISTER - SHPC_MSI_CHTRL (Offset 12)	704
		14.6.2.8 SHPC Command Register - SHPC, CMD (Offset 14)	705
		14.6.2.9 SHPC Command Status Register - SHPC CMDSTS (Offset 16)	705
		14.6.2.10 Interrupt Locator Register - INT_LOC (Offset 18)	706
		14.6.2.11 SERR Locator Register - SERR LOC (Offset 1C)	706
		14.6.2.12 SERR Interrupt Register - SERR-INT (Offset 20)	707
		14.6.2.13 Logical Slot Register - (Offset 24)	708
Intel	XScale®	Core and Core Performance Monitoring	711
15 1	High-L	avel Overview of Intel XScale <sup>®</sup> Core	711
13.1	15 1 1	APM Compatibility	711
	15.1.1	Fostures	712
	13.1.2	15.1.2.1 Multiply/ACcumulate (MAC)	712
		15.1.2.2 Memory Management	712
		15.1.2.2 Instruction Cache	713
		15 1 2 4 Branch Target Buffer	713
		15.1.2.5 Data Cache	713
15.2	CP14 F	Registers	714
•	15.2.1	Registers 0-3: Performance Monitoring	714
	15.2.2	Register 1: Clock Count Register CCNT	715
	15.2.3	Register 6: Core Clock Configuration Register CCLKCFG	
	15.2.4	Registers 8-15: Software Debug	716
	•		

15



	15.3	CP15 Registers	717
		15.3.1 Register 0: ID and Cache Type Registers	718
		15.3.2 Register 1: Control and Auxiliary Control Registers	719
		15.3.3 Register 2: Translation Table Base Register	721
		15.3.4 Register 3: Domain Access Control Register	721
		15.3.5 Register 5: Fault Status Register	722
		15.3.6 Register 6: Fault Address Register	722
		15.3.7 Register 7: Cache Functions	723
		15.3.8 Register 8: TLB Operations	725
		15.3.9 Register 9: Cache Lock Down	725
		15.3.10 Register 10: TLB Lock Down	726
		15.3.11 Register 13: Process ID	727
		15.3.11.1 The PID Register Affect On Addresses	727
		15.3.12 Register 14: Breakpoint Registers	728
		15.3.13 Register 15: Coprocessor Access Register	729
	15.4	Core Performance Monitoring Unit (CPMON)	730
40	<b>T</b> :	- · · ·	704
16	Ilmer	S	731
	16.1	Timer Operation	733
		16.1.1 Basic Programmable Timer Operation	733
		16.1.2 Watch Dog Timer Operation	734
		16.1.3 Load/Store Access Latency for Timer Registers	735
	16.2	Timer Interrupts	736
	16.3	Timer State Diagram	737
	16.4	Timer Registers	738
		16.4.1 Power Up/Reset Initialization	738
		16.4.2 Timer Mode Registers – TMR0:1	739
		16.4.2.1 Bit 0 - Terminal Count Status Bit (TMRx.tc)	739
		16.4.2.2 Bit 1 - Timer Enable (TMRx.enable)	740
		16.4.2.3 Bit 2 - Timer Auto Reload Enable (TMRx.reload)	740
		16.4.2.4 Bit 3 - Timer Register Privileged Read/Write Control (TMRx.pri)	741
		16.4.2.5 Bits 4, 5 - Timer Input Clock Select (TMRX.csel1:0)	741
		16.4.3 Timer Count Register – TCR0:1	742
		16.4.4 Timer Reload Register – TRRU:1	
		16.4.5 Timer Interrupt Status Register – TISR	744
	40 -	16.4.6 Watch Dog Timer Control Register – WDTCR	745
	16.5	Uncommon TCRX and TRRX Conditions	746
17	Interr	upt Controller Unit and IOAPIC	747
	17.1	Overview	747
		17.1.1 IOAPIC	
	172	Theory of Operation	749
		17.2.1 Interrupt Controller Unit	749
			750
		17 2 2 1 PCI IRQ# Mechanism	750
		17.2.2.2 PCI MSI Mechanism	
		17.2.2.3 PCI Virtual Wire Mechanism	751
	17.3	The Intel XScale <sup>®</sup> Core Exceptions Architecture	752
		17.3.1 CPSR and SPSR	752
		17.3.2 The Exception Process	752
		17.3.3 Exception Priorities and Vectors	753

	17.3.4	Software Requirements For Exception Handling	753
	0	17.3.4.1 Nesting FIQ and IRQ Exceptions	753
17.4	Intel <sup>®</sup> 8	30333 I/O Processor	
	Externa	al Interrupt Interface	754
	17.4.1	Interrupt Inputs	754
	17.4.2	Outbound Interrupts	756
	17.4.3	Interrupt Routing	757
17.5	Intel <sup>®</sup> 8	30333 I/O Processor	
	Interrup	ot Controller Unit	758
	17.5.1	Programmer Model	759
		17.5.1.1 Active Interrupt Source Control and Status	759
		17.5.1.2 Prioritization and Vector Generation for Active Interrupt Sources	760
	17.5.2	Operational Blocks	762
	17.5.3	80333: Internal Peripheral Interrupt	763
		17.5.3.1 Normal Interrupt Sources	763
		17.5.3.2 Error Interrupt Sources	765
	17.5.4	High-Priority Interrupt ( <b>HPI#</b> )	766
	17.5.5	limer Interrupts	766
	17.5.6	Software Interrupts	766
17.6	The Int	el <sup>®</sup> 80333 I/O Processor IOAPIC	767
	17.6.1	PCI IRQ# Interrupts	767
	17.6.2	PCI MSI Interrupts	768
	17.6.3	PCI Virtual Wire Interrupts	768
	17.6.4	PCI Express Legacy INTx Support and Boot Interrupt	768
	17.6.5	Buffer Flushing	769
	17.6.6	EOI Special Cycles	769
4 7 7	17.6.7	Upstream Interrupt Message Format	769
17.7	Default	Status	/ / 0
17.8	Interrup	bt Control Unit Registers	//1
	17.8.1	Interrupt Control Register 0 - INICILU.	112
	17.8.2	Interrupt Control Register 1 - INICIL1	/ / 5
	17.8.3	Interrupt Steering Register 0 - INTSTRU	///
	17.8.4	Interrupt Steering Register 1 - INTSTR1	780
	17.8.5	IRQ Interrupt Source Register 0 - IINTSRC0	782
	17.8.6	IRQ Interrupt Source Register 1 - IINTSRC1	785
	17.8.7	FIQ Interrupt Source Register 0 - FINTSRC0	787
	17.8.8	FIQ Interrupt Source Register 1 - FIN I SRC1	790
	17.8.9	Interrupt Priority Register 0 - IPR0	793
	17.8.10	) Interrupt Priority Register 1 - IPR1	794
	17.8.11	I Interrupt Priority Register 2 - IPR2	795
	17.8.12	2 Interrupt Priority Register 3 - IPR3	796
	17.8.13	3 Interrupt Base Register - IN I BASE	797
	17.8.14	Interrupt Size Register - INTSIZE	798
	17.8.15	5 IRQ Interrupt Vector Register - IINTVEC	799
	17.8.16	S FIQ Interrupt Vector Register - FINTVEC	800
	17.8.17	YPCI Interrupt Routing Select Register - PIRSR	801
17.9	IOAPIC	Registers	803
	17.9.1	Contiguration Space Registers	803
		17.9.1.1 IOAPIC Vendor ID - APIC_VID (Offset 00)	805
		17.9.1.2 IUAPIC Device ID - APIC_DID (Offset 02)	805
		17.9.1.3 IOAPIC Device Command Register - APIC_CMD (Offset 04)	805



	17.9.1.4	IOAPIC Status Register - APIC STS (Offset 06)	806
	17.9.1.5	IOAPIC Revision ID - APIC RID (Offset 08)	806
	17.9.1.6	IOAPIC Class Code - APIC CC (Offset 09)	807
	17.9.1.7	IOAPIC Header Register - APIC BIST/HT/MLT/CLS (Offset 0C)	807
	17.9.1.8	IOAPIC Memory Base Register - APIC_MBAR (Offset 10)	807
	17.9.1.9	IOAPIC Subsystem Vendor ID - APIC_SSVID (Offset 2C)	808
	17.9.1.10	IOAPIC Subsystem ID - APIC_SSID (Offset 2E)	808
	17.9.1.11	IOAPIC Capability Pointer - APIC CAPPTR (Offset 34)	808
	17.9.1.12	lOAPIC Alternate Base Address Register _APIC_ABAR (Offset 40)	808
	17.9.1.13	IOAPIC PCI Express Capability Identifier -	
		APIC_EXP_CAPID (Offset 44)	809
	17.9.1.14	IOAPIC PCI Express Next Item Pointer -	
		APIC_EXP_NXTP (Offset 45)	809
	17.9.1.15	i IOAPIC PCI Express Capability -	
		APIC_EXP_CAP (Offset 46)	809
	17.9.1.16	IOAPIC PCI Express Device Capabilities Register -	
		APIC_EXP_DCAP (Offset 48)	810
	17.9.1.17	IOAPIC PCI Express Device Control Register -	
		APIC_EXP_DCTL (Offset 4C)	811
	17.9.1.18	IOAPIC PCI Express Device Status Register -	
		APIC_EXP_DSTS (Offset 4E)	812
	17.9.1.19	IOAPIC PCI Express Link Capabilities Register -	
		APIC_EXP_LCAP (Offset 50)	812
	17.9.1.20	IOAPIC PCI Express Link Control Register -	
		APIC_EXP_LCTL (Offset 54)	813
	17.9.1.21	IOAPIC PCI Express Link Status Register -	~ 4 4
	470400	APIC_EXP_LSTS (Offset 56)	814
	17.9.1.22		044
	170100	APIC_PIM_CAPID (UIISel 60)	014
	17.9.1.23		011
	170104	AFIC_FIVI_INATE (Oliset oD)	014
	17.9.1.24		<b>815</b>
	170125	INAPIC Power Management Control/Status Register -	015
	17.3.1.25	APIC PM CSR (Offset 70)	815
	17 9 1 26	CIOAPIC Power Management Bridge Support Extensions -	010
	17.5.1.20	APIC PM BSE (Offset 72)	816
	179127	/IOAPIC Power Management Data Field -	010
	17.0.1.27	APIC PM DATA (Offset 73)	816
1792	Direct Me	emory Space Registers	817
17.0.2	17021	Access Disposition	817
	17.0.2.1	Register Summary	817
	17923	IOAPIC Index Register APIC IDX (Offset 00)	818
	17924	IOAPIC Window Register - APIC, WIND (Offset 10)	818
	17.9.2.5	IOAPIC IRQ Pin Assertion Register - APIC PAR (Offset 20)	818
	17.9.2.6	IOAPIC EOI Register - APIC EOI (Offset 40)	818
1793		ndirect Memory Space Registers	819
	17931	Register Summary	819
	17932	IOAPIC ID - APIC ID (Offset 00)	819
	17933	IOAPIC Version - APIC VS (Offset 04)	820
	17.9.3.4	IOAPIC Arbitration ID - APIC ARBID (Offset 08)	820
	17.9.3.5	IOAPIC Boot Configuration - APIC BCFG (Offset 0C)	820
	17.9.3.6	IOAPIC Redirection Table Low DWORD IRQxx -	0
		APIC RDLxx (Offset 10-3E)	821
		(	

		17.9.3.7 IOAPIC Redirection Table High DWORD IRQxx - APIC_RDHxx (Offset 11 - 3F)	822
18	Gene	eral Purpose I/O Unit	823
	18 1	General Purpose Input Output Support	823
	10.1	18 1 1 General Purpose Inputs	823
		18.1.2 General Purpose Outputs	
		18.1.3 Reset Initialization of General Purpose Input Output Function	824
		18.1.4 GPIO Pin Multiplexing	824
	18.2	Register Definitions	825
		18.2.1 GPIO Output Enable Register - GPOE	826
		18.2.2 GPIO Input Data Register - GPID	827
		18.2.3 GPIO Output Data Register - GPOD	828
19	Perip	oheral Registers	829
	19.1	Overview	829
	19.2	Accessing the PCI Configuration Registers	830
	19.3	Accessing Peripheral Memory-Mapped Registers	831
	19.4	Accessing Peripheral Registers Using	
		the Core Coprocessor Register Interface	831
	19.5	Architecturally Reserved Memory Space	831
	19.6	PCI Configuration Register Address Space	833
	19.7	Derinherel Memory Menned Register Address Space	845
	19.0	Coprocessor Pagister Space	047 862
	19.9		
20	Clock	king and Reset	865
20		king and Reset	865
20	<b>Clock</b> 20.1	king and Reset	865
20	<b>Clock</b> 20.1	king and Reset	865 865 866
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	865 865 866 866
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	865 865 866 866 866 866
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	865 865 866 866 866 866 866
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1 20.2	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1	king and Reset Clocking Overview	
20	<b>Clock</b> 20.1	king and Reset Clocking Overview 20.1.1 Clocking Theory of Operation 20.1.2 Clocking Region 1 20.1.3 Clocking Region 2 20.1.4 Clocking Region 3 20.1.5 Clocking Region 4 20.1.6 Clocking Region 5 20.1.7 Clocking Region 6 20.1.8 Clocking Region 7 20.1.9 Clocking Region Summary Reset Overview 20.2.1 <b>PWRGD</b> Reset Mechanism 20.2.2 <b>RSTIN#</b> Reset Mechanism 20.2.3 PCI Express Reset Mechanism 20.2.4 Software PCI Reset Mechanism	
20	<b>Clock</b> 20.1	king and Reset       20.1.1     Clocking Theory of Operation       20.1.2     Clocking Region 1       20.1.3     Clocking Region 2       20.1.4     Clocking Region 3       20.1.5     Clocking Region 4       20.1.6     Clocking Region 5       20.1.7     Clocking Region 6       20.1.8     Clocking Region 7       20.1.9     Clocking Region Summary       Reset Overview     20.2.1       20.2.1     PWRGD Reset Mechanism       20.2.2     RSTIN# Reset Mechanism       20.2.3     PCI Express Reset Mechanism       20.2.4     Software PCI Reset Mechanism       20.2.5     Hot-Plug Reset Mechanism	865 866 866 866 866 866 867 867 867 867 867
20	<b>Clock</b> 20.1	king and Reset       Clocking Overview       20.1.1 Clocking Theory of Operation       20.1.2 Clocking Region 1       20.1.3 Clocking Region 2       20.1.4 Clocking Region 3       20.1.5 Clocking Region 4       20.1.6 Clocking Region 5       20.1.7 Clocking Region 6       20.1.8 Clocking Region 7       20.1.9 Clocking Region Summary       Reset Overview       20.2.1 PWRGD Reset Mechanism       20.2.2 RSTIN# Reset Mechanism       20.2.3 PCI Express Reset Mechanism       20.2.4 Software PCI Reset Mechanism       20.2.5 Hot-Plug Reset Mechanism       20.2.6 I/O Processor Reset	
20	<b>Clock</b> 20.1	king and Reset       Clocking Overview       20.1.1     Clocking Theory of Operation       20.1.2     Clocking Region 1       20.1.3     Clocking Region 2       20.1.4     Clocking Region 3       20.1.5     Clocking Region 4       20.1.6     Clocking Region 5       20.1.7     Clocking Region 6       20.1.8     Clocking Region 7       20.1.9     Clocking Region Summary       Reset Overview     20.2.1       20.2.1     PWRGD Reset Mechanism       20.2.2     RSTIN# Reset Mechanism       20.2.3     PCI Express Reset Mechanism       20.2.4     Software PCI Reset Mechanism       20.2.5     Hot-Plug Reset Mechanism       20.2.6     I/O Processor Reset       20.2.7     Internal Bus Reset	
20	<b>Clock</b> 20.1	king and Reset       Clocking Overview       20.1.1 Clocking Theory of Operation       20.1.2 Clocking Region 1       20.1.3 Clocking Region 2       20.1.4 Clocking Region 3       20.1.5 Clocking Region 4       20.1.6 Clocking Region 5       20.1.7 Clocking Region 6       20.1.8 Clocking Region 7       20.1.9 Clocking Region Summary       Reset Overview       20.2.1 PWRGD Reset Mechanism       20.2.2 RSTIN# Reset Mechanism       20.2.3 PCI Express Reset Mechanism       20.2.4 Software PCI Reset Mechanism       20.2.5 Hot-Plug Reset Mechanism       20.2.6 I/O Processor Reset       20.2.7 Internal Bus Reset       20.2.8 Intel XScale <sup>®</sup> core Reset Mechanism	
20	<b>Clock</b> 20.1	king and Reset       Clocking Overview       20.1.1     Clocking Theory of Operation       20.1.2     Clocking Region 1       20.1.3     Clocking Region 2       20.1.4     Clocking Region 3       20.1.5     Clocking Region 4       20.1.6     Clocking Region 5       20.1.7     Clocking Region 6       20.1.8     Clocking Region 7       20.1.9     Clocking Region Summary       Reset Overview     20       20.2.1     PWRGD Reset Mechanism       20.2.2     RSTIN# Reset Mechanism       20.2.3     PCI Express Reset Mechanism       20.2.4     Software PCI Reset Mechanism       20.2.5     Hot-Plug Reset Mechanism       20.2.6     I/O Processor Reset       20.2.7     Internal Bus Reset       20.2.8     Intel XScale <sup>®</sup> core Reset Mechanism       20.2.9     Reset Summary	
20	20.2 20.3	king and Reset       20.1.1     Clocking Theory of Operation       20.1.2     Clocking Region 1       20.1.3     Clocking Region 2       20.1.4     Clocking Region 3       20.1.5     Clocking Region 4       20.1.6     Clocking Region 5       20.1.7     Clocking Region 6       20.1.8     Clocking Region 7       20.1.9     Clocking Region Summary       Reset Overview     20.2.1       20.2.1     PWRGD Reset Mechanism       20.2.2     RSTIN# Reset Mechanism       20.2.3     PCI Express Reset Mechanism       20.2.4     Software PCI Reset Mechanism       20.2.5     Hot-Plug Reset Mechanism       20.2.6     I/O Processor Reset       20.2.7     Internal Bus Reset       20.2.8     Intel XScale <sup>®</sup> core Reset Mechanism       20.2.9     Reset Sequencing	



21	Test	Logic Ui	nit and Testability	881
	21.1	Overvie	eM.	
	21.2	Test Co	ontrol/Observe Pins	
	21.3	IFFF 1	149 1 Standard Test Access Port (TAP)	882
	2110	21.3.1	TAP Pin Description	883
		21.0.1	21311 Test Clock ( <b>TCK</b> )	883
			21.3.1.2 Test Mode Select ( <b>TMS</b> )	883
			21.3.1.3 Test Data Input (TDI)	
			21.3.1.4 Test Data Output (TDO)	883
			21.3.1.5 Asynchronous Reset (TRST#)	883
		21.3.2	TAP Controller	884
			21.3.2.1 Test-Logic-Reset State	885
			21.3.2.2 Run-Test/Idle State	885
			21.3.2.3 Select-DR-Scan State	885
			21.3.2.4 Capture-DR State	885
			21.3.2.5 Shift-DR State	886
			21.3.2.6 Exit1-DR State	886
			21.3.2.7 Pause-DR State	886
			21.3.2.8 Exit2-DR State	886
			21.3.2.9 Update-DR State	887
			21.3.2.10 Select-IR-Scan State	887
			21.3.2.11 Capture-IR State	887
			21.3.2.12 SHIII-IR State	007
			21.3.2.13 EXILI-IN SIGLE	007 888
			21.3.2.14 Fause-IN State	888
			21 3 2 16 Undate-IR State	888
		2133	TAP Controller Configuration	889
		21.0.0	TAP Controller Begisters	890
		21.5.4	21.3.4.1 Instruction Register	800
			21.3.4.2 Instructions	891
			21.3.4.2.1High-Z	
			21.3.4.3 Boundary-Scan Register.	
			21.3.4.4 Bypass Register	893
			21.3.4.5 Device Identification Register	893

#### Figures

1	Intel <sup>®</sup> 80333 I/O Processor Functional Block Diagram	43
2	80333 Programming Model	58
3	Chipset Configuration Model	60
4	Type 1 to Type 0 Translation, PCI and PCI-X	62
5	I/O Forwarding	65
6	Memory Forwarding	67
7	Internal Arbitration Scheme	81
8	80333 PME# Messaging	
9	80333 Capabilities	102
10	ATU Block Diagram	160
11	ATU Queue Architecture Block Diagram	161
12	Inbound Address Detection	166
13	Inbound Translation Example	167
14	Internal Bus Memory Map	178
15	Outbound Address Translation Windows	
16	Direct Addressing Window	181
17	Private Device Example	199
18	ATU Interface Configuration Header Format	230
19	ATU Interface Extended Configuration Header Format (Power Management)	231
20	ATU Interface Extended Configuration Header Format (MSI Capability)	231
21	ATU Interface Extended Configuration Header Format (PCI-X Capability)	231
22	ATU Interface Extended Configuration Header Format (VPD Capability)	232
23	PCI Memory Map	319
24	Internal Bus Memory Map	
25	Overview of Circular Queue Operation	
26	Circular Queue Operation	
27	Intel <sup>®</sup> 80333 I/O Processor BIU and MCU Architecture	
28	DDR SDRAM 64-bit Memory Address Map	
29	DMA Controller	
30	DMA Channel Block Diagram	
31	DMA Chain Descriptor	
32	DMA Chaining Operation	
33	Example of Gather Chaining	
34	Synchronizing to Chained Transfers	
35	Optimization of an Unaligned DMA.	
36	Optimization of an Unaligned DMA	
37	CRC-32C Generator Polynomial	
38	DMA Programming Model State Diagram	
39	Software Example for Channel Initialization	
40	Software Example for POI-to-Local DMA Transfer	
41	Software Example for Local Memory-to-Local Memory DMA Transfer	
42	Application Accelerator Plack Diagram	
43	Principle / Four Source Descriptor Format	
44	Chain Descriptor Format for Fight Source Addresses (YOP Function)	
40	Chain Descriptor Format for Sixteen Source Addresses (XOR Function)	413 //16
40 ⊿7	Chain Descriptor Format for Thirty Two Source Addresses (XOR Function)	/10
<u>۲</u> ۲	Chain Descriptor Format for Dual-XOR-transfer	
40 ⊿0	P+O Base Chain Descriptor Format	۲24 ۸۵۸
-5		



50	P+Q Chain Descriptor Format for Six Source Addresses (XOR Function)	425
51	P+Q Chain Descriptor Format for Twelve Source Addresses (XOR Function)	426
52	P+Q Chain Descriptor Format for Sixteen Source Addresses (XOR Function)	429
53	XOR Chaining Operation	433
54	Example of Gather Chaining for Four Source Blocks	434
55	Synchronizing to Chained AA Operation	437
56	The Bit-wise XOR Algorithm	
57	Hardware Assist XOR Unit	
58	The Bit-wise XOR Algorithm including the P+Q RAID-6 Mode	444
59	GF Multiply Function	445
60	Galois Field Logarithm Transformation Table	
61	Galois Field Inverse Logarithm Transformation table	446
62	P+Q RAID-6 Generation Equation	
63	The Bit-wise Dual-XOR Algorithm	
64	An example of Zero Result Buffer Check	
65	An example of Zero Result Buffer Check with P+Q RAID-6	453
66	Example of a Memory Block Fill Operation	
67	Application Accelerator Programming Model State Diagram	
68	Optimization of an Unaligned Data Transfer	457
69	Pseudo Code: Application Accelerator Initialization	459
70	Pseudo Code: Application Accelerator Chain Resume Initialization	459
71	Pseudo Code: Suspend Application Accelerator	459
72	Pseudo Code: XOR Transfer Operation	460
73	Pseudo Code: Dual-XOR Transfer Operation	
74	Pseudo Code: Memory Block Fill Operation	
75	Pseudo Code: Zero Result Buffer Check Operation	462
76	Memory Controller Block Diagram	491
77	Dual-Bank DDR SDRAM Memory Subsystem	503
78	DDR SDRAM 64-bit Memory Address Map	506
79	FDDR SDRAM 64-bit Physical Map	506
80	Logical Memory Image of a DDR SDRAM Memory Subsystem	511
81	Supported DDR SDRAM Extended Mode Register Settings	513
82	Supported DDR-II SDRAM Extended Mode Register Settings	514
83	Supported DDR SDRAM Mode Register Settings	514
84	DDR SDRAM Initialization Sequence (controlled with software)	515
85	DDR SDRAM Pinelined Reads	517
86	DDR SDRAM Read 36 Bytes ECC Enabled BI –4	518
87	DDR SDRAM Write 36 Bytes, ECC Enabled, BL=4	520
88	DDR SDRAM Pinelined Writes	522
89	Refresh While the Memory Bus is Not Busy	523
90	FCC Write Flow	526
91	Intel <sup>®</sup> 80333 I/O Processor G-Matrix (generates the ECC)	527
92	Sub 64-bit DDR SDRAM Write (D <sub>2</sub> )	520
02	ECC Read Data Flow	531
9 <u>7</u>	Intel <sup>®</sup> $80333 I/O$ Processor H-Matrix (indicates the single-bit error location)	532
95	Power Failure Sequence	522
96	Power Failure State Machine	530 530
97	Power Failure Sequence	505 540
98	The Perinheral Rus Interface   Init	5+0 583
aa	Data Width and Low Order Address Lines	586

100 Four Mbyte Flash Memory System	588
101 120 ns Flash Read Cycle	589
102 120 ns Flash Write Cycle	590
103 I <sup>2</sup> C Bus Configuration Example	602
104 I <sup>2</sup> C Bus Interface Unit Block Diagram	603
105 Start and Stop Conditions	606
106 START and STOP Conditions	607
107 Data Format of First Byte in Master Transaction	610
108 Acknowledge on the I <sup>2</sup> C Bus	611
109 Clock Synchronization During the Arbitration Procedure	612
110 Arbitration Procedure of Two Masters	613
111 Master-Receiver Read from Slave-Transmitter	617
112 Master-Receiver Read from Slave-Transmitter / Repeated Start	
/ Master-Transmitter Write to Slave-Receiver	617
113 A Complete Data Transfer	
114 Master-Transmitter Write to Slave-Receiver	619
115 Master-Receiver Read to Slave-Transmitter	619
116 Master-Receiver Read to Slave-Transmitter / Repeated START	
/ Master-Transmitter Write to Slave-Receiver	
117 General Call Address	620
118 Basic SMBus Transfer Waveform	637
119 Start (S) / Repeat Start (Sr) Signaling	637
120 Stop (P) Signaling	638
121 ACK (A) Signaling	638
122 NACK (N) Signaling	638
122 Intel <sup>®</sup> 80333 I/O Processor Busy	630
124 DW/ORD Configuration Read Protocol (SMRus Block Write/Block Read, PEC Enabled)	640
125 DWORD Memory Read Protocol (SMBus Block Write/Block Read, PEC Enabled).	640
126 DWORD Configuration Road Protocol (SMBus Mord Write/Mord Road, PEC Enabled)	6/1
127 DWORD Configuration Read Protocol (SMBus Rlock Write/Rlock Read, PEC Disabled).	6/1
127 DWORD Connigulation Read Protocol (Simbus Diock White/Diock Read, FEC Disabled).	6/1
120 DWORD Memory Read Protocol (SMBus Block White/Block Read, FEC Disabled)	041 6/1
129 DWORD Configuration Write Protocol (SMBus Word Write, DEC Enchad)	041 642
130 DWORD Configuration while Protocol (Sivibus block while, FEC Enabled)	
131 DWORD Memory White Protocol (SMBus Word White, PEC Enabled)	
132 DWORD Conliguration write Protocol (SiMBus Byte write, PEC Enabled)	
133 DWORD Memory Write Protocol (SMBus Word Write/(Word, Byte) Read, PEC Enabled)	0.10
134 DWORD Memory Read Protocol (SMBus Word Write/Byte Read, PEC Enabled)	
135 Example UART Data Frame	654
136 NRZ Bit Encoding Example – (0100 1011)	654
137 Intel® 80333 I/O Processor Arbitration Block Diagram	681
138 Arbitration Example	683
139 Arbitration Between Three Initiators	686
140 Intel XScale <sup>®</sup> Core Back-to-Back Transactions with MTT1 Enabled	688
141 The Intel XScale <sup>®</sup> Core Architecture Features	712
142 Programmable Timer Functional Diagram	731
143 Timer Unit State Diagram	737
144 System Interrupt Architecture	748
145 Interrupt Controller Block Diagram (Active Interrupt Source Registers)	759
146 Interrupt Controller Block Diagram (FIQ/IRQ Interrupt Vector Generation)	761
147 80333 Interrupt Controller Connections	762



148 Intel <sup>®</sup> 80333 I/O Processor IOAPIC Interrupt Support	767
149 Memory Address Space	832
150 Intel <sup>®</sup> 80333 I/O Processor Clocking Regions Diagram	865
151 Intel <sup>®</sup> 80333 I/O Processor Reset Block Diagram	871
152 Intel <sup>®</sup> 80333 I/O Processor System Reset	878
153 IEEE 1149.1 Standard. Block Diagram	882
154 Timing of Actions in a TAP Controller State	884
155 TAP Controller State Diagram	884
156 TAP Controller Configuration	889
157 Device ID Register	893

#### Tables

	l erminology	48
2	PCI Express Interface Pins	52
3	PCI Interface Pins	53
4	PCI Interface Pins: 64-Bit Extensions	55
5	PCI Clock and Reset Pins	56
6	Reset Strap Pins	57
7	Configuration Spaces Visible to Software	59
8	Addressable 80333 Spaces	59
9	Secondary PCI Device Addressing	61
10	Inbound Transaction Ordering	70
11	Outbound Transaction Ordering	
12	PCI Mode Pin/Strap Encoding	72
13	PCI-X Initialization Pattern	72
14	PCI Transactions Supported	73
15	PCI-X Transactions Supported	74
16	80333 Implementation of Requester Attribute Fields	82
17	80333 Implementation of Completer Attribute Fields	82
10	Split Completion Abort Registers	02 
10	LOCK Transaction Handling in 80333	02
20	Transaction Termination Translation on Immediate Terminations of	05
20	Completion Required Cycles to PCI/X	01
21	Transaction Termination Translation on Split Terminations of	
21	Completion Required Cycles to PCI/X	91
22	Transaction Termination Translation on Immediate Terminations of	
	Completion Required Cycles to Internal Switch	92
23	Transaction Termination Translation on Terminations of	
_0	Completion Required Cycles to PCI Express	~~
24	Summary of Error Reporting	92
24 25	Summary of Error Reporting	92 94 99
24 25 26	Summary of Error Reporting Bit Attribute Definitions	92 94 99 99
24 25 26 27	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space	92 94 99 .103 .104
24 25 26 27 28	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID	92 94 99 .103 .104 .105
24 25 26 27 28 29	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD	92 94 99 .103 .104 .105 .106
24 25 26 27 28 29 30	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS.	92 94 .103 .104 .105 .106 .108
24 25 26 27 28 29 30 31	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD. Primary Device Status - PSTS Class Code - CC	92 94 .103 .104 .105 .106 .108 .110
24 25 26 27 28 29 30 31 32	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS Class Code - CC Cache Line Size - CLS	92 94 .103 .104 .105 .106 .108 .110 .111
24 25 26 27 28 29 30 31 32 33	Summary of Error Reporting. Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS. Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT	92 94 .103 .104 .105 .106 .108 .110 .111
24 25 26 27 28 29 30 31 32 33 34	Summary of Error Reporting. Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD. Primary Device Status - PSTS. Class Code - CC. Cache Line Size - CLS Primary Master Latency Timer - PMLT. Header Type - HEADTYP	92 94 .103 .104 .105 .106 .106 .108 .110 .111 .111
24 25 26 27 28 29 30 31 32 33 34 35	Summary of Error Reporting. Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT Header Type - HEADTYP SHPC 64-Bit Base Address Register - SHPC BAR	92 94 .103 .104 .105 .106 .108 .110 .111 .111 .111 .111
24 25 26 27 28 29 30 31 32 33 34 35 36	Summary of Error Reporting. Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS. Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT. Header Type - HEADTYP. SHPC 64-Bit Base Address Register - SHPC_BAR. Bus Numbers - BNUM	92 94 .103 .104 .105 .106 .108 .110 .111 .111 .111 .112 .112
24 25 26 27 28 29 30 31 32 33 34 35 36 37	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT Header Type - HEADTYP. SHPC 64-Bit Base Address Register - SHPC_BAR Bus Numbers - BNUM Secondary Master Latency Timer - SMI T	92 94 .103 .104 .105 .106 .108 .110 .111 .111 .111 .112 .112 .112 .113
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38	Summary of Error Reporting. Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT Header Type - HEADTYP SHPC 64-Bit Base Address Register - SHPC_BAR Bus Numbers - BNUM Secondary Master Latency Timer - SMLT I/O Base and Limit - IOBI	92 94 .103 .104 .105 .106 .106 .110 .111 .111 .111 .112 .112 .113 .113
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39	Summary of Error Reporting	92 94 99 .103 .104 .105 .106 .108 .110 .111 .111 .111 .111 .112 .112 .113 .113
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40	Summary of Error Reporting	92 94 99 .103 .104 .105 .106 .108 .110 .111 .111 .111 .111 .111 .112 .113 .113
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41	Summary of Error Reporting. Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS. Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT. Header Type - HEADTYP. SHPC 64-Bit Base Address Register - SHPC_BAR Bus Numbers - BNUM Secondary Master Latency Timer - SMLT I/O Base and Limit - IOBL Secondary Status - SSTS Memory Base and Limit - MBL Prefetchable Memory Base and Limit - PMBI	
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42	Summary of Error Reporting	92 94 .103 .104 .105 .106 .106 .108 .110 .111 .111 .111 .112 .112 .113 .113 .114 .115 .116 .116
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT Header Type - HEADTYP SHPC 64-Bit Base Address Register - SHPC_BAR Bus Numbers - BNUM Secondary Master Latency Timer - SMLT //O Base and Limit - IOBL Secondary Status - SSTS Memory Base and Limit - MBL Prefetchable Memory Base upper 32 Bits - PMBU32 Prefetchable Memory Limit Linper 32 Bits - PMBU32 Prefetchable Memory Limit Linper 32 Bits - PMBU32	92 94 .103 .104 .105 .106 .106 .108 .110 .111 .111 .111 .112 .112 .113 .113 .114 .115 .116 .116 .116
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT Header Type - HEADTYP SHPC 64-Bit Base Address Register - SHPC_BAR Bus Numbers - BNUM Secondary Master Latency Timer - SMLT I/O Base and Limit - IOBL Secondary Status - SSTS Memory Base and Limit - MBL Prefetchable Memory Base Upper 32 Bits - PMBU32 Prefetchable Memory Limit Upper 32 Bits - PMLU32 I/O Base and Limit - IOBE	92 94 .103 .104 .105 .106 .106 .108 .110 .111 .111 .111 .112 .113 .113 .114 .115 .116 .116 .116 .116
24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45	Summary of Error Reporting Bit Attribute Definitions Configuration Space – Legacy Region PCI Express Extended Configuration Space Identifiers - ID Command - PCICMD Primary Device Status - PSTS Class Code - CC Cache Line Size - CLS Primary Master Latency Timer - PMLT Header Type - HEADTYP SHPC 64-Bit Base Address Register - SHPC_BAR Bus Numbers - BNUM Secondary Master Latency Timer - SMLT I/O Base and Limit - IOBL Secondary Status - SSTS Memory Base and Limit - MBL Prefetchable Memory Base and Limit - PMBL Prefetchable Memory Base Upper 32 Bits - PMBU32 Prefetchable Memory Limit Upper 32 Bits - PMLU32 I/O Base and Limit Upper 16 Bits - IOBLU16 Capabilities List Pointer - CCAP	92 94 99 .103 .104 .105 .106 .108 .110 .111 .111 .111 .112 .113 .113 .114 .115 .116 .116 .116 .117 .117



46	Interrupt Information - INTR[A]	117
47	Interrupt Information - INTR[B]	117
48	Bridge Control - BCTRL	118
49	Bridge Configuration Register - BCNF	120
50	Multi-Transaction Timer - MTT	122
51	PCI Clock Control - PCLKC	122
52	PCI Express Capability Identifier - EXP_CAPID	122
53	PCI Express Next Item Pointer _ EXP_NXTP	122
54	PCI Express Capability - EXP_CAP	123
55	PCI Express Device Capabilities Register - EXP_DCAP	123
56	PCI Express Device Control Register - EXP_DCTL	124
57	PCI Express Device Status Register - EXP_DSTS	125
58	PCI Express Link Capabilities Register - EXP_LCAP	125
59	PCI Express Link Control Register - EXP_LCTL	126
60	PCI Express Link Status Register - EXP_LSTS	127
61	MSI Capability Identifier - MSI_CAPID	127
62	Next Item Pointer - MSI_NXTP	127
63	MSI Message Control - MSI_MC	128
64	MSI Message Address - MSI_MA	128
65	MSI Message Data - MSI_MD	128
66	Power Management Capabilities Identifier - PM_CAPID	128
67	Power Management Next Item Pointer - PM_NXTP	129
68	Power Management Capabilities - PM_CAP	129
69	Power Management Control/Status Register - PM_CSR	130
70	Power Management Bridge Support Extensions - PM_BSE	130
71	Power Management Data Field - PM_DATA	131
72	SHPC Capability Identifier - SHPC_CAPID	131
73	SHPC Next Item Pointer - SHPC_NXTP	131
74	SHPC DWORD Select Register - SHPC_DWSEL	132
75	SHPC Status - SHPC_STS	132
76	SHPC DWORD - SHPC_DWORD	133
77	PCI-X Capabilities Identifier - PX_CAPID	133
78	PCI-X Next Item Pointer - PX_NXTP	133
79	PCI-X Secondary Status - PX_SSTS	134
80	PCI-X Bridge Status - PX_BSTS	135
81	PCI-X Upstream Split Transaction Control - PX_USTC	135
82	PCI-X Downstream Split Transaction Control - PX_DSTC	136
83	Bridge Initialization Register A-Segment - BINIT[A]	137
84	Bridge Initialization Register B-Segment - BINIT[B]	138
85	PCI Express Advanced Error Capability Identifier - EXPAERR_CAPID	139
86	PCI Express Uncorrectable Error Status - ERRUNC_STS	140
87	PCI Express Uncorrectable Error Mask - ERRUNC_MSK	141
88	PCI Express Uncorrectable Error Severity - ERRUNC_SEV	142
89	PCI Express Correctable Error Status - ERRCOR_STS	143
90	PCI Express Correctable Error Mask - ERRCOR_MSK	143
91	Advanced Error Control and Capability Register - ADVERR_CTL	144
92	PCI Express Transaction Header Log - HDR_LOG	144
93	Uncorrectable PCI/X Status Register - PCI-XERRUNC_STS	145
94	Uncorrectable PCI/X Error Mask Register - PCI-XERRUNC_MSK	147
95	Uncorrectable PCI/X Error Severity Register - PCI-XERRUNC_SEV	149

96	Uncorrectable PCI/X Error Pointer Register - PCI-XERRUNC_PTR	151
97	Uncorrectable PCI/X Header Log - PCI-XHDR_LOG	151
98	Uncorrectable PCI/X Data Error Log - PCI-XD_LOG	151
99	Other PCI/X Error Logs and Control - PCI-XERRLOGCTL	152
100	Internal Arbiter Control A-Segment - ARB CNTRL[A]	154
101	Internal Arbiter Control B-Segment - ARB CNTRL[B]	154
102	Strap Status Register - SSR	155
103	Offset 300: PWRBGT HDR - Power Budgeting Enhanced Capability Header	156
104	Power Budgeting Data Select Register - PWRBGT_DSEL	156
105	Power Budgeting Data Register - PWRBGT_DATA	
106	Power Budgeting Capability Register - PWRBGT CAP	
107	Power Budgeting Information Registers 0:23 - PWRBGT_INFO[0:23]	
108	ATU Command Support	163
109	Outbound Address Translation Control	175
110	Internal Bus-to-PCI Command Translation for Two Memory Windows/DMA Channels	176
111	Internal Bus-to-PCI Command Translation for I/O Window	177
112	Inhound Queues	189
113	Inbound Read Prefetch Data Sizes	190
11/	PCI to Internal Rus Command Translation for All Inbound Transactions	101
115	Authound Aueues	102
116	ATU Inhound Data Flow Ordering Rules	102
117	ATU Authound Data Flow Ordering Rules	10/
110	Inhound Transaction Ordering Summany	106
110	Authound Transaction Ordering Summary	107
119	ATU Error Departing Summany, DCU Interface	
120	ATU Error Departing Summary Internal Publisher	222
121	Aldress Translation Unit Degisters	225
122	Address Translation Unit Registers	233
123	ATU YOL Configuration Register Space	236
124	ATU Vendor ID Register - ATUVID	238
125	ATU Device ID Register - ATUDID	239
126	ATU Command Register - ATUCMD	240
127	ATU Status Register - ATUSR	241
128	ATU Revision ID Register - ATURID	243
129	ATU Class Code Register - ATUCCR	244
130	ATU Cacheline Size Register - ATUCLSR	245
131	ATU Latency Timer Register - ATULT	246
132	ATU Header Type Register - ATUHTR	247
133	ATU BIST Register - ATUBISTR	248
134	Inbound ATU Base Address Register 0 - IABAR0	249
135	Inbound ATU Upper Base Address Register 0 - IAUBAR0	250
136	Inbound ATU Base Address Register 1 - IABAR1	251
137	Inbound ATU Upper Base Address Register 1 - IAUBAR1	252
138	Inbound ATU Base Address Register 2 - IABAR2	253
139	Inbound ATU Upper Base Address Register 2 - IAUBAR2	254
140	ATU Subsystem Vendor ID Register - ASVIR	255
141	ATU Subsystem ID Register - ASIR	256
142	Expansion ROM Base Address Register -ERBAR	257
143	ATU Capabilities Pointer Register - ATU_Cap_Ptr	258
144	Memory Block Size Read Response	259
145	ATU Base Registers and Associated Limit Registers	260



146 ATU Interrupt Line Register - ATUILR	
147 ATU Interrupt Pin Register - ATUIPR	
148 ATU Minimum Grant Register - ATUMGNT	
149 ATU Maximum Latency Register - ATUMLAT	
150 Inbound ATU Limit Register 0 - IALR0	
151 Inbound ATU Translate Value Register 0 - IATVR0	
152 Expansion ROM Limit Register - ERLR	
153 Expansion ROM Translate Value Register - ERTVR	
154 Inbound ATU Limit Register 1 - IALR1	
155 Inbound ATU Limit Register 2 - IALR2	
156 Inbound ATU Translate Value Register 2 - IATVR2	
157 Outbound I/O Window Translate Value Register - OIOWTVR	
158 Outbound Memory Window Translate Value Register 0- OMWTVR0	
159 Outbound Upper 32-bit Memory Window Translate Value Register 0- OUMWTVR0	
160 Outbound Memory Window Translate Value Register 1- OMWTVR1	
161 Outbound Upper 32-bit Memory Window Translate Value Register 1- OUMWTVR1	
162 Outbound Upper 32-bit Direct Window Translate Value Register - OUDWTVR	
163 PCI Express-to-PCI Bridge Secondary A Segment Bus Number Register - PEBSABN	R278
164 PCI Express-to-PCI Bridge Secondary B-Segment	-
Bus Number Register - PEBSBBNR	
165 PCI Express-to-PCI Bridge Primary Bus Number Register - PEBPBNR	
166 PCI Express-to-PCI Bridge Device Number Register - PEBDNUM	
167 ATU Configuration Register - ATUCR	
168 PCI Configuration and Status Register - PCSR	
169 ATU Interrupt Status Register - ATUISR	
170 ATU Interrupt Mask Register - ATUIMR	
171 Inbound ATU Base Address Register 3 - IABAR3	
172 Inbound ATU Upper Base Address Register 3 - IAUBAR3	
173 Inbound ATU Limit Register 3 - IALR3	
174 Inbound ATU Translate Value Register 3 - IATVR3	
175 Outbound Configuration Cycle Address Register - OCCAR	
176 Outbound Configuration Cycle Data Register - OCCDR	
177 VPD Capability Identifier Register - VPD_CAPID	
178 VPD Next Item Pointer Register - VPD_NXTP	
179 VPD Address Register - VPD_AR	
180 VPD Data Register - VPD_DR	
181 Power Management Capability Identifier Register - PM_CAPID	
182 Power Management Next Item Pointer Register - PM_NXTP	
183 Power Management Capabilities Register - PM_CAP	
184 Power Management Control/Status Register - PM_CSR	
185 PCI-X_Capability Identifier Register - PX_CAPID	
186 PCI-X Next Item Pointer Register - PX_NXTP	
187 PCI-X Command Register - PX_CMD	
188 PCI-X Status Register - PX_SR	
189 PCI-A Drive Strength Control Register - PADSCR	
190 PCI-A Drive Strength Value Register - PADSVR	
191 PCI-B Drive Strength Control Register - PBDSCR	
192 PCI-B Drive Strength Value Register - PBDSVR	
193 MU Summary	
194 Circular Queue Ordering Requirements	

195	Circular Queue Summary	324
196	Queue Starting Addresses	326
197	Circular Queue Summary	331
198	Circular Queue Status Summary	331
199	Message Unit Register	335
200	Inbound Message Register - IMRx	336
201	Outbound Message Register - OMRx	337
202	Inbound Doorbell Register - IDR	338
203	Inbound Interrupt Status Register - IISR	339
204	Inbound Interrupt Mask Register - IIMR	340
205	Outbound Doorbell Register - ODR	341
206	Outbound Interrupt Status Register - OISR	342
207	Outbound Interrupt Mask Register - OIMR	343
208	MU Configuration Register - MUCR	344
209	Queue Base Address Register - QBAR	345
210	Inbound Free Head Pointer Register - IFHPR	346
211	Inbound Free Tail Pointer Register - IFTPR	347
212	Inbound Post Head Pointer Register - IPHPR	348
213	Inbound Post Tail Pointer Register - IPTPR	349
214	Outbound Free Head Pointer Register - OFHPR	350
215	Outbound Free Tail Pointer Register - OFTPR	351
216	Outbound Post Head Pointer Register - OPHPR	352
217	Outbound Post Tail Pointer Register - OPTPR	353
218	Index Address Register - IAR	354
219	MSI Capability Identifier Register - MSI CAPID	355
220	MSI Next Item Pointer Register - MSI NXTP	356
221	MSI Message Control Register - MSI MCR	357
222	MSI Message Address Register - MSI MAR	358
223	MSI Message Upper Address Register - MSI MUAR	359
224	MSI Message Data Register - MSI MDR	360
225	Contiguous Byte Enable Encodings	364
226	Internal Bus Command Summary	365
227	Bus Interface Unit Register Tables	370
228	BIU Status Register - BIUSR	371
229	BIU Error Address Register - BEAR	373
230	BIU Control Register - BIUCR	374
231	DMA Registers	379
232	DMA Interrupt Summary	395
233	DMA Controller Unit Registers	398
234	Channel Control Register x - CCRx	399
235	Channel Status Register x - CSRx	400
236	Descriptor Address Register x - DARx	401
237	Next Descriptor Address Register x - NDARs	402
238	PCI Address Register x - PADRx	403
239	PCI Upper Address Register x - PU ADRx	404
240	Local Address Register x - LADRx	405
241	Byte Count Register x - BCRx	406
242	Descriptor Control Register x - DCRx	407
243	· · · · · · · · · · · · · · · · · · ·	408
244	Register Description	412



245	Descriptor Summary	.431
246	AA Operation and Command Combination Summary	. 439
247	Typical AA Operation and Addressing Summary	.440
248	AA Interrupts	.463
249	Application Accelerator Unit Registers	. 465
250	Accelerator Control Register - ACR	. 466
251	Accelerator Status Register - ASR	.467
252	Accelerator Descriptor Address Register - ADAR	.468
253	Accelerator Next Descriptor Address Register - ANDAR	.469
254	Data / Source Address Register - SAR1/PQSAR1	.470
255	Source Address Register232 - SAR232	.472
256	P+Q RAID-6 Source Address Registers 216 - PQSAR216	.473
257	Galois Field Multiplier Registers 1.5- GFMR1.5	474
258	Destination Address Register - DAR	476
259	Accelerator Byte Count Register - ABCR	477
260	Accelerator Descriptor Control Register - ADCR	478
261	Extended Descriptor Control Register 0 - FDCR0	482
262	Extended Descriptor Control Register 1 - EDCR1	484
263	Extended Descriptor Control Register 2 - EDCR2	486
264	Commonly Lised Terms	490
265	DDR SDRAM Memory Configuration Ontions	500
266	DDR SDRAM Interface Signals	501
267	DDR-II SDRAM Interface Signals	502
268	Supported DDR SDRAM Bank and Page Sizes	504
269	Bank Address Decode	504
270	DDR SDRAM Address Decode Summary	504
270	DDR SDRAM Address Decode #1	505
272	DDR SDRAM Address Decode #2	505
273	DDR SDRAM Address Decode #2	505
274	DDR SDRAM Address Begister Summary	506
275	Address Decoding for DDR SDRAM Memory Banks	507
276	Programming Codes for the DDR SDRAM Bank Size	507
270	Programming Values for the DDR SDRAM 32-bit Size Register (S32SRI29:201)	508
278	DDR SDRAM Commands	512
270	Typical Refresh Frequency Register Values	524
280	Syndrome Decoding	530
200	Overlanning Address Priorities	536
201	MOLI Error Response	5/2
202	Momory Controllor Pagistar	542
200	DDP SDPAM Initialization Provider SDIP	545
204	DDR SDRAM Initialization Register - SDIR	540
200	DDR SDRAM Control Register 0 - SDCR0	5/0
200	SDRAM Control Register SDRA	551
201	SDRAM Base Register - SDBR	552
200	SDRAM Boundary Register - SBR0	552
209	DDR SDRAM 32-hit Region Size Projetor - S22SP	551
290	ECC Control Pogistor - ECCP	554
291	ECC Log Pagistore - ELOCA ELOC1	555
29Z	ECC Address Projectore - ECAPO ECAP1	550
293	ECC Test Pagister - ECTST	550
294		. 550



295	Memory Controller Interrupt Status Register - MCISR	559
296	MCU Port Transaction Count Register - MPTCR	560
297	MCU Preemption Control Register - MPCR	561
298	Refresh Frequency Register - RFR	562
299	DCAL Control and Status Register - DCALCSR	563
300	DCAL Address Register - DCALADDR	565
301	DCAL Data Registers 17-0 - DCALDATA[17:0]	566
302	OCD Adjust Field Encoding	567
303	OCD Definition of DCALDATA[17:0] Registers	568
304	Receive Enable Calibration of DCALDATA[9] Register	568
305	Receive Enable Calibration Definition of DCALDATA[17:0] Registers	569
306	DQS Calibration Field Encodings for DCALDATA[17:0] Registers	570
307	DQS Definition of DCALDATA[17:0] Registers	570
308	Receive Enabled Delay Register - RCVDLY	571
309	Slave Low Mix 0 - SLVLMIX0	572
310	Slave Low Mix 1 - SLVLMIX1	573
311	Slave High Mix 0 - SLVHMIX0	574
312	Slave High Mix 1 - SLVHMIX1	575
313	Slave Length - SLVLEN	576
314	Master Mix - MASTMIX	577
315	Master Length- MASTLEN	578
316	DDR Drive Strength Status Register - DDRDSSR	579
317	DDR Drive Strength Control Register - DDRDSCR	580
318	DDR Miscellaneous Pad Control Register - DDRMPCR	581
319	Bus Signal Descriptions	587
320	Flash Wait State Profile Programming	589
321	Peripheral Bus Interface Register	592
322	PBI Control Register - PBCR	593
323	Memory Block Size Limit Register Value	594
324	PBI Base Address Register 0 - PBBAR0	595
325	PBI Limit Register 0 - PBLR0	596
326	PBI Base Address Register 1- PBBAR1	597
327	PBI Limit Register 1 - PBLR1	598
328	PBI Memory-less Boot Registers 2:0 - PMBR[2:0]	599
329	PBI Drive Strength Control Register - PBDSCR	600
330	I <sup>2</sup> C Interface Pins	601
331	I <sup>2</sup> C Bus Definitions	602
332	Modes of Operation	605
333	START and STOP Bit Definitions	606
334	Master Transactions	614
335	Slave Transactions	618
336	General Call Address Second Byte Definitions	620
337	I <sup>2</sup> C Register Summary	625
338	I <sup>2</sup> C Control Register x - ICRx	626
339	I <sup>2</sup> C Status Register x - ISRx	628
340	I <sup>2</sup> C Slave Address Register x - ISARx	630
341	I <sup>2</sup> C Data Buffer Register x - IDBRx	631
342	I <sup>2</sup> C Bus Monitor Register x - IBMRx	632
343	SMBus Interface Pins	633
344	SMBus Command Encoding	635



345	SMBus Status Byte Encoding	.640
346	SMBus Register Summary Table	645
347	Register Summary	645
348	SMBus Controller Command Register - SM_CMD	646
349	SMBus Controller Byte Count Register - SM_BC	646
350	SMBus Controller Address Register 3- SM ADDR3	647
351	SMBus Controller Address Register 2 SM ADDR2	647
352	SMBus Controller Address Register 1 - SM ADDR1	647
353	SMBus Controller Address Register 0 - SM ADDR0	648
354	SMBus Controller Data Registers - SM DATA	648
355	SMBus Controller Status Register - SM STS	649
356	SMBus Enable Register - SMBER	650
357	UART Signal Descriptions	653
358	Divisor Values for Typical Baud Rates	658
359	UART Register Addresses as Offsets of a Base	660
360	UART Unit Registers	660
361	UART Register MMR Addresses	661
362	UART x Receive Buffer Register - (UxRBR)	662
363	UART x Transmit Holding Register - (UxTHR)	.662
364	UART x Interrupt Enable Register - (UxIER)	663
365	UART x Interrupt Identification Register - (UxIIR)	.664
366	Interrupt Identification Register Decode	665
367	UART x FIFO Control Register - (UxFCR)	666
368	UART x Line Control Register - (UxLCR)	668
369	UART x Modem Control Register - (UxMCR)	670
370	UART x Line Status Register - (UxLSR).	.672
371	UART x Modem Status Register - (UxMSR)	675
372	UART x Scratchpad Register - (UxSCR)	676
373	UART x Divisor Latch Low Register - (UxDLL).	.677
374	UART x Divisor Latch High Register - (UxDLH)	677
375	UART x FIFO Occupancy Register - (UxFOR)	678
376	UART x Auto-Baud Control Register - (UxABR)	679
377	UART x Auto-Baud Count Register - (UxACR)	680
378	Priority Programming Example	.684
379	Bus Arbitration Example – Three Bus Initiators	684
380	Bus Arbitration Example – Six Bus Initiators	685
381	Arbitration Flow	687
382	Arbitration Reset	689
383	Arbiter Register	689
384	Internal Arbitration Control Register - IACR	690
385	Programmed Priority Control	690
386	Multi-Transaction Timer Register - MTTR1	.691
387	Multi-Transaction Timer Register - MTTR2	692
388	Parallel 1-slot No-Glue Mode Hot-Plug Pins	695
389	SHPC Working Register Set Layout	702
390	Standard Hot-Plug Controller Base Offset Register - SHPC BASEOFF	702
391	Slots Available Register 1 - SLOTS AVAIL1	702
392	Slots Available Register 2 - SLOTS AVAIL2	703
393	Slot Configuration Register - SLOT CONFIG	703
394	Secondary Bus Configuration Register - SBUS_CFG	704
# intel®

395	SHPC MSI Control Register - SHPC_MSI_CNTRL	704
396	SHPC Programming Interface Register - SHPC_PIF	705
397	SHPC Command Register - SHPC_CMD	705
398	SHPC Command Status Register - SHPC_CMDSTS	705
399	Interrupt Locator Register - INT_LOC	706
400	SERR Locator Register - SERR_LOC	706
401	SERR Interrupt Register - SERR-INT	707
402	Logical Slot Register- Slot SERR-INT Mask Field	708
403	Logical Slot Register- Slot Event Latch Field	709
404	Logical Slot Register- Slot Status Field	710
405	CP14 Registers	714
406	Accessing the Performance Monitoring Registers	714
407	Clock Count Register CCNT	715
408	Core Clock Count Register CCNT	715
409	Accessing the Debug Registers	716
410	CP15 Registers	717
411	ID Register	718
412	Cache Type Register	718
413	ARM* Control Register	719
414	Auxiliary Control Register	720
415	Translation Table Base Register	721
416	Domain Access Control Register	721
417	Fault Status Register	722
418	Fault Address Register	722
419	Cache Functions	723
420	TLB Functions	725
421	Cache Lockdown Functions	725
422	Data Cache Lock Register	725
423	TLB Lockdown Functions	726
424	Accessing Process ID.	727
425	Process ID Register	727
426	Accessing the Debug Registers	728
427	Coprocessor Access Register	729
428	Timer Performance Ranges	732
429	Timer Mode Register Control Bit Summary	733
430	Timer Responses to Register Bit Settings	735
431	Timer Registers	738
432	Timer Powerup Mode Settings	738
433	Timer Mode Register – TMRx	739
434	Timer Input Clock (TCLOCK) Frequency Selection	741
435	Timer Count Register – TCRx	742
436	Timer Reload Register – TRRx	743
437	Timer Interrupt Status Register – TISR	744
438	Watch Dog Timer Control Register WDTCR	745
439	Uncommon TMRx Control Bit Settings	746
440	Exception Priorities And Vectors	753
441	Interrupt Input Pin Descriptions	754
442	Interrupt Routing Summary	757
443	Normal Interrupt Sources	763
444	Error Interrupt Sources	765



445	Exception Priorities and Vectors	.768
446	Front Side Bus Delivery Address Format	.769
447	Front Side Bus Delivery Data Format	.769
448	Default Interrupt Routing and Status Values	.770
449	Interrupt Controller Register Addresses	.771
450	Interrupt Control Register 0 - INTCTL0	.772
451	Interrupt Control Register 1 - INTCTL1	.775
452	Interrupt Steering Register 0 - INTSTR0	.777
453	Interrupt Steering Register 1 - INTSTR1	.780
454	IRQ Interrupt Source Register 0 - IINTSRC0	.782
455	IRQ Interrupt Source Register 1 - IINTSRC1	.785
456	FIQ Interrupt Source Register 0 - FINTSRC0.	.787
457	FIQ Interrupt Source Register 1 - FINTSRC1	.790
458	Interrupt Priority Register 0 - IPR0	793
459	Interrupt Priority Register 1 - IPR1	.794
460	Interrupt Priority Register 2 - IPR2	795
461	Interrupt Priority Register 3 - IPR3	796
462	Interrupt Base Register - INTRASE	797
463	Interrupt Size Register - INTSIZE	798
464	IRQ Interrupt Vector Register- IINTVEC	799
465	FIQ Interrupt Vector Register- FINTVEC	800
466	PCI Interrupt Routing Select Register - PIRSR	801
467	IOAPIC Configuration Space Layout	803
468	IOAPIC Configuration Register Summary	804
469	IOAPIC Vendor ID - APIC VID	805
470	IOAPIC Device ID - APIC DID	.805
471	IOAPIC Device Command Register - APIC CMD	.805
472	IOAPIC Status Register - APIC STS	.806
473	IOAPIC Revision ID - APIC RID	.806
474	IOAPIC Class Code - APIC CC	.807
475	IOAPIC Header Register - APIC BIST/HT/MLT/CLS.	.807
476	IOAPIC Memory Base Register - APIC MBAR	.807
477	IOAPIC Subsystem Vendor ID - APIC SSVID	.808
478	IOAPIC Subsystem ID - APIC SSID	.808
479	IOAPIC Capability Pointer - APIC CAPPTR	.808
480	IOAPIC Alternate Base Address Register - APIC ABAR	.808
481	IOAPIC PCI Express Capability Identifier - APIC EXP CAPID	.809
482	IOAPIC PCI Express Next Item Pointer - APIC EXP NXTP	.809
483	IOAPIC PCI Express Capability - APIC EXP CAP	.809
484	IOAPIC PCI Express Device Capabilities Register - APIC EXP DCAP	.810
485	IOAPIC PCI Express Device Control Register - APIC EXP DCTL	.811
486	IOAPIC PCI Express Device Status Register - APIC EXP DSTS	
487	IOAPIC PCI Express Link Capabilities Register - APIC EXP LCAP	.812
488	IOAPIC PCI Express Link Control Register - APIC EXP LCTL	.813
489	IOAPIC PCI Express Link Status Register - APIC EXP LSTS	. 814
490	IOAPIC Power Management Capability Identifier - APIC PM CAPID (Offset 6C)	.814
491	IOAPIC Power Management Next Item Pointer - APIC PM NXTP (Offset 6D)	.814
492	IOAPIC Power Management Capabilities - APIC PM CAP - (Offset 6E)	.815
493	IOAPIC Power Management Control/Status Register - APIC PM CSR (Offset 70)	.815
494	IOAPIC Power Management Bridge Support Extensions - APIC_PM_BSE (Offset 72)	.816

# intel®

495	IOAPIC Power Management Data Field - APIC_PM_DATA (Offset 73)	816
496	IOAPIC Direct Memory Registers	817
497	IOAPIC Direct Memory Space Register Summary	817
498	IOAPIC Index Register - APIC_IDX	818
499	IOAPIC Window Register - APIC_WIND	818
500	IOAPIC IRQ Pin Assertion Register - APIC_PAR	818
501	IOAPIC EOI Register - APIC_EOI	818
502	IOAPIC Indirect Memory Space Register Summary	
503	IOAPIC ID - APIC_ID	
504	IOAPIC Version - APIC_VS	
505	IOAPIC Arbitration ID - APIC_ARBID	
506	IOAPIC Boot Configuration - APIC_BCFG	
507	IOAPIC Redirection Table Low DWORD IRQxx - APIC_RDLxx	
508	IOAPIC Redirection Table High DWORD IRQxx - APIC_RDHxx	
509	GPIO Pin Multiplexing	
510	General Purpose I/O Registers Addresses	
511	GPIO Output Enable Register - GPOE	
512	GPIO Input Data Register - GPID	
513	Output Data Register - GPOD	
514	Intel <sup>®</sup> 80333 I/O Processor PCI Programming Model	830
515	PCI Configuration Register Locations	
516	PCI Configuration Register Locations	
517	Intel XScale <sup>®</sup> Core Local Addresses Assigned to Integrated Peripherals	
518	Peripheral Memory-Mapped Register Locations	
519	Intel XScale <sup>®</sup> Core Coprocessor Registers Assigned to Integrated Peripherals	
520	Coprocessor Register Locations	
521	Clock Pin Summary	
522	80333 Clock Region Summary	
523	Internal Bus Reset Summary	
524	Reset Summary	
525	Reset Strap Signals	
526	Fest Control/Observe Pins	
527	TLU TAP Controller Instruction Set	
528	Intel XScale <sup>®</sup> Core TAP Controller Instruction Set	
529	Device ID Register Field Definitions	
530	Intel XScale <sup>®</sup> Core Device ID Register Settings	
531	TLU Device ID Register Settings	



# **Revision History**

Date	Revision	Description
March 2005	001	Initial Release.



# Introduction

# 1.1 About This Document

This document is the authoritative and definitive reference for the external architecture of the Intel<sup>®</sup> 80333 I/O processor<sup>1</sup> (80333).

Intel Corporation assumes no responsibility for any errors which may appear in this document nor does it make a commitment to update the information contained herein.

Intel retains the right to make changes to these specifications at any time, without notice. In particular, descriptions of features, timings, packaging, and pin-outs does not imply a commitment to implement them. In fact, this specification does not imply a commitment by Intel to design, manufacture, or sell the product described herein.

# 1.1.1 How To Read This Document

This document describes the product-specific features of the 80333. Each chapter describes a different feature and starts with an overview followed by the theory of operation.

The reader should have a working understanding of the *PCI Local Bus Specification*, Revision 2.3 and *PCI Express Specification*, Revision 1.0a.

# 1.1.2 Other Relevant Documents

- 1. Intel XScale<sup>®</sup> Core Developer's Manual (Order Number: 273473), Intel Corporation
- 2. *Intel<sup>®</sup> XScale™ Microarchitecture Programmer's Reference Manual* (Order Number: 273436), Intel Corporation
- 3. ARM Architecture Reference Manual ARM Limited. Order number: ARM DDI 0100E.
- 4. Intel<sup>®</sup> 80331 I/O Processor Developer's Manual (Order Number: 273942), Intel Corporation
- 5. PCI Local Bus Specification, Revision 2.3 PCI Special Interest Group
- 6. PCI-to-PCI Bridge Architecture Specification, Revision 1.1 PCI Special Interest Group
- 7. PCI Bus Power Management Interface Specification, Revision 1.1 PCI Special Interest Group
- 8. PCI-X Addendum to the PCI Local Bus Specification, Revision 1.0a PCI Special Interest Group
- 9. PCI Hot-Plug Specification, Revision 1.0 PCI Special Interest Group
- 10. PCI Express Specification, Revision 1.0a PCI Special Interest GroupIEEE Standard Test Access Port and Boundary Scan Architecture 1149.1a - IEEE
- 11. System Management Bus Specification, Revision 2.0

<sup>1.</sup> ARM\* architecture compliant.



# 1.2 About the Intel<sup>®</sup> 80333 I/O Processor

The 80333 is a multi- function device that integrates the Intel XScale<sup>®</sup> core (ARM\* architecture compliant) with intelligent peripherals and dual PCI Express-to-PCI Bridges. The 80333 consolidates, into a single system:

- Intel XScale<sup>®</sup> core
- x8 PCI Express Upstream Link
- Two PCI Express-to-PCI Bridges supporting PCI-X interface
- PCI Standard Hot-Plug Controller (segment B)
- Address Translation Unit (PCI-to-Internal Bus Application Bridge) interfaced to segment A
- High-Performance, Dual-Ported Memory Controller
- Interrupt Controller with 17 external interrupt inputs
- Two Direct Memory Access (DMA) Controller
- Application Accelerator with RAID 6 support
- Messaging Unit
- Peripheral Bus Interface Unit
- Two I<sup>2</sup>C Bus Interface Units
- Two 16550 compatible UARTs with flow control (4 pins)
- Eight General Purpose Input Output (GPIO) ports

It is an integrated processor that addresses the needs of intelligent I/O applications and helps reduce intelligent I/O system costs.

PCI Express\* is an industry standard, high performance, low latency system interconnect. The 80333 PCI Express upstream link is capable of x8 lane widths at 2.5GHz operation as defined by the *PCI Express Specification*, Revision 1.0a. The addition of the Intel XScale<sup>®</sup> core brings intelligence to the PCI Express-to-PCI Bridges.

The 80333 integrates dual PCI Express-to-PCI Bridges with the ATU as an integrated secondary PCI device. The Upstream PCI Express port implements the PCI-to-PCI Bridge programming model according to the *PCI Express Specification*, Revision 1.0a. The Primary Address Translation Unit is compliant with the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a definitions of an 'application bridge'.



Figure 1 is a block diagram of the 80333.





# **1.3 Features**

The 80333 combines the Intel XScale<sup>®</sup> core with powerful new features to create an intelligent I/O processor. This multi-device I/O Processor is fully compliant with the *PCI Local Bus Specification*, Revision 2.3 and the *PCI Express Specification*, Revision 1.0a. 80333-specific features include:

- Intel XScale<sup>®</sup> Core
- PCI Express 2.5GHz x8 link
- Application Accelerator Unit
- Address Translation Units
- Memory Controller
- Peripheral Bus Interface

- Two I<sup>2</sup>C Bus Interface Units
- Two DMA Controllers
- Messaging Unit
- Internal Bus
- Two UARTs
- Interrupt Controller and GPIOs
- Two PCI Express-to-PCI Bridges to secondary PCI-X 133MHz Bus interfaces

The subsections that follow briefly overview each feature. Refer to the appropriate chapter for full technical descriptions.

# 1.3.1 Intel XScale<sup>®</sup> Core

The 80333 core processor is based upon the Intel XScale<sup>®</sup> core. The core processor operates at a maximum frequency of 800 MHz. The instruction cache is 32 Kbytes in size and is 32-way set associative. Also, the core processor includes a data cache that is 32 Kbytes and is 32-way set associative and a mini data cache that is 2 Kbytes and is 2-way set associative.

# 1.3.2 PCI Express-to-PCI Bridge Units

80333 provides dual PCI Express-to-PCI Bridge units. These bridge units share a common upstream PCI Express interface compliant to the *PCI Express Specification*, Revision 1.0a. The PCI Express interface supports a port lane width of 8, for up to 2GB/s per direction (4GB/s total) at 2.5Gbps bit rate. The PCI-X secondary interfaces support 64-bit 133MHz compliant to the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a. These two secondary PCI bus interface are referred to the "A" and "B" segment, where the 80333 ATU resides on "A" segment.

# intel

# 1.3.3 Address Translation Units

An Address Translation Unit (ATU) allows PCI transactions direct access to the 80333 local memory. The Address Translation Unit supports transactions between PCI address space and 80333 address space. Address translation for the ATU is controlled through programmable registers accessible from both the PCI interface and the Intel XScale<sup>®</sup> core. The PCI interface of the ATU is connected to the 80333 "A" Secondary PCI interface of the bridge. Upstream access to the PCI Express interface is controlled by inverse decode with the address windows of the bridge. Dual access to registers allows flexibility in mapping the two address spaces. The ATU also supports the power management extended capability configuration header that as defined by the *PCI Bus Power Management Interface Specification*, Revision 1.1.

# 1.3.4 Memory Controller

The Memory Controller allows direct control of a DDR SDRAM memory subsystem. It features programmable chip selects and support for error correction codes (ECC). The memory controller can be configured for DDR SDRAM at 333 MHz and DDR-II at 400 MHz. The memory controller supports pipelined access and arbitration control to maximize performance. The memory controller is dual-ported, with a dedicated interface for the Intel XScale<sup>®</sup> core Bus Interface Unit and a second interface to the Internal Bus. The memory controller interface configuration support includes Unbuffered DIMMs, Registered DIMMs, and discrete DDR SDRAM devices.

External memory can be configured as host addressable memory or private 80333 memory utilizing the Address Translation Unit and Bridges.

## 1.3.5 Application Accelerator Unit

The Application Accelerator Unit (AA) provides low-latency, high-throughput data transfer capability between the AA unit, the 80333 local memory and the PCI bus. It executes data transfers from and to the 80333 local memory, from the PCI bus to the 80333 local memory, or from the 80333 local memory to the PCI bus. The AA unit performs XOR operations, computes parity, generates and verifies an eight byte Data Integrity field, performs memory block fills, and provides the necessary programming interface. The AA unit in the 80333 has been enhanced to support RAID 6 functionality.

## 1.3.6 Peripheral Bus Interface

The Peripheral Bus Interface Unit is a data communication path to the flash memory components or other peripherals of a 80333 hardware system. The PBI includes support for either 8-/16-bit devices. To perform these tasks at high bandwidth, the bus features a burst transfer capability which allows successive 8-/16-bit data transfers.

## 1.3.7 DMA Controller

The DMA Controller allows low-latency, high-throughput data transfers between PCI bus agents and the local memory. Two separate DMA channels accommodate data transfers to the PCI bus. Both channels include a local memory to local memory transfer mode. The DMA Controller supports chaining and unaligned data transfers. It is programmable through the Intel XScale<sup>®</sup> core only.



# 1.3.8 I<sup>2</sup>C Bus Interface Unit

The I<sup>2</sup>C (Inter-Integrated Circuit) Bus Interface Unit allows the Intel XScale<sup>®</sup> core to serve as a master and slave device residing on the I<sup>2</sup>C bus. The I<sup>2</sup>C unit uses a serial bus developed by Philips Semiconductor consisting of a two-pin interface. The bus allows the 80333 to interface to other I<sup>2</sup>C peripherals and microcontrollers for system management functions. It requires a minimum of hardware for an economical system to relay status and reliability information on the I/O subsystem to an external device. Also refer to I<sup>2</sup>C *Peripherals for Microcontrollers* (Philips Semiconductor).

The 80333 includes two I<sup>2</sup>C bus interface units.

## 1.3.9 Messaging Unit

The Messaging Unit (MU) provides data transfer between the 80333 and PCI-Express system. It uses interrupts to notify each system when new data arrives. The MU has four messaging mechanisms: Message Registers, Doorbell Registers, Circular Queues and Index Registers. Each allows a host processor or external PCI device and the 80333 to communicate through message passing and interrupt generation.

## 1.3.10 Internal Bus

The Internal Bus is a high-speed interconnect between internal units and Intel XScale<sup>®</sup> core processor. The Internal Bus operates at 333 MHz and is 64 bits wide.

## 1.3.11 UART Unit

The 80333 includes two UART unit. The UART Unit allows the Intel XScale<sup>®</sup> core to serve as a master and slave device residing on the UART bus. The UART unit uses a serial bus consisting of a two-pin interface. The bus allows the 80333 to interface to other peripherals and microcontrollers. Also refer to *16550 Device Specification* (National Semiconductor\*).

## 1.3.12 Interrupt Controller Unit

The Interrupt Controller Unit (ICU) aggregates interrupt sources both external and internal of 80333 to the Intel XScale<sup>®</sup> core processor. The ICU supports high performance interrupt processing with direct interrupt service routine vector generation on a per source basis. Each source has programmability for masking, core processor interrupt input, and priority.

## 1.3.13 GPIO

The 80333 includes 8 General Purpose I/O (GPIO) pins.

## 1.3.14 SMBus Unit

The SMBus (System Management Bus) Interface Unit allows the 80333 to serve as a slave device on the SMBus. SMBus is based on the principles of the  $I^2C$  bus and allows the 80333 to interface to system SMBus for external access and control of internal registers.

# intel®

# **1.4 Terminology and Conventions**

# 1.4.1 Representing Numbers

All numbers in this document can be assumed to be Base10 unless designated otherwise. In text, numbers in Base16 are represented as "nnnH", where the "H" signifies hexadecimal. In pseudo code descriptions, hexadecimal numbers are represented in the form 0x1234ABCD. Binary numbers are not explicitly identified but are assumed when bit operations or bit ranges are used.

# 1.4.2 Fields

A *reserved* field is a read-only field that may be used by a future implementation. Software should not modify or depend on any values in reserved fields.

A *preserved* field is a read-write field that may be used by a future implementation. Software should not modify or depend on any values in preserved fields.

A *read/write* field can written to a new value following initialization. This field can always be read to return the current value.

A *read only* field can be read to return the current value. Writes to *read only* fields are treated as no-op operations and does not change the current value nor result in an error condition.

A *read/clear* field can also be read to return the current value. A write to a *read/clear* field with the data value of 0 causes no change to the field. A write to a *read/clear* field with a data value of 1 causes the field to be cleared (reset to the value of 0). For example, when a *read/clear* field has a value of F0H, and a data value of 55H is written, the resultant field is A0H.

A *read/set* field can also be read to return the current value. A write to a *read/set* field with the data value of 0 causes no change to the field. A write to a *read/set* field with a data value of 1 causes the field to be set (set to the value of 1). For example, when a *read/set* field has a value of F0H, and a data value of 55H is written, the resultant field is F5H.

A *writeonce/readonly* field can be written to a new value **once** following initialization. After the this write has occurred, the *writeonce/readonly* field treats all subsequent writes as no-op operations and does not change the current value or result in an error condition. The field can always be read to return the current value.

# 1.4.3 Specifying Bit and Signal Values

The terms *set* and *clear* in this specification refer to bit values in register and data structures. When a bit is set, its value is 1; when the bit is clear, its value is 0. Likewise, *setting* a bit means giving it a value of 1 and *clearing* a bit means giving it a value of 0.

The terms *assert* and *deassert* refer to the logically active or inactive value of a signal or bit, respectively.



# 1.4.4 Signal Name Conventions

All signal names use the signal name convention of using the "#" symbol at the end of a signal name to indicate that the signal's active state occurs when it is at a low voltage. The absence of the "#" symbol indicates that the signal's active state occurs when it is at a high voltage.

# 1.4.5 Terminology

To aid the discussion of the 80333 architecture, the following terminology is used:

#### Table 1.Terminology

Term	Description			
Core processor	Intel XScale <sup>®</sup> core within the 80333.			
Downstream	At or toward a PCI Express port directed away from the root complex (to a bus with a higher number).			
DWORD	32-bit data quantity			
Host processor	Processor located upstream from the 80333.			
Inbound	At or toward the Internal Bus of the 80333 from the PCI interface of the ATU.			
Local bus	80333 Internal Bus.			
Local memory	Memory subsystem on the Intel XScale <sup>®</sup> core DDR SDRAM or Peripheral Bus Interface busses.			
Local processor	Intel XScale <sup>®</sup> core within the 80333.			
Outbound	At or toward the PCI interface of the 80333 ATU from the Internal Bus.			
QWORD	64-bit data quantity.			
Short	16-bit quantity.			
Upstream	At or toward a PCI Express port directed to the PCI Express root complex (to a bus with a lower number).			
word	16-bit quantity.			



# PCI Express-to-PCI Bridges

# 2

# 2.1 Overview

80333 integrates two PCI Express-to-PCI bridges. Each bridge follows the PCI-to-PCI Bridge programming model. The PCI Express port of the 80333 compliant to the *PCI Express Specification*, Revision 1.0a. The two PCI bus interfaces are fully compliant to the *PCI Local Bus Specification*, Revision 2.3 and *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0b.

# 2.1.1 PCI Express Interface Features

- PCI Express Specification, Revision 1.0a compliant
- Support for single x8 or single x4 PCI Express operation.
- 64-bit addressing support
- 32-bit CRC (cyclic redundancy checking) covering all transmitted data packets.
- 16-bit CRC on all link message information.



## 2.1.2 PCI-X Interface Features

- PCI Local Bus Specification, Revision 2.3 compliant.
- PCI-to-PCI Bridge Specification, Revision 1.1 compliant.
- *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0b compliant.64-bit 66 MHz, 3.3 V, NOT 5 V tolerant.
- Four and five external REQ/GNT Pairs for internal arbiter on segment A and B respectively. Only two pairs are available when operating SHPC is used in parallel mode.
- Programmable bus parking on either the last agent or always on 80333.External PCI clock-feed support for asynchronous primary and secondary domain operation.On-Die Termination (ODT) of 8.33 Kohms @ 40%.
- 64-bit addressing for inbound and outbound transactions
- Outbound LOCK# support.
- No inbound LOCK# support.
- PCI fast Back-to-Back capable as target.
- Up to four active and four pending inbound delayed transactions (memory read commands only) and up to two outbound delayed transaction.
- Private Device support for PCI devices to be controlled by the Intel XScale<sup>®</sup> core processor.
- Private Memory support for 80333 local memory access by Private Devices.
- Full peer-to-peer read write capability between the two PCI segments of 80333.

## 2.1.3 Power Management

- Support for PCI Express Active State Power Management (ASPM) L0s link state.
- Support for *PCI Bus Power Management Interface Specification*, Revision 1.1 compatibility D0, D3hot and D3cold device power states.
- Support PME# event generation on behalf of downstream PCI devices.

## 2.1.4 80333 in Server Systems

When 80333 is used as part of the Lindenhurst chipset and Twin-Castle chipset in server systems, 80333 PCI Express-to-PCI Bridges would behave like a traditional host bridge component with a transparent PCI-to-PCI Bridge programming model and touting a standard Hot-Plug capability (when enabled). Essentially, 80333 would serve as an electrical bridge between PCI Express and PCI and is transparent to software. Similarly, in PCI Express add-in card implementations, the 80333 provides an electrical bridge between PCI Express on the server motherboard, to PCI bus on the add-in card.

# 2.2 Signal Description

intel

The "#" symbol at the end of a signal name indicates that the active, or asserted state occurs when the signal is at a low voltage level. When "#" is not present after the signal name the signal is asserted when at the high voltage level. The following notations are used to describe the signal type:

- I: Input pin
- O: Output pin
- OD: Open-drain Output pin
- I/O: Bi-directional Input/Output pin
- I/OD: Bi-directional Input/Open-drain Output pin

intel

# 2.2.1 PCI Express Interface

### Table 2. PCI Express Interface Pins

Signal	I/O	Description	
REFCLK+/ REFCLK-	I	PCI Express Clock in: This pin receives a 100 MHz clock input from an external high quality crystal. This clock is used as the reference clock for the PCI Express Tx/Rx circuitry and by the PCI Express core PLL to generate a 250 MHz clock for the PCI Express core logic.	
PE0Tp[7:0]/ PE0Tn[7:0]       O       PCI Express Serial Data Transmit: There are eight serial differential output links in the Finterface each running at 2.5 Gb/s.         • X8 Mode: All PE0Tp[7:0]/ PE0Tn[7:0] are used       • X8 Mode: Only PE0Tp[3:0]/ PE0Tn[3:0] are used		<ul> <li>PCI Express Serial Data Transmit: There are eight serial differential output links in the PCI Express interface each running at 2.5 Gb/s.</li> <li>X8 Mode: All PE0Tp[7:0]/ PE0Tn[7:0] are used</li> <li>X4 Mode: Only PE0Tp[3:0]/ PE0Tn[3:0] are used</li> </ul>	
PE0Rp[7:0]/ PE0Rn[7:0]	I	<ul> <li>PCI Express Serial Data Receive: There are eight serial differential input links in the PCI Express interface each running at 2.5 Gb/s.</li> <li>X8 Mode: All PE0Rp[7:0]/ PE0Rn[7:0] are used</li> <li>X4 Mode: Only PE0Rp[3:0]/ PE0Rn[3:0] are used</li> </ul>	
PE_VSSBG	Ι	Ground for part of the BandGap, separated from the rest of the $V_{SS}$	
PE_RCOMPO	Ι	Connected to external resistor.	
PE_ICOMPI	I	Connected to same external resister as PE_RCOMPO on the board. This is an input sensing not for current compensation.	
PE_VCCBG I Ground for part of the BandGap, separated from the reset of the V <sub>CC</sub> s.		Ground for part of the BandGap, separated from the reset of the $V_{CC}s$ .	
Total			

## 2.2.2 PCI Bus Interface (Two Instances)

Each interface is marked by either the letter "A" or "B" to signify the interface. Therefore, A\_AD refers to the AD bus on PCI bus A, and B\_AD refers to the AD bus on PCI bus B. All pin names described in the following sections, an 'X' in the name indicates either A or B, for the PCI bus A and PCI bus B sides. For example, X\_PAR signal would be called A\_PAR on the PCI bus A and B\_PAR on the PCI bus B.

#### 1/0 **Description** Signal PCI Address/Data: These signals are a multiplexed address and data bus. During the address phase X\_AD[31::0] I/O or phases of a transaction, the initiator drives a physical address on AD31:00. During the data phases of a transaction, the initiator drives write data, or the target drives read data. Bus Command and Byte Enables: These signals are a multiplexed command field and byte enable field. During the address phase or phases of a transaction, the initiator drives the transaction type on I/O C/BE#[3:0]. When there are two address phases, the first address phase carries the dual address X\_C/BE#[3::0] command and the second address phase carries the transaction type. For both read and write transactions, the initiator drives byte enables on C/BE#[3:0] during the data phases. Parity: Even parity calculated on 36 bits - AD[31:0] plus C/BE[3:0]#. It is calculated on all 36 bits regardless of the valid byte enables. It is generated for address and data phases. It is driven identically to the AD[31:0] lines, except it is delayed by exactly one PCI clock. It is an output during the address phase for all P64H2 initiated transactions and all data phases when the 80333 is the X\_PAR I/O initiator of a PCI write transaction, and when it is the target of a read transaction. 80333 checks parity when it is the initiator of PCI read transactions and when it is the target of PCI write transactions. Device Select: The 80333 asserts DEVSEL# to claim a PCI transaction. As a target, the 80333 asserts DEVSEL# when a PCI master peripheral attempts an access to an internal address or an X DEVSEL# I/O address destined for PCI Express. As an initiator, DEVSEL# indicates the response to a 80333 initiated transaction on the PCI bus. DEVSEL# is tri-stated from the leading edge of PCIRST#. DEVSEL# remains tri-stated by the 80333 until driven as a target. Frame: FRAME# is driven by the Initiator to indicate the beginning and duration of an access. While X\_FRAME# FRAME# is asserted data transfers continue. When FRAME# is negated the transaction is in the final I/O data phase. Initiator Ready: IRDY# indicates the ability of the initiator to complete the current data phase of the X IRDY# I/O transaction. A data phase is completed when both IRDY# and TRDY# are sampled asserted. Target Ready: Indicates the ability of the target to complete the current data phase of the transaction. X TRDY# I/O A data phase is completed when both TRDY# and IRDY# are sampled asserted. TRDY# is tri-stated from the leading edge of PCIRST#. TRDY# remains tri-stated by the 80333 until driven as a target. X\_STOP# I/O Stop: Indicates that the target is requesting an initiator to stop the current transaction. Parity Error: Driven by an external PCI device when it receives data that has a parity error. Driven by X PERR# 80333 when, as a initiator it detects a parity error during a read transaction and as a target during I/O write transactions. System Error: The 80333 samples SERR# as an input and conditionally forwards it to the PCI I/OD X\_SERR# Express. PCI Requests: Request inputs into the 80333 internal arbiter. These inputs must be tied to VCC when A\_REQ#[3:0] I the internal arbiter is disabled. These are four requests on the A-Segment, and five requests on the B\_REQ#[4:0] B-Segment. A\_GNT#[3:0] PCI Grants: This signal indicates that an initiator can start a transaction on the PCI Bus. These are 0 four grants on the A-Segment, and five grants on the B-Segment. B\_GNT#[4:0] 66 MHz Enable: This input signal from the PCI Bus indicates the speed of the PCI Bus. When it is high then the Bus speed is 66 MHz and when it is low then the bus speed is 33 MHz. This signal will X\_M66EN I/OD be used to generate appropriate clock (33 or 66 MHz) on the PCI Bus. This pin does not have on-die 8.2K pull-up. This pull-up must be provided externally.

### Table 3.PCI Interface Pins (Sheet 1 of 2)



## Table 3.PCI Interface Pins (Sheet 2 of 2)

Signal	I/O	Description			
X_PCI-XCAP	I	PCI-X Capable: Indicates whether all devices on the PCI bus are PCI-X devices, so that the 80333 can switch into PCI-X mode. Note that this pin is also used by the Standard Hot-Plug Controller when in the 1-slot-no-glue mode to sample the card's bus/mode capability.			
X_LOCK#	0	PCI Lock: Indicates an exclusive bus operation and may require multiple transactions to complete. This signal is an output from 80333 when it is initiating exclusive transactions on PCI. PLOCK# is ignored when PCI masters are granted the bus. The 80333 does not propagate locked transaction upstream.			
Total	61				



# 2.2.3 PCI Bus Interface 64-Bit Extension (Two Interfaces)

### Table 4. PCI Interface Pins: 64-Bit Extensions

Signal	I/O	Description		
X_AD[63::32]	I/O	PCI Address/Data: These signals are a multiplexed address and data bus. This bus provides an additional 32 bits to the PCI bus. During the data phases of a transaction, the initiator drives the upper 32 bits of 64-bit write data, or the target drives the upper 32 bits of 64-bit read data, when REQ64# and ACK64# are both asserted.		
X_C/BE#[7::4]         I/O         Bus Command and Byte enables upper 4 bits: These signals are a multiplexed co byte enable field. For both reads and write transactions, the initiator will drive byte AD[63:32] data bits on C/BE7:4] during the data phases when REQ64# and ACK6 asserted.		Bus Command and Byte enables upper 4 bits: These signals are a multiplexed command field and byte enable field. For both reads and write transactions, the initiator will drive byte enables for the AD[63:32] data bits on C/BE7:4] during the data phases when REQ64# and ACK64# are both asserted.		
X_PAR64	I/O	PCI interface upper 32 bits parity: This carries the even parity of the 36 bits of AD[63:32] and C/BE#[7:4] for both address and data phases.		
X_REQ64#I/OPCI interface request 64-bit transfer: This is asserted by the initiator to indicate requesting a 64-bit data transfer. It has the same timing as FRAME#. When the this signal is an output. When the 80333 is the target this signal is an input.		PCI interface request 64-bit transfer: This is asserted by the initiator to indicate that the initiator is requesting a 64-bit data transfer. It has the same timing as FRAME#. When the 80333 is the initiator, this signal is an output. When the 80333 is the target this signal is an input.		
X_ACK64#         I/O         PCI interface acknowledge 64-bit transfer: This is asserted by the target only when R asserted by the initiator, to indicate the target ability to transfer data using 64 bits. It h timing as DEVSEL#.		PCI interface acknowledge 64-bit transfer: This is asserted by the target only when REQ64# is asserted by the initiator, to indicate the target ability to transfer data using 64 bits. It has the same timing as DEVSEL#.		
Total	39			



# 2.2.4 PCI Bus Interface Clocks and, Reset and Power Management (Two Interfaces)

Signal	I/O	Description			
A_CLKO[3::0], B_CLKO[4::0]	0	CI Clock Output: 33/66/100/133 MHz clock for a PCI device.			
X_CLKIN	I	PCI Clock In: This signal is PCI clock feedback input. The PCI PLL synchronizes this input to the clock that feeds the internal PCI logic.			
X_RST#	0	PCI Reset: 80333 asserts RST# to reset devices that reside on the secondary PCI bus.			
X_CLKOUT O Feedback routing. The		Feedback Clock Output: Clock feedback OUT for the PCI PLL for compensating for the board routing. This should be connected to the X_CLKIN input for feeding the PCI interface logic.			
X_PME# I		PCI Power Management Event: PCI bus power management event signal. This is a shared open drain input to 80333 from all the PCI cards on the corresponding PCI bus segment. This is a level sensitive signal that will be converted to a PME event on PCI Express.			
Total	14				

### Table 5. PCI Clock and Reset Pins



# 2.2.5 Reset Straps

### Table 6.Reset Strap Pins

Signal I/O		Description			
X_133EN	I	Enable PCI-X at 133 MHz for PCI Bus A/B: This pin, when high, allows the PCI-X segment to run at 133 MHz when X_PCI-XCAP is sampled high. When low, the PCI-X segment will only run at 100 MHz when X_PCI-XCAP is sampled high.			
Total	1				



# 2.3 **Programming Model and Addressing**

## 2.3.1 Programming Model

For software running on PCI Express, 80333 Primary PCI (Upstream PCI Express port) appears as a multi-function device. The notion of a multi-function device on the upstream port (rather than a device with multiple device numbers) is required to support PCI Express Hot-Plug. There are multiple PCI devices in 80333, with the second device being the ATU which exists on the Secondary PCI A-Segment. There are four functions of the upstream device as seen from PCI Express – two IOAPIC functions, two PCI bridge functions. All the functions have the same device number (Device=0) but with different function numbers. The Address Translation Unit (ATU) is an integrated secondary PCI A-segment device and is a single function device.

The 80333 PCI Express-to-PCI Bridge is a transparent bridge where primary and secondary buses share the same address space. A PCI bus segment in 80333 appears to software running on the primary side as a standard PCI-to-PCI Bridge. Figure 2 shows 80333 with both the PCI bus segments.



#### Figure 2. 80333 Programming Model

The PCI Express Interface becomes the logical primary PCI bus and each physical PCI bus segment becomes a separate secondary PCI bus with a PCI-to-PCI Bridge model. The standard Hot-Plug controller (SHPC) associated with PCI bus Segment-B appears to software as a capability of the PCI-to-PCI Bridge. The IOAPIC controllers reside as separate PCI functions (function numbers 1 and 3) on the primary bus. Function 1 is intended to be used to service interrupts from PCI bus A and function 3 is intended to be used to service interrupts from PCI bus B. IOAPIC is only visible to software running on the primary side and not the secondary side. Table 7 summarizes the configuration information of both the PCI bus segments. The devices mentioned and their function/device numbers are as viewed from the primary interface.



# 2.3.2 Addressing

### Table 7. Configuration Spaces Visible to Software

Function	Bus Number	PCI Express Device Number	PCI Express Function Number
PCI Bus A	Primary #	0	0
PCI Bus B	Primary #	0	2
IOAPIC A	Primary #	0	1
IOAPIC B	Primary #	0	3
Address Translation Unit (ATU)	Secondary A-Segment #	0xE	0

## 2.3.2.1 Addressable 80333 Spaces

Before discussing all the addressing/configuration aspects of 80333, here is a brief summary of all addressable spaces within the 80333 PCI Express-to-PCI Bridges (functions 0 and 2), their access mechanism and a description of when they are applicable. A detailed description of each of these spaces will follow in the later chapters. The address spaces of the IOAPICs, SHPCand ATU are detailed in the corresponding chapters.

#### Table 8.Addressable 80333 Spaces

Addrossable Space	Access			
Audressable Space	PCI Express	PCIA	PCIB	ATU
PCI to PCI Pridge Configuration Space A	Yes	No	No	Yes <sup>a</sup>
FCI-IO-FCI BILIQE CONINGULATION SPACE A	(Type 0)			(Type 0)
DCI to DCI Bridge Configuration Space B	Yes	Nia	NL-	Yes <sup>a,b</sup>
PCI-IO-PCI Bridge Configuration Space B	(Type 0)	INO	INO	(Type 1)
IOARIC Configuration Space A	Yes	No	No	No
TOAFIC Configuration Space A	(Type 0)			NO
IOARIC Configuration Space R	Yes	No	No	No
IOAFIC Configuration Space B	(Type 0)	NO	NO	NO
ATU Configuration Space A	Yes	Yes	No	NI/A
ATO Configuration Space A	(Type 1)	(Type 0)	NO	19/75
IOAPIC Memory Space A	Yes	Yes	Yes	Yes
IOAPIC Memory Space B	Yes	Yes	Yes	Yes
SHPC Memory Space B	Yes	Yes	Yes	Yes
ATU Memory Space A	Yes	Yes	Yes	N/A

a. Configuration cycles to the PCI Express-to-PCI Bridge Configuration Space from the Intel XScale<sup>®</sup> core are performed by Configuration transactions from the ATU to Device Number 0. Refer to Section 2.9.5.56, "Bridge Initialization Register -BINIT (Offset FC)" for details.

b. Configuration transactions to the B-segment bridge are necessary to clear the Configure Retry status in the "Bridge Initialization Register - BINIT (Offset FC)" during initialization.



## 2.3.2.2 Chipset Configuration Model

The configuration access mechanism in a PCI Express-based chipset is built around the PCI-to-PCI Bridge model to reuse the existing PCI software base and for ease of system scalability reasons. In this model, devices in MCH (Memory controller, PCI-to-PCI Bridges to extended I/O interconnects etc.) and ICH (USB, PCI-to-PCI Bridge to physical PCI bus, etc.) are mapped on the logical PCI bus 0. Each PCI Express port in MCH is bridged from the logical PCI bus 0 through a PCI-to-PCI Bridge with the primary side of the PCI-to-PCI Bridge being the logical PCI bus 0 and the secondary side of the PCI-to-PCI Bridge being the PCI Express bus. 80333 sitting on a PCI Express bus is mapped on to a logical PCI bus number greater than 0.

### Figure 3. Chipset Configuration Model





## 2.3.2.3 Secondary PCI Devices

PCI devices under control of the Intel XScale<sup>®</sup> core processor, can be configured as private devices, and hidden from BIOS and host software. Devices are hidden by inhibiting the assertion of the IDSEL input of the device during configuration cycles. This is configured through the BINIT register. 80333 supports public and private devices according to the following table.

#### Table 9. Secondary PCI Device Addressing

Device Number	Signal used for IDSEL input	Public/Private	Notes		
0	AD16	Reserved	Dedicated for Bridge		
1	AD17				
2	AD18		Based on BINIT Register Device Hiding Enable bit (bit 2) Available for secondary PCI devices.		
3	AD19				
4	AD20	Public or Private			
5	AD21				
6	AD22				
7	AD23				
8	AD24				
9	AD25	Pacanyod	Used by ATU to address Extended Configure space of Bridge		
10 (0xA)	AD26	Reserved			
11 (0xB)	AD27				
12 (0xC)	AD28		Available for accordant BCI devises		
13 (0xD)	AD29		Available for secondary PCI devices.		
14 (0xE)	AD30	Public	Dedicated for ATU on A-Segment PCI Bus (wired internally to ATU IDSEL)		
15 (0xF)	AD31		Available for secondary PCI devices.		

## 2.3.2.4 Configuration Space Access

80333 supports configuration space accesses from PCI Express using both the legacy *PCI-to-PCI Bridge Specification*, Revision 1.1 access mechanism and enhanced PCI Express configuration space access mechanism. 80333 supports configuration space accesses from the Intel XScale<sup>®</sup> core via the ATU bus also.

#### 2.3.2.4.1 PCI Express Configuration Access

• PCI Express Enhanced Configuration Access Mechanism

*PCI-to-PCI Bridge Specification*, Revision 1.1 defines the configuration space region of a PCI function to be up to 256 B. PCI Expressextends the PCI configuration space from 256 B to 4 K. The region up to 256 B can be accessed using the *PCI-to-PCI Bridge Specification*, Revision 1.1 defined mechanism for configuration accesses (CFC/CF8). The region above 256 B is accessible only via the enhanced configuration access mechanism defined in PCI Express. This mechanism utilizes a flat memory-mapped address region to access the configuration space. The core-logic chipset converts the legacy *PCI-to-PCI Bridge Specification*, Revision 1.1 or the enhanced PCI Express configuration space accesses into PCI Express configuration cycles.

The extended address bits used to access the configuration region above 256 B will be all 0s when the *PCI-to-PCI Bridge Specification*, Revision 1.1 compatible access mechanism is used or when accessing devices on PCI. Note that 80333, when it translates Type 1 configuration transactions from PCI Express-to-PCI and finds the extended address bits to be non-zero, will



terminate the transaction with an unsupported request response on PCI Express. All configuration accesses on PCI Express are aligned DWORD only.

— Type 0 Accesses to the 80333:

The bridge configuration space and the IOAPIC configuration space are accessed from PCI Express by a Type 0 configuration transaction. Type 0 transactions from PCI Express (not peer PCI) to the bridge segments return a configuration retry response on PCI Express provided the "Configure Retry" bit in the BINIT register is set.

**Error Response**: Type 0 configuration transactions that do not match function number 0-3 and a device number of 0 result in a UR<sup>1</sup> response on PCI Express.

— Type 1 Accesses to the 80333:

Type 1 accesses from PCI Express are used to access the PCI bus segments and the Address Translation Unit (ATU). Type 1 configuration transactions to PCI that do not complete within 40  $\mu$ s from the time they are received on PCI Express, will receive a configuration retry response on PCI Express. 80333 will continue to retry the transaction on PCI even when a retry response had been signaled on PCI Express, and when complete the transaction is discarded.

— Type 1 to Type 0 Translation:

A bridge performs a Type 1 to Type 0 translation when the Type 1 transaction is generated on PCI Express and is intended for a device attached directly to its secondary bus. The 80333 must convert the configuration command to a Type 0 format so that the secondary bus device can respond to it. The resulting Type 0 address to be driven on PCI is shown in Figure 4. Device numbers are decoded to generate a single '1' in address bits 31:16. When the device number is greater than 16, then all bits are '0'. A Device number of 0 converts to PCI AD(16) being a '1' and so on. The internal ATU on secondary PCI A-Segment is addressed in this manner when the Type 0 transaction addresses Device number 14 (0xE) on secondary PCI A-segment bus number.

*Note:* When the device-hiding bit in the BINIT register is set, 80333 will master abort all Type 0 transactions to the PCI bus targeting device numbers 0 to 9. Device numbers 15 to 10 are never hidden i.e. master aborted by 80333.

#### Figure 4. Type 1 to Type 0 Translation, PCI and PCI-X



<sup>1.</sup> Refer to PCI Express Specification, Revision 1.0a for details of the UR response type.

# intel

— Type 1-to-Type 1 Forwarding:

The 80333 forwards a Type 1 configuration cycle unchanged to the PCI bus under the following conditions:

- The Type field is 1010 - a configuration read or a configuration write transaction.

- The Et field in the request header is a 11 and the Format is x01.

- The bus number falls in the range defined by the lower limit (exclusive) in the secondary bus number register and the upper limit (inclusive) in the subordinate bus number register (i.e., Secondary Bus Number < Type 1 Bus Number  $\leq$  Sub-ordinate Bus Number).

— Type 1 to Special Cycle Forwarding:

The 80333 translates a Type 1 configuration write transaction on PCI Express into a special cycle on PCI, but does not translate a Type 1 configuration access on PCI to a special cycle on PCI Express. A cycle to be translated has the following attributes in the address:

- The Type (bus command) is 1010.
- The Et field in the request header is a 11 and the Format is 101.
- The device number is equal to 11111b.
- The function number is equal to 111b.
- The register number is equal to 000000b.

- The bus number is equal to the value in the secondary bus number register in configuration space.

The address and data are forwarded unchanged. Devices ignore the address and decode only the bus command. The data phase contains the special cycle message. The transaction will master abort, but results in a normal completion on the opposite bus (normal completion status on PCI Express, no DEVSEL# on PCI).

**Error Response**: Type 1 configuration transactions from PCI Express with bus numbers that do not match the range of the secondary bus number up to the sub-ordinate bus number (both inclusive), for either PCI buses will result in a master abort response (UR) on the PCI Express.

Type 1 PCI Express configuration transactions that have the extended address bits to be non-zero are terminated on PCI Express with a UR response.

### 2.3.2.4.2 PCI Configuration Access

• Type 0 Accesses from Intel XScale<sup>®</sup> core to the 80333 Bridge configuration Space.

The 80333 supports access of the configuration space from the Intel XScale<sup>®</sup> core. The Intel XScale<sup>®</sup> core can access the bridge configuration space with a Type 0 configuration transaction when the following conditions are met:

- The transaction is a Type 0 configuration read or a configuration write transaction.
- The address bit AD[16] is asserted to address Device number is zero and the function number is zero (A-segment bridge is function #0).

Inbound configuration transactions are enabled in the "Bridge Initialization Register - BINIT (Offset FC)" (enabled by default).



## 2.3.2.5 I/O Space Access Mechanism

I/O accesses always target the PCI bus from PCI Express. No I/O accesses are allowed from PCI-to-PCI Express and nor are any I/O accesses to internal devices (APIC, SHPC) allowed.

One I/O window can be set up in the PCI-to-PCI Bridge space for forwarding I/O transactions from PCI Express-to-PCI. No I/O transactions can be forwarded from PCI-to-PCI Express. Refer to Section 2.3.2.7 to see how I/O cycles in the VGA range are handled. The registers and register bits listed below define the setup and control of this I/O window:

- I/O Base and Limit (IOBL) Registers in the PCI-to-PCI Bridge configuration space
- I/O Enable bit (IOSE) in the Command Register in the PCI-to-PCI Bridge configuration space
- Enable 1 K (EN1K) granularity in the 80333 PCI-to-PCI Bridge configuration Register

To enable outbound I/O transactions, the I/O enable bit must be set in the command register in 80333 configuration space (bit 0 at offset 04-05h). When the I/O enable bit is not set, all I/O transactions initiated on PCI Express will receive a master abort completion. No inbound I/O transactions may cross the bridge and are therefore never claimed on the PCI bus. The 80333 implements one set of I/O base and limit address registers in configuration space that define an I/O address range for the bridge. PCI Express I/O transactions with addresses that fall inside the range defined by the I/O base and limit registers are forwarded to PCI, and PCI I/O transactions with addresses that fall outside this range are master aborted.

Setting the base address to a value greater than that of the limit address turns off the I/O range. When the I/O range is turned off, no I/O transactions are forwarded to PCI even when the I/O enable bit is set.

The I/O range has a minimum granularity of 4 KB and is aligned on a 4 KB boundary. The maximum I/O range is 64 K. This range may be lowered to 1 K granularity by setting the EN1K bit in the Bridge Configuration register (BCNF) at offset 40H.

The base register consists of an 8-bit field at configuration address 1Ch, and a 16-bit field at address 30H. The top four bits of the 8-bit field define bits <15:12> of the I/O base address. The bottom four bits read only as 0H to indicate that the 80333 supports 16-bit I/O addressing only. Bits <11:0> of the base address are assumed to be 0,which naturally aligns the base address to a 4 KB boundary. The I/O base upper 16 bits register at offset 30H is reserved.

After chip reset, the value of the I/O base address is initialized to 0000H. The I/O limit register consists of an 8-bit field at offset 1Dh and a 16-bit field at offset 32h. The top four bits of the 8-bit field define bits <15:12> of the I/O limit address. The bottom four bits read only as 0H to indicate that 16-bit I/O addressing is supported. Bits <11:0> of the limit address are assumed to be FFFh, which naturally aligns the limit address to the top of a 4 KB I/O address block. The 16 bits contained in the I/O limit upper 16 bits register at offset 32h are reserved. After chip reset, the value of the I/O limit address is reset to 0FFFh (i.e., the lower 4 K in the 64 K space).



#### Figure 5. I/O Forwarding



*Note:* When the EN1K bit is set in the Bridge Configuration register (BCNF), the base and limit registers are changed such that the top 6 bits of the 8-bit field define bits <15:10> of the I/O base/limit address, and the bottom 2 bits read only as 0H to indicate support for 16-bit I/O addressing. Bits <9:0> are assumed to be 0 (for the base register) and 1 (for the limit register), which naturally aligns the address to a 1 KB boundary.

**Error Response**: I/O transactions from PCI Express that do not match the I/O address forwarding window of either PCI-to-PCI Bridges will result in a UR response. Note that software is responsible for making sure that the I/O window programmed into the registers of the two PCI-to-PCI Bridges do not overlap.

## 2.3.2.6 Memory Space Access Mechanism

80333 supports 64 bits of memory addressing on both interfaces.

Memory Accesses from (to) PCI to (from) PCI Express

Two memory windows can be setup for forwarding memory transactions from PCI Express-to-PCI. These windows are defined as part of the standard PCI-to-PCI Bridge configuration space. Inverse decoding is used for forwarding transactions from PCI-to-PCI Express. Refer to Section 2.3.2.7 to see how memory cycles in the VGA range are handled. The registers and register bits listed below control the setup and operation of these memory windows:

- Memory-mapped I/O Base and Limit (MBL) registers.
- Prefetchable Memory Base and Limit (PMBL) registers.
- Prefetchable Memory Base and Limit upper 32 bits (PMBLU32) register.
- Memory Enable (MSE) bit in the Command register.
- Master Enable bit (BME) in the Command register.

To enable outbound memory transactions, the memory space enable bit in the command register must be set (bit 1 of offset 04-05h). To enable inbound memory transactions, the master enable bit in the command register must be set (bit 2 of offset 04-05h). 80333 will not prefetch data from PCI devices. Inbound prefetching is controllable and discussed later in Section 2.5, "PCI Express and PCI-X Bridge Interfaces".

- Memory Mapped I/O Window

BIOS software uses the memory mapped I/O window to map all non-prefetchable (i.e., have side effects on reads or non-memory like devices) memory space in PCI. The memory mapped I/O base address and memory mapped I/O limit address registers define an address range that the 80333 uses to determine when to forward memory commands. The 80333 forwards a memory transaction from PCI Express-to-PCI when the address falls within the range, and forwards it from PCI-to-PCI Express (or the peer PCI segment) when the address is outside the range (provided that they do not fall into the prefetchable memory range. See Memory Accesses from (to) PCI to (from) PCI Express bullet. This memory range supports 32-bit addressing only (addresses 4 GB). It has a granularity and alignment of 1 MB. This range is defined by a 16-bit base address register at offset 20H in configuration space and a 16-bit limit address register at offset 22h. The top 12 bits of each of these registers correspond to bits <31:20> of the memory address.

# intel





The low four bits are hardwired to 0. The low 20 bits of the base address are assumed to be all '0', which results in a natural alignment to a 1 MB boundary. The low 20 bits of the limit address are assumed to be all '1', which results in an alignment to the top of a 1 MB block. Setting the base to a value greater than that of the limit turns off the memory range.

Prefetchable Memory Window

The prefetchable memory base and address registers, along with their upper 32-bit counterparts, define an additional address range that the 80333 uses to forward accesses. BIOS software uses this window to map all the prefetchable memory spaces in the PCI. But 80333 treats the memory reads in this region to be still non-prefetchable. The 80333 forwards a memory transaction from PCI Express-to-PCI when the address falls within the range, and forwards transactions from PCI-to-PCI Express (or the peer PCI segment) when the address is outside the range and do not fall into the regular memory range (see Memory Accesses from (to) PCI to (from) PCI Express bullet). This memory range supports 64-bit addressing, has a granularity and alignment of 1 MB. This lower 32-bits of the range are defined by a 16-bit base register at offset 24h in configuration space and a 16-bit limit register at offset 28h. The top 12 bits of each of these registers correspond to bits <31:20> of the memory address. The low 4 bits are hardwired to 1h, indicating 64-bit address support. The low 20 bits of the base address are assumed to be all '0', which results in a natural alignment to a 1 MB boundary. The low 20 bits of the limit address are assumed to be all '1', which results in an alignment to the top of a 1 MB block. The upper 32-bits of the range are defined by a 32-bit base register at offset 28h in configuration space, and a 32-bit limit register at offset 2Ch. Setting the entire base (with upper 32-bits) to a value greater than that of the limit turns off the memory range.



- Private A-Segment Memory Window

The private memory window for the A-segment define an address range that the 80333 uses to map private devices to, and to locate local memory for private device access. This memory range is enabled through the Bridge Initialization Register and is fixed at address range where A[63:62]=10<sub>2</sub>. Accesses to this range will not be claimed by the A-segment PCI Express-to-PCI Bridge. This range is intended to be mapped to the ATUs private BAR windows (ATUBAR3) and the private device BARs. When disabled, all addresses in this range are forwarded upstream from PCI-to-PCI Express.

• Memory Accesses to SHPC Memory Space (B-segment only)

Memory accesses to SHPC memory space are handled through a 64-bit BAR in the PCI-to-PCI configuration space and an access enable bit:

- A 64-bit BAR (SHPC\_BAR) in PCI-to-PCI Bridge configuration space
- Memory space enable bit (MSE) in the command register in the PCI-to-PCI Bridge configuration space.
- Memory Accesses to IOAPIC Memory Space

Memory accesses to IOAPIC memory space are handled through two address ranges and an access enable bit:

- A 32-bit BAR (MBAR) in the IOAPIC configuration space.
- An alternate 32-bit BAR (ABAR), also in the IOAPIC configuration space.
- Memory space enable bit (MSE) in the command register in the IOAPIC configuration space.

# intel

## 2.3.2.7 VGA Addressing

When a VGA-compatible device exists behind a 80333 bridge, the VGA enable bit in the bridge control register is set (offset 3 at 3E-3Fh). When set, the 80333 forwards all transactions addressing the VGA frame buffer memory and VGA I/O registers from PCI Express-to-PCI, regardless of the values of the 80333 base and limit address registers. When set, 80333 does not forward VGA frame buffer memory accesses to PCI Express regardless of the values of the memory address ranges. However, the I/O enable and memory enable bit in the command register must still be set. When cleared, 80333 forwards transactions addressing the VGA frame buffer memory and VGA I/O registers from PCI Express-to-PCI when the defined memory and I/O address ranges enable forwarding. When cleared, accesses to the VGA frame buffer memory are forwarded from PCI-to-PCI Express when the defined memory address ranges enable forwarding. However, the master enable bit must still be set. The VGA I/O addresses are never forwarded to PCI Express. when the inbound I/O enable bit in BINIT register is cleared. When this bit is set and also the VGA enable bit is set, 80333 will not forward the VGA I/O addresses from PCI-to-PCI-Express.

- The VGA frame buffer consists of the following memory address range 000A 0000H–00B FFFFh
- The VGA I/O addresses consist of the I/O addresses 3B0H–3BBh and 3C0H–3DFh. These I/O addresses are aliased every 1 KB throughout the first 64 KB of I/O space, when the VGA 16-bit decode bit in the bridge control register (bit 4) is cleared. This means that address bits <9:0> (3B0H-3BBh and 3C0H-3DFh) are decoded, <15:10> are not decoded and can be any value, and address bits <31:16> must be all 0s. When the VGA 16-bit decode bit is set, then 80333 does the entire 16 bit decode on the VGA I/O addresses. When software sets the VGA enable bit in one bridge, the ISA enable bit must be set in the other bridge.

# 2.4 Transaction Ordering

## 2.4.1 80333 Transaction Ordering

80333 follows the producer-consumer model of a standard PCI Express-to-PCI bridge. Based on this model, 80333 implements a set of ordering rules in the inbound and outbound directions. The ordering plane covered by these rules spans the transaction domain covered by PCI Express and either of the two PCI segments. 80333 uses a single PCI Express virtual channel, for both PCI segments, to communicate with the host chip-set root complex. This would mean that, though in general, there is no ordering requirements between the traffic flowing to/from the two PCI segments (except those imposed by peer-to-peer read and write traffic between the two PCI segments), they appear ordered on PCI Express. As for the peer-to-peer traffic between the two PCI segments, 80333 internally tracks these transactions in its PCI Express interface, through a fence mechanism since it maintains its inbound PCI Express posted buffers independent for the two PCI segments.

Accesses to the internal 80333 configuration registers, follow no ordering relationship with respect to transactions moving to and from PCI and PCI Express buses. What this means is that inbound memory and configuration reads and writes (when enabled via the BINIT register) would be completed out of order with the transactions pending in the 80333 inbound queues towards PCI Express. Outbound memory/configuration transactions to the internal register space could complete out of order with respect to transactions pending in the outbound queues towards PCI. Software must be aware that any semaphore mechanism implemented through the internal 80333 register space would require a dummy read to PCI or PCI Express space to push the writes that could be pending in the 80333 queues in either direction. The ordering tables in the next two sections do not consider these transactions.

## 2.4.1.1 Inbound Transaction Ordering

Table 10 lists the combined set of ordering rules in the inbound path of 80333 for PCI transactions.

Row Pass Column	Posted Write	Delayed /Split Rd Request	Delayed /Split Rd Completion	Delayed /Split Write Completion	
Posted Write	No	Yes	No	No	
Delayed/Split Rd Request	No	Yes <sup>a</sup>	No	No	
Delayed/Split Rd Completion	No	Yes	Yes	Yes	
Delayed/Split Write Completion	No	Yes	Yes	Yes	

#### Table 10. Inbound Transaction Ordering

a. Subsequent requests only (prefetches). All inbound initial requests are in order.



## 2.4.1.2 Outbound Transaction Ordering

Table 11 lists the combined set of ordering rules in the outbound path of 80333.

### Table 11.Outbound Transaction Ordering

Row pass Column	Posted Write	Delayed /Split Rd Request	Delayed /Split Write Request	Delayed /Split Rd Completion	
Posted Write	No	Yes	Yes	Yes	
Delayed/Split Rd Request	No	Yes <sup>a</sup>	Yes	Yes	
Delayed/Split Write Request	No	Yes	Yes	Yes	
Delayed/Split Rd Completion	No	Yes	Yes	Yes	

a. 80333 supports two outbound completion required requests per PCI segment. Outbound delayed/split read requests can pass each other when issued on the PCI bus.



# 2.5 PCI Express and PCI-X Bridge Interfaces

This section deals with the specifics of the operation of the PCI and PCI Express interfaces and the transaction flow details at each of these interfaces.

## 2.5.1 PCI-X Interface

### 2.5.1.1 Initialization

80333 is the source bridge for the PCI bus and hence senses the M66EN and PCI-XCAP pins to decide the mode and frequency of operation. Encoding on the M66EN and PCI-XCAP pins along with the PCI reset straps is shown below that identifies the capabilities of the system.

Table 12. PCI Mode Pin/Strap Encoding

M133EN <sup>a</sup>	M66EN	PCI-XCAP	PCI Bus Mode	PCI Frequency	
N/A	Ground	Ground	PCI Conventional	33 MHz	
N/A	Not Connected	Ground	PCI Conventional	66 MHz	
N/A	N/A	Pull-down	PCI-X	66 MHz	
Ground	Ground N/A		PCI-X	100 MHz	
Not Connected	N/A	Not Connected	PCI-X	133 MHz	

a. This pin is used by the motherboard designer to run the PCI-X bus at 100 MHz even when all the PCI-X cards on the bus are 133MHz capable, so as to accommodate the board routing limitation on frequency. Note that to accommodate running a PCI-X bus at 66MHz, when all cards are capable of 133 MHz, the motherboard has to drive the PCI-XCAP pin to VCC/2.

M66EN/M133EN Straps have internal pull-ups, thereby making Not Connected condition a pull-up.

Once 80333 identifies the capabilities of the PCI bus devices, it drives the initialization pattern on DEVSEL#, STOP#, TRDY#, FRAME# and IRDY# pins as per Table 13 to initialize the PCI bus devices to the proper mode and frequency. The patterns shown in Table 13 below are guaranteed to be stable on the rising edge of the **A\_RST#** and **B\_RST#** pins. Refer to the PCI specifications for details of the timing of these patterns with respect to the **A\_RST#** and **B\_RST#** pins.

#### Table 13. PCI-X Initialization Pattern

FRAME# and IRDY# are Deasserted							
DEVSEL#	STOP#	TRDY#	Mode	Clock Period (ns)		Clock Frequency (MHz)	
				Maximum	Minimum	Minimum	Maximum
Deasserted	Deasserted	Deasserted	PCI 33	•	30	0	33
			PCI 66	30	15	33	66
Deasserted	Deasserted	Asserted	PCI-X	20	15	50	66
Deasserted	Asserted	Deasserted	PCI-X	15	10	66	100
Deasserted	Asserted	Asserted	PCI-X	10	7.5	100	133
Asserted	Deasserted	Deasserted	PCI-X	Reserved			
Asserted	Deasserted	Asserted	PCI-X				
Asserted	Asserted	Deasserted	PCI-X				
Asserted	Asserted	Asserted	PCI-X				

In summary:

- PCIRST# is an output from 80333.
- PCI clocks are actively driven out from 80333.
- 80333 drives X\_AD[31:0], X\_BE[3:0], X\_PAR low during PCI bus reset.
- 80333 drives X\_REQ64# low during reset.


# 2.5.1.2 Transactions Supported

## 2.5.1.2.1 PCI Mode

Table 14 lists all the transactions supported by 80333 on the PCI bus. Note that PCI command encodings do not correspond to the transactions listed below are ignored by 80333. 80333 supports full 64-bit addressing inbound and outbound. This implies that 80333 as a target can accept dual address cycles and as a master can generate dual address cycles.

## Table 14. PCI Transactions Supported

Transaction	Encoding	Master	Target
Interrupt acknowledge	0000	No	No
Special cycle (PCI Express Type1-to-PCI Special Cycle)	0001	Yes	No
I/O read	0010	Yes	No
I/O write	0011	Yes	No
Memory read	0110	Yes	Yes
Memory write	0111	Yes	Yes
Configuration Read	1010	Yes	No
Configuration Write	1011	Yes	No
Memory Read Multiple	1100	No	Yes
Dual Address Cycle	1101	Yes	Yes
Memory Read Line	1110	No	Yes
Memory Write and Invalidate	1111	No	Yes
LOCK Transaction	-	Yes	No



## 2.5.1.2.2 PCI-X Mode

Table 15 lists the transactions that 80333 supports when the PCI interface is in the PCI-X mode. PCI-X commands that are not any of the commands listed in the table are ignored by 80333. 80333 supports the memory write block command for writes that are multiples of cache-line (from DMA engine in MCH).

## Table 15. PCI-X Transactions Supported

Transaction	Encoding	Master	Target
Interrupt acknowledge	0000	No	No
Special cycle (PCI Express Type1-to-PCI Special Cycle)	0001	Yes	No
I/O read	0010	Yes	
I/O write	0011	Yes	
Reserved	0100	No	No
Reserved	0101	No	No
Memory read DWORD	0110	Yes	Yes
Memory Write	0111	Yes	Yes
Alias to Memory Read Block	1000	No	Yes
Alias to Memory Write Block	1001	No	Yes
Configuration Read	1010	Yes	
Configuration Write	1011	Yes	
Split Completion	1100	Yes	Yes
Dual Address Cycle	1101	Yes	Yes
Memory Read Block	1110	Yes	Yes
Memory Write Block	1111	Yes	Yes
LOCK Transaction	-	Yes	No



## 2.5.1.3 Read Transactions

## 2.5.1.3.1 Prefetchable

Any memory read line or memory read multiple commands on PCI that are decoded by the 80333 are prefetched on the PCI Express interface. The prefetchability of a given PCI read request is determined by the prefetch policy (PP) bits [4:3] in the "Bridge Configuration Register - BCNF (Offset 40)" on page 120. The amount of data prefetched depends on the clock frequency, x\_REQ64#, and command type. The 80333 does not prefetch past a 4 KB page boundary.

## 2.5.1.3.2 Delayed

All memory read transactions are delayed read transactions. When the 80333 accepts a delayed read request, it samples the address, command, and address parity. This information is entered into the delayed transaction queue. When in PCI-X mode, transactions follow the split transaction model of PCI-X. Read data returned from PCI Express for an active delayed transaction entry is forwarded to the PCI-X master as a split completion.

# 2.5.1.4 Configuration Transactions

Type 0 configuration transactions are issued when the intended target resides on the same PCI bus as the initiator. A Type 0 configuration transaction is identified by the configuration command and the lowest 2 bits of the address set to  $00_2$ .

Type 1 configuration transactions are issued when the intended target resides on another PCI bus, or when a special cycle is to be generated on another PCI bus. A Type 1 configuration transaction is identified by the configuration command and the lowest 2 bits of the address set to  $01_2$ .

The register number is found in both Type 0 and Type 1 formats and gives the Dword address of the configuration register to be accessed. The function number is also included in both Type 0 and Type 1 formats and indicates which function of a multifunction device is to be accessed. For single-function devices, this value is not decoded. Type 1 configuration transaction addresses also include a 5-bit field designating the device number that identifies the device on the target PCI bus that is to be accessed. In addition, the bus number in Type 1 transactions specifies the PCI bus to which the transaction is targeted.

## 2.5.1.5 Decoding

In the PCI mode, the 80333 supports the linear increment address mode only for bursting memory transfers (indicated when the low 2 address bits are equal to 0). When either of these address bits is nonzero, the 80333 disconnects the transaction after the first data transfer. The 80333 decodes all PCI cycles with medium DEVSEL# timing. In the PCI-X mode, 80333 always decodes as a Type A target. Also, in PCI-X mode, 80333 decodes split completions using the primary bus number field.

Refer to Section 2.3 for a general description of the addressing and decoding.

## 2.5.1.6 Transaction Termination

### 2.5.1.6.1 PCI Mode Transaction Termination

- 80333 as Master
  - Normal Termination

As a PCI master, the 80333 uses normal termination when DEVSEL# is returned by the target within five clock cycles of FRAME# assertion. It terminates a transaction when one of the following conditions are met:

- All write data for a write transaction is transferred from 80333 data buffers to the target (80333 does not generate back-to-back transactions)

- All read data for a read transaction have been transferred from the target to 80333
- The master latency timer expires and the 80333s bus grant is de-asserted
- Master Abort

When a 80333 initiated transaction is not responded to with DEVSEL# within five clocks of FRAME# assertion, the 80333 terminates the transaction with a master abort. The 80333 sets the received master abort bit in the secondary status register. Read requests (configuration, I/O, memory) that receive master abort termination are sent back to PCI Express/peer PCI with a master abort status.

*Note:* When 80333 performs a Type 1 to special cycle translation, a master abort is the expected termination for the special cycle on the target bus. In this case, the master abort received bit is not set, and the Type 1 configuration transaction is disconnected after the first data phase.

- Target Abort

When the 80333 receives a target abort, and the cycle requires completion on PCI Express, the 80333 will return the target abort status to PCI Express. 80333 sets the received target abort status bit in the secondary status register for all target aborts it receives on the PCI bus. Target abort can happen on any data phase of a PCI-X transaction, and a read completion packet to PCI Express/peer PCI, incurring a target abort in the middle of the packet would return valid data to the point of target abort and a target abort completion status for the reminder.

- Disconnect and Retry

When 80333 receives a disconnect response from a target, it will re-initiate the transfer with the remaining length. When 80333 receives a retry, it will wait at least two PCI clocks before it retries the transaction. When the retried transaction is a write, 80333 will retry the write till it completes normally or with a target or master abort. When the retried transaction is a delayed read or delayed write transaction, 80333 will allow memory reads and writes to pass the transaction. Refer to Section 2.4 for details on the kinds of reordering allowed. Retry is not considered an error condition and so there is no error logging or reporting done on a retry.

# intel

- 80333 as Target
  - Retry

The 80333 terminates a transaction with retry to an initiator when one of the following conditions is met:

- A new memory read transaction and 80333 delayed transaction queue is full.
- A memory read that has already been queued, but has not completed on PCI Express.

- A memory read that has been queued and completed on PCI Express but ordering rules require an outbound posted write to complete ahead of it.

- A LOCK transaction has been established from PCI Express-to-PCI
- A memory write transaction and 80333 has no free buffer space to accept the write.

- A memory write is from a master other than the master that was previously retried (starvation prevention mechanism).

— Disconnect

The 80333 disconnects an initiator when one of the following conditions is met.

- The 80333 cannot accept any more write data.
- The 80333 has no more read data to deliver.
- When the memory address is non-linear.
- Configuration reads and writes after the first data phase.
- When the inverse decode window ends
- Target Abort

The 80333 returns a target abort to PCI when one of the following conditions is met.

- The cycle master aborted or target aborted on PCI Express/peer PCI.
- Configuration read or write transaction has address or data parity errors.

### 2.5.1.6.2 PCI-X Mode Transaction Termination

- 80333 as Master
  - Initiator Disconnect or Satisfaction of Byte Count

As a PCI-X master, the 80333 uses normal termination (initiator disconnect or satisfaction of byte count) when DEVSEL# is returned by the target within six clock cycles after address phase. 80333 terminates a transaction when one of the following conditions are met:

- All write data indicated in the byte count of the write transaction is transferred from 80333 data buffers to the target. 80333 never does an initiator disconnect on a write before the byte count size has been satisfied

- Initiator disconnect at the next ADB on a split read completion because of 80333 data buffer running dry

- Initiator disconnect at the next ADB when the master latency timer expires and 80333s bus grant is de-asserted

- Master Abort

When a 80333 initiated transaction is not responded to with DEVSEL# within six clocks after address phase, the 80333 terminates the transaction with a master abort. The 80333 sets the received master abort bit in the secondary status register. Read requests (configuration, I/O, memory) that receive master abort termination are sent back to PCI Express/peer PCI with a master abort status. Delayed write requests that receive master abort are sent back to PCI Express with a master abort status.

- *Note:* When 80333 performs a Type 1 to special cycle translation, a master abort is the expected termination for the special cycle on the target bus. In this case, the master abort received bit is not set, and the Type 1 configuration transaction is disconnected after the first data phase.
  - Target Abort

When the 80333 receives a target abort, and the cycle requires completion on PCI Express, the 80333 will return the target abort status to PCI Express. 80333 sets the received target abort status bit in the secondary status register for all target aborts it receives on the PCI bus. Target abort can happen on any data phase of a PCI-X transaction, and a read completion packet to PCI Express/peer PCI, incurring a target abort in the middle of the packet would return valid data to the point of target abort and all 1s for the reminder of the length and a target abort completion status for the entire packet.

— Disconnect and Retry

When 80333 receives a disconnect response (single data phase or at next ADB) from a target, it will re-initiate the transfer with the remaining length. When 80333 receives a retry, it will wait at least two PCI clocks before it retries the transaction. When the retried transaction is a write, 80333 will retry the write till it completes normally or with a target or master abort. When the retried transaction is a delayed read or delayed write transaction, 80333 will allow memory reads, split completions and writes to pass the transaction. Refer to Section 2.4 for details on the kinds of reordering allowed. Retry is not considered an error condition and so there is no error logging or reporting done on a retry.

— Split Response

 $80333\ could\ receive\ split\ response\ for\ memory\ reads,\ I/O\ and\ configuration\ read\ and\ write\ transactions.$ 

March 2005

# intel

- 80333 as Target
  - Retry

The 80333 responds with a retry to PCI-X when one of the following conditions is met.

- A memory read transaction and 80333 delayed transaction queue is full.
- -- A LOCK transaction has been established from PCI Express-to-PCI.
- A memory write transaction and 80333 has no free buffer space to accept the write.
- A memory write is from a master other than the master that was previously retried (starvation prevention mechanism.
- (starvation prevention mechanism.
- *Note*:80333 never retries a completion since it always has enough buffer space for all split requests it sends out. No transaction information is retained on any writes.
  - Disconnect

80333 disconnects a transaction on PCI-X when one of the following conditions is met.

- The 80333 cannot accept any more write data and an ADB is reached.
- Split completion packet being sent, ADB is reached and 80333 read buffers running dry.
- Inversed decode window ends and an inbound write is in progress, irrespective of write buffer availability.
- Target Abort

The 80333 returns a target abort to PCI-X when one of the following conditions is met. - 80333 returns a target abort response in a split completion packet to PCI-X agent and split cycle target aborted on PCI Express/peer PCI.

- Configuration reads and writes with address, data parity or attribute phase parity errors.

- Master Abort

The 80333 master aborts all memory transactions on PCI-X when Bus master enable (BME) bit is unset.

The 80333 returns a master abort response in a split completion when the split cycle master aborted on PCI Express/peer PCI.

Split Response

All memory read cycles that cross 80333 receive this termination, when they are not retried.

• Termination on Device Boundary Crossing

On PCI-X any split request that crosses BAR boundary (initial address + length > BAR limit) will result in a normal response up to the BAR range and a "byte count out of range" response for the remainder of the length.

# 2.5.1.7 Arbitration

80333 supports a high-performance internal PCI arbiter that supports up to five external masters on the B-Segment and four external masters on the A-Segment. The request inputs into the internal arbiter include: 5 external request inputs (B-segment, 4 on A-Segment), the 80333 Address Translation Unit (A segment only), a PCI Express request input and a standard Hot-Plug controller request input (B-Segment only). The internal standard Hot-Plug controller always has the highest priority. The internal standard Hot-Plug controller requests for ownership of the PCI bus under the following conditions:

- A PCI card is being enabled on to the bus and bus needs to be quiescent.
- A PCI-X card is to be initialized and the host has to drive the initialization pattern to the card.

The internal arbiter, once it gives the grant to the standard Hot-Plug controller, is required to keep the grant asserted as long as the request from the standard Hot-Plug controller is active. In the case of PCI, during this bus ownership period by the Hot-Plug controller, the bus is idle and parked on 80333 (i.e., the 80333 is not generating any valid transactions on the bus). In the case of PCI-X, during this period, the 80333 could be driving an initialization pattern to the Hot-Plug card being powered up. Apart from the standard Hot-Plug controller, all other request inputs to the internal arbiter are split into two groups, a high priority group and a low priority group. Any master, including the PCI Express interface, can be programmed to be in either of the two groups. This could also mean that all the request inputs into the arbiter (excluding the standard Hot-Plug controller) could be in one single group. Within a group, priority is round-robin. The entire low-priority group represents one slot in the high-priority group. 80333 provides a 16-bit arbiter control register to control two aspects of the internal arbiter behavior:

- Priority group for a master (i.e., whether a master is in low priority group or high priority group).
- Bus parking on last PCI agent or PCI Express.

For controlling the priority level, there is one bit for each of the PCI REQ# inputs, one bit for the ATU, and one bit for the PCI Express request input. Since the standard Hot-Plug is always the highest priority, it has no control bit. Bit 7 in the control register is for PCI Express (PCI Express-to-PCI Bridge) bit 5 is for ATU on the A-Segment, bit 4 is for REQ[4]# (B-Segment only), bit 3 is for REQ[3]# and so on. A value of '1' in a bit position puts the corresponding master in the high-priority group. Figure 7 represents the arbiter scheme with bits 7 to 0 in the arbiter control register being "1100011". In Figure 7, M0 represents master 0 (REQ[0]#), M1 represents master 1 (REQ[1]#) and so on. Bit 8 in the arbiter control register controls the bus parking behavior of the internal arbiter. A value of 0 instructs the internal arbiter to always park the bus on PCI Express. A value of 1 instructs the internal arbiter to park the bus on the last PCI master. 80333 also supports an 8-bit MTT (Multi-Transaction Timer) register that influences the behavior of the internal arbiter. This register controls the amount of time that the 80333s arbiter allows a PCI initiator to perform multiple back-to-back transactions on the PCI bus. The number of clocks programmed in the MTT represents the guaranteed time slice (measured in PCI clocks) allotted to the current agent, after which the arbiter will grant another agent that is requesting the bus.



Figure 7. Internal Arbitration Scheme





# 2.5.1.8 PCI-X Protocol Specifics

## 2.5.1.8.1 Attributes

Table 16 describes how the 80333 fills in attribute fields where the *PCI-to-PCI Bridge Specification*, Revision 1.1, leaves some implementation leeway.

#### Table 16. 80333 Implementation of Requester Attribute Fields

Attribute	Function
No Snoop (NS)	80333 just forwards this attribute in both directions and does nothing with it internally.
Relaxed Ordering (RO)	This bit allows relaxed ordering of transactions, which the 80333 does not permit. This bit is simply forwarded in the 80333, and is never generated on PCI-X from an PCI Express packet or vice-versa.
Tag	Since the 80333 can have two outstanding requests on PCI-X at a time, this field can be either 0 or 1
Byte Counts	From PCI Express, this is based upon the length field from PCI Express, which is DWORD based.

### 2.5.1.8.2 4 GB and 4 K Page Crossover

The *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0b, allows burst transactions to cross page (in the 80333 case, this is 4 K) and 4 GB address boundaries. As a PCI-X master, the 80333 will always end the transaction at a 4 K boundary. As a PCI-X target, the 80333 will allow a burst past a 4 K page boundary. Note that on PCI Express, requests never cross a 4 K boundary, reads or writes.

#### 2.5.1.8.3 Wait States

The 80333 will never generate wait states as a target except in the case of configuration read and write accesses, which are handled with immediate completions.

#### 2.5.1.8.4 Split Transactions

• Completer attributes

#### Table 17. 80333 Implementation of Completer Attribute Fields

Attribute	Function
Byte Count Modified (BCM)	80333 will not set this bit.
Split Completion Error (SCE)	<ul> <li>The 80333 will only set this bit when:</li> <li>a memory read command from PCI-X master or target aborted on PCI Express.</li> <li>When 80333 does a queue discard operation of inbound queues.</li> </ul>
Split Completion Message (SCM)	This bit shadows the SCE bit.

• Unexpected Split Completions

The 80333 will assert **DEVSEL**# and discard the data when the Requester ID matches the bridge, but the Tag does not match that of any outstanding requests (0 or 1) from this device.

• Split Completion Messages

The 80333 can only generate error messages for cycles that cross the bridge that master or target abort. Now Dword cycles will cross the bridge that requires completion (i.e. I/O cycles). Therefore, the 80333 can only generate a "PCI-X Bridge Error" completion message for the memory read commands as indicated in Table 18.

#### Table 18. Split Completion Abort Registers

Index	Message
00h	Master-Abort: The 80333 encountered a Master-Abort on the destination bus.
01h	Target-Abort: The 80333 encountered a Master-Abort on the destination bus

82

# 2.5.2 LOCK Cycles

A lock is established when a memory read from PCI Express that targets PCI with the lock bit set is responded to with a **TRDY#** by a PCI target. The bus is unlocked when the Unlock Special Cycle is received on PCI Express. When the PCI bus is locked, all inbound memory transactions from that bus are retried. The 80333 inbound read prefetch engine stops issuing any more requests on the PCI Express bus. Note though that read completions for inbound read requests issued ahead of the lock being established on the PCI bus could return on PCI Express after the PCI lock has been established and 80333 should be able to accept them.

Once the bus is locked, any PCI Express cycle to PCI will be driven with the **LOCK**# pin asserted, even when that particular cycle is not locked. This should not occur, because under lock, peer-to-peer accesses will be internally blocked and the MCH should not be sending any non-locked transactions downstream.

When one PCI bus segment is locked, the other is still free to accept cycles (i.e., that bus is not locked. However, these are not allowed to proceed on PCI Express or the locked PCI segment). Therefore, once the PCI bus is locked, no more cycles will proceed onto PCI Express from the non-locked PCI segment, or from the IOAPICs.

When during the LOCK sequence, any of the locked read commands results in a master or target abort (either on the PCI bus or the internal switch interconnect), then 80333 loses lock after sending a completion packet on PCI Express. In case of a memory write receiving a target or master abort during a LOCK sequence, 80333 only unlocks after the unlock message is received on PCI Express.

- Outbound LOCK is supported by the 80333.
- Inbound LOCK transactions are treated with the LOCK signal ignored. Also locks to internal devices, the SHPC and the IOAPIC are not supported by 80333. See the summary in the table below for summary of 80333 responses to LOCK transactions.

## Table 19.LOCK Transaction Handling in 80333

End Point	Source		
Litter offic	PCI	PCI Express	
SHPC Memory	Ignore <sup>a</sup>	Error Reported <sup>b</sup>	
IOAPIC	Ignore <sup>a</sup>	Error Reported <sup>b</sup>	
PCI	N/A	Forward to PCI w/ LOCK#	
PCI Express	Ignore <sup>a</sup> N/A		

a. Transaction is treated as when it is a normal read or write transaction.

b. For locked reads a response of UR is reported on PCI Express.



## 2.5.2.1 Error Logging and Escalation

80333 supports the PCI Express-advanced error logging capability with bridge extensions and this capability is located in the enhanced PCI Express configuration space. All errors on PCI Express and PCI/X are classified as either correctable or uncorrectable fatal or uncorrectable non-fatal. Refer to PCI Express and bridge specifications for details of the classification. Each of the uncorrectable errors is individually tracked with separate status bits in the capability and can be programmed to generate either a ERR\_FATAL or ERR\_NONFATAL message on PCI Express. Also the first uncorrectable error on each interface (which has its corresponding mask bit cleared) causes a header log and subsequent errors cause only a status bit to be set. Once a header is logged on a given interface, further header logging on that interface is reenabled only after the error status bit corresponding to the first error is cleared by software. On the PCI interface, 80333 also logs the data on the transaction cycle on which a data parity error was detected.

#### 2.5.2.1.1 Bridge and APIC Error Escalation Rules on PCI Express

All uncorrectable errors detected by 80333 on PCI Express, PCI/X and internally are escalated to ERR\_FATAL or ERR\_NONFATAL message on PCI Express only when the corresponding mask bit for that error is cleared in the advanced error register and the ERR\_FATAL or ERR\_NONFATAL message is enabled. ERR\_FATAL and ERR\_NONFATAL messages are both enabled when the SERR enable bit in the PCICMD register is set. ERR\_FATAL/NON\_FATAL messages are also enabled when the corresponding enable bits in the PCI Express capability structure are set.

All uncorrectable data errors forwarded from PCI Express or from internal RAMs to PCI-X would cause the appropriate status bits to be set in the Secondary Status (master data parity error bit for writes/delayed PCI reads) and PCI-XERRUNC\_STS (PCI/X PERR# detected bit) registers (when PERR# is observed on PCI bus), and also cause the appropriate header to be logged on the PCI/X side of the advanced error capability (when enabled per the rules for header logging), but would not cause any error messages to be sent on PCI Express. Error message is only generated on PCI Express on observance of PERR# asserted on PCI, when it was a bus error rather than a forwarded error from 80333.

Note that for error logging and escalation for errors from the PCI Express interface follow the rules per the PCI Express spec.

The following additional rules apply for all bridge error escalation:

- SERR assertion forwarding from PCI/X to PCI Express is enabled either when the SERR enable bit in the bridge control register is set or the mask bit for the SERR detected, in the advanced error reporting register is cleared.
- Posted Write Master aborts on PCI/X can be escalated on to PCI Express either when the Master Abort Mode bit in the Bridge control register is set. (i.e. when the master abort mode bit is clear, posted write master aborts do not generate any PCI Express error message) or the "PCI/X detected master abort" mask bit in the advanced capability register is clear.
- Posted Write Target Aborts and Posted Write Data Parity errors on PCI/X when 80333 is master on PCI/X bus, are escalated to PCI Express when the SERR enable bit is set, even when the advanced error mask bit is set.
- Secondary discard timer expiry will generate an error message on PCI Express either when the "SERR enable on discard timer expiry" bit in the bridge control register is set or the corresponding mask bit in the PCI Express advanced error capability register is clear.

# intel

- Uncorrectable address parity errors (PCI/X interface only and not PCI Express) and uncorrectable data parity errors on PCI Express and PCI/X interfaces generate error messages on PCI Express either when the parity error response enable bit in the PCICMD and BCTL register are set respectively (or) the appropriate mask bits in the PCI Express capability register is clear.
  - Uncorrectable data parity errors on PCI Express include the received poisoned TLP and any data parity errors in the J unit queues.
  - Uncorrectable data parity errors on PCI/X includes data parity errors detected by 80333 as target on PCI/X bus.
- Signaled system error bit is set in PCISTS register anytime an ERR\_FATAL or ERR\_NONFATAL message is signaled with the SERR enable bit is set:

In all cases, 80333 uses the severity bits in the advanced error capability to escalate between FATAL and NONFATAL message.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual PCI Express-to-PCI Bridges

أNt

# 2.6 **Power Management**

# 2.6.1 PCI Express Link Power Management

## 2.6.1.1 Hardware Controlled Active State Power Management

PCI Express defines a Hardware-initiated power management of the PCI Express link called the active state power management. Under hardware control, the link can be put into a low-power L0s link state or an even lower power L1 link state. Latency to recover from L0s back to L0 (active link state) is smaller than that of recovery from L1 to L0. 80333 only supports the PCI Express active state power management link state L0s. 80333 does NOT support optional active state power management link state L1. Active state power management is totally traffic dependent and is not initiated by software. Refer to the PCI Express specifications for more details on the PCI Express active state power management.

## 2.6.1.2 OS-Driven PCI-PM Compatible Power Management

80333 supports PCI Express link states required to implement PCI-PM 1.1 compatible 80333 device states (D0, D3hot, D3cold). These states are PCI Express L0, L1 and L3. When both bridge segments and APICs in 80333 are programmed to D3hot state, PCI Express link in 80333 enters the link state L1. Note that when the APIC functions are hidden via the respective hide bits in the BCNF register, they appear as functions in D3hot state and also appear as functions with L0s active state power management enabled, from the perspective of the PCI Express link power management logic. In legacy OSes where the APIC function could be hidden from the OS, 80333 PCI Express link power state activation logic assumes that when a bridge segment is in D3hot, the associated APIC is also in a D3hot state. When the PCI Express link in 80333 is in L1 because of software driven power management, the only message that can cause the link to come out of L1 is a PME message. Refer to the PCI Express specification for details of the protocol involved in transitioning the link to L1 state.

80333 also supports the PM\_TUN\_OFF/PM\_TO\_ACK protocol to support D3cold/L3 device/link states.

# 2.6.2 PCI Bus Power Management

80333 supports bus state B0 and B3 corresponding to bridge device states of D0 and D3cold. 80333 does stopping the PCI bus clocks when in D3hot state and hence does NOT support the bus state B2.

# 2.6.3 80333 Device Power Management

Each 80333 Bridge and IOAPIC function supports *PCI Bus Power Management Interface Specification*, Revision 1.1 power managements states D0, D3hot and D3cold states. Each function, when in the D3hot state, behaves as follows:

- The function responds to Configuration cycles on PCI Express.
- The function does **NOT** respond to Memory cycles on PCI Express except split completions.
- The function does NOT respond to I/O cycles on PCI Express except split completions.
- The function does **NOT** initiate PCI Express transactions other than split completions (bridge functions only).
- The function does not reset its registers, when programmed to D0 from D3hot. Also, 80333 does not assert **A\_RST#** or **B\_RST#** when in D3hot.
- *Note:* When a bridge segment is in D3hot state, a Hot-Plug event on that segment also creates a PME event to be generated. Software needs to setup the Hot-Plug to generate wake event instead of interrupts when the bridge segment is in D3hot. Refer to *PCI Hot-Plug Specification*, Revision 1.0 section 6.4.1 for details.



# 2.6.4 **Power Management Event Signaling**

80333 has to work in a cabled PCI Express system (split chassis solution) and hence it will support conveying PCI power management events (PME#) over PCI Express via an in-band mechanism. Power management events are generated on behalf of PCI devices that require a change in their power state or internally by the standard Hot-Plug controller in 80333, when it detects a Hot-Plug event with the bridge segment in D3hot state. 80333 does not support any method to "wake" the PCI Express hierarchy before it signals a PME message (i.e., 80333 does not support either the WAKE# side band signal or the in-band PCI Express tone generation mechanism). Waking is needed when the PCI Express link upstream of 80333 is in a non-communicative state with clock and/or power removed. The expectation from 80333 is both ends of the link are fully powered and clocked (i.e., the link is fully communicative) when signaling the PCI power management events.

80333 will support a PME# event pin for conveying power management events on the PCI bus sourced directly by 80333. The PME output from all the PCI devices on the segment are wire-ORed to obtain a composite PME signal which is routed to the 80333. 80333 converts the level-sensitive PME# signal into a PME\_MSG packet on PCI Express. This message will carry the bus number of the PCI bus that caused the PME# assertion. The power manager software needs the bus number information when invoked. Since the bus number of the PCI bus needs to be passed in the PME\_MSG, this scheme only works for waking from the PCI buses directly below 80333. The method envisioned by legacy PME# reporting is illustrated Figure 8 below.



#### Figure 8. 80333 PME# Messaging



The exact mechanism for generating the PME\_MSG packet in 80333 involves sending a message over PCI Express whenever the PME input pin is asserted. Note that this packet would have to carry the bus number of the PCI bus generating the PME#. This means that 80333 would have to construct the requestor ID of the PME\_MSG packet with the secondary bus number register in the corresponding PCI-to-PCI Bridge header space. For PME generated by the standard hot-plug controller, the Requester ID would be: {Primary bus number of 80333:Primary Device Number:2 (function # of B-segment as only B-segment has SHPC)}. There is a chance that PME messages could be lost in PCI Express. Hence 80333 will implement a counter to periodically sample PME# pin and generate a message when PME# is asserted. Refer to the PCI Express spec for the service time-out value (aka sample rate). Current speed value for the sample rate is 100ms.

80333 does not support PME# assertion from the D3cold state and also 80333 does not support Vaux.

intel

# 2.7 **RAS Features and Error Handling**

# 2.7.1 RAS Features

# 2.7.1.1 PCI Express Error Handling

PCI Express link in 80333 is 32-bit CRC protected providing for very high reliability. With a target bit error rate of  $10^{-12}$  the 32-bit CRC combined with the 8b/10b encoding on the serial link, provides for greater than 10 years in MTBF (mean time between failure). The smaller link packets will utilize a 16-bit CRC scheme. PCI Express also provides for a software-transparent recovery from temporary link failures. When received packets are in error, hardware could automatically retransmit the packet.

## 2.7.1.2 PCI Express Hot-Plug

80333 would support a cabled PCI Express environment where 80333 would be part of a remote I/O box and cabled to the system board. 80333 would support the PCI Express Hot-Plug protocol for hot add/removal from the system board, providing for very high serviceability.

## 2.7.1.3 Internal RAM Error Protection

All RAMs within 80333 PCI Express-to-PCI Bridge are parity protected. There is single bit of parity for every 64 bits.

# 2.7.1.4 PCI Error Protection

PCI buses are parity protected. Upper and lower 32 bits on the PCI bus are separately parity protected.

# 2.7.1.5 Advanced Error Reporting

80333 supports the PCI Express advanced error reporting feature, which allows for OS-level error recovery for system errors and also for system debug. 80333 supports the base PCI Express advanced error reporting feature and also supports the bridge-specific extensions for reporting PCI and PCI-X errors.



# 2.7.2 Error Handling

## 2.7.2.1 Completion Required Transaction Termination Translation Between Interfaces

The following sections describe the behavior of the 80333 on both PCI Express and PCI/PCI-X under various termination conditions on PCI/X and PCI Express for completion required requests.

## 2.7.2.1.1 Outbound Transactions – PCI Express to PCI/X

• Immediate Termination on PCI/X

Immediate termination corresponds to a condition where 80333 masters a transaction on PCI or PCI-X and receives an immediate termination response for that transaction. The behavior described for completion required cycles is independent of the setting of the Master Abort Mode bit, and is independent of whether the cycle is exclusive (locked) or not.

### Table 20. Transaction Termination Translation on Immediate Terminations of Completion Required Cycles to PCI/X

PCI/PCI-X Termination	PCI Express Completion
Normal Completion	Successful
Configure and I/O Writes Only; Normal Completion w/ Data Parity Error <sup>a</sup>	UR
Master Abort	UR
Target Abort	CA

a. The data parity error could also be because of internal queue error in 80333. Also in PCI mode PERR# would be sampled asserted on a data parity error.

• Split Termination on PCI/X

Split termination error translations happen when a completion required transaction is originally mastered by 80333 on the PCI/X bus, receives a split termination response and a split completion error message is later received by 80333 on behalf of the original request. Table 21 describes independent behavior of the Master Abort Mode bit and whether or not the cycle is exclusive (locked) or not. Note, when a target or master abort is returned on PCI Express for the first read of an exclusive access, the attached PCI/X bus is not locked. This is of special importance to the completion messages of "byte count out of range, "device specific", and "reserved/illegal" codes. 80333 must not lock its bus on these errors, even though they are not explicitly master or target aborts on the PCI-X interface.

### Table 21. Transaction Termination Translation on Split Terminations of Completion Required Cycles to PCI/X

PCI/PCLX Termination	Message		PCI Express Completion	
	Class	Index		
Successful	0	00H	Successful	
Master Abort	1	00H	UR	
Target Abort	1	01H	CA	
Write Data Parity Error	1	02H	UR	
Byte Count Out of Range <sup>a</sup>	2	00H	UR	
Write Data Parity Error	2	01H	UR	
Device Specific	2 8xH		CA	
Reserved/Illegal	Others		CA	

a. 80333 does not receive this response since it splits outbound memory read transactions on an ADB. 80333 aliases this response to a UR.



• Internal Switch Terminations

### Table 22. Transaction Termination Translation on Immediate Terminations of Completion Required Cycles to Internal Switch

Switch Termination	PCI Express Completion
Successful	Completions are generated from the end devices
Master Abort	UR

### 2.7.2.1.2 Inbound Transactions – PCI/X-to-PCI Express

#### • Split Termination on PCI Express

Table 23 provides the translation for split completion errors received on PCI Express and how they are translated to PCI/X.

### Table 23. Transaction Termination Translation on Terminations of Completion Required Cycles to PCI Express

PCI Express Termination	PCI/X Completion		
Successful (SC)		Successful	
		Memory Reads: PCI Target Abort <sup>a</sup>	
		I/O Reads: PCI Target Abort	
	MAM = 1	I/O Writes: PCI Target Abort	
		Configuration Reads: PCI Target Abort	
Unsupported request (UR)/Config-		Configuration Writes: PCI Target Abort	
uration Retry (Configure transac-	MAM = 0	Memory Reads: PCI Return all Fs	
tions only)		I/O reads: PCI Return all Fs	
		I/O writes: Normal Completion	
		Configure reads: PCI Return all Fs	
		Configure writes: Normal Completion	
	PCI-X Split Master Abort <sup>b</sup>		
СА	PCI Target Abort <sup>a</sup>		
0.1	PCI-X Split Target Abort <sup>b</sup>		

a. 80333 streams data to the point of error and then signals a target abort.

b. The 80333 will issue a Split Completion Error Message with either Master Abort or Target Abort for the remaining completion sequence when an abort is detected on PCI Express/PCI. When several bytes of data returned successfully from PCI Express/PCI, and have not yet been sent back on PCI-X, when the abort is detected on from PCI Express/PCI the 80333 will stream data to the point of error (rounded to the ADB boundary), and generate the Split Completion Error Message after that.

## 2.7.2.1.3 Error Logging and Escalation

80333 supports the PCI Express-advanced error logging capability with bridge extensions and this capability is located in the enhanced PCI Express configuration space. All errors on PCI Express and PCI/X are classified as either correctable or uncorrectable fatal or uncorrectable non-fatal. Refer to PCI Express and bridge specifications for details of the classification. Each of the uncorrectable errors is individually tracked with separate status bits in the capability and can be programmed to generate either a ERR\_FATAL or ERR\_NONFATAL message on PCI Express. Also the first uncorrectable error on each interface causes a header log and subsequent errors cause only a status bit to be set. Once a header is logged on a given interface, further header logging on that interface is reenabled only after the error status bit corresponding to the first error is cleared by software. On the PCI interface, 80333 also logs the data on the transaction cycle on which a data parity error was detected.

- Bridge and APIC Error Escalation Rules on PCI Express:
  - All uncorrectable errors detected by 80333 on PCI Express, PCI/X and internally are escalated to ERR\_FATAL or ERR\_NONFATAL message on PCI Express only when the corresponding mask bit for that error is cleared in the advanced error register and the ERR\_FATAL or ERR\_NONFATAL message is enabled. ERR\_FATAL and ERR\_NONFATAL messages are both enabled when the SERR enable bit in the PCICMD register is set.
  - ERR\_FATAL/NON\_FATAL messages are also enabled when the corresponding enable bits in the PCI Express capability structure are set.
  - The following additional rules apply for all bridge error escalation:
  - SERR assertion forwarding from PCI/X to PCI Express is enabled either when the SERR enable bit in the bridge control register is set or the mask bit for the SERR detected, in the advanced error reporting register is cleared.
  - Posted Write Master aborts on PCI/X can be escalated on to PCI Express either when the Master Abort Mode bit in the Bridge control register is set. (i.e. when the master abort mode bit is clear, posted write master aborts do not generate any PCI Express error message) or the "PCI/X detected master abort" mask bit in the advanced capability register is clear.
  - Secondary discard timer expiry will generate an error message on PCI Express either when the "SERR enable on discard timer expiry" bit in the bridge control register is set or the corresponding mask bit in the PCI Express advanced error capability register is clear.
  - Uncorrectable address parity errors (PCI/X interface only and not PCI Express) and uncorrectable data parity errors on PCI Express and PCI/X interfaces generate error messages on PCI Express either when the parity error response enable bit in the PCICMD and BCTL/Sc\_PCICMD register are set respectively (or) the appropriate mask bits in the PCI Express capability register is clear.
  - Uncorrectable data parity errors on PCI Express include the received poisoned TLP and any data parity errors in the PCI Express interface unit queues.
  - Uncorrectable data parity errors on PCI/X includes data parity errors detected by 80333 as target on PCI/X bus.

The signaled system error bit is set in the PCISTS register anytime an ERR\_FATAL or ERR\_NONFATAL message is signaled with the SERR enable bit is set.



## 2.7.2.1.4 Summary of Error Reporting

## Table 24.Summary of Error Reporting (Sheet 1 of 3)

	PCI Express-PCI Bridge Uncorrectable Error Logging/Escalation				
#	Error	Comments	Base Error Severity	Advanced Error Log	
Mas	ter and Target Aborts		·		
	Primary - Secondary	Forwarding			
1	Received Immediate Master Abort	80333 was master on PCI/X bus and received an master abort (both posted and non-posted transactions)	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	
2	Received Split Completion Master Abort Message	80333 received a split completion error message with master abort status (non-posted transactions only)	Uncorrectable: Non-Fatal	None	
3	Immediate Target Abort	80333 was master on PCI/X bus and received an target abort (both posted and non-posted transactions)	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	
4	Received Split Completion Target Abort Message	80333 received a split completion error message with Target abort was master on PCI/X bus and received an master abort (both posted and non-posted transactions)	Uncorrectable: Non-Fatal	None	
5	Other Split Completion Error Message Errors	Includes Device-Specific, Write data parity error, Byte Count-out-of-Range (non-posted transactions only)	Uncorrectable: Non-Fatal	None	
6	Split Completion Master Abort	80333 was master on PCI/X bus for a completion and received a master abort	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	
7	Split Completion Target Abort	80333 was master on PCI/X bus for a completion and received a Target abort	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	
	Secondary - Primary	Forwarding			
8	Received Split Completion Master Abort Status	80333 received a split completion status of master abort from PCI Express (non-posted transactions only)	Uncorrectable: Non-Fatal	None	
9	Received Split Completion Completion Abort Status	80333 received a split completion status of target abort from PCI Express (non-posted transactions only)	Uncorrectable: Non-Fatal	None	
10	Received Unexpected Split Completions	80333 received a split completion status from PCI Express that does not match any outstanding inbound requests in 80333	Uncorrectable: Non-Fatal	Error Log in PCI Express portion of the advanced error log	
	Internal 80333 Terminations				
11	Internal Switch Master Abort (both posted and non-posted transactions)	80333 received transaction that did not match the address ranges of any function within 80333 (includes memory, I/O) or did not match the device/function of any function within 80333 or the bus number ranges covered by the bridge segments.	Uncorrectable: Non-Fatal	Error Log in PCI Express portion of the advanced error log	
12	CA Signaled by SHPC	SHPC received a memory read that was greater than a DWORD in length.	Uncorrectable: Non-Fatal	Error Log in PCI Express portion of the advanced error log	
13	CA Signaled APIC	APIC received a memory read or write (optional) that was greater than a DWORD in length	Uncorrectable: Non-Fatal	None	

March 2005

94



# Table 24.Summary of Error Reporting (Sheet 2 of 3)

	PCI Express-PCI Bridge Uncorrectable Error Logging/Escalation				
#	Error	Comments	Base Error Severity	Advanced Error Log	
Misc	ellaneous Errors				
14	PCI SERR# Detected	Bridge detected SERR# asserted on PCI/X (Asserted by some other agent. 80333 could even be driving the transaction that caused this error)	Uncorrectable: Fatal	Error Log in Bridge extension fields	
15	Delayed Transaction Timer Expiry	Valid for secondary interface in PCI mode only. When so, when the secondary read discard timer expires an error is signaled on the PCI Express, when enabled	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	
Add	ress Parity Errors				
16	Uncorrectable PCI Address Parity Error	Bridge is target of a transaction where an uncorrectable address parity error happened i.e. the decode logic of the bridge indicates a HIT and there is an uncorrectable address parity error. In PCI mode, 80333 will claim this transaction and forward it to PCI Express as when no error happened (regardless of the parity error response bit). In PCI-X mode, 80333 will issue a target abort. Error escalation on this is dependent on the parity error response bit in the bridge control register and also the address parity error severity bit in the PCI Express advanced error bridge extension fields.	Uncorrectable: Fatal	Error Log in Bridge extension fields	
		Bridge is not target of a transaction where an uncorrectable address parity error happened i.e. the decode logic of the bridge indicates a MISS and there is an uncorrectable address parity error. 80333 will not claim this transaction though error logging/escalation proceeds as when it was claimed.	Uncorrectable: Fatal	Error Log in Bridge extension fields	
17	Uncorrectable PCI Address Parity Error	Bridge is target of a transaction where an uncorrectable attribute parity error happened i.e. the decode logic of the bridge indicates a HIT and there is an uncorrectable attribute parity error. Error escalation on this is dependent on the parity error response bit in the bridge control register and also the address parity error severity bit in the PCI Express advanced error bridge extension fields.	Uncorrectable: Fatal	Error Log in Bridge extension fields	
Data	Parity Errors				
	Primary - Secondary	Data Forwarding			
		Data was received from PCI Express bad (includes requests and completions) and forwarded to PCI with bad parity	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	
17	PCI Express Poisoned TLP received	Data was received from PCI Express bad and forwarded to SHPC unit with bad parity. Error escalation is handled per received poisoned TLP condition	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	
		Data was received from PCI Express bad and forwarded to APIC unit with bad parity. Error escalation is handled per received poisoned TLP condition	Uncorrectable: Non-Fatal	No Advanced Error Log	
18	PCI Express Poisoned TLP received	Parity error detected when bridge acting as a master on PCI while forwarding a good PCI Express transaction or when forwarding a transaction not corrupted by an internal 80333 queues.	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields	



## Table 24.Summary of Error Reporting (Sheet 3 of 3)

	PCI Express-PCI Bridge Uncorrectable Error Logging/Escalation							
#	Error	Comments	Base Error Severity	Advanced Error Log				
	Secondary - Primary	Data Forwarding						
19	PCI PERR# signaled	80333 is the target of the memory, I/O or Configure write or 80333 is the master of a read that is completed immediately by target and 80333 detects a parity error in the data. In both cases, 80333 forwards data to PCI Express or peer PCI with bad parity signal.	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields				
	Internal 80333 Data Queue Parity Error							
20	Data Parity Errors in bridge Queues in either direction (Includes internal queue errors towards SHPC but not APIC)	Bridge queues errors found on data that are forwarded from PCI Express to PCI or vice-versa. These are captured with the "bridge data error" bit in the Bridge extension field. Note that this bridge data error bit captures internal queue errors towards SHPC also.	Uncorrectable: Non-Fatal	Error Log in Bridge extension fields				
21	Data Parity Errors in bridge Queues to APIC)	Bridge queue data errors on transactions towards APIC. These errors are escalated using the UR control bits in the APIC	Uncorrectable: Non-Fatal	No Advanced Error Log				

# 2.7.2.2 Data Poisoning

80333 provides for bad data to be propagated from one interface to the other with a poison bit. There is one poison bit for every 64 bits inside 80333. Bad data received on PCI Express will be forwarded to PCI where it will be driven on PCI with the parity bit flipped. Similarly bad data from 80333 will be forwarded to PCI Express as a poisoned TLP. A similar scheme applies when the source of the bad parity is the internal SRAMs.

# 2.8 **BIOS/Device Driver Recommendations**

This section is intended to collect in one place all requirements on BIOS and device driver that are imposed by the 80333 architecture. This section would also collect the recommended settings for various control registers in 80333 for various modes of 80333 operation internally and on the bus.

# 2.8.1 Standard Hot-Plug Controller Initialization Requirements

The standard Hot-Plug controller in 80333 relies on software-based (BIOS/firmware/driver) initialization. 80333 does not provide a way to do hardware-based SHPC initialization (i.e., 80333 does not support pin straps or a PROM solution for initializing the SHPC registers). Nor does 80333 SHPC support an automatic initialization of the PCI bus at reset. The standard Hot-Plug controller in 80333 is reset and requires an initialization whenever the PCI Express link in 80333 goes through a reset. This initialization applies to two aspects of Hot-Plug initialization

- Initializing the hardware-init registers in the SHPC
- Initializing the PCI bus

The PCI Express link in 80333 goes through a reset under three conditions

- One of 80333 reset pins gets asserted (PWRGD, RSTIN#).
- A reset is detected on PCI Express link because the north component is in going through a reset.
- A reset is detected on PCI Express when software initiates a reset of PCI Express link from the north component.

Under each of these conditions, it is expected that the SHPC (registers and the bus) is initialized by the host software. During power-on, this initialization routine is executed as part of the POST code. In scenarios where 80333 is hot-inserted onto PCI Express (80333 in an add-in card or in a split-chassis RemotePCI solution) or when 80333 enters into reset because of a software initiated reset in the superior bridge, running in the OS context, the SHPC registers need to be initialized. For ACPI-enabled systems and OS, this initialization routine can be included as part of the platform firmware and the OSHP method, described in the *PCI Standard Hot-Plug Controller and Subsystem Specification*, Revision 1.0, could be used to initialize the SHPC registers. For Non-ACPI systems and OSes, the initialization needs to be handled by the platform Hot-Plug driver appropriately.



# 2.8.2 80333-Specific Bridge Register Initialization

80333 includes a device-specific registers that allow for control of 80333 behavior, both internally and externally. Examples of these device-specific registers are – arbiter control register, prefetch control register etc. Depending on the usage context of 80333, these registers might need programming to a different value, than the reset-default, every time 80333 goes through a component reset. 80333 goes through a component reset when

- One of 80333s reset pins gets asserted (PWRGD, RSTIN#).
- A reset is detected on PCI Express link because the north component is in going through a reset.
- A reset is detected on PCI Express when software initiates a reset of PCI Express link from the north component.

Under each of these conditions, it is expected that the host software initializes these device-specific registers. During power-on, this initialization routine gets executed as part of the POST code. In scenarios where 80333 is hot-inserted onto PCI Express (80333 in an add-in card or in a split-chassis RemotePCI solution) or when 80333 enters into reset because of a software initiated reset in the superior bridge, running in the OS context, the 80333-specific registers need to be initialized. For ACPI-enabled systems and OS, this initialization routine can be included as part of the platform firmware and the OSHP method, described in the *PCI Standard Hot-Plug Controller and Subsystem Specification*, Revision 1.0, could be used to initialize the 80333-specific device registers along with the SHPC registers. For Non-ACPI systems and OSes or when the SHPC is disabled, the initialization needs to be handled by the platform Hot-Plug driver appropriately.

# 2.8.3 Peer-to-Peer Memory Read Support

When peer-to-peer memory reads are enabled in 80333 via the peer memory read enable bit in the Bridge Configuration register, the software needs to turn-off prefetching on memory read commands by setting the PP bits in Prefetch Control Register to 01.

# 2.8.4 Lockout Bit Usage Guidelines

The secondary lockout bit is set coming out of a component reset (cold boot or hot-insertion). When set, the lockout bit causes retries of all configuration transactions on the PCI bus. The primary side software clears this bit when it has completed the 80333 configuration process. Since retries happen on the PCI bus when this is set, this could hang the bus in hot-insertion cases where 80333 is being polled by software on behalf of a wake event from another card. So the primary board software may have to complete configuration of 80333 early in the wake-up process and clear the lockout bit, to avoid holding the backplane.

# 2.8.5 Performance Considerations

# 2.8.5.1 Non-QWORD Aligned Accesses Outbound

80333 does not handle Non-QWORD-aligned accesses (both posted and non-posted) outbound in a performance-optimal way. When possible, drivers need to be written to access QWORD-aligned addresses downstream to get best performance.

*Note:* Only one segment is active. PCI/X requests are 1KByte in length for both reads and writes.



# 2.9 Configuration Registers

This section will discuss all the 80333 registers pertaining to the PCI Express-to-PCI bridge operation.

# 2.9.1 Bit Attribute Definition

The nomenclature used for describing the bit attributes in Table 25 is used throughout this section.

## Table 25.Bit Attribute Definitions

Mnemonic	Attribute
RO	Read-Only. This bit cannot be altered by software. This bit could be hardwired to return a fixed value all the time or set by hardware on an event.
PR	Preserved: Reserved for future RW implementations; software must preserve value read for writes to bits. 80333 hardware implements these bits as read-only 0s.
RV	Reserved: Reserved for future implementations; software must use 0 for writes to bits. 80333 hardware implements these bits as read-only zeros
ROS	Read-Only Sticky. Register bits are read-only and cannot be altered by software. Bits are not cleared by reset and can only be reset with the <b>PWRGD</b> reset condition.
RZSet	Read 0 to Set. Reading this bit when the current value of the bit is a 0, will flip the bit to a 1. Software has to write a 1 to clear this bit. Writing a zero has no affect.
RW	Read-Write. Software can do a full read and write of this bit.
RW1Set	Read and write one to set. Software needs to write a 1 to set this bit. Writing a 0 has no effect on this bit. Software can clear this bit through separate RWC bit or hardware will reset it.
RWS	Read-Write and Sticky. Software can read and write from this bit and only a <b>PWRGD</b> reset can reset this bit.
RWC	Read and Write One to Clear. Software needs to write a 1 to this bit to clear it when set. Write of 0 has no effect on this bit.
RWCS	Read and Write One to Clear and Sticky through reset. Software needs to write a 1 to this bit to clear it when set. Write of 0 has no effect on this bit. Only a <b>PWRGD</b> reset can reset this bit.
RWO	Read and Write Once. Software can read this bit but can write to it only once. Once written, this bit retains its content and only a reset can clear its contents. Software must update all RWO (and RWOS) bits within a DWORD address with a single DWORD write.
RWOS	Read and Write Once Sticky. Software can read this bit but can write to it only once. Once written, this bit retains its content and only a <b>PWRGD</b> reset can clear its contents. Software must update all RWO (and RWOS) bits within a DWORD address with a single DWORD write.
Strap	Strap: Read-only register whose power on default is based on sampling a strap pin at the rising edge of <b>PWRGD</b>
WT	Write Transient. This bit is always read as a zero. Write of 1 to this bit cause other side effects that are specific to every WT bit.
RWT	Read Write Transient. Each RWT bit is associated with a WT bit and the RWT bit is writable when the associated WT bit is a 1.

# 2.9.2 Accesses to Reserved Registers

Software must not attempt to write to the registers that are marked 'reserved'. Writing to these registers would yield undetermined results. Reads of these registers could yield either value. Note that individual register bits that are marked 'reserved' can be treated as such i.e. behavior of those bits is as per the attribute definition for that bit.



# 2.9.3 80333-Specific Bridge Register Initialization

80333 includes a device-specific registers that allow for control of 80333 behavior, both internally and externally. Examples of these device-specific registers are - arbiter control register, prefetch control register etc. Depending on the usage context of 80333, these registers might need programming to a different value, than the reset-default, every time 80333 goes through a component reset. 80333 goes through a component reset when

- One of 80333 reset pins gets asserted (PWRGD, RSTIN#)
- A reset is detected on PCI Express link because the north component is in going through a reset
- A reset is detected on PCI Express when software initiates a reset of PCI Express link from the north component

Under each of these conditions, it is expected that the host software initializes these device-specific registers. During power-on, this initialization routine gets executed as part of the POST code. In scenarios where 80333 is hot-inserted onto PCI Express (80333 in an add-in card or in a split-chassis RemotePCI solution) or when 80333 enters into reset because of a software initiated reset in the superior bridge, running in the OS context, the 80333-specific registers need to be initialized. For ACPI-enabled systems and OS, this initialization routine can be included as part of the platform firmware and the OSHP method, described in the SHPC spec, could be used to initialize the 80333-specific device registers along with the SHPC registers. For Non-ACPI systems and Oses or when the SHPC is disabled, the initialization needs to be handled by the platform's Hot-Plug driver appropriately.



# 2.9.4 80333 - Specific Bridge Register Initialization

80333 includes a device-specific registers that allow for control of the 80333 behavior, both internally and externally. Examples of these device-specific registers are:

- arbiter control register.
- prefetch control register.
- etc.

Depending on the usage context of 80333, these registers might need programming to a different value, than the reset-default, every time 80333 goes through a component reset. 80333 goes through a component reset when

- One of the 80333 reset pins gets asserted (PWRGD, RSTIN#).
- A reset is detected on PCI Express link because the north component is in going through a reset.
- A reset is detected on PCI Express when software initiates a reset of PCI Express link from the north component.

Under each of these conditions, it is expected that the host software initializes these device-specific registers. During power-on, this initialization routine gets executed as part of the POST code. In scenarios where 80333 is hot-inserted onto PCI Express (80333 in an add-in card or in a split-chassis RemotePCI solution) or when 80333 enters into reset because of a software initiated reset in the superior bridge, running in the OS context, the 80333-specific registers need to be initialized.

For ACPI-enabled systems and OS, this initialization routine can be included as part of the platform firmware and the OSHP method, described in the SHPC specification, could be used to initialize the 80333-specific device registers along with the SHPC registers. For Non-ACPI systems and OSes or when the SHPC is disabled, the initialization needs to be handled by the platform Hot-Plug driver appropriately.



# 2.9.5 Configuration Registers

The bridge configuration space follows the standard PCI Express-to-PCI Bridge configuration space format. Refer to the *PCI Express-to-PCI Bridge Specification*, Revision 1.0 for details on the format. Each bridge contains an identical set of registers as described in this section for its respective PCI segment. Table 26 shows the 80333 configuration registers and their address byte offset values. Figure 9 gives the capabilities that 80333 supports.<sup>1</sup>

## Figure 9. 80333 Capabilities



<sup>1.</sup> The SHPC Capability is only visible when Hot-Plug is enabled in 80333 with **B\_HSLOT[3]** reset strap sampled high at **PWRGD**.



				Byte Offset		Byte Offset
DID VID						80h
PSTS PCICMD		04h		84h		
C	lass Code (CO	C)	REVID	08h		88h
	HDR	PMLT	CLS	0Ch		8Ch
	SUDC			10h		90h
	SHEC	_DAN		14h		94h
SMLT		BNUM		18h		98h
SS	TS	IO	BL	1Ch		9Ch
	Μ	BL		20h		A0h
	PN	1BL		24h		A4h
	PME	3U32		28h		A8h
	PME	3L32		2Ch	Reserved	ACh
	IOBI	_U16		30h		B0h
	Reserved		CAPPTR	34h		B4h
	Rese	erved		38h		B8h
BCI	RL	IN		3Ch		BCh
PCLKC	MII	BC		40h		CON
EXP_	CAP	EXP_NXTP	_CAPID	44h		C4h
	EXP_	DCAP		48h		C8h
EXP_	DSTS	EXP_	DCTL	4Ch		CCh
	EXP_	LCAP		50h		D0h
EXPL	DSTS	EXP_	LCTL	54h		D4h
	Rese	erved		58h	PX_SSIS PX_NXIP PX_CAPID	D8h
MSI_	_MC	MSI_NXTP	MSI_CAPID	5Ch	PX_BSTS	DCh
	MSI	_MA		60h		
				64h	PX_DSTC	E4h
				68n		E8n
PM_CAP PM_NXTP PM_CAPID						
	- IVI_DOE		000	7/h	Reserved	
	SHPC	SHPC	SHPC	/ 411		
SHPC_STS	_DWSEL	_NXTP	_CAPID	78h		F8h
SHPC_DWORD					BINIT	FCh

# Table 26. Configuration Space – Legacy Region



# Table 27. PCI Express Extended Configuration Space

Registe	r DWORD	Byte Offset		
EXP_AE	RR_CAPID	100		
ERRU	NC_STS	104		
ERRUI	NC_MSK	108		
ERRU	NC_SEV	10C		
ERRC	OR_STS	110		
ERRC	OR_MSK	114		
ADVE	RR_CTL	118		
		11C		
HDR	2106	120		
	(	124		
		128		
	PCI-XERRUNC_STS	12C		
Reserved	PCI-XERRUNC_MSK	130		
Reserved	PCI-XERRUNC_SEV	134		
	PCI-XERRUNC_PTR	138		
		13C		
PCI-XH		140		
		144		
	148			
PCI-X	14C			
PCI-XER	RLOGCTL	154		
		158		
Res	served	15C		
		160		
		164		
ARB_CNTRL	Reserved	168		
Res	served	16C		
Reserved	SSR	170		
Res	erved	174 - 313		
PWRBGT_INFO[0:23]				
Reserved				



# 2.9.5.1 Identifiers - ID (Offset 00)

Contains the vendor and device identifiers for software.

### Table 28. Identifiers - ID

Bits	Туре	Reset		Description
21.10	RO	Α	В	Device ID (DID): Indicates what device number was assigned by the latel for 80222
51.10		0370H	0372H	Device in (DD). Indicates what device number was assigned by the interior 60555
15:0	RO	8086H		Vendor ID (VID): 16-bit field that indicates that Intel is the vendor.

# 2.9.5.2 Command - PCICMD (Offset 04)

This controls how the device behaves on the primary interface (PCI Express). As this component is a bridge, additional command information is located in a separate register called "Bridge Control" located at offset 3E.

Table 29.	Command - PCICMD	(Sheet 1 of 2	)
-----------	------------------	---------------	---

Bits	Туре	Reset	Description			
15:11	RV	00H	Reserved			
10	RW	RW	0	INTx Mask: This bit disables the SHPC from asserting IRQ[23]# wire to the I/OxAPIC. This bit is valid only when the MSI is disabled i.e. the MSI enable bit in the MSI_MC register is a zero. 0 = enables the assertion of its IRQ[23]# signal to the IOAPIC		
			1 = disables the assertion of its IRQ[23]# signal.			
			When IRQ[23]# is already asserted when this bit is set, it must be deasserted.			
9	RO	0	Fast Back-to-back enable (FBE): This bit has no meaning on PCI Express. It is hardwired to '0'.			
	RW		SERR# Enable (SEE): Controls the enable for PCI-compatible SERR# logging/reporting on PCI Express (along with the PCISTS[14] bit).			
8		RW	0	<ul> <li>0 = Disable reporting errors</li> <li>1 = Enable reporting of non-fatal and fatal errors to the Root Complex</li> <li>NOTE: Errors are reported when enabled either through this bit or through the PCI Express specific bits in the Device Control Register.</li> </ul>		
7	RO	0	Wait Cycle Control (WCC): Reserved.			
6		0	Parity Error Response Enable (PERE): Controls the response to data parity errors forwarded from the PCI Express interface (either internal queue errors in the PCI Express cluster or a poisoned TLP received from PCI Express)/peer PCI on read completions and inbound Writes with poisoned data.			
0	NVV	RVV	RVV		0	<ul> <li>0 = Ignore these errors on PCI Express/peer-PCI interface and NOT set the MDPD bit (but still follow the rules for poisoned tlp received/Bridge internal error reporting rules, etc.)</li> <li>1 = Reports read completion data parity errors on PCI Express and sets the MDPD bit in the status register.</li> </ul>
5	RO	0	VGA Palette Snoop Enable (VGA_PSE): Reserved.			
4	RO	0	Memory Write and Invalidate Enable (MWIE): Memory write and invalidate transactions are not generated, as PCI Express does not have a corresponding transfer type.			
3	RO	0	Special Cycle Enable (SCE): Reserved			



Table 29. C	ommand - PCICMD	(Sheet 2 of 2)
-------------	-----------------	----------------

Bits	Туре	Reset	Description
	RW	0	Bus Master Enable (BME): Controls the ability to act as a master on PCI Express when forwarding memory transactions from PCI or when generating MSI transactions on behalf of the SHPC.
2			<ul> <li>Do not respond to any memory transactions on the PCI interface and stops issuing new requests on PCI Express (includes both internally generated MSI or any forwarded transaction from the PCI interface).</li> <li>Process transactions normally.</li> </ul>
			<b>NOTE:</b> This bit does not stop completions on PCI Express from being issued. Software must guarantee that when this bit is set that all inbound posted transactions are flushed in the bridge segment. Otherwise delayed completions (e.g. configuration read completions) could be stuck behind a posted write and cannot proceed from PCI to PCI Express.
	RW	RW 0	Memory Space Enable (MSE): Controls the response as a target to memory transactions from primary or secondary.
1			<ul> <li>0 = All memory transactions targeting secondary are master aborted and all memory transactions from secondary to primary are claimed.</li> <li>1 = Primary to secondary forwarding and vice-versa happen per normal rules for memory forwarding.</li> </ul>
0	RW		I/O Space Enable (IOSE): Controls the response as a target to I/O transactions from primary or secondary.
		RW	0



# 2.9.5.3 Primary Device Status - PSTS (Offset 06)

For the writable bits in this register, writing a '1' will clear the bit. Writing a '0' to the bit will have no effect.

## Table 30.Primary Device Status - PSTS (Sheet 1 of 2)

Bits	Туре	Reset	Description
15	RWC	0	<ul> <li>Detected Parity Error (DPE): indicates a data parity error (tailer received) was detected from PCI Express or peer PCI segment (writes or read completions). This bit gets set even when the Parity Error Response Enable bit (bit 6 of the command register) is not set.</li> <li>0 = No error.</li> <li>1 = Trailer received or Data Parity Error detected.</li> <li>NOTE: SRAM soft errors on the data path from the PCI Express interface to the individual PCI bridge interfaces within 80333 are also tracked using this bit.</li> </ul>
14	RWC	0	Signaled System Error (SSE): Indicates when ERR_FATAL or ERR_NONFATAL messages are sent to the Root Complex (and the SERR enable bit (bit 8 in PCICMD register) is set) for errors that would have caused an SERR# assertion on conventional PCI. 0 = No error. 1 = ERR_FATAL or ERR_NONFATAL message sent.
13	RWC	0	Received Master Abort (RMA): Indicates when acting as master on PCI Express a completion packet with UR-EC status is received. 0 = No error. 1 = UR-EC status received in completion packet.
12	RWC	0	<ul> <li>Received Target Abort (RTA): Indicates when acting as master on PCI Express a completion packet with CA status is received.</li> <li>0 = No error.</li> <li>1 = CA Status received in completion packet.</li> </ul>
11	RWC	0	Signaled Target Abort (STA): Indicates a completion packet with CA status was generated on PCI Express (either forwarded from PCI interface or internally signaled by SHPC). 0 = No error. 1 = CA Status generated in completion packet.
10:9	RO	00	DEVSEL# Timing (DVT): These bits have no meaning on PCI Express. Fast decode timing is reported.
8	RWC	0	Master Data Parity Error Detected (MDPD): Indicates a completion packet from PCI Express/peer PCI/X from a previous read request has been received, and a data parity error (either from a bus interface error or internal data path error in 80333) is detected, and the Parity Error Response Enable bit in the Command Register (offset 04h, bit 6) is set. This bit is also set when a bridge segment forwards a write with poisoned parity from PCI/X to PCI Express/peer PCI/X and the Parity Error Response Enable bit in the Command (offset 04h, bit 6) is set. 0 = No error. 1 = Data Parity Error in completion packet or Tailer received. NOTE: SRAM soft errors on the data path from the internal PCI Express interface to the individual PCI bridge interfaces within 80333 are also tracked using this bit.
7	RO	0	Fast Back-to-Back Capable (FBC): This bit has no meaning on PCI Express.
6	RV	0	Reserved
5	RO	0	66 MHz Capable (C66): This bit has no meaning on PCI Express.


## Table 30. Primary Device Status - PSTS (Sheet 2 of 2)

Bits	Туре	Reset	Description
4	RO	1	Capabilities List Enable (CAPE): Indicates the capabilities pointer in the bridge. Offset 34H indicates the offset for the first entry in the linked list of capabilities. 0 = No capabilities enabled. 1 = Capabilities enabled and accessible from 34H.
3	RO	0	<b>Interrupt Status:</b> Read-only bit reflects the SHPC interrupt state, when the interrupt is generated via the IRQ[23]# wire (not via MSI). Only when the INTx mask bit in the command register is a 0 and this Interrupt Status bit is a 1, and MSI is disabled does SHPC assert the IRQ[23]# signal to the IOAPIC. Setting the INTx mask bit to a 1 has no effect on this bit setting.
2:0	RV	0H	Reserved.



## 2.9.5.4 Revision ID - REVID (Offset 08)

Revision ID Register.

## 2.9.5.5 Class Code - CC (Offset 09)

This contains the class code, sub class code, and programming interface for the device.

Bits	Туре	Reset	Description
23:16	RO	06h	Base Class Code (BCC): The value of "06h" indicates that this is a bridge device.
15:8	ROS	04h	Sub Class Code (SCC): 8-bit value that indicates this is a PCI-to-PCI Bridge.
7:0	RO	00H	Programming Interface (PIF): Indicates that this is standard (non-subtractive) PCI-to-PCI Bridge.

#### Table 31.Class Code - CC



## 2.9.5.6 Cache Line Size - CLS (Offset 0C)

This indicates the cache line size of the system.

Bits	Туре	Reset	Description
7:0	RW	00H	<ul> <li>Cache Line Size (CLS): Specifies the system cacheline size in units of DWORDS.</li> <li>'08H', represents a 32-byte line (8 DWORDs)</li> <li>'10H' represents a 64-byte line</li> <li>'20H' represents a 128-byte line.</li> <li>Any value outside this range will default to a 64-byte line. When creating read requests to PCI Express, this value is used to partition speculative PCI read requests on cache line aligned boundaries. This register has no other effect on 80333.</li> </ul>

#### Table 32.Cache Line Size - CLS

## 2.9.5.7 Primary Master Latency Timer - PMLT (Offset 0D)

This register does not apply to PCI Express, and is maintained as R/W for software compatibility.

#### Table 33. Primary Master Latency Timer - PMLT

Bits	Туре	Reset	Description
7:3	RO	0 0000	Time Value (TV): not applicable for PCI Express.
2:0	RV	000	Reserved.

## 2.9.5.8 Header Type - HEADTYP (Offset 0E)

This register determines how the rest of the configuration space is laid out.

#### Table 34.Header Type - HEADTYP

Bits	Туре	Reset	Description
7	RO	1	Multi-function device (MFD): Reserved as '1' to indicate the bridge is a multi-function device.
6:0	RO	000 0001	Header Type (HTYPE): Defines the layout of addresses 10H through 3FH in configuration space. Reads as 01h to indicate that the register layout conforms to the standard PCI-to-PCI Bridge layout.



## 2.9.5.9 SHPC 64-Bit Base Address Register - SHPC\_BAR (Offset 10)

64-bit SHPC base address register, used to access the SHPC working register set. This BAR appears in the configuration space only when the standard Hot-Plug controller is enabled (the B\_HSLOT[3] is sampled as 1 (high) at **PWRGD** asserting edge).

#### Table 35. SHPC 64-Bit Base Address Register - SHPC\_BAR

Bits	Туре	Reset	Description
SHPC Ena	abled: B_H	SLOT[3] is sa	mpled as 1 (high) at PWRGD asserting edge
63:8	RW	0	Base Address: These bits are read write and are used by BIOS to understand that SHPC needs 4 KBytes of memory space and then write a valid 4 KByte aligned base address.
7:4	RV	0	Reserved
3	RO	0	Prefetchable: This bit is a read-only 0 to indicate that this register needs to be mapped into the non-prefetchable space.
2:1	RO	10	Type: These bits are read-only with a reset default of 10, indicating that this register can map anywhere in the 64-bit memory space.
0	RO	0	Memory Space Indicator (MSI): This bit is a read-only 0 indicating that this BAR maps into memory space.
SHPC Disabled: B_HSLOT[3] is sampled as 0 (low) at PWRGD asserting edge			
63:0	RV	0000 0000 0000 0000H	Reserved.

#### 2.9.5.10 Bus Numbers - BNUM (Offset 18)

This contains the primary, secondary, and maximum subordinate bus number registers.

#### Table 36.Bus Numbers - BNUM

Bits	Туре	Reset	Description
23:16	RW	00H	Subordinate Bus Number (SBBN): Indicates the highest PCI bus number downstream of this bridge. Any type one configuration cycle on PCI Express with bus number greater than the secondary bus number and less than or equal to the subordinate bus number will be forwarded as a type one configuration cycle to the secondary PCI bus.
15:8	RW	00H	Secondary Bus Number (SCBN): Indicates the bus number of PCI to which the secondary interface is connected. Any type one configuration cycle matching this bus number will be translated to a Type 0 configuration cycle and run on the PCI bus.
7:0	RW	00H	Primary Bus Number (PBN): Indicates the bus number of PCI Express. Any Type 1-configuration cycle with a bus number less than this number will not be accepted by this bridge (in other words, it still may match the other bridge). 80333 uses this register for forming the requester ID on PCI Express.



## 2.9.5.11 Secondary Master Latency Timer - SMLT (Offset 1B)

This timer controls the amount of time that the 80333 will continue to burst data on its secondary interface. The counter starts counting down from the assertion of FRAME#. When the grant is removed, then the expiration of this counter will result in the de-assertion of FRAME#. When the grant has not been removed, then the 80333 may continue ownership of the bus. Secondary master latency timer default value should be 64 in PCI-X mode (Section 8.6.1 of the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0b).

#### Table 37. Secondary Master Latency Timer - SMLT

Bits	Туре	Reset	Description
7:3	RW	PCI 0 0000 <sub>2</sub> PCI-X 0 1001 <sub>2</sub>	<ul> <li>Secondary Latency Timer (TV): 5-bit value that indicates the number of PCI clocks, in</li> <li>8-clock increments, that the bridge will remain as a master of the PCI bus when another master is requesting use of the PCI bus.</li> <li>NOTE: In scenarios like Hot-Plug when SW changes the mode of the PCI bus from conventional to PCI-X or vice-versa, this register changes its default value also appropriately.</li> </ul>
2:0	RV	0002	Reserved.

## 2.9.5.12 I/O Base and Limit - IOBL (Offset 1C)

Defines the base and limit, aligned to a 4 KB boundary, of the I/O area of the bridge. Accesses from PCI Express that are within the ranges specified in this register will be sent to PCI when the I/O space enable bit is set. Accesses from PCI that are outside the ranges specified will master abort (UR-EC Completion Status).

Bits	Туре	Reset	Description
15:12	RW	0H	I/O Limit Address Bits [15:12] (IOLA): Defines the top address of an address range to determine when to forward I/O transactions from PCI Express-to-PCI. These bits correspond to address lines 15:12 for 4 KB aligned window. Bits 11:0 are assumed to be FFFh.
11:10	RO - EN1K=0 RW - EN1K=1	00 <sub>2</sub>	I/O Limit Address Bits [11:10] (IOLA1K): When the EN1K bit is set in the Bridge Configuration register (BCNF), these bits become read/write and are compared with I/O address bits 11:10 to determine the 1 K limit address. When the EN1K bit is cleared, this field becomes Read Only.
9:8	RO	002	I/O Limit Addressing Capability (IOLC): This is hard-wired to $00_2$ , indicating support for only 16-bit I/O addressing.
7:4	RW	0 0000 <sub>2</sub>	I/O Base Address Bits [15:12] (IOBA): Defines the bottom address of an address range to determine when to forward I/O transactions from one interface to the other. I/O These bits correspond to address lines 15:12 for 4 KB alignment. Bits 11:0 are assumed to be 000H.
3:2	RO - EN1K=0 RW - EN1K=1	00 <sub>2</sub>	I/O Base Address Bits [11:10] (IOBA1K): When the EN1K bit is set in the Bridge Configuration register (BCNF), these bits become read/write and are compared with I/O address bits 11:10 to determine the 1 K base address. When the EN1K bit is cleared, this field becomes Read Only.
1:0	RO	002	I/O Base Addressing Capability (IOBC): This is hard-wired to $00_2$ , indicating support for only 16-bit I/O addressing.

#### Table 38. I/O Base and Limit - IOBL

## 2.9.5.13 Secondary Status - SSTS (Offset 1E)

For the writable bits in this register, writing a '1' will clear the bit. Writing a '0' to the bit will have no effect.

Bits	Туре	Reset	Description
15	RWC	0	Detected Parity Error (DPE): This bit is set to a '1' whenever the bridge detects an address, attribute or data parity error on the PCI bus. This bit gets set even when the Parity Error Response Enable bit (bit 0 of offset 3E-3F) is not set.
14	RWC	0	Received System Error (RSE): This bit is set to a '1' when a SERR# assertion is received on PCI.
13	RWC	0	Received Master Abort (RMA): This bit is set to a '1' whenever the bridge, as an initiator on the PCI bus, receives a master-abort or when the bridge receives a PCI-X split completion packet with a master abort.
12	RWC	0	Received Target Abort (RTA): This bit is set to a '1' whenever the bridge, as an initiator on PCI, receives target-abort on PCI. For "completion required" PCI Express packets, this event should force a completion status of "target abort" on PCI Express, and set the Signaled Target Abort in the Primary Status Register.
11	RWC	0	Signaled Target Abort (STA): This bit is set to 1 when the bridge, as a target on the PCI Bus, signals a target abort.
10:9	RO	01	DEVSEL# Timing (DVT): Indicates that the bridge responds in medium decode time to transactions on the PCI interface (secondary bus).
8	RWC	0	<ul> <li>Master Data Parity Error Detected. (MDPD): This bit is set to a '1' when all of the following are true:</li> <li>PCI Rules: <ul> <li>The bridge is the initiator on PCI.</li> <li>PERR# is detected asserted (writes) or the bridge asserted PERR# (reads).</li> <li>The Parity Error Response Enable bit in the Bridge Control Register (bit 0, offset 3Eh) is set.</li> </ul> </li> <li>PCI-X Rules: <ul> <li>The bridge receives a split completion message from PCI-X which indicates write data parity error (or) the bridge receives a split completion and it detected an uncorrectable data parity error in either the data or the split completion message (or) the bridge saw the PERR# pin asserted when mastering a write transaction (posted or non-posted) on PCI-X (terminated with either a split response or immediate completion) (or) the bridge detected an uncorrectable data parity error in the Split Response and calculates an uncorrectable data parity error in the Split Response.</li> <li>The Parity Error Response Enable bit in the Bridge Control Register (bit 0, offset 3Eh) is set.</li> </ul> </li> </ul>
7	RO	1	Fast Back-to-Back Capable (FBC): Indicates that the secondary interface can receive fast back-to-back cycles.
6	RV	0	Reserved
5	RO	1	66 MHz Capable (C66): Indicates the secondary interface of the bridge is 66 MHz capable.
4:0	RV	00H	Reserved

#### Table 39.Secondary Status - SSTS



## 2.9.5.14 Memory Base and Limit - MBL (Offset 20)

Defines the base and limit, aligned to a 1 MB boundary, of the non-prefetchable memory area of the bridge. Accesses from PCI Express that are within the ranges specified in this register will be sent to PCI when the memory space enable bit is set. Accesses from PCI that are outside the ranges specified will be forwarded to PCI Express when the bus master enable bit is set. Note that even though this region is non-prefetchable, peer reads from PCI could potentially prefetch through this window. This prefetching can be turned off with the prefetch policy bits (PP bits 42:41) in the PREFCTL register.

Reset default is with the memory base and limit registers all 0s. Note that this register must be programmed appropriately to enable or disable the space.

Bits	Туре	Reset	Description
31:20	RW	000H	Memory Limit (ML): These bits are compared with bits 31:20 of the incoming address to determine the upper 1 MB aligned value (exclusive) of the range. The incoming address must be less than or equal to ML:FFFFF.
19:16	RV	0H	Reserved.
15:4	RW	000H	Memory Base (MB): These bits are compared with bits 31:20 of the incoming address to determine the lower 1 MB aligned value (inclusive) of the range. The incoming address must be greater than or equal to MB:00000.
3:0	RV	ОH	Reserved.

#### Table 40. Memory Base and Limit - MBL



## 2.9.5.15 **Prefetchable Memory Base and Limit - PMBL (Offset 24)**

Defines the base and limit, aligned to a 1 MB boundary, of the prefetchable memory area of the bridge. Accesses from PCI Express that are within the ranges specified in this register will be sent to PCI when the memory space enable bit is set. Accesses from PCI that are outside the ranges specified will be forwarded to PCI Express when the bus master enable bit is set. Note that even though this register specifies a valid prefetchable memory window, the bridge will never Prefetch through this window in the outbound direction (reads from PCI Express-to-PCI). Nor does the bridge do any byte-merging in this window.

Note though that peer reads from PCI could prefetch through this window. This prefetching can be turned off with the prefetch policy bits (PP bits 42:41) in the PREFCTL register.

Reset default is with the prefetchable memory base and limit registers all 0s. Note that this register must be programmed appropriately to enable or disable the space.

#### Table 41. Prefetchable Memory Base and Limit - PMBL

Bits	Туре	Reset	Description
31:20	RW	000H	Prefetchable Memory Limit (PML): These bits are compared with bits 31:20 of the incoming address to determine the upper 1 MB aligned value (inclusive) of the range. The incoming address must be less than or equal to PMLU:PML:FFFFF.
19:16	RO	1H	64-bit Indicator (IS64L): Indicates that 64-bit addressing is supported for the limit. This value must be in agreement with the IS64B field.
15:4	RW	000H	Prefetchable Memory Base (PMB): These bits are compared with bits 31:20 of the incoming address to determine the lower 1 MB aligned value (inclusive) of the range. The incoming address must be greater than or equal to PMBU:PMB:00000.
03:0	RO	1H	64-bit Indicator (IS64B): Indicates that 64-bit addressing is supported for the limit. This value must be in agreement with the IS64L field.

#### 2.9.5.16 **Prefetchable Memory Base Upper 32 Bits - PMBU32 (Offset 28)**

This defines the upper 32 bits of the prefetchable address base register.

#### Table 42. Prefetchable Memory Base Upper 32 Bits - PMBU32

Bits	Туре	Reset	Description
31:0	RW	0000 0000H	Prefetchable Memory Base Upper Portion (PMBU): Full 64-bit addressing supported.

#### 2.9.5.17 Prefetchable Memory Limit Upper 32 Bits - PMLU32 (Offset 2C)

This defines the upper 32 bits of the prefetchable address base register.

#### Table 43. Prefetchable Memory Limit Upper 32 Bits - PMLU32

Bits	Туре	Reset	Description
31:0	RW	00000000H	Prefetchable Memory Limit Upper Portion (PMLU): Full 64-bit addressing supported.



## 2.9.5.18 I/O Base and Limit Upper 16 Bits - IOBLU16 (Offset 30)

Since I/O is limited to 64 KB, this register is reserved and not used.

#### Table 44. I/O Base and Limit Upper 16 Bits - IOBLU16

Bits	Туре	Reset	Description
31:16	RO	0000H	I/O Base High 16 Bits (IOBH): Reserved.
15:0	RO	0000H	I/O Limit High 16 Bits (IOLH): Reserved.

#### 2.9.5.19 Capabilities List Pointer - CCAP (Offset 34)

Contains the pointer for the first entry in the capabilities list.

#### Table 45. Capabilities List Pointer - CCAP

Bits	Туре	Reset	Description
7:0	RO	44H	Capabilities Pointer (PTR): Indicates that the pointer for the first entry in the capabilities list is at 44H (PCI Express capability) in configuration space.

## 2.9.5.20 Interrupt Information - INTR (Offset 3C)

This register contains information on interrupts on the bridge. The two bridges (A-segment Bridge or function #0, and B-segment Bridge or function #2) have different definitions since the A-segment does not support a Standard Hot-Plug controller.

#### Table 46. Interrupt Information - INTR[A]

Bits	Туре	Reset	Description	
15:8	RO	00H	Interrupt Pin (PIN): Indicates no interrupt is used by the Bridge-A segment.	
7:0	RW	00H	Interrupt Line (LINE): This register is used to convey the interrupt line routing information between the initialization code and the device driver. This is not as such used by 80333.	

#### Table 47. Interrupt Information - INTR[B]

Bits	Туре	Reset	Description		
	RO	SHPC Enabled (B_HSLOT[3]=1 at PWRGD asserting edge)			
15:8		02H	Interrupt Pin (PIN): 80333 B-segment bridge has an integrated standard Hot-Plug controller, which is a source of interrupt. The logical primary bus interrupt pin is INTB# for function 2 (B-segment bridge). Note that the Hot-Plug interrupt is routed internally to IRQ#[23] of the B-segment's IOAPIC.		
		SHPC Disab	led (B_HSLOT[3]=0 at PWRGD asserting edge)		
		00H	Interrupt Pin (PIN): Indicates no interrupt is used by the Bridge-B segment.		
7:0	RW	00H Interrupt Line (LINE): This register is used to convey the interrupt line routing information between the initialization code and the device driver. This is not as such used by 80333.			

## 2.9.5.21 Bridge Control - BCTRL (Offset 3E)

This register provides extensions to the Command register that are specific to a bridge. The Bridge Control register provides many of the same controls for the secondary interface that are provided by the Command register for the primary interface. Some bits affect operation of both interfaces of the bridge.

#### Table 48.Bridge Control - BCTRL (Sheet 1 of 2)

Bits	Туре	Reset	Description	
15:12	RV	0H	Reserved.	
11	RW	0	<ul> <li>Discard Timer SERR# Enable (DTSE): Controls the generation of ERR_NONFATAL/ERR_FATAL on the primary interface in response to a timer discard on the secondary interface.</li> <li>0 = Do not generate ERR_NONFATAL/ERR_FATAL on a secondary timer discard (but the corresponding mask bit in AER could still enable a message when clear).</li> <li>1 = Generate ERR_NONFATAL/ERR_FATAL in response to a secondary timer discard provided the appropriate mask bit in the advanced capability register is clear.</li> </ul>	
10	RWC	0	Discard Timer Status (DTS): This bit is set to a '1' when the secondary discard timer expires (there is no discard timer for the primary interface).	
9	9 RW 0 Secondary Discard Timer (SDT): Sets the maximum number of PCI of bridge waits for an initiator on PCI to repeat a delayed transaction red starts once the delayed transaction completion is at the head of the q master has not repeated the transaction at least once before the cound discards the transaction from its queues. 0 = The PCI master time-out value is between 2^15 and 2^16 PCI of 1 = The PCI master time out value is between 2010 and 2011 PCI of 1		Secondary Discard Timer (SDT): Sets the maximum number of PCI clock cycles that the bridge waits for an initiator on PCI to repeat a delayed transaction request. The counter starts once the delayed transaction completion is at the head of the queue. When the master has not repeated the transaction at least once before the counter expires, the bridge discards the transaction from its queues. 0 = The PCI master time-out value is between 2^15 and 2^16 PCI clocks. 1 = The PCI master time-out value is between 2^10 and 2^11 PCI clocks	
8	RW	0	Primary Discard Timer (PDT): Not relevant to PCI Express.	
7	RO	0	Fast Back-to-Back Enable (FBE): The bridge cannot generate fast back-to-back cycles or the PCI bus from PCI Express initiated transactions.	
6	RW	0	<ul> <li>Secondary Bus Reset (SBR): Controls x_RST# assertion on PCI.</li> <li>0 = The bridge deasserts x_RST#. Note that the onus is on the software to guarantee the proper resetting of this bit to guarantee the secondary bus x_RST# timing requirements.</li> <li>1 = The bridge asserts PCIRST#. Bridge configuration registers are not reset on this bit being set.</li> <li>Once this bit is set, the bridge will complete the currently running transaction on PCI and then reset the bus. Note that it is the responsibility of software to make sure that all pending transactions with the bus segment are complete before setting this bit. Failing which, transactions could be lost.</li> </ul>	
5	5 RW 0		Master Abort Mode (MAM): Controls the bridges behavior when a master abort occurs on either interface.         Master Abort on PCI Express/peer PCI:         Memory reads         0 = Bridge asserts TRDY# on PCI. It drives all '1's for reads.         1 = Bridge returns a target abort on PCI.         This bit has no effect on inbound posted transactions master aborting on the PCI Express upstream device or when the PCI segment is in the PCI-X mode.         I/O and Configuration Writes         0 = Bridge completes transaction normally.         1 = Bridge returns a target abort on PCI.         Master Abort on PCI/PCI-X: (posted transactions only)         0 = No action on PCI Express.         1 = ERR_NONFATAL/ERR_FATAL response (independent of advanced error mask bit), when enabled. Refer to Section 2.7for details.	



Table 48.	. Bri	dge Contro	ol - BCTRL (Sheet 2 of 2)	
Bits	Туре	Reset	Description	
4	RW	0	VGA 16-bit Decode: Enables the bridge to provide 16-bits decoding precluding the decode of VGA alias addresses every 1 KB. This bit enable bit (bit 3 of this register) to be set '1'.	
			VGA Enable (VGAE): Modifies the response to VGA compatible ad	

Table 48.	Bridge Contro	DI - BCTRL	(Sheet 2 of 2)
	Dridge Contro		

4	RW	0	VGA 16-bit Decode: Enables the bridge to provide 16-bits decoding of VGA I/O address precluding the decode of VGA alias addresses every 1 KB. This bit requires the VGA enable bit (bit 3 of this register) to be set '1'.	
			VGA Enable (VGAE): Modifies the response to VGA compatible addresses. When set to a '1', the bridge forwards the following transactions from PCI Express-to-PCI regardless of the value of the I/O base and limit registers. The transactions are qualified by the memory enable and I/O enable in the command register.	
2		0	Memory addresses: 000A0000H-000BFFFFH	
3	RVV	0	I/O addresses: 3B0H-3BBH and 3C0H-3DFH. For the I/O addresses, bits [63:16] of the address must be '0', and bits [15:10] of the address are ignored (i.e., aliased).	
			The same holds true from secondary accesses to the primary interface in reverse for memory accesses and also for I/O when the inbound I/O enable bit is set in the BINIT register, from secondary to primary.	
2	RW	0	ISA Enable (IE): Modifies the response by the bridge to ISA I/O addresses. This only applies to I/O addresses that are enabled by the I/O Base and I/O Limit registers and are in the first 64 KB of PCI I/O space. When this bit is set, the bridge will block any forwarding from primary to secondary of I/O transactions addressing the last 768 bytes in each 1 KB block (offsets 100H to 3FFh). This bit has reverse effect on I/O transfers originating on the secondary bus when the inbound I/O enable bit is set in the BINIT register.	
1	RW	0 SERR# Enable (SE): Controls the forwarding of secondary interface SERR# asserti the primary interface. When set, the 80333 will send a PCI Express ERR_NONFATAL/ERR_FATAL cycle (based on Advanced Error capability's PCI SE detected severity bit) when all of the following are true. SERR# is asserted on the secondary interface. This bit is set or the SERR# detected mask bit in the Advanced Error capability ERR_NONFATAL/ERR_FATAL messages are enabled to be sent.		
0	RW	0	Parity Error Response Enable (PERE): Controls the response to address and data parity errors on the secondary interface. When the bit is cleared, the bridge must ignore any par errors that it detects and continue normal operation. The bridge must generate parity every when parity error reporting is disabled.	



## 2.9.5.22 Bridge Configuration Register - BCNF (Offset 40)

80333-specific bridge control bits.

## Table 49. Bridge Configuration Register - BCNF (Sheet 1 of 2)

Bits	Туре	Reset	Description		
15:14	RW	-	<ul> <li>PCI Mode (PMODE): Determines the mode of operation of the PCI bus. These bits both reflect the status of the current PCI bus mode and also let software change the mode by writing to these bits. The power up value of this register is written based upon the Table 12, "PCI Mode Pin/Strap Encoding" on page 72.</li> <li>Bits Mode</li> <li>00 Conventional PCI</li> <li>01 PCI-X</li> <li>10, 11 Reserved (alias to 00 @ 33 MHz)</li> <li>When Reserved codes are written, 80333 runs the bus at PCI 33MHz and the PMODE and PFREQ fields reflect that mode/speed.</li> <li>NOTE: These bits provided for debug only. With Hot-Plug enabled, software uses Standard Hot-Plug commands to change PCI bus mode and frequency, instead of these bits. Using these bits when Hot-Plug is enabled could have undesirable results.</li> <li>With Hot-Plug disabled, 80333 checks (via PMODE/PFREQ fields, when SBR bit is set) requested software speed and mode for consistent slot and bus segments capabilities (Expressed via pins/straps indicated in ). When the requested speed/mode is greater than the capabilities of the slot/bug-segment. then 80333aliases the command to PCI 33 MHz.</li> </ul>		
13	RWS	1	APIC Configure Space Disable: When 1 80333 disables all configuration accesses to IOAPIC configuration space on the appropriate segment from PCI Express and PCI. When 1 all configuration access from PCI Express, PCI to the IOAPIC of the appropriate segment are master aborted. This bit has no affect on the SMBus accesses to the IOAPIC configuration space and also this bit has no effect on memory accesses to IOAPIC.		
			By default the IOAPIC configuration space is hidden.		
12	RW	0	decoded to 1 K, down from the 4 K limit that currently exists in the I/O base and limit registers. It does this by redefining bits [11:10] and bits [3:2] of the IOBL register at offset 1C to be read/write, and enables them to be compared with I/O address bits [11:10] to determine when they are within the bridge I/O range.		
			segments with two IOHs and the I/O space is only 64 K.		
11	RV	1	Reserved.		
10:09	RW	он	<ul> <li>PCI Frequency (PFREQ): Determines the frequency the PCI bus operates. The power up value of this register is written based upon Table 13. After software determines the busses capabilities, it sets this value, and the PMODE bits to the desired frequency and resets the PCI bus. The values are encoded as follows:</li> <li>Bit FrequencyComments</li> <li>00 33Only valid when PMODE is '00'</li> <li>01 66Valid for PMODE being '10' or '01' or '00'</li> <li>10 100Only valid when PMODE is '10' or '01'</li> <li>11 133Only valid when PMODE is '10' or '01'</li> <li>Invalid combinations should not be written by software. Results will be indeterminate.</li> <li>NOTE: The PFREQ bits can not be used to change the actual PCI bus frequency. The only</li> </ul>		
			way to change the PCI bus frequency is to reset the bus or use the Hot Plug controller.		
8	RV	0	Reserved		
7	RW	1	Peer Memory Read Enable (PMRE): When '1', peer memory reads from one internal device to another are supported (PCI-PCI, PCI-APIC, PCI-SHPC). When '0' (normal operation), peer memory reads are not allowed and all memory reads from a PCI bridge will be sent to PCI Express regardless of address.		



Bits	Туре	Reset	Description		
6	RO	0	Enable Fence Special Cycle (EFSC): 80333 does not support the Requestor ID fence cycle on PCI Express and hence this is 0.		
5	RW	0	SHPC GPE Message Enable (SGME): Enable Redirection of Hot-Plug interrupts to Assert/Deassert GPE Messages on PCI Express.		
4	RV	0	Reserved.		
2	RW	0	Outbound Delayed Transaction Resource Partitioning (ODTP): Determines how the two outbound delayed transaction entries are partitioned between PCI Express and peer PCI traffic.		
5			<ul> <li>0 = Max 1 entry for PCI Express and Max 1 entry for peer PCI.</li> <li>1 = Max 2 entries for both PCI Express and peer PCI. First come first served.</li> </ul>		
2	RV	0	Reserved		
1:0		RW 00	Maximum Inbound Delayed Transactions (MDT): Controls the maximum number of inbound delayed transactions the 80333 is allowed to have:		
	D\//		00: 4 active, 4 pending		
			01: 1 active, 1 pending		
			10: 2 active, 2 pending		
			• 11: Reserved		

## Table 49.Bridge Configuration Register - BCNF (Sheet 2 of 2)



## 2.9.5.23 Multi-Transaction Timer - MTT (Offset 42)

This register controls the amount of time that the 80333 arbiter allows a PCI initiator to perform multiple back-to-back transactions on the PCI bus. The number of clocks programmed in the MTT represents the guaranteed time slice (measured in PCI clocks) allotted to the current agent, after which the arbiter will grant another agent that is requesting the bus.

#### Table 50. Multi-Transaction Timer - MTT

Bits	Туре	Reset	Description
7:3	RW	00H	Timer Count Value (MTC): This field specifies the amount of time that grant will remain asserted to a master continuously asserting its request for multiple transfers. This field specifies the count in an 8-clock (PCI clock) granularity.
2:0	RV	000b	Reserved.

## 2.9.5.24 PCI Clock Control - PCLKC (Offset 43)

This register controls the enable or disable of the 80333 PCI clock outputs.

#### Table 51. PCI Clock Control - PCLKC

Bits	Туре	Reset	Description
7	PR	1b	Preserved.
6	RW	1b	PCI Feedback Clock Control: These bit enables the PCI Feedback clock output buffer, when 1. Otherwise the buffer is tristated.
5	PR	1b	Preserved.
			PCI Clock Control: These bits enable the PCI clock output buffers, when 1. Otherwise the buffers are tristated. Bit 0 corresponds to X_CLKO[0], Bit 1 corresponds to X_CLKO[1], etc.
4:0	RW	1 1111b	A-side implements bits 3:0
			B-side implements bits 4:0
			The tristating of the clock is asynchronous to the output clocks.

#### 2.9.5.25 PCI Express Capability Identifier - EXP\_CAPID (Offset 44)

This register stores the PCI Express capability ID value.

#### Table 52. PCI Express Capability Identifier - EXP\_CAPID

Bits	Туре	Reset	Description
7:0	RO	10H	PCI Express Capability ID: Indicates PCI Express capability

## 2.9.5.26 PCI Express Next Item Pointer - EXP\_NXTP (Offset 45)

This register stores the byte offset pointer to the next capability list item of the bridge.

#### Table 53. PCI Express Next Item Pointer \_ EXP\_NXTP

Bits	Туре	Reset	Description
7:0	RO	5Ch	Next Capability Pointer: The offset of the next capabilities list item, which is the MSI capability.



## 2.9.5.27 PCI Express Capability - EXP\_CAP (Offset 46)

This register carries the version number of the capability item and other base information contained in the capability structure.

## Table 54. PCI Express Capability - EXP\_CAP

Bits	Туре	Default	Description
15:14	RV	002	Reserved
13:9	RO	0 0000 <sub>2</sub>	Interrupt Message Number: Not relevant for 80333
8	RO	02	Slot Implemented: Not relevant for 80333
7:4	RO	7H	Device/Port Number: Indicates Bridge Association value for PCI Express-to-PCI Bridge
3:0	RO	1H	Version Number: Indicates PCI Express capability structure version number

## 2.9.5.28 PCI Express Device Capabilities Register - EXP\_DCAP (Offset 48)

This register carries information on the PCI Express link capabilities.

Bits	Туре	Default	Description
31:28	RV	0H	Reserved
27:26	RO	00 <sub>2</sub>	Slot Power Limit Scale: In combination with the Slot Power Limit value, specifies the upper limit on power supplied by slot. Power limit (in Watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Value field. This value is set by the Set_Slot_Power_Limit message.
25:18	RO	00H	Slot Power Limit Value: In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in Watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. This value is set by the Set_Slot_Power_Limit message.
17:15	RV	002	Reserved.
14	RO	02	Power Indicator Present: Not Supported.
13	RO	02	Attention Indicator Present: Not Supported.
12	RO	02	Attention Button Present: Not Supported.
11:9	RO	0002	Endpoint L1 Acceptable Latency: L1 ASPM is not supported.
8:6	RO	0002	Endpoint L0s Acceptable Latency: The least latency possible out of L0s is supported.
5	RO	02	Extended Tag Field Supported: Ignore. Only a 5 bit tag is supported.
4:3	RO	002	Phantom Functions Supported: Not Supported.
2:0	RO	001 <sub>2</sub>	Supported Max Payload sizes: 256 byte packets is the max supported.

#### Table 55. PCI Express Device Capabilities Register - EXP\_DCAP



## 2.9.5.29 PCI Express Device Control Register - EXP\_DCTL (Offset 4C)

This register carries command bits that control 80333 behavior on PCI Express.

Table 56.	PCI Express	<b>Device Control</b>	Register - EXP	_DCTL
-----------	-------------	-----------------------	----------------	-------

Bits	Туре	Default	Description
15	RW	02	Bridge Configure Retry Enable: When set, 80333 is enabled to return a configuration retry response on PCI Express for a Configuration transaction to PCI/X.
14:12	RW	010 <sub>2</sub>	Max_Read_Request_Size: 80333 cannot send requests greater than the size indicated by this field. Encodings are Bits Max_Read_Request_Size 000b 128B max read request size 001b 256B max read request size 010b 512B max read request size 011b 1024B max read request size 100b 2048B max read request size 101b 4096B max read request size 110b Reserved (Bridge defaults to 512B) 111b Reserved (Bridge defaults to 512B)
11	RO	02	Enable No Snoop: Does not apply to 80333 since it does not set the NS bit on MSI transactions it generates.
10	RO	02	Auxiliary (AUX) Power PM Enable: Not Supported
9	RO	02	Phantom Function Enable: Not Supported
8	RO	02	Extended Tag Field Enable: Ignored since only 5 bit tag is supported
7:5	RW	000 <sub>2</sub>	Maximum Payload Size: Indicates the maximum payload size supported for TLPs. Supported encodings are: Bits Max Payload Size 000b 128B 001b 256B All other values default to 128B
4	RO	02	Enable Relaxed Ordering: Does not apply to 80333 since it does not set the RO bit on MSI requests it generates.
3	RW	02	Unsupported Request Reporting Enable: This bit controls enabling of ERR_NONFATAL or ERR_FATAL messages on PCI Express for reporting Unsupported Request errors. Note that for requests from PCI Express that are unsupported or which master abort on the internal switch use this enable bit.
2	RW	02	Report Fatal Errors: When this bit is set generation of the ERR_FATAL message is enabled
1	RW	02	Report NonFatal Errors: When this bit is set generation of the ERR_NONFATAL message is enabled.
0	RW	02	Report Correctable Errors: When this bit is set generation of ERR_CORR message is enabled.



## 2.9.5.30 PCI Express Device Status Register - EXP\_DSTS (Offset 4E)

This register carries information on the PCI Express device status.

Bits	Туре	Default	Description
15:6	RV	000H	Reserved.
5	RO	02	Transactions pending: Reserved
4	RO	02	Auxiliary Power Detected: Auxiliary Power is not support.
3	RWC	0 <sub>2</sub>	Unsupported Request Detected: This bit is set when any unsupported request from PCI Express is received. This includes requests that are not claimed by any function in 80333 (i.e. MA-ed on the internal switch) but does NOT include any request that are forwarded to the PCI interface with completions returned with an unsupported request status.
2	RWC	02	Detected Fatal Error: When set, a fatal error has been detected (regardless of whether an error message was generated) either on the PCI Express or PCI/X interface or internally. This bit is set till software writes a 1 to clear it.
1	RWC	02	Detected NonFatal Error: When set, an nonfatal error has been detected (regardless of whether an error message was generated) either on the PCI Express or PCI/X interface or internally. This bit is set till software writes a 1 to clear it.
0	RWC	02	Detected Correctable Error: When set, a correctable error has been detected (regardless of whether an error message was generated) either on the PCI Express or PCI/X interface or internally. This bit is set till software writes a 1 to clear it.

#### Table 57. PCI Express Device Status Register - EXP\_DSTS

## 2.9.5.31 PCI Express Link Capabilities Register - EXP\_LCAP (Offset 50)

#### Table 58. PCI Express Link Capabilities Register - EXP\_LCAP

Bits	Туре	Default	Description
31:24	RO	02	Port Number: Not applicable.
23:18	RV	02	Reserved.
17:15	RO	111 <sub>2</sub>	L1 Exit Latency: L1 transition is not supported.
14:12	RO	111 <sub>2</sub>	L0s Exit Latency: The value in these bits is influenced by bit 6 in link control register. Note that software could write the bit 6 in link control register to either a 1 or 0 and these bits should change accordingly. The mapping is shown below: Bit 6 in LCTLL0s Exit Latency 0 = DEM[5:3] (because currently L0s cannot work with different ref clocks) 1 = DEM[2:0]
11:10	RO	01 <sub>2</sub>	ASPM Support: Only L0s is supported.
9:4	RO	00 1000 <sub>2</sub>	Maximum Link Width: X8 link width supported.
3:0	RO	1H	Maximum Link Speed: 2.5 Gb/s link speed supported.



## 2.9.5.32 PCI Express Link Control Register - EXP\_LCTL (Offset 54)

Bits	Туре	Default	Description
15:8	RV	00H	Reserved.
7	RW	02	Extended Synchronize – This bit when set forces extended transmission of 4096 FTS ordered sets in FTS and an extra 1024 TS1 at exit from L1 prior to entering L0. This mode provides external devices monitoring the link time to achieve bit and symbol lock before the link enters L0 state and resumes communication. Default value for this bit is 0
6	RW	0 <sub>2</sub>	Common Clock Configuration: Indicates the relationship of the reference clock between 80333 and the component at the opposite end of 80333 Upstream PCI Express interface. 0 = clock is asynchronous. 1 = clock is common. NOTE: This bit is used to reflect the proper L0s exit latency value in the EXP_LCAP register.
5	RO	02	Retrain Link: Not Applicable.
4	RO	02	Disable Link: Not Applicable.
3	RO	02	Read Completion Boundary Control: Not Applicable.
2	RV	02	Reserved.
1:0	RW	02	<ul> <li>ASPM Control: Enables bridge upstream interface to enter L0s</li> <li>00 L0s entry disabled.</li> <li>01 80333 enters L0s per the Specification requirement for L0s entry.</li> <li>10 L0s entry disabled.</li> <li>11 80333 enters L0s per the Specification requirement for L0s entry.</li> </ul>

## Table 59. PCI Express Link Control Register - EXP\_LCTL



## 2.9.5.33 PCI Express Link Status Register - EXP\_LSTS (Offset 56)

Bits	Туре	Default	Description	
15:13	RV	0002	Reserved.	
12	ROS	1 <sub>2</sub>	Slot Clock Configuration – Indicates that when 80333 is on a PCI Express connector, that it is using the same reference clock as is provided at the connector. 0 = indicates independent reference clock. 1 = indicates same reference clock.	
11	RO	02	Link Training: Not Applicable.	
10	RO	02	Link Width Negotiation Error: Not Applicable.	
9:4	RO	Set by PCI Express Link Layer after training is complete	Negotiated Link Width: This field indicates the negotiated width of PCI Express Link.         00 00012       X1 - Supported         00 00102       X2 - Not Supported         00 01002       X4 - Supported         00 10002       X8 - Supported         00 11002       X12 - Not Supported         00 11002       X12 - Not Supported         01 00002       X16 - Not Supported         10 00002       X32 - Not Supported         Valid values for 80333 are x1, x4, and x8. All other values are reserved.	
3:0	RO	1H	Link Speed: The only speed supported is 2.5 Gbps.	

#### Table 60. PCI Express Link Status Register - EXP\_LSTS

## 2.9.5.34 MSI Capability Identifier - MSI\_CAPID (Offset 5C)

PCI Express MSI capability identifier register.

#### Table 61. MSI Capability Identifier - MSI\_CAPID

Bits	Туре	Reset	Description
7:0	RO	05h	Capability ID (MCID): Capabilities ID indicates MSI.

## 2.9.5.35 Next Item Pointer - MSI\_NXTP (Offset 5D)

#### Table 62. Next Item Pointer - MSI\_NXTP

Bits	Туре	Reset	Description
7:0	RO	6CH	Next Pointer (MNPTR): Points to the next capabilities list pointer and this is the last PCI Express Power Management capability item.



## 2.9.5.36 MSI Message Control - MSI\_MC (Offset 5E)

Bits	Туре	Reset	Description
15:8	RV	0	Reserved.
7	RO	1	64-Bit Address Capable: 80333 is capable of generating a 64-bit message address
6:4	RW	000	Multiple Message Enable: Only one message is supported. These bits are R/W for software compatibility,
3:1	RO	000	Multiple Message Capable: Only one message is supported.
0	RW	0	MSI Enable: When set, MSI is enabled and traditional interrupt pins are not used to generate interrupts. (When set SHPC does not use IRQ[23]# wire to the internal IOAPIC to generate interrupts)

#### Table 63. MSI Message Control - MSI\_MC

## 2.9.5.37 MSI Message Address - MSI\_MA (Offset 60)

#### Table 64. MSI Message Address - MSI\_MA

Bits	Туре	Reset	Description
63:2	RW	0	Address (ADDR): System specified message address, always DWORD aligned.
1:0	RV	00	Reserved.

## 2.9.5.38 MSI Message Data - MSI\_MD (Offset 68)

#### Table 65. MSI Message Data - MSI\_MD

Bits	Туре	Reset	Description
15:0	RW	0000H	Data (DATA): This 16-bit field is programmed by system software when MSI is enabled. Its content is driven onto the lower word (D[15:0]) of the MSI memory write transaction.

## 2.9.5.39 **Power Management Capabilities Identifier - PM\_CAPID (Offset 6C)**

#### Table 66. Power Management Capabilities Identifier - PM\_CAPID

Bits	Туре	Reset	Description
7:0	RO	01H	Identifier (ID): Indicates this is a PCI Compatible PM.



## 2.9.5.40 Power Management Next Item Pointer - PM\_NXTP (Offset 6D)

#### Table 67. Power Management Next Item Pointer - PM\_NXTP

Bits	Туре	Reset	Description		
SHPC En	SHPC Enabled (B_HSLOT[3]=1 at the rising edge of PWRGD)				
7:0	RO	78h	Next Pointer: When SHPC is enabled, this points to the SHPC as the next capability.		
SHPC Disabled (B_HSLOT[3]=0 at the rising edge of PWRGD)					
7:0	RO	D8h	Next Pointer: When SHPC is disabled, points to the PCI-X capability as the next capability.		

## 2.9.5.41 Power Management Capabilities - PM\_CAP (Offset 6E)

#### Table 68. Power Management Capabilities - PM\_CAP

Bits	Туре	Reset	Description
15:11	RO	19H	PME_Support: PME assertion is supported on behalf of SHPC when in D3hot. PME assertion from D3cold is not supported.
10	RO	0	D2 Support: Not Supported.
9	RO	0	D1 Support: Not Supported.
8:6	RO	000	Auxiliary Current: Auxiliary Power is not Supported.
5	RO	0	DSI: Device specific initialization is not required when transitioning to D0 from D3hot state. This bit is zero.
4	RV	0	Reserved.
3	RO	0	PME Clock: Not relevant.
2:0	RO	010	Version: PM implementation is compliant with <i>PCI Bus Power Management Interface Specification</i> , Revision 1.1.

## 2.9.5.42 Power Management Control/Status Register - PM\_CSR (Offset 70)

#### Table 69. Power Management Control/Status Register - PM\_CSR

Bits	Туре	Reset	Description
15	RWCS	0	PME Status: This bit is set when a PME request would normally be sent on behalf of SHPC, independent of the state of the PME_En bit. SHPC requests a PME message when 80333 is in D3hot state and a Hot-Plug operation is requested. Refer to the Chapter 14, "Standard Hot-Plug Controller" for the details of PME generation.
14:13	RO	00	Data Scale: Not Supported.
12:9	RO	ОH	Data Select: Not Supported
8	RWS	0	PME En: Gates assertion of the PME message on behalf of SHPC
7:2	RV	00 0000	Reserved.
1:0	RW	00	PowerState – This 2-bit field is used both to determine the current power state of a function and to set the function into a new power state. Supported field values are given below. 00b – D0 01b – Reserved 10b – Reserved 11b – D3 hot When software attempts to write an unsupported, optional state to this field, the write operation must complete normally on the bus; however, the data is discarded and no state change occurs.

## 2.9.5.43 Power Management Bridge Support Extensions -PM\_BSE (Offset 72)

#### Table 70. Power Management Bridge Support Extensions - PM\_BSE

Bits	Туре	Reset	Description
7	RO	0	BPCC_En (Bus Power/Clock Control Enable): Neither bus or clock control of PCI is supported when in D3hot state. This bit is wired to a 0.
6	RO	0	B2/B3#: Not Supported. This bit has no meaning since BPCC_En bit is a 0.
5:0	RV	00 0000	Reserved.



## 2.9.5.44 **Power Management Data Field - PM\_DATA (Offset 73)**

#### Table 71. Power Management Data Field - PM\_DATA

Bits	Туре	Reset	Description
7:0	RV	0	Reserved: the data register is not reported by 80333

## 2.9.5.45 SHPC Capability Identifier - SHPC\_CAPID (Offset 78)

This register stores the SHPC capability ID value.

#### Table 72. SHPC Capability Identifier - SHPC\_CAPID

Bits	Туре	Reset	Description		
SHPC En	SHPC Enabled (B_HSLOT[3]=1 at the rising edge of PWRGD)				
7:0	RO	0CH	SHPC Capability ID: Used to detect the presence of an SHPC integrated with a PCI-to-PCI Bridge.		
SHPC Disabled (B_HSLOT[3]=0 at the rising edge of PWRGD)					
7:0	RV	00H	Reserved.		

## 2.9.5.46 SHPC Next Item Pointer - SHPC\_NXTP (Offset 79)

This register stores the byte offset pointer to the next capability list item of the bridge.

#### Table 73. SHPC Next Item Pointer - SHPC\_NXTP

Bits	Туре	Reset	Description			
SHPC En	SHPC Enabled (B_HSLOT[3]=1 at the rising edge of PWRGD)					
7:0	RO	D8H	Vext Pointer: Points to the PCI-X Capability.			
SHPC Disabled (B_HSLOT[3]=0 at the rising edge of PWRGD)						
7:0	RV	00H	Reserved.			



## 2.9.5.47 SHPC DWORD Select Register - SHPC\_DWSEL (Offset 7A)

This register is used to select the DWORD offset in the SHPC working register set for read and write by the SHPC software. Valid only when SHPC is enabled.

#### Table 74. SHPC DWORD Select Register - SHPC\_DWSEL

Bits	Туре	Reset	Description			
SHPC En	SHPC Enabled (B_HSLOT[3]=1 at the rising edge of PWRGD)					
7:0	RW	00H	DWORD Select: Selects the DWORD from the SHPC Working Register set that is accessible through the DWORD Data register. Accesses to the DWORD Data register have no effect on the DWORD Select field. The value of 0 selects the first DWORD of the SHPC Working set. A value of 1 selects the second DWORD, and so on. This field has a default value of 0.			
SHPC Disabled (B_HSLOT[3]=0 at the rising edge of PWRGD)						
7:0	RV	00H	Reserved.			

## 2.9.5.48 SHPC Status - SHPC\_STS (Offset 7B)

SHPC Status Register.

#### Table 75. SHPC Status - SHPC\_STS

Bits	Туре	Reset	Description		
SHPC En	SHPC Enabled (B_HSLOT[3]=1 at the rising edge of PWRGD)				
7	RO	-	Controller Interrupt Pending (CIP): This bit is set when one or more bits are set in the Interrupt Locator register in the SHPC working register set. This bit is cleared when no bits are set in the Interrupt Locator register.		
6	RO	-	Controller System Error Pending (CSP): This bit is set when one or more bits are set in the SERR Locator register in the SHPC working register set. This bit is cleared when no bits are set in the SERR Locator register.		
5:0	RV	00H	Reserved.		
SHPC Disabled (B_HSLOT[3]=0 at the rising edge of PWRGD)					
7:0	RV	00H	Reserved.		



## 2.9.5.49 SHPC DWORD - SHPC\_DWORD - (Offset 7C)

SHPC Data Register.

#### Table 76. SHPC DWORD - SHPC\_DWORD

Bits	Туре	Reset	Description			
SHPC En	SHPC Enabled (B_HSLOT[3]=1 at the rising edge of PWRGD)					
31:0	RW	00H	DWORD Data – This field allows software to access the SHPC Working Register set via the Capabilities List Item in Configuration Space. The DWORD Select field selects the SHPC Working Register set DWORD that is accessed by reads and writes to this register. Accessing SHPC Working Register set registers through this field behaves the same as accessing them through memory-mapped accesses (see Section 14.6, "SHPC Working Register Set"). Multiple accesses to the DWORD Data register continue to affect the same DWORD when the DWORD Select field is unchanged. When the PCI-to-PCI Bridge integrated with the SHPC is not in the D0 power management state, reads from this register must complete successfully but the returned value is undefined and the behavior of writes is undefined.			
SHPC Dis	SHPC Disabled (B_HSLOT[3]=0 at the rising edge of PWRGD)					
31:0	RV	0000 0000H	Reserved.			

## 2.9.5.50 PCI-X Capabilities Identifier - PX\_CAPID (Offset D8)

Identifies this item in the Capabilities list as a PCI-X register set. It returns 07h when read.

#### Table 77. PCI-X Capabilities Identifier - PX\_CAPID

Bits	Туре	Reset	Description
7:0	RO	07H	Identifier (ID): Indicates this is a PCI-X capabilities list.

## 2.9.5.51 **PCI-X Next Item Pointer - PX\_NXTP (Offset D9)**

Indicates where the next item in the capabilities list resides. This is the end of the list, so "00H" is returned.

#### Table 78. PCI-X Next Item Pointer - PX\_NXTP

Bits	Туре	Reset	Description
7:0	RO	00H	PCI-X is the last capability list item and hence these bits are all 0s.



## 2.9.5.52 PCI-X Secondary Status - PX\_SSTS (Offset DA)

This is the PCI-X command register, which controls various modes of the bridge.

#### Table 79.PCI-X Secondary Status - PX\_SSTS

Bits	Туре	Reset	Description
15:9	RV	00H	Reserved.
			Secondary Clock Frequency (SCF): This field is set with the frequency of the secondary bus. The values are:
			Bits Max Frequency Clock Period
			000 PCI Mode N/A
8:6	RO	-	001 66 15
			010 100 10
			011 133 7.5
			1xx Reserved Reserved
			The default value for this register is given in Table 13.
5	RO	0	Split Request Delayed. (SRD): This bit is supposed to be set by a bridge when it cannot forward a transaction on the secondary bus to the primary bus because there is not enough room within the limit specified in the Split Transaction Commitment Limit field in the Downstream Split Transaction Control register. The 80333 will never set this bit.
4	RO	0	Split Completion Overrun (SCO): This bit is supposed to be set when a bridge terminates a Split Completion on the secondary bus with retry or Disconnect at next ADB because its buffers are full. The 80333 will never set this bit.
3	RWC	0	Unexpected Split Completion (USC): This bit is set when an unexpected split completion is received with a bus number that matches a bus number on the primary side of 80333, but with a requester ID:tag not equal to the requester ID:tag of any of the two split requests that 80333 is tracking on PCI-X. This bit is cleared by software writing a '1'.
2	RWC	0	Split Completion Discarded (SCD): This bit is set when the 80333 discards a split completion moving toward the secondary bus because the requester would not accept it. This bit cleared by software writing a '1'.
1	RO	1	133 MHz Capable (C133): This bit indicates that the 80333 secondary interface is capable of 133 MHz operation in PCI-X mode.
0	RO	1	64-bit Device (D64): Indicates the width of the secondary bus as 64-bits.



## 2.9.5.53 PCI-X Bridge Status - PX\_BSTS (Offset DC)

Identifies PCI-X status register for the bridge primary side.

Bits	Туре	Reset		Description
31:22	RV	(	0	Reserved.
21	RO	0		Split Request Delayed (SRD): This bit is supposed to be set by a bridge when it cannot forward a transaction on to the primary bus from the secondary bus, because there is not enough room within the limit specified in the Split Transaction Commitment Limit field in the upstream Split Transaction Control register. The 80333 does not set this bit.
20	RO	0		Split Completion Overrun (SCO): This bit will not be set by the 80333 because the 80333 will never request on PCI Express more data than it has room to receive.
19	RO	0		Unexpected Split Completion (USC): sets this field when a completion on PCI Express is destined to a specific bridge segment (either A or B) but tag does not match.
18	RO	0		Split Completion Discarded (SCD): This does not apply to PCI Express.
17	RO	0		133 MHz Capable (C133): Not applicable.
16	RO	0		64-bit Device (D64): Not applicable.
15:8	RO	00H		Bus Number (BNUM): This register is an alias to the PBN field of the BNUM register at offset 18h.
7:3	RO	00H		Device Number (DNUM): Device number as captured on Type-0 configuration writes to the bridge segment.
2.0	RO	А	В	Eunction Number (ENLIM): Read only hits for PCLX diagnostic software
2.0		000	010	Tunction runnber (Friow). Read only bits for POPA diagnostic software.

#### Table 80. PCI-X Bridge Status - PX\_BSTS

## 2.9.5.54 PCI-X Upstream Split Transaction Control - PX\_USTC (Offset E0)

This register identifies controls behavior of the 80333 buffers for forwarding Split Transactions from the secondary bus to PCI Express.

#### Table 81. PCI-X Upstream Split Transaction Control - PX\_USTC

Bits	Туре	Reset	Description
31:16	RW	FFFFH	Split Transaction Limit (STL): This field is R/W in case some diagnostic software wants to use it. It is not used by the 80333 for modifying its "commitment" level. 80333 internal launch algorithms keep buffers from being over allocated.
15:00	RO	FFFFH	Split Transaction Capacity (STC): 80333 essentially has infinite capacity, because its launch algorithm keeps buffers from overrunning.



## 2.9.5.55 PCI-X Downstream Split Transaction Control - PX\_DSTC (Offset E4)

This register controls behavior of the 80333 buffers for forwarding Split Transactions from PCI Express to the secondary bus.

#### Table 82. PCI-X Downstream Split Transaction Control - PX\_DSTC

Bits	Туре	Reset	Description
31:16	RW	FFFFH	Split Transaction Limit (STL): This field is R/W in case some diagnostic software wants to use it. It is not used by the 80333 for modifying its "commitment" level. 80333 internal launch algorithms keep buffers from being over-allocated.
15:0	RO	FFFFH	Split Transaction Capacity (STC): 80333 essentially has infinite capacity, because its launch algorithm keeps buffers from overrunning.



## 2.9.5.56 Bridge Initialization Register - BINIT (Offset FC)

Controls bridge for special addressing and transaction response behavior for the Intel XScale<sup>®</sup> core control of I/O controllers on the A-segment PCI-X bus. The A-segment and B-segment bridges have different content for this register as described below.

Table 83	<b>Bridge Initialization</b>	Register A-Sec	ment - <b>BINITIA</b> 1
	Dridge mitianzation	Register A-Ocy	

Bits	Туре	Reset	Description
31:7	RV	000 0000H	Reserved
			Device Hide Lockout: Disables Requester with Requestor ID of zero (Bus#:Dev#:Fct# = 0:0:0) during device hiding.
6	RWS	0	<ul> <li>0 = Requester ID zero access to hidden devices Enabled.</li> <li>1 = Requester ID zero access to hidden devices Disabled.</li> <li>NOTE: This bit is only valid when Device Hiding is enabled (bit 2 is set).</li> </ul>
5	RWS	0	<b>ATU Hide:</b> This bit, when set will hide the ATU from host configuration cycles by not asserting PCI IDSEL# signal on configuration transactions to the ATU.
		•	0 = ATU Visible. 1 = ATU Hidden.
			Private Addressing Enable: When this bit is set, 80333 will hard-code certain address ranges to the <i>secondary</i> segment of the A-segment bridge. This address range is intended to be used for private devices and the private memory space of the ATU (ATUBAR3). The hard coded address range is
			A[63:62] = 10 Secondary side of A-Segment
4	RWS	0	This address range will be forwarded from the primary to the A-segment secondary side and also never forwarded from the secondary to the primary side, irrespective of the setting of the prefetchable base and limit registers. Note that even when the private addressing is enabled, the normal 80333 behavior defined for BME, MSE, IOSE bits in the PCICMD register are still true.
			<ul><li>0 = Private addressing disabled.</li><li>1 = Private addressing enabled.</li></ul>
3	RW	Varies with external state of	Configuration Cycle Retry: This bit, when set will result in a configure retry response on 80333 for all Type 0 configuration transactions from PCI Express to the internal 80333 bridge registers. When clear, Type 0 transactions will be completed normally to the internal 80333 bridge registers.
			on rising edge of <b>PWRGD</b>
			Device Hiding Enable: Enables hiding of devices on the secondary PCI bus in support of I/O processor control.
2	RWS	S 0	<ul> <li>0 = All downstream devices are visible to all requestors.</li> <li>1 = Device numbers 0 to 9 on the immediate secondary bus are hidden from downstream configuration transactions. See "Secondary PCI Devices" on page 61.</li> <li>NOTE: See also bit 6.</li> </ul>
1	RWOS	1	Inbound Configuration Enable: Controls the behavior for Inbound Configuration transactions allowing S/W access by the Intel XScale <sup>®</sup> core processors. Decoding rules are specified in Chapter 3, "Address Translation Unit". This is a write-once register.
			<ul> <li>0 = Inbound Configurations transactions are not claimed from ATU.</li> <li>1 = Inbound Configurations transactions are claimed from the ATU.</li> </ul>
0	RV	1	Reserved.



Bridge millanzation Register B Cognotic BritingBr	Table 84.	<b>Bridge Initialization</b>	<b>Register B-Segment</b>	- BINIT[B]
---	-----------	------------------------------	---------------------------	------------

Bits	Туре	Reset	Description
31:4	RV	000 0000H	Reserved.
3	RW	Varies with external state of <b>RETRY</b> pin on rising edge of <b>PWRGD</b>	<ul> <li>Configuration Cycle Retry: This bit, when set will result in a configure retry response on 80333 for all Type 0 configuration transactions from PCI Express to the internal 80333 bridge registers. When clear, Type 0 transactions will be completed normally to the internal 80333 bridge registers.</li> <li>0 = Normal response to Type 0 Configuration transactions.</li> <li>1 = Response to Type 0 Configuration transaction is the Configure Retry Response.</li> <li>NOTE: Configuration Cycle Retry bits in the BINIT[A] and the ATU PCSR registers must also be cleared to 0 for normal response to all configuration cycles.</li> </ul>
2:0	RV	011	Reserved.

# intel

## 2.9.5.57 PCI Express Advanced Error Capability Identifier -EXPAERR\_CAPID (Offset 100)

This register stores the PCI Express extended capability ID value.

#### Table 85. PCI Express Advanced Error Capability Identifier - EXPAERR\_CAPID

Bits	Туре	Reset	Description
31:20	RO	300H	Next PCI Express Extended Capability Pointer: Points to the PCI Express Power Budgeting Capability as the next capability.
19:16	RO	1H	Advanced Error Capability Version Number: PCI Express Advanced Error Reporting Extended Capability Version Number.
15:0	RO	0001H	Advanced Error Capability ID: PCI Express Extended Capability ID indicating Advanced Error Reporting Capability.



## 2.9.5.58 PCI Express Uncorrectable Error Status -ERRUNC\_STS (Offset 104)

The Uncorrectable Error Status register reports error status of individual uncorrectable error sources. An individual error status bit that is set to "1" indicates that a particular error occurred; software may clear an error status by writing a 1 to the respective bit.

#### Table 86. PCI Express Uncorrectable Error Status - ERRUNC\_STS

Bits	Туре	Reset	Description
31:21	RV	0	Reserved.
20	RWCS	0	Unsupported Request Error Status: Set by 80333 whenever an unsupported request is detected on PCI Express including those that master abort on the switch or signaled by the SHPC (on write data parity errors - configuration cycles and memory cycles).
19	RO	0	ECRC Check: 80333 does not do ECRC check and this bit is never set.
18	RWCS	0	Malformed TLP: 80333 sets this bit when it receives a malformed TLP. Header logging is done.
17	RWCS	0	Receiver Overflow: 80333 would set this when PCI Express interface unit receive buffers overflow.
16	RWCS	0	Unexpected Completion: 80333 sets this bit whenever it receives a completion with a requestor id that does not match side A or side B or when it receives a completion with a matching requestor id but an unexpected tag field. 80333 logs the header of the unexpected completion.
15	RWCS	0	Completer Abort: 80333 sets this bit and logs the header associated with the request when SHPC signals a completer abort. 80333 logs the header.
14	RWCS	0	Completion Time-out: 80333 sets this bit when inbound memory/configure/I/O reads do not receive completions within 16-32ms.
13	RWCS	0	Flow Control Protocol Error Status: 80333 sets this bit when there is a flow control protocol error detected.
12	RWCS	0	Poisoned TLP Received: 80333 sets this bit when a poisoned TLP is received from PCI Express. Note that internal queue errors are not covered by this bit. 80333 logs the header of the poisoned TLP packet.
11:5	RV	0	Reserved.
4	RWCS	0	Data Link Protocol Error: 80333 sets this bit when there is a data link protocol error detected.
3:1	RV	0	Reserved.
0	RWCS	0	Training Error: 80333 does not set this bit.



## 2.9.5.59 PCI Express Uncorrectable Error Mask -ERRUNC\_MSK (Offset 108)

The Uncorrectable Error Mask register controls reporting of individual uncorrectable errors by device to the host bridge via a PCI Express error message and also logging of the header. Refer to PCI Express specifications for details of how the mask bits function. A masked error (respective bit set in mask register) is not reported to the host bridge by 80333. nor is the header logged (status bits unaffected by the mask bit) or the pointer updated. There is a mask bit per bit of the Uncorrectable Error Status register.

Bits	Туре	Reset	Description
31:21	RV	0	Reserved.
20	RWS	0	Unsupported Request Error Status Error Mask: 0 = Not Masked. 1 = Masked.
19	RO	0	ECRC Check Error Mask: 0 = Not Masked. 1 = Masked.
18	RWS	0	Malformed TLP Error Mask: 0 = Not Masked. 1 = Masked.
17	RWS	0	Receiver Overflow Error Mask: 0 = Not Masked. 1 = Masked.
16	RWS	0	Unexpected Completion Error Mask: 0 = Not Masked. 1 = Masked.
15	RWS	0	Completer Abort Error Mask: 0 = Not Masked. 1 = Masked.
14	RWS	0	Completion Time Out Error Mask: 0 = Not Masked. 1 = Masked.
13	RWS	0	Flow Control Protocol Error Status Error Mask: 0 = Not Masked. 1 = Masked.
12	RWS	0	Poisoned TLP Received Error Mask: 0 = Not Masked. 1 = Masked.
11:5	RV	0	Reserved.
4	RWS	0	Data Link Protocol Error Mask: 0 = Not Masked. 1 = Masked.
3:1	RV	0	Reserved.
0	RO	0	Training Error Mask: N/A for 80333.

#### Table 87. PCI Express Uncorrectable Error Mask - ERRUNC\_MSK

## 2.9.5.60 PCI Express Uncorrectable Error Severity -ERRUNC\_SEV (Offset 10C)

The Uncorrectable Error Severity register controls whether an individual uncorrectable error is reported as a fatal error. An uncorrectable error is reported as fatal when the corresponding error bit in the severity register is set. When the bit is cleared, the corresponding error is considered non-fatal.

Bits	Туре	Reset	Description
31:21	RV	0	Reserved.
20	RWC	0	Unsupported Request Error Status Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
19	RO	0	ECRC Check Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
18	RWC	1	Malformed TLP Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
17	RWC	1	Receiver Overflow Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
16	RWC	0	Unexpected Completion Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
15	RWC	0	Completer Abort Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
14	RWC	0	Completion Time Out Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
13	RWC	0	Flow Control Protocol Error Status Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
12	RWC	0	Poisoned TLP Received Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
11:5	RV	0	Reserved.
4	RWC	1	Data Link Protocol Error Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
3:1	RV	0	Reserved.
0	RO	0	Training Error Mask: N/A for 80333.

#### Table 88. PCI Express Uncorrectable Error Severity - ERRUNC\_SEV



## 2.9.5.61 PCI Express Correctable Error Status -ERRCOR\_STS (Offset 110)

The Correctable Error Status register reports error status of individual correctable error sources on a PCI Express device. When an individual error status bit is set to "1" it indicates that a particular error occurred; software may clear an error status by writing a 1 to the respective bit.

Bits	Туре	Reset	Description
31:13	RV	0	Reserved.
12	RWCS	0	Replay Timer Time-out Status: 80333 sets this bit when replay timer time-out happened.
11:9	RV	0	Reserved.
8	RWCS	0	Replay Number Rollover Status: 80333 sets this bit when the replay number rolls over from 11 to 00.
7	RWCS	0	Bad DLLP Status: 80333 sets this bit on CRC errors on DLLP.
6	RWCS	0	Bad TLP Status: 80333 sets this bit on CRC errors on TLP.
5:1	RV	0	Reserved.
0	RWCS	0	Receiver Error: 80333 sets this bit when the physical layer detects a receiver error.

#### Table 89. PCI Express Correctable Error Status - ERRCOR\_STS

## 2.9.5.62 PCI Express Correctable Error Mask -ERRCOR\_MSK (Offset 114)

The Correctable Error Mask register controls reporting of individual correctable errors via ERR\_COR message. A masked error (respective bit set in mask register) is not reported to the host bridge by 80333. There is a mask bit per bit in the Correctable Error Status register.

#### Table 90. PCI Express Correctable Error Mask - ERRCOR\_MSK

Bits	Туре	Reset	Description
31:13	RV	0	Reserved.
12	RWS	0	Replay Timer Time-out Mask: 0 = Not Masked. 1 = Masked.
11:9	RV	0	Reserved.
8	RWS	0	Replay Number Rollover Mask: 0 = Not Masked. 1 = Masked.
7	RWS	0	Bad DLLP Mask: 0 = Not Masked. 1 = Masked.
6	RWS	0	Bad TLP Mask: 0 = Not Masked. 1 = Masked.
5:1	RV	0	Reserved.
0	RWS	0	Receiver Error Mask: 0 = Not Masked. 1 = Masked.



## 2.9.5.63 Advanced Error Control and Capability Register -ADVERR\_CTL (Offset 118)

The register gives the status and control for ECRC checks and also the pointer to the first uncorrectable error that happened.

#### Table 91. Advanced Error Control and Capability Register - ADVERR\_CTL

Bits	Туре	Reset	Description
31:9	PR	0	Preserved.
8	RO	0	ECRC Check Enable: 80333 does not support ECRC check and this bit is reserved.
7	RO	0	ECRC Check Capable – 80333 is not ECRC check capable.
6	RO	0	ECRC Generation Enable – 80333 cannot generated ECRC and this bit is ignored by 80333.
5	RO	0	ECRC Generation Capable – 80333 cannot generated ECRC.
4:0	ROS	0 0000	The First Error Pointer - Identifies the bit position of the first error reported in the Uncorrectable Error Status register. This register rearms itself (which does not change its current value) once the error status bit pointed to by the pointer is cleared by software by writing a 1 to that status bit.

#### 2.9.5.64 PCI Express Transaction Header Log -HDR\_LOG (Offset 11C-12B)

Transaction header log for PCI Express error.

#### Table 92. PCI Express Transaction Header Log - HDR\_LOG

Bits	Туре	Reset	Description
127:0	ROS	0	Header of the PCI Express packet in error. Once an error is logged in this register, it remains locked for further error loggings until the time the software clears the status bit that cause the header log i.e. the error pointer is rearmed to log again.
# intel®

### 2.9.5.65 Uncorrectable PCI/X Status Register -PCI-XERRUNC\_STS (Offset 12C)

### Table 93. Uncorrectable PCI/X Status Register - PCI-XERRUNC\_STS (Sheet 1 of 2)

Bits	Туре	Reset	Description
15:14	RV	00	Reserved.
13	RWCS	0	Internal Bridge Data Error: Accounts internal data errors in 80333 data queues in either direction. 80333 does NOT log any headers for this error.
12	RWCS	0	PCI/X SERR# Detected: 80333 sets this bit whenever it detects the PCI SERR# pin is asserted. There is no header logging associated with the setting of this bit.
11	RWCS	0	PCI/X PERR# Detected: 80333 sets this bit whenever it detects the PCI bus PERR# pin asserted when it is mastering a write (memory, I/O or configuration) or sourcing data during a split/delayed read completion on the PCI/X bus. 80333 logs the header of the transaction in which the PERR# was detected (regardless of the data phase in which it is detected), in the PCI/X header log register. This bit is also set when 80333 receives a PCI-X Split Completion Message with Write data parity error status. The header log under this condition is the command, address and attribute portion during the Split Completion Message.
			Note that this status bit and also the associated header log are always done irrespective of whether the PERR# detected was because of a PCI bus error or because of a forwarded poisoned data. But error message escalation to PCI Express is done only when the PERR# detected was NOT because of forwarded poisoned data.
10	RWCS	0	PCI Delayed Transaction Timer Expiry: This bit is set by 80333 when it detects that a DT time-out has happened on a hardDT read stream or on an inbound I/O or configuration transaction. No header is logged.
9	RWCS	0	PCI/X Uncorrectable Address Parity Error Detected: 80333 sets this bit when it is the target of an inbound transaction and an address parity error was detected by 80333 (regardless of whether the bus mode is PCI or PCI-X). 80333 logs the header of the transaction in which it detected the address/attribute parity error in the PCI/X header log register.
8	RWCS	0	PCI-X Uncorrectable Attribute Parity Error Detected: 80333 sets this bit when it is the target of an inbound transaction and an attribute parity error was detected by 80333). 80333 logs the header of the transaction in which it detected the address/attribute parity error in the PCI/X header log register.
7	RWCS	0	PCI-X Uncorrectable Data Parity Error Detected: 80333 sets this bit in all PCI modes (PCI, PCI-X) when it is the target of an inbound transaction or when it is mastering a PCI delayed read with target sourcing data to 80333, and data parity error was detected by 80333. 80333 logs the header of the transaction in which it detected the data parity error in the PCI/X header log register.
6	RWCS	0	Split Completion Message Data Error: This bit is set when a split completion message is received with an uncorrectable data parity error. <sup>a</sup>
5	RWCS	0	Unexpected Split Completion: This bit is set when a completion is received from pCI-X that matches the bus number range on the primary side of 80333 but the Requester ID:tag combination does not match one of the NP tx that 80333 has outstanding on PCI-X
4	RV	0	Reserved.
3	RWCS	0	PCI/X Detected Master Abort: 80333 set this bit when it is the master of a request transaction on the PCI/X bus and received a master abort. Header is logged for that transaction. This bit is also set when 80333 receives a PCI-X Split Completion Message with Master Abort Status. The header log under this condition is the command, address and attribute portion during the Split Completion Message.



### Table 93. Uncorrectable PCI/X Status Register - PCI-XERRUNC\_STS (Sheet 2 of 2)

Bits	Туре	Reset	Description
2	RWCS	0	PCI/X Detected Target Abort (optional in specification): 80333 set this bit when it is the master of a request transaction on the PCI/X bus and received a target abort. Header is logged for that transaction. This bit is also set when 80333 receives a PCI-X Split Completion Message with Target Abort Status. The header log under this condition is the command, address and attribute portion during the Split Completion Message.
1	RWCS	0	PCI-X Detected Split Completion Master Abort: 80333 sets this bit when a split completion it sends on PCI-X bus master aborts. 80333 logs the header of the split completion.
0	RWCS	0	PCI-X Detected Split Completion Target Abort (optional in specification): 80333 sets this bit when a split completion it sends on PCI-X bus target aborts. 80333 logs the header.

a. 80333 aliases the completion status to CA when it detects this condition.



### 2.9.5.66 Uncorrectable PCI/X Error Mask Register -CIXERRUNC\_MSK (Offset 130)

This register masks the reporting of PCI/X uncorrectable errors. There is one mask bit per error. Note that the status bits are set in the status register irrespective of whether the mask bit is on or off. The mask bit also affects the header log for the PCI/X transaction. When the mask bit is on, the header is not logged and no error message is generated on PCI Express and the PCI/X pointer is not updated. Refer to Chapter 11, "SMBus Interface Unit" for full details of error escalation via the advanced error capability or the legacy mode SERR bits.

### Table 94. Uncorrectable PCI/X Error Mask Register - PCI-XERRUNC\_MSK (Sheet 1 of 2)

Bits	Туре	Reset	Description
15:14	RV	00	Reserved.
13	RWC	0	Internal Bridge Data Error Mask: 0 = Not Masked. 1 = Masked.
12	RWC	1	PCI/X SERR# Detected Mask: 0 = Not Masked. 1 = Masked.
11	RWC	0	PCI/X PERR# Detected Mask: 0 = Not Masked. 1 = Masked.
10	RWC	1	PCI Delayed Transaction Timer Expiry Mask: 0 = Not Masked. 1 = Masked.
9	RWC	1	PCI/X Uncorrectable Address Parity Error Detected Mask: 0 = Not Masked. 1 = Masked.
8	RWC	1	PCI-X Uncorrectable Attribute Parity Error Detected Mask: 0 = Not Masked. 1 = Masked.
7	RWC	1	PCI-X Uncorrectable Data Parity Error Detected Mask: 0 = Not Masked. 1 = Masked.
6	RWS	0	Split Completion Message Data Error Mask 0 = Not Masked. 1 = Masked.
5	RWS	1	Unexpected Split Completion Error Mask 0 = Not Masked. 1 = Masked.
4	PR	0	Preserved.
3	RWC	1	PCI/X Detected Master Abort Mask: 0 = Not Masked. 1 = Masked.



Table 94. Unco	rrectable PCI/X Error Mask Regist	ter - PCI-XERRUNC_MSK (Sheet 2 of 2)
----------------	-----------------------------------	--------------------------------------

Bits	Туре	Reset	Description
2	RWC	0	PCI/X Detected Target Abort (optional in specification) Mask: 0 = Not Masked. 1 = Masked.
1	RWC	0	PCI-X Detected Split Completion Master Abort Mask: 0 = Not Masked. 1 = Masked.
0	RWC	0	PCI-X Detected Split Completion Target Abort (optional in specification) Mask: 0 = Not Masked. 1 = Masked.



### 2.9.5.67 Uncorrectable PCI/X Error Severity Register -PCI-XERRUNC\_SEV (Offset 134)

This register controls the severity of reporting of PCI-X uncorrectable errors. There is one mask bit per error. When a bit is set to a 1, the corresponding error, when enabled, will generate an ERR\_FATAL message on PCI Express. When a bit is cleared to 0, it will cause a ERR\_NONFATAL on PCI Express.

Table 95. Uncorrectal	ole PCI/X Error Severity	Register - PCI-XERRUNC	_SEV (Sheet 1 of 2)
-----------------------	--------------------------	------------------------	---------------------

Bits	Туре	Reset	Description
15:14	RV	00	Reserved.
13	RWC	0	Internal Bridge Data Error Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
12	RWC	1	PCI/X SERR# Detected Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
11	RWC	0	PCI/X PERR# Detected Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
10	RWC	0	PCI Delayed Transaction Timer Expiry Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
9	RWC	1	PCI/X Uncorrectable Address Parity Error Detected Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
8	RWC	1	PCI-X Uncorrectable Attribute Parity Error Detected Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
7	RWC	0	PCI-X Uncorrectable Data Parity Error Detected Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
6	RWS	1	Split Completion Message Data Error Severity 0 = ERR_NONFATAL. 1 = ERR_FATAL.
5	RWS	0	Unexpected Split Completion Error Severity 0 = ERR_NONFATAL. 1 = ERR_FATAL.
4	PR	0	Preserved.
3	RWC	0	PCI/X Detected Master Abort Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.



## Table 95. Uncorrectable PCI/X Error Severity Register - PCI-XERRUNC\_SEV (Sheet 2 of 2)

Bits	Туре	Reset	Description
2	RWC	0	PCI/X Detected Target Abort (optional in specification) Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
1	RWC	0	PCI-X Detected Split Completion Master Abort Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.
0	RWC	0	PCI-X Detected Split Completion Target Abort (optional in specification) Severity: 0 = ERR_NONFATAL. 1 = ERR_FATAL.



### 2.9.5.68 Uncorrectable PCI/X Error Pointer -PCI-XERRUNC\_PTR (Offset 138)

This register points to the first error that occurred. This register is rearmed (which does not change its current value) when the error status register corresponding to the error that this register is pointing to is cleared by software writing a 1 to the bit.

### Table 96. Uncorrectable PCI/X Error Pointer Register - PCI-XERRUNC\_PTR

Bits	Туре	Reset	Description
15:4	PR	000H	Preserved.
3:0	ROS	0Н	PCI First Error Pointer: This register points to the first error that got logged in the status register (providing the corresponding mask bit is clear and the pointer is rearmed). This register rearms itself when the status bit corresponding to the error that this register is pointing to is cleared by software writing a 1 to the bit. The PCI/X Uncorrectable error pointer behavior is valid in both secondary PCI modes: PCI and PCI-X.

### 2.9.5.69 Uncorrectable PCI/X Error Transaction Header Log -PCI-XHDR\_LOG (Offset 13C-14B)

Transaction header log for PCI errors. The log in this register corresponds to one of the status bits set in the PCI/X uncorrectable status register. Once an error is logged in this register, it remains locked for further error loggings until the time the software clears the status bit that cause the header log (i.e. the error pointer is rearmed to log again). This register operation is valid in both secondary PCI modes: PCI and PCI-X.

### Table 97. Uncorrectable PCI/X Header Log - PCI-XHDR\_LOG

Bits	Туре	Reset	Description
127:64	ROS	0	Address: These bits capture the 64-bit address of the transaction in which an error was detected. In case of a 32-bit address, the upper address bits are all zeros. Address is logged on all error conditions.
63:44	RV	0	Reserved.
43:40	ROS	ОH	DAC Bus Command: This captures the value of C/BE[3:0] during the second address phase of a DAC transaction.
39:36	ROS	ОH	Bus Command: This captures the value of C/BE[3:0] during the 1st address phase of the transaction.
35:0	ROS	0	Attribute: This carries the attribute of the transaction. The field is arranged as <cbe[3:0]#::ad[31:0]> during the attribute phase of the transaction. When the bus is in PCI mode, this register is all zeros.</cbe[3:0]#::ad[31:0]>

### 2.9.5.70 Uncorrectable PCI/X Data Error Log -PCI-XD\_LOG (Offset 14C-153)

This register is logged for all correctable or uncorrectable data parity errors.

### Table 98. Uncorrectable PCI/X Data Error Log - PCI-XD\_LOG

Bits	Туре	Reset	Description
63:0	ROS	0	PCI-X Data Log: This register is logged with the PCI data bus value whenever 80333 is the target of a data transfer and it detects a data parity error. This register is not defined when the log valid bit in the error log and control register is not set. This register re-arms itself for loading again, when software clears the log valid bit by writing a one that bit. For 32-bit data transfers, only the lower 32 bits are logged.



### 2.9.5.71 Other PCI/X Error Logs and Control -PCI-XERRLOGCTL(Offset 154)

This register is contains bits logged for uncorrectable data parity errors (in PCI or PCI-X), uncorrectable address/attribute parity errors and PCI REQ# line of failure

### Table 99. Other PCI/X Error Logs and Control - PCI-XERRLOGCTL (Sheet 1 of 2)

Bits	Туре	Reset	Description	
31:19	RV	0	Reserved.	
18	ROS	0	Data Log: When 0 indicates the data log is from a parity error. This bit is logged along with the DLOG register and is rearmed when the log valid bit cleared. This is also only valid when the log valid bit is set by 80333, and will not be set by 80333.	
			PCI-X Attribute Parity (PP): This bit represents the parity detected in the attribute phase of a request and completion. When the 80333 is driving, it is the value driven. When the 80333 is receiving, it is the value captured.	
17	ROS	0	This bit is logged along with the PCI-XHDRLOG register when there is an attribute or address parity error. This bit is not loaded for any other error conditions. Once an error is logged in this bit, this bit is rearmed to load again until software clears the corresponding status bit in the PCI-X uncorrectable status register.	
			PCI Address High (PAH): This bit represents the parity detected in the second phase (upper 32-bits) of a dual address cycle. It is forced to '0' when the address was a single address cycle. When 80333 is driving, it is the value driven. When the 80333 is receiving, it is the value captured.	
16	ROS	0	This bit is logged along with the PCI-XHDRLOG register when there is an address or attribute parity error (this bit is never loaded independently of the PCI-XHDRLOG register). This bit is not loaded for any other error conditions. Once an error is logged in this bit, this bit is NOT rearmed to load again until software clears the corresponding status bit in the PCI-X uncorrectable status register.	
15	ROS	POS 0	PCI Address Low (PAL): For PCI-X requests and all PCI cycles, this bit represents the parity detected in the first phase (lower 32-bits) of a dual address cycle, or just the address of a regular address cycle. For PCI-X completions, this bit represents the first clock (requester attributes) driven in the completion cycle. When 80333 is driving, it is the value driven. When the 80333 is receiving, it is the value captured.	
		Koo	Ŭ	This bit is logged along with the PCI-XHDRLOG register when there is an address or attribute parity error (this bit is never loaded independently of the PCI-XHDRLOG register). This bit is not loaded for any other error conditions. Once an error is logged in this bit, this bit is NOT rearmed to load again until software clears the corresponding status bit in the PCI-X uncorrectable status register.
14	RWCS	0	REQ# Log Valid: Set when REQ# log bits are valid. Clearing this bit will re-enable logging into the REQ# log register bits.	
			REQ# Log: These bits capture the REQ# of the PCI agent mastering the transaction when 80333 detected an uncorrectable or correctable address, attribute or data parity error. That is, the REQ# log is valid when either of the three conditions occur that cause either of bits 9:7 to be set. Once a log is made in the REQ# log, further logging is of the REQ# log bits is stopped till the REQ# log valid bit is cleared. (note that this register is not dependent on the clearing of status bits in PCIXERRUNC Status Register or the PX ECCSTS register, to rearm itself).	
13:11	KOS	ROS	0	000 = REQ0 001 = REQ1 010 = REQ2 011 = REQ3 100 = REQ4 101 = REQ5 110 = REQ6 111 = Reserved
10	RV	0	Reserved	



### Table 99. Other PCI/X Error Logs and Control - PCI-XERRLOGCTL (Sheet 2 of 2)

Bits	Туре	Reset	Description	
9	RWCS	0	Log Valid: This is set by 80333 whenever it logs a value in the DLOG register and the byte enable log bits in this register. Software clears this register by writing a 1 to 1, and this will rearm the DLOG register and the byte enable log register bits to start loading again.	
8	ROS	0	Data Bus Width: This bit is set when the data logged in the DLOG register is 64 bits. Otherwise this bit is clear. When clear the upper 32 bits of the DLOG register are invalid.	
7:0	ROS	0	PCI-X Byte Enable Log: This error is logged whenever 80333 is the target of a data transfer and it detects a data parity error. This register is logged along with the DLOG register. This register is not defined when the log valid bit (bit 9 above) is not set. This register re-arms itself for loading again, when software clears the log valid bit by writing a one that bit.	



### 2.9.5.72 Internal Arbiter Control Register -ARB\_CNTRL (Offset 16A)

### Table 100. Internal Arbiter Control A-Segment - ARB\_CNTRL[A]

Bits	Туре	Reset	Description	
15:9	RV	0	Reserved.	
8	RW	0	<ul> <li>Bus Parking Control: Controls the bus parking behavior of the internal arbiter:</li> <li>0 = Bus is parked on the last PCI agent using the bus.</li> <li>1 = Bus is always parked on 80333.</li> </ul>	
7	RW	1	<ul> <li>Bridge Priority Ring Allocation: Priority ring allocation for 80333 bridge requests:</li> <li>0 = 80333 is in the low priority ring.</li> <li>1 = 80333 is in the high priority ring.</li> </ul>	
6	RV	0	Reserved.	
5	RW	1	ATU Ring Allocation: Priority ring allocation for 80333 ATU requests: 0 = The ATU is in the low priority ring of the internal arbiter. 1 = The ATU is in the high priority ring of the internal arbiter.	
4	PR	1	Preserved.	
3:0	RW	1111	<ul> <li>PCI Master Priority Ring Allocation: Bit 0 corresponds to REQ#[0], 1 corresponds to REQ#[1] and so on:</li> <li>0 = The corresponding master is in the low priority ring of the internal arbiter.</li> <li>1 = The corresponding master is in the high priority ring of the internal arbiter.</li> </ul>	

### Table 101. Internal Arbiter Control B-Segment - ARB\_CNTRL[B]

Bits	Туре	Reset	Description	
15:9	RV	0	Reserved.	
8	RW	0	<ul> <li>Bus Parking Control: Controls the bus parking behavior of the internal arbiter:</li> <li>0 = Bus is parked on the last PCI agent using the bus.</li> <li>1 = Bus is always parked on 80333.</li> </ul>	
7	RW	1	<ul> <li>Bridge Priority Ring Allocation: Priority ring allocation for 80333 bridge requests:</li> <li>0 = 80333 is in the low priority ring.</li> <li>1 = 80333 is in the high priority ring.</li> </ul>	
6	RV	0	Reserved.	
5	PR	1	Preserved.	
4:0	RW	1 1111	<ul> <li>PCI Master Priority Ring Allocation: Bit 0 corresponds to REQ#[0], 1 corresponds to REQ#[1] and so on:</li> <li>0 = The corresponding master is in the low priority ring of the internal arbiter.</li> <li>1 = The corresponding master is in the high priority ring of the internal arbiter.</li> </ul>	



## 2.9.5.73 Strap Status Register - SSR (Offset 170)

Indicates the status of various reset straps in 80333.

Bits	Туре	Reset	Description		
15	RO	0	Configure Retry Strap: Reflects the RETRY strap value at the rising edge of <b>PWRGD</b> .		
14:12	RV	0	Reserved.		
11:8	RO	Strap	PCI Slot Count: B_HSLOT[3:0]# values that was sampled on the rising edge of <b>PWRGD</b> .		
7:1	RO	Strap	Manageability Address (MA): These 7 bits represent the address the SMBus slave port will respond to when an access is attempted. This register will have the following value:         Bit       Value         7       '1'         6       '1'         5       PBI A[19]         4       '0'         3       PBI A[18]         2       PBI A[17]         1       PBI A[16]         Only the value from function 0 is valid.		
0	RO	0	P133EN Status: Reflects the status of the X_133EN pin sampled at rising edge of <b>PWRGD</b> . This serves as a dashboard for the BIOS.		

 Table 102.
 Strap Status Register - SSR



### 2.9.5.74 Power Budgeting Enhanced Capability Header -PWRBGT\_CAPID (Offset 300)

This register defines the capability identifier.

### Table 103. Offset 300: PWRBGT\_HDR - Power Budgeting Enhanced Capability Header

Bits	Туре	Reset	Description	
31:20	RO	000H	Next PCI Express Extended Capability Pointer: This is the last capability.	
19:16	RO	1H	Power Budgeting Capability Version Number: PCI Express Power Budgeting Capability Version Number.	
15:0	RO	0004H	Power Budgeting Capability ID: PCI Express Power Budgeting Capability ID indicating Advanced Error Reporting Capability.	

### 2.9.5.75 Power Budgeting Data Select Register -PWRBGT\_DSEL (Offset 304)

### Table 104. Power Budgeting Data Select Register - PWRBGT\_DSEL

Bits	Туре	Reset	Description	
7:0	RW	0	Data Select: This read-write register indexes the Power Budgeting Data reported through the Data register and selects the DWORD of Power Budgeting Data that should appear in the Data Register. Index values for this register start at 0 to select the first DWORD of Power Budgeting Data; subsequent DWORDs of Power Budgeting Data are selected by increasing index values. A value of 0 selects the DWORD data starting at address 0x314 to appear in the data register at offset 0x308, a value of 1 selects the DWORD data starting at address 0x318 to appear in the in the data registers at offset 0x308 and so on. Values greater than 23 for this register will report all zeros in the data register.	

### 2.9.5.76 Power Budgeting Data Register -PWRBGT\_DATA (Offset 308)

### Table 105. Power Budgeting Data Register - PWRBGT\_DATA

Bits	Туре	Reset	Description
31:0	RO	0000 0000H	Data: This read-only register returns the DWORD of Power Budgeting Data selected by the Data Select Register. 80333 reports its power consumption for PM states D0, D3; Types idle, sustained, Max; Power rails 12V, 3.3V, 1.8V, Thermal.



### 2.9.5.77 Power Budgeting Capability Register -PWRBGT\_CAP (Offset 30C)

### Table 106. Power Budgeting Capability Register - PWRBGT\_CAP

Bits	Туре	Reset	Description
7:1	RV	000 0000	Reserved.
0	RWO	0	System Allocated – This bit when set indicates that the power budget for the device is included within the system power budget. Reported Power Budgeting Data for this device should be ignored by software for power budgeting decisions when this bit is set.

### 2.9.5.78 Power Budgeting Information Registers 23:0 -PWRBGT\_INFO[23:0] (Offset 314-370)

80333 implements a set of 24 Power Budgeting information registers.

### Table 107. Power Budgeting Information Registers 0:23 - PWRBGT\_INFO[0:23]

Bits	Туре	Reset	Description	
31:21	RV	0000H	Reserved.	
31:21	RV	0000H	Reserved. Init software can program this field to report various power consumption values in various power states. Refer to <i>PCI Express Specification</i> , Revision 1.0a for the format of this field for reporting the various power consumption values. The default values are chose to represent the 80333 power consumption with both segments in PCI-X 133MHz. Default Values for PWRBGT_INFO[23:0] Registers (bits 20:0) REGISTER Fn#0 Fn#1 PWRBGT_INFO[0]010130010110 PWRBGT_INFO[1]018135 018111 PWRBGT_INFO[2]038136 038111	
20:0	RW	Values in description.	PWRBGT_INFO[3]016130 016110         PWRBGT_INFO[4]01E130 01E130         PWRBGT_INFO[5]03E130 03E130         PWRBGT_INFO[6]050125 050125         PWRBGT_INFO[6]050125 050125         PWRBGT_INFO[7]058126 058126         PWRBGT_INFO[8]078134 078134         PWRBGT_INFO[9]056125 056125         PWRBGT_INFO[10]05E125 05E125         PWRBGT_INFO[10]05E125 07E127         PWRBGT_INFO[23:12]000000 000000	



**This Page Left Intentionally Blank** 

# intel

# **Address Translation Unit**

# 3

This chapter describes the operation modes, setup, and implementation of the module which interfaces between the PCI bus and the Intel<sup>®</sup> 80333 I/O processor (80333) internal bus.

# 3.1 Overview

As indicated in Figure 10, the Address Translation Unit (ATU) — the interface between the PCI bus and the on-chip internal bus — consists of the Address Translation Unit (ATU), the Expansion ROM Unit and the Messaging Unit (MU) described in Chapter 4, "Messaging Unit"

The ATU supports both inbound and outbound address translation. The ATU provides access between the PCI bus and the 80333 internal bus. The ATU and the MU share PCI address space.

Transactions initiated on the PCI bus and targeted at the 80333 internal bus are referred to as *inbound transactions* (PCI to internal bus). Transactions initiated on the 80333 internal bus and targeted at the PCI bus are referred to as *outbound transactions* (internal bus to PCI). The ATU accepts multiple inbound or outbound transactions and processes them simultaneously.

During inbound transactions, the ATU converts PCI addresses (initiated by a PCI bus master) to internal bus addresses and initiates the data transfer on the 80333 internal bus. During outbound transactions, the ATU converts internal bus addresses to PCI addresses and initiates the data transfer on the PCI bus.

The Messaging Unit provides a mechanism for the system processor and the 80333 to transfer control information. The Messaging Unit occupies the first 4 Kbytes in Memory Window 0 of the ATU address space. PCI masters on the PCI interface of the 80333 access the MU by addressing the ATU anywhere in the first 4 KB offset from the ATU Base Address Register 0.

The Expansion ROM provides the PCI mechanism for downloading device/board driver code during system boot sequence. It consists of a separate inbound address range which accesses a Flash EPROM device connected through the 80333 memory controller. Refer to the *PCI Local Bus Specification*, Revision 2.3 for details of Expansion ROM usage.

The Address Translation Unit, the Expansion ROM Translation Unit, and the Messaging Unit appear as a single PCI device on the PCI bus. That is, the 80333 is a single-function device.

The ATU supports the PCI 64-bit and 66 MHz extensions providing up to 532 Mbytes/sec of PCI bandwidth as well as the PCI-X 64-bit and 133 MHz extensions providing up to 1064 Mbytes/sec of PCI bandwidth. On the internal interface, the ATU implements the 80333 internal bus protocol which provides for a maximum of 2.7 Gbytes/sec using 64-bit/333 MHz signaling.

The ATU includes three extended capability headers that implement Power Management capability as defined by the *PCI Bus Power Management Interface Specification*, Revision 1.1, MSI capability as defined by *PCI Local Bus Specification*, Revision 2.3, and PCI-X capability as defined by *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a.

The functionality of the ATU is described in the following sections. The ATU and the MU have a memory-mapped register interface that is visible from either the PCI interface, the internal bus interface, or both.



### Figure 10. ATU Block Diagram



intel



Figure 11. ATU Queue Architecture Block Diagram

# 3.2 ATU Address Translation

The ATU allows PCI masters on the PCI bus to initiate transactions to the 80333 internal bus and allows the Intel XScale<sup>®</sup> core (ARM\* architecture compliant) to initiate transactions to the PCI bus.

The ATU implements an address windowing scheme to determine which addresses to claim and translate to the destination bus.

- The address windowing mechanism for inbound translation is described in Section 3.2.1.1, "Inbound Address Translation" on page 165
- The address windowing mechanism for outbound translation is described in Section 3.2.2, "Outbound Transactions- Single Address Cycle (SAC) Internal Bus Transactions" on page 175 and Section 3.2.3, "Outbound Write Transaction" on page 183

The ATU has the ability to accept up to eight inbound PCI read transactions and four inbound PCI write transactions simultaneously. Also, the ATU has the ability to accept up to eight outbound internal bus read transactions and four outbound internal bus write transactions simultaneously. Refer to Figure 11 and Section 3.5 for details of the ATU queue architecture.

The ATU unit allows for recognition and generation of multiple PCI cycle types. Table 108 shows the PCI and PCI-X commands supported for both inbound and outbound ATU transactions. The type of operation seen by the ATU on inbound transactions is determined by the PCI master who initiates the transaction. Claiming an inbound transaction depends on the address range programmed within the inbound translation window. The type of transaction used by the ATU on outbound transactions generated by the core processor is determined by the internal bus address and the fixed outbound windowing scheme.

ATU supports the 64-bit addressing specified by the *PCI Local Bus Specification*, Revision 2.3. This 64-bit addressing extension is supported for both inbound and outbound data transactions. This is in addition to the 64-bit data extensions supported by the 80333.

ATU does not support exclusive access using the PCI LOCK# signal. Also, the ATU does not guarantee atomicity for outbound transactions.

# intel

### Table 108. ATU Command Support

PCI Command Encoding	PCI Command Type	PCI-X Command Type	Claimed on Inbound Transactions on PCI Bus	Generated by Outbound Transactions on PCI Bus	Valid Internal Bus Command
0000	Interrupt Acknowledge	Interrupt Acknowledge	No	No	Interrupt Acknowledge
0001	Special Cycle	Special Cycle	No	No	Special Cycle
0010	I/O Read	I/O Read	No	Yes	I/O Read
0011	I/O Write	I/O Write	No	Yes	I/O Write
0100	reserved	reserved	No	No	reserved
0101	reserved	reserved	No	No	reserved
0110	Memory Read	Memory Read DWORD	Yes	Yes	Memory Read DWORD
0111	Memory Write	Memory Write	Yes	Yes	Memory Write
1000	reserved <sup>a</sup>	Alias to Memory Read Block	Yes	No	Alias to Memory Read Block
1001	reserved <sup>a</sup>	Alias to Memory Write Block	Yes	No	Alias to Memory Write Block
1010	Configuration Read	Configuration Read	Yes	Yes	Configuration Read
1011	Configuration Write	Configuration Write	Yes	Yes	Configuration Write
1100	Memory Read Multiple	Split Completion	Yes	Yes	Split Completion
1101	Dual Address Cycle	Dual Address Cycle	Yes	Yes	Dual Address Cycle
1110	Memory Read Line	Memory Read Block	Yes	Yes <sup>b</sup>	Memory Read Block
1111	Memory Write and Invalidate	Memory Write Block	Yes	Yes	Memory Write Block

a. The ATU claims PCI commands 8 and 9 when issued as Dual Address Cycle (DAC). See Intel<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #19 for more details.

b. PCI-X mode only.

Inbound and outbound ATU transactions are best described by the data flows used on the PCI bus and the 80333 internal bus during read and write operations. The following sections describe read and write operations for inbound ATU transactions (PCI to internal bus) and outbound transactions (internal bus to PCI).



# 3.2.1 Inbound Transactions

Inbound transactions which target the ATU are translated and executed on the 80333 internal bus. As a PCI target, the ATU is capable of accepting all PCI memory read and write operations as either a 32-bit or a 64-bit PCI target. In the conventional PCI mode *Memory Write* and *Memory Write and Invalidate* operations are performed as posted operations and all memory read operations are performed as delayed reads. In the PCI-X mode *Memory Write, Memory Write Block,* and *Alias to Memory Read DWORD, Memory Read Block,* and *Alias to Memory Read Block*, operations are executed as split transactions. The ATU is capable of accepting configuration read and write cycles. In the conventional PCI mode, *Configuration Writes* are performed as delayed read operations and *Configuration Reads* are performed as delayed read operations. In the PCI-X mode, both *Configuration Writes* and *Configuration Reads* are performed as gent as gent operations.

Inbound memory write transactions have their addresses entered into the inbound write address queue (IWADQ) and data entered into the inbound write data queue (IWQ). The IWQ/IWADQ pair are capable of holding up to 4 write operations up to the size of the data queue. Inbound configuration writes use the inbound delayed write queue (IDWQ) for address and data. Refer to Section 3.5 for details of queue operation. Inbound read operations (memory and configuration) have their address entered into the inbound transaction queue (ITQ) and the data is returned to the PCI master in the inbound read queue (IRQ). The ITQ is capable of holding up to 8 delayed read requests (split read requests when operating in the PCI-X mode).

In the conventional PCI mode, for inbound transactions, the ATU is a slave on the PCI bus and is a requester on the internal bus. PCI slave operation is defined in the *PCI Local Bus Specification*, Revision 2.3. In the PCI-X mode, for inbound transactions, the ATU initially is a target on the PCI bus and becomes an initiator when performing split completion transactions, and is an initiator on the internal bus.

164



### 3.2.1.1 Inbound Address Translation

The ATU allows external PCI bus initiators to directly access the internal bus. These PCI bus initiators can read or write 80333 memory-mapped registers or 80333 local memory space. The process of inbound address translation involves two steps:

- 1. Address Detection.
  - Determine when the 32-bit PCI address (64-bit PCI address during DACs) is within the address windows defined for the inbound ATU.
  - Claim the PCI transaction with medium DEVSEL# timing in the conventional PCI mode and with Decode A DEVSEL# timing in the PCI-X mode.
- 2. Address Translation.
  - Translate the 32-bit PCI address (lower 32-bit PCI address during DACs) to a 32-bit 80333 internal bus address.

The ATU uses the following registers in inbound address window 0 translation:

- Inbound ATU Base Address Register 0
- Inbound ATU Limit Register 0
- Inbound ATU Translate Value Register 0

The ATU uses the following registers in inbound address window 2 translation:

- Inbound ATU Base Address Register 2
- Inbound ATU Limit Register 2
- Inbound ATU Translate Value Register 2

The ATU uses the following registers in inbound address window 3 translation:

- Inbound ATU Base Address Register 3
- Inbound ATU Limit Register 3
- Inbound ATU Translate Value Register 3
- *Note:* Inbound Address window 1 is not a translate window. Instead, window 1 may be used to allocate host memory for Private Devices. Inbound Address window 3 does not reside in the standard section of the configuration header (offsets 00H 3CH), thus the host BIOS does not configure window 3. Window 3 is intended to be used as a special window into local memory for private PCI agents controlled by the 80333 in conjunction with the Private Memory Space of the bridge. PCI Express-to-PCI Bridge in 80333.

Inbound address detection is determined from the 32-bit PCI address, (64-bit PCI address during DACs) the base address register and the limit register. In the case of DACs none of the upper 32-bits of the address is masked during address comparison. The algorithm for detection is:

### Equation 1. Inbound Address Detection

When PCI\_Address [31:0] and Limit\_Register[31:0] == Base\_Register[31:0] and PCI\_Address [63:32] == Base\_Register[63:32] (for DACs only) the PCI Address is claimed by the Inbound ATU.



Figure 12 shows an example of inbound address detection.

### Figure 12. Inbound Address Detection



The incoming 32-bit PCI address (lower 32-bits of the address in case of DACs) is bitwise ANDed with the associated inbound limit register. When the result matches the base register (and upper base address matches upper PCI address in case of DACs), the inbound PCI address is detected as being within the inbound translation window and is claimed by the ATU.

*Note:* The first 4 Kbytes of the ATU inbound address translation window 0 are reserved for the Messaging Unit. See Section 3.3, "Messaging Unit" on page 187.

Once the transaction is claimed, the address must be translated from a PCI address to a 32-bit internal bus address. In case of DACs upper 32-bits of the address is simply discarded and only the lower 32-bits are used during address translation. The algorithm is:

### **Equation 2.** Inbound Translation

Intel<sup>®</sup> I/O processor Internal Bus Address = (PCI\_Address[31:0] and ~Limit\_Register[31:0]) | ATU\_Translate\_Value\_Register[31:0].

The incoming 32-bit PCI address (lower 32-bits in case of DACs) is first bitwise ANDed with the bitwise inverse of the limit register. This result is bitwise ORed with the ATU Translate Value and the result is the internal bus address. This translation mechanism is used for all inbound memory read and write commands excluding inbound configuration read and writes. Inbound configuration cycle translation is described in Section 3.2.1.4, "Inbound Configuration Cycle Translation" on page 173.

In the PCI mode for inbound memory transactions, the only burst order supported is Linear Incrementing. For any other burst order, the ATU signals a Disconnect after the first data phase. The PCI-X supports linear incrementing only, and hence above situation is not encountered in the PCI-X mode.



Figure 13 shows an inbound translation example for 32-bit addressing. This example would hold true for an inbound transaction from PCI bus.

Figure 13. Inbound Translation Example





### 3.2.1.2 Inbound Write Transaction

An inbound write transaction is initiated by a PCI master and is targeted at either 80333 local memory or a 80333 memory-mapped register.

Data flow for an inbound write transaction on the PCI bus is summarized as:

- The ATU claims the PCI write transaction when the PCI address is within the inbound translation window defined by the ATU Inbound Base Address Register (and Inbound Upper Base Address Register during DACs) and Inbound Limit Register.
- When the IWADQ has at least one address entry available and the IWQ has at least one buffer available, the address is captured and the first data phase is accepted.
- The PCI interface continues to accept write data until one of the following is true:
  - The initiator performs a disconnect.
  - The transaction crosses a buffer boundary.
- When an address parity error is detected during the address phase of the transaction, the address parity error mechanisms are used. Refer to Section 3.7.1 for details of the address parity error response.
- When operating in the PCI-X mode when an attribute parity error is detected, the attribute parity error mechanism described in Section 3.7.1 is used.
- When a data parity error is detected while accepting data, the slave interface sets the appropriate bits based on PCI specifications. No other action is taken. Refer to Section 3.7.2.6 for details of the inbound write data parity error response.

Once the PCI interface places a PCI address in the IWADQ, when IWQ has received data sufficient to cross a buffer boundary or the master disconnects on the PCI bus, the ATUs internal bus interface becomes aware of the inbound write. When there are additional write transactions ahead in the IWQ/IWADQ, the current transaction remains posted until ordering and priority have been satisfied (Refer to Section 3.5.3) and the transaction is attempted on the internal bus by the ATU internal master interface. The ATU does not insert target wait states nor do data merging on the PCI interface, when operating in the PCI mode.

In the PCI-X mode memory writes are always executed as immediate transactions, while configuration write transactions are processed as split transactions. The ATU generates a Split Completion Message, (with Message class = 0h - Write Completion Class and Message index = 00h - Write Completion Message) once a configuration write is successfully executed.

Also, when operating in the PCI-X mode a write sequence may contain multiple write transactions. The ATU handles such transactions as independent transactions.

Data flow for the inbound write transaction on the internal bus is summarized as:

- The ATU internal bus master requests the internal bus when IWADQ has at least one entry with associated data in the IWQ.
- When the internal bus is granted, the internal bus master interface initiates the write transaction by driving the translated address onto the internal bus. For details on inbound address translation, see Section 3.2, "ATU Address Translation" on page 162.
- When **IB\_DEVSEL**# is not returned, a master abort condition is signaled on the internal bus. The current transaction is flushed from the queue and **SERR**# may be asserted on the PCI interface.
- The ATU initiator interface asserts **IB\_REQ64**# to attempt a 64-bit transfer. When **IB\_ACK64**# is not returned, a 32-bit transfer is used. Transfers of less than 64-bits use the **IB\_C/BE[7:0]**# to mask the bytes not written in the 64-bit data phase. Write data is transferred from the IWQ to the internal bus when data is available and the internal bus interface retains internal bus ownership.
- The internal bus interface stops transferring data from the current transaction to the internal bus when one of the following conditions becomes true:
  - The internal bus initiator interface loses bus ownership. The ATU internal initiator terminates the transfer (initiator disconnection) at the next ADB (for the internal bus ADB is defined as a naturally aligned 128-byte boundary) and attempt to reacquire the bus to complete the delivery of remaining data using the same sequence ID but with the modified starting address and byte count.
  - A Disconnect at Next ADB is signaled on the internal bus from the internal target. When the transaction in the IWQ completes at that ADB, the initiator returns to idle. When the transaction in the IWQ is not complete, the initiator attempts to reacquire the bus to complete the delivery of remaining data using the same sequence ID but with the modified starting address and byte count.
  - A Single Data Phase Disconnect is signaled on the internal bus from the internal target. When the transaction in the IWQ needs only a single data phase, the master returns to idle. When the transaction in the IWQ is not complete, the initiator attempts to reacquire the bus to complete the delivery of remaining data using the same sequence ID but with the modified starting address and byte count.
  - The data from the current transaction has completed (satisfaction of byte count). An initiator termination is performed and the bus returns to idle.
  - A Master Abort is signaled on the internal bus. SERR# may be asserted on the PCI bus. Data is flushed from the IWQ.



### 3.2.1.3 Inbound Read Transaction

An inbound read transaction is initiated by a PCI initiator and is targeted at either 80333 local memory or a 80333 memory-mapped register space. The read transaction is propagated through the inbound transaction queue (ITQ) and read data is returned through the inbound read queue (IRQ).

When operating in the conventional PCI mode, all inbound read transactions are processed as delayed read transactions. When operating in the PCI-X mode, all inbound read transactions are processed as split transactions. The ATUs PCI interface claims the read transaction and forwards the read request through to the internal bus and returns the read data to the PCI bus. Data flow for an inbound read transaction on the PCI bus is summarized in the following statements:

- The ATU claims the PCI read transaction when the PCI address is within the inbound translation window defined by ATU Inbound Base Address Register (and Inbound Upper Base Address Register during DACs) and Inbound Limit Register.
- When operating in the conventional PCI mode, when the ITQ is currently holding transaction information from a previous delayed read, the current transaction information is compared to the previous transaction information (based on the setting of the DRC Alias bit in Section 3.10.43, "ATU Configuration Register ATUCR" on page 282). When there is a match and the data is in the IRQ, return the data to the master on the PCI bus. When there is a match and the data is not available, a Retry is signaled with no other action taken. When there is not a match and when the ITQ has less than eight entries, capture the transaction information, signal a Retry and initiate a delayed transaction. When there is not a match and when the ITQ is full, then signal a Retry with no other action taken.
  - When an address parity error is detected, the address parity response defined in Section 3.7 is used.
- When operating in the conventional PCI mode, once read data is driven onto the PCI bus from the IRQ, it continues until one of the following is true:
  - The initiator completes the PCI transaction. When there is data left unread in the IRQ, the data is flushed.
  - An internal bus Target Abort was detected. In this case, the QWORD associated with the Target Abort is never entered into the IRQ, and therefore is never returned.
  - Target Abort or a Disconnect with Data is returned in response to the Internal Bus Error.
  - The IRQ becomes empty. In this case, the PCI interface signals a Disconnect with data to the initiator on the last data word available.
- When operating in the PCI-X mode, when ITQ is not full, the PCI address, attribute and command are latched into the available ITQ and a Split Response Termination is signalled to the initiator.
- When operating in the PCI-X mode, when the transaction does not cross a 1024 byte aligned boundary, then the ATU waits until it receives the full byte count from the internal bus target before returning read data by generating the split completion transaction on the PCI-X bus. When the read requested crosses at least one 1024 byte boundary, then ATU completes the transfer by returning data in 1024 byte aligned chunks.

# intel

- When operating in the PCI-X mode, once a split completion transaction has started, it continues until one of the following is true:
  - The requester (now the target) generates a Retry Termination, or a Disconnection at Next ADB (when the requester is a bridge)
  - The byte count is satisfied.
  - An internal bus Target Abort was detected. The ATU generates a Split Completion Message (message class=2h - completer error, and message index=81h - target abort) to inform the requester about the abnormal condition. The ITQ for this transaction is flushed. Refer to Section 3.7.1.
  - An internal bus Master Abort was detected. The ATU generates a Split Completion Message (message class=2h - completer error, and message index=80h - Master abort) to inform the requester about the abnormal condition. The ITQ for this transaction is flushed. Refer to Section 3.7.1
- When operating in the conventional PCI mode, when the master inserts wait states on the PCI bus, the ATU PCI slave interface waits with no premature disconnects.
- When a data parity error occurs signified by **PERR**# asserted from the initiator, no action is taken by the target interface. Refer to Section 3.7.2.5.
- When operating in the conventional PCI mode, when the read on the internal bus is target-aborted, either a target-abort or a disconnect with data is signaled to the initiator. This is based on the ATU ECC Target Abort Enable bit (bit 0 of the ATUIMR for ATU). When set, a target abort is used, when clear, a disconnect is used.
- When operating in the PCI-X mode (with the exception of the MU queue ports at offsets 40h and 44h), when the transaction on the internal bus resulted in a target abort, the ATU generates a Split Completion Message (message class=2h completer error, and message index=81h internal bus target abort) to inform the requester about the abnormal condition. For the MU queue ports, the ATU returns either a target abort or a single data phase disconnect depending on the ATU ECC Target Abort Enable bit (bit 0 of the ATUIMR for ATU). The ITQ for this transaction is flushed. Refer to Section 3.7.1.
- When operating in the conventional PCI mode, when the transaction on the internal bus resulted in a master abort, the ATU returns a target abort to inform the requester about the abnormal condition. The ITQ for this transaction is flushed. Refer to Section 3.7.1
- When operating in the PCI-X mode, when the transaction on the internal bus resulted in a master abort, the ATU generates a Split Completion Message (message class=2h completer error, and message index=80h internal bus master abort) to inform the requester about the abnormal condition. The ITQ for this transaction is flushed. Refer to Section 3.7.1.
- When operating in the PCI-X mode, when the Split Completion transaction completes with either Master-Abort or Target-Abort, the requester is indicating a failure condition that prevents it from accepting the completion it requested. In this case, since the Split Request addresses a location that has no read side effects, the completer must discard the Split Completion and take no further action.



The data flow for an inbound read transaction on the internal bus is summarized in the following statements:

- The ATU internal bus master interface requests the internal bus when a PCI address appears in an ITQ and transaction ordering has been satisfied. When operating in the PCI-X mode the ATU does not use the information provided by the Relax Ordering Attribute bit. That is, ATU always uses conventional PCI ordering rules.
- Once the internal bus is granted, the internal bus master interface drives the translated address onto the bus and wait for **IB\_DEVSEL#**. When a Retry is signaled, the request is repeated. When a master abort occurs, the transaction is considered complete and a target abort is loaded into the associated IRQ for return to the PCI initiator (transaction is flushed once the PCI master has been delivered the target abort).
- Once the translated address is on the bus and the transaction has been accepted, the internal bus target starts returning data with the assertion of **IB\_TRDY#**. Read data is continuously received by the IRQ until one of the following is true:
  - The full byte count requested by the ATU read request is received. The ATU internal bus initiator interface performs a initiator completion in this case.
  - When operating in the conventional PCI mode, a Target Abort is received on the internal bus from the internal bus target. In this case, the transaction is aborted and the PCI side is informed.
  - When operating in the PCI-X mode, a Target Abort is received on the internal bus from the internal bus target. In this case, the transaction is aborted. The ATU generates a Split Completion Message (message class=2h - completer error, and message index=81h target abort) on the PCI bus to inform the requester about the abnormal condition. The ITQ for this transaction is flushed.
  - When operating in the conventional PCI mode, a single data phase disconnection is received from the internal bus target. When the data has not been received up to the next QWORD boundary, the ATU internal bus master interface attempts to reacquire the bus. When not, the bus returns to idle.
  - When operating in the PCI-X mode, a single data phase disconnection is received from the internal bus target. The ATU IB initiator interface attempts to reacquire the bus to obtain remaining data.
  - When operating in the conventional PCI mode, a disconnection at Next ADB is received from the internal bus target. The bus returns to idle.
  - When operating in the PCI-X mode, a disconnection at Next ADB is received from the internal bus target. The ATU IB initiator interface attempts to reacquire the bus to obtain remaining data.

To support *PCI Local Bus Specification*, Revision 2.0 devices, the ATU can be programmed to ignore the memory read command (Memory Read, Memory Read Line, and Memory Read Multiple) when trying to match the current inbound read transaction with data in a DRC queue which was read previously (DRC on target bus). When the Read Command Alias Bit in the ATUCR register is set, the ATU does not distinguish the read commands on transactions. For example, the ATU enqueues a DRR with a Memory Read Multiple command and performs the read on the internal bus. Some time later, a PCI master attempts a Memory Read with the same address as the previous Memory Read Multiple. When the Read Command Bit is set, the ATU would return the read data from the DRC queue and consider the Delayed Read transaction complete. When the Read Command bit in the ATUCR was clear, the ATU would not return data since the PCI read commands did not match, only the address.

## 3.2.1.4 Inbound Configuration Cycle Translation

The ATU only accepts Type 0 configuration cycles with a function number of zero.

The ATU is configured through the PCI bus. When operating in the conventional PCI mode, all inbound configuration cycles are processed as delayed transactions. When operating in the PCI-X mode, all inbound configuration cycles are processed as split transactions. The translation mechanism for inbound configuration cycles is defined by the *PCI Local Bus Specification*, Revision 2.3.

The ATU configuration space is selected by a PCI configuration command and claims the access (by asserting DEVSEL#) when the IDSEL pin is asserted, the PCI command indicates a configuration read or write, and address bits AD[1:0] are  $00_2$  all during the address phase. The ATU interface ignores any configuration command (IDSEL active) where AD[1:0] are  $not 00_2$  (e.g. Type 1 commands). During the configuration access address phase, the PCI address is divided into a number of fields to determine the actual configuration register access. These fields, in combination with the byte enables during the data phase create the unique encoding necessary to access the individual registers of the configuration address space:

- AD[7:2] Register Number. Selects one of 64 DWORD registers in the ATU PCI configuration address space.
- **C/BE[3:0]#** Used during the data phase. Selects which actual configuration register is used within the DWORD address. Creates byte addressability of the register space.
- **AD**[10:8] Function Number. Used to select which function of a multi-function device is being accessed. The ATU is function 0 and therefore it only responds to 000<sub>2</sub> in this bit field and ignore all other bit combinations.

The ATU configuration address space starts at internal address FFFF.E100H. Therefore **AD**[7:2] equal to  $000000_2$  equates to address FFFF.E100H and **AD**[7:2] equal to  $000001_2$  results in address FFFF.E104H and so on.

For inbound configuration reads, the IRQ and ITQ are used in the same manner as inbound memory read operations. The internal bus cycle that results is a 32-bit transaction where **I\_REQ64#** is not asserted.

For inbound configuration writes, the ATU adds a delayed write data queue, IDWQ, which holds data in the same manner as the IWQ. The transaction information from the configuration write operation on the PCI interface is captured into the IDWQ (when full, a Retry is signaled). The data from the delayed write (split write when operating in the PCI-X mode) request cycle is latched into the IDWQ and forwarded to the internal bus interface. Once transaction ordering and priority have been satisfied, the internal bus master interface requests the internal bus and deliver the write data to the target as defined in Section 3.2.1.2.

The status of the transaction on the internal bus is returned to the PCI initiator on the PCI bus. When operating in the conventional PCI mode, the retry cycle from the initiator is accepted once the write has been completed on the internal bus and the status has been captured for return to the PCI master. When operating in the PCI-X mode, a Split Completion Message (message class=0h and message index= 00h) is generated on the PCI bus, once the write has been completed on the internal bus and the status has been latched for return to the PCI master. Since Master Aborts and Target Aborts cannot occur during configuration cycles on the internal bus, normal completion status is returned. The data from PCI completion transaction is discarded.



### 3.2.1.5 Discard Timers

The ATU implements discard timers for inbound delayed transactions. These timers prevent deadlocks when the initiator of a retried delayed transaction fails to complete the transaction within  $2^{10}$  or  $2^{15}$  PCI clock cycles on the initiating bus when operating in the conventional PCI mode. The timer starts counting when the delayed request becomes a delayed completion by completing on the internal bus and all passing rules are satisfied. When the originating master on the PCI bus has not retried the transaction before the timer expires, the completion transaction is discarded.

Discard timer values are controlled by the ATU Configuration Register's Discard Timer Value bit. The ATU queues covered by discard timers are the IRQ and the IDWQ. After discarding a transaction, the ATU must set the Discard Timer Status bit in the ATU Configuration Register. The ATU does *not* assert the **SERR#** signal after discarding a transaction.

# 3.2.2 Outbound Transactions- Single Address Cycle (SAC) Internal Bus Transactions

Outbound transactions initiated by the 80333 core processor are directed to the PCI interface through the ATU. The core processor always generates Single Address Cycles on the internal bus. As a PCI master, the ATU is capable of PCI I/O transactions, PCI memory reads (in case of conventional PCI), memory read DWORD (in case of PCI-X), PCI memory writes, configuration reads and writes, and DAC cycles. Outbound memory transactions are always attempted as 64-bit PCI transactions. Outbound memory write operations are performed as posted operations and outbound memory read operations are all performed as split read operations.

Outbound transactions use a separate set of queues from inbound transactions. Outbound write operations have their address entered into the outbound write address queue (OWADQ) and their data into the outbound write queue (OWQ). Outbound read transactions, performed as split transactions, use the Outbound Transaction Queue (OTQ) to store address, and get data returned into the outbound read queue (ORQ). Refer to Section 3.5.2 for details of outbound queue architecture. Outbound configuration transactions use a special outbound port structure. Refer to Section 3.2.3 for details.

For outbound write transactions, the ATU is a target on the internal bus and initiator on the PCI bus. For outbound read transactions, the ATU is a completer on the internal bus (initially accepts the split read as a target and then provides read data by initiating a split completion).

### 3.2.2.1 Outbound Address Translation - Single Address Cycle (SAC) Internal Bus Transactions

In addition to providing the mechanism for inbound translation, the ATU translates Intel XScale<sup>®</sup> core-initiated cycles to the PCI bus. This is known as *outbound address translation*. Outbound transactions are processor or DMA transactions targeted at the PCI bus. The ATU internal bus target interface claims internal bus cycles and completes the cycle on the PCI bus on behalf of the Intel XScale<sup>®</sup> core or DMAs. The ATU supports two different outbound windowing modes:

- Address Translation Windowing
- Direct Address Windowing (No translation)

Figure 14 shows a 80333 memory map with all reserved address locations highlighted. The two 64 MByte outbound translation memory windows exist from 8000.0000H to 83FF.FFFFH and 8400.0000H to 87FF.FFFFH. The direct addressing window is from 0000.0040H to 7FFF.FFFFH. The 64KByte outbound I/O window is from 9000.0000H to 90000.FFFFH.

The response of the ATU to Outbound Transactions is controlled by bits in the ATU Command Register and the ATU Configuration Register. The Outbound ATUs behavior for the different combinations of these control bits is described in Table 109.

### Table 109. Outbound Address Translation Control

Outbound Response	Bus Master Enable (ATUCMD - bit2)	Outbound Enable (ATUCR - bit1)
Master-Abort	0	0
Retry	0	1
Master-Abort	1	0
Claim <sup>a</sup>	1	1

a. The ATU may respond with a Retry in this case when the Outbound Transaction Queues are full.



### 3.2.2.2 Outbound Address Translation Windows- Single Address Cycle (SAC) Internal Bus Transactions

Inbound translation involves a programmable inbound translation window consisting of a base and limit register and a translate value register for PCI to internal bus translation. The outbound address translation windows use a similar methodology except that the outbound translation window base addresses and limit sizes are fixed in 80333 internal bus local address space; this removes the need for separate base and limit registers.

Figure 15 on page 180 illustrates the outbound address translation windows. The ATU has three windows: two are 64 Mbyte and one is 64 Kbyte. The two outbound memory translation windows range from 8000.0000H to 87FF.FFFFH. After these two windows, the outbound I/O window range from 9000.0000H to 9000.FFFFH.

The two Memory windows are 64 Mbytes and I/O window is 64 Kbytes. An internal bus cycle with an address within one of the outbound windows initiates a read or write cycle on the PCI bus. The PCI cycle type depends on which translation window the local bus cycle "hits". The read or write decision is based on the internal bus cycle type.

ATU has a window dedicated to the following outbound PCI/PCI-X transaction types:

- Memory reads and Memory writes Memory Window
- I/O reads and writes I/O Window

### Table 110. Internal Bus-to-PCI Command Translation for Two Memory Windows/DMA Channels

Internal Bus Command	Conventional PCI Command	PCI-X Command
Memory Write	Memory Write	Memory Write
Memory Write Block	Memory Write and Invalidate <sup>a</sup> or Memory Write	Memory Write Block
Memory Read Block	Memory Read Multiple	Memory Read Block
Memory Read DWORD	Memory Read	Memory Read DWORD
Alias to Memory Write Block	Memory Write and Invalidate or Memory Write	Memory Write Block
Alias to Memory Read Block	Memory Read Multiple	Memory Read Block

a. The ATU converts (Alias to) Memory Write Block to Memory Write and Invalidate when the following four conditions are met, otherwise the Memory Write Block is converted to a Memory Write:

1.) Memory Write and Invalidate transactions are enabled in the "ATU Command Register - ATUCMD".

2.) The Cache Line size is set to 8 or 16 DWORDS in the "ATU Cacheline Size Register - ATUCLSR".

3.) Starting address of the Outbound PCI bus request is aligned to the cache line size.

4.) Byte count intended for the Outbound PCI bus request is a multiple of the cache line size.



### Table 111. Internal Bus-to-PCI Command Translation for I/O Window

Internal Bus Command <sup>a</sup>	Conventional PCI Command	PCI-X Command
Memory Write	I/O Write	I/O Write
Memory Write Block	I/O Write	I/O Write
Memory Read Block	I/O Read	I/O Read
Memory Read DWORD	I/O Read	I/O Read
Alias to Memory Write Block	I/O Write	I/O Write
Alias to Memory Read Block	I/O Read	I/O Read

a. User should designate memory region containing I/O Window (9000.000H to 90000.FFFFH) as non-cacheable and non-bufferable from Intel XScale<sup>®</sup> core. This guarantees all load/stores to I/O Window are of DWORD quantities. In event the user inadvertently issues a read to the I/O Window that crosses a DWORD address boundary, the ATU target aborts transaction. All writes are terminated with a Single-Phase-Disconnect and only bytes 3:0 is relevant dependent on the Byte Enables.

The windowing scheme refers to:

- An Internal Bus read cycle that addresses a Memory Window is translated to a Memory Read on the PCI bus
- An Internal Bus write cycle that addresses a Memory Window is translated to a Memory Write on the PCI bus
- An Internal Bus read cycle that addresses the I/O Window is an I/O Read on the PCI bus
- An Internal Bus write cycle that addresses the I/O Window is an I/O Write on the PCI bus



### Figure 14. Internal Bus Memory Map





The translation portion of outbound ATU transactions is accomplished with a translate value register in the same manner as inbound translations. Each outbound memory window is associated with two translation registers which provide lower and upper translation addresses (OMWTVR0-1, OUMWTVR-1). When the corresponding OUMWTVRx register is all-zero a SAC transaction is generated on the PCI bus. Otherwise, a DAC is generated on the PCI bus. ATU uses the following registers during outbound address translation:

- Outbound Memory Window Translate Value Register 0 (OMWTVR0)
- Outbound Upper 32-bit Memory Window Translate Value Register 0 (OUMWTVR0)
- Outbound Memory Window Translate Value Register 1 (OMWTVR1)
- Outbound Upper 32-bit Memory Window Translate Value Register 1 (OUMWTVR1)
- Outbound I/O Window Translate Value Register (OIOWTVR)
- Outbound Configuration Cycle Address Register (OCCA)

See Section 3.8 for details on outbound translation register definition and programming constraints.

The translation algorithm used, as stated, is very similar to inbound translation. For memory transactions, the algorithm is:

### Equation 3. Outbound Address Translation

PCI Address = (Internal\_Bus\_Address and 03FF.FFFFH) | Window\_Translate\_Value\_Register

For memory transactions, the internal bus address is bitwise ANDed with the inverse of 64 Mbytes which clears the upper 6 bits of address. The result is bitwise ORed with the outbound window translate value register to create the lower 32-bits of the PCI address.

For I/O transactions, the algorithm is:

#### Equation 4. I/O Transactions

PCI Address = (Internal\_Bus\_Address and 0000.FFFFH) | Window\_Translate\_Value\_Register

For I/O transactions, the internal bus address is bitwise ANDed with the inverse of 64 Kbytes which clears the upper 16 bits of address. Address aliasing is prevented by the outbound window translate value registers which only allow values on boundaries equivalent to the window length. PCI I/O addresses are byte addresses and not word addresses. The PCI I/O address two least significant bits are determined by byte enables that the processor issues. For example, when the Intel XScale<sup>®</sup> core performs a 2-byte write and generates byte enables of  $0011_2$ , the ATU sets the two least significant bits of PCI I/O address to  $10_2$ .





### Figure 15. Outbound Address Translation Windows
## 3.2.2.3 Direct Addressing Window - Single Address Cycle (SAC) Transactions

The second method used by outbound cycles from the Intel XScale<sup>®</sup> core to the PCI bus is the direct addressing window. This is a window of addresses in Intel XScale<sup>®</sup> core address space that act in the same manner as the outbound translation windows either without any translation or with the translation of address bit 31 only. This allows the Direct Addressing window to translate to different address ranges on the PCI bus (0000.0040H to 7FFF.FFFFH or 8000.0000H to FFFF.FFFFH). A Intel XScale<sup>®</sup> core read or write to an internal bus address within the direct addressing window initiates a read or write on the PCI bus with the same address (with the possible exception of address bit 31) as used on the internal bus. When the Outbound Upper 32-bit Direct Window Translate Value Register (OUDWTVR) is not all-zero, then a DAC transaction is generated on the PCI bus using the upper-32 bit address provided in OUDWTVR. Otherwise a SAC transaction is generated on the PCI bus. Figure 16 shows two examples of outbound writes that are through the direct addressing window.

Direct Addressing is limited to PCI memory read and write commands only. I/O cycles are not supported with direct addressing.

#### Figure 16. Direct Addressing Window



The internal bus side of the direct addressing window address range is fixed in the lower 2 Gbytes of the 80333 local address space (except for the first 32 Bytes which are reserved for the Reset and Exception vectors). Internal bus cycles with an address from 0000.0040H to 7FFF.FFFFH are forwarded to the PCI bus, when enabled. The following bits within the ATUCR affect direct addressing operation:

- ATUCR Direct Addressing Enable bit when set, enables the direct addressing window. When clear, addresses within the direct addressing window are not forwarded to the PCI bus.
- ATUCR Direct Addressing Upper 2G Translation Enable when set, the ATU forwards internal bus cycles with an address between 0000.0040H and 7FFF.FFFFH to the PCI bus with bit 31 of the address set (8000.0000H FFFF.FFFFH). When clear, no translation occurs.



### 3.2.2.4 Outbound DMA Transactions

The ATU provides each DMA channel with its own dedicated decode window that claims all transactions generated by that channel that are destined for the PCI bus.

Depending on the contents of the PCI Upper Address field of a particular DMA descriptor, the ATU may initiate a SAC or a DAC transaction on the PCI bus in response to the DMA channel's request. As the DMA channel transfers data to the ATU, the lower 32-bits of the Internal Bus address are passed to the PCI bus unaltered (see Section 6.3.1, "Chain Descriptors" on page 380 for more details).

# intel

## 3.2.3 Outbound Write Transaction

An outbound write transaction is initiated by the Intel XScale<sup>®</sup> core<sup>1</sup> or by one of the DMAs and is targeted at a PCI target on the PCI bus. The outbound write address and write data are propagated from the 80333 internal bus to a PCI bus through the OWADQ and OWQ, respectively.

The ATUs internal bus target interface claims the write transaction and forwards the write data through to the targeted PCI bus. The data flow for an outbound write transaction on the internal bus is summarized in the following statements:

- For Single Address Cycles (SACs), ATU internal bus target interface latches the address from the internal bus into the OWADQ when that address is inside one of the outbound translate windows (see Section 3.5) and the OWQ is not full.
- For Dual Address Cycles (DACs), ATU internal bus target interface latches address from the internal bus into the OWADQ when the OWQ is not full and OWADQ is not full.
- Once outbound address is latched, internal bus target interface stores write data into the OWQ until the internal bus transaction completes or the reaches a buffer boundary. The initiator of the transaction is disconnected at an ADB when the transaction reaches a buffer boundary.
- When the OWADQ is full, the target interface signals a Retry on the internal bus to the outbound cycle initiator.
- When OWADQ latches the address and corresponding data is latched in a buffer in OWQ, the outbound cycle is enabled for transmission on the PCI Bus and PCI interface requests PCI bus.

The PCI interface is responsible for completing the outbound write transaction with the PCI address translated from the OWADQ and the data in the OWQ. The data flow for an outbound write transaction on the PCI bus is summarized in the following statements:

- ATU PCI interface requests PCI bus when completed internal bus transaction is in OWADQ and data associated with transfer in OWQ. Once bus is granted, PCI master interface writes PCI translated address from OWADQ to the PCI bus and wait for the transaction to be claimed.
- When Master Abort seen during address phase, transaction flushed and OWADQ/OWQ are cleared. Refer Section 3.7.3 for full details on PCI master abort conditions during outbound transactions.
- In the conventional PCI mode, once the PCI write transaction is claimed, the PCI interface transfers data from the OWQ to the PCI bus until one of the following is true:
  - The PCI target signals a Retry or Disconnect. The ATU PCI master attempts to reacquire the PCI bus to complete the write transaction.
  - The **GNT#** signal is deasserted and the master latency timer has expired. In this case, the master interface attempts to reacquire the PCI bus and complete the write transaction.
  - The PCI target signals a Target-Abort. In this case, the OWQ and OTQ are cleared and the transaction is aborted. The appropriate error bits are set as defined in Section 3.7.4.
  - The transaction terminates normally by transferring all data (full byte count) associated with it. The write address is removed from the OWADQ and the interface returns to idle.

<sup>1.</sup> For best performance, the user should designate the two Outbound Memory Windows (8000.0000H to 83FF.FFFFH, 8400.0000H to 87FF.FFFFH) as non-cacheable and bufferable from the Intel XScale<sup>®</sup> core. This assignment enables the Intel XScale<sup>®</sup> core to issue multiple outstanding transactions to the Outbound Memory Windows, thereby, taking full advantage of the ATU outbound queue architecture. However, the user needs to be aware that the Outbound ATU queue architecture does not maintain strict ordering between read and write requests as described in Table 117, "ATU Outbound Data Flow Ordering Rules" on page 194. In the event that the user requires strict ordering to be maintained, the user can either change the designation of this region of memory to be non-cacheable and non-bufferable or enforce the requirement in software.



- In the PCI-X mode, once the PCI memory write transaction is claimed, the PCI interface transfers data from the OWQ to the PCI bus until one of the following is true:
  - The PCI target signals a Retry or Single Data Phase Disconnect. The ATU PCI initiator attempts to reacquire the PCI bus to complete the write transaction.
  - reacquire the PCI bus to complete the write transaction.
  - The **GNT#** signal is deasserted and the master latency timer has expired. In this case, the master interface attempts to reacquire the PCI bus and complete the write transaction.
  - The PCI target signals a Target-Abort. In this case, the OWQ and OWADQ are cleared and the transaction is aborted. The appropriate error bits are set as defined in Section 3.7.4.
  - The transaction terminates normally with Satisfaction of Byte Count. The write address is removed from the OWADQ and the interface returns to idle.
- In the PCI-X mode, once the PCI I/O write transaction is claimed, the PCI interface transfers data from the OWQ to the PCI bus until one of the following is true:
  - The PCI target signals a Retry. The ATU PCI initiator attempts to reacquire the PCI bus to complete the write transaction.
  - The transaction terminates normally with Satisfaction of Byte Count or with Single Data Phase Disconnect. The write address is removed from the OWADQ and the interface returns to idle.
  - The PCI target signals a Target-Abort. In this case, the OWQ and OWADQ are cleared and the transaction is aborted. The appropriate error bits are set as defined in Section 3.7.4.
  - The transaction terminates with Split Response Termination. The write address is removed from the OWADQ and the interface returns to idle only when it receives the corresponding Split Completion Message.

When a data parity error is encountered (**PERR**# detected), the master interface continues writing data to clear the queue.

In the conventional PCI mode when the PCI target deasserts TRDY#, no action is taken by the ATU master other than inserting wait states.

# intel

## 3.2.4 Outbound Read Transaction

An outbound read transaction is initiated by the Intel XScale<sup>®</sup> core<sup>1</sup> or one of the DMAs and is targeted at a PCI slave on the PCI bus. The read transaction is propagated through the outbound transaction queue (OTQ) and read data is returned through the outbound read queue (ORQ).

The ATUs internal bus target interface claims the Memory Read Dword and Memory Read Block and Alias to Memory Read Block transaction and forwards the read request through to the PCI bus and returns the read data to the internal bus.

Data flow for an outbound read transaction on the internal bus is summarized as follows:

- For Single Address Cycles (SACs), the ATU internal bus interface latches the internal bus address when the address is inside an outbound address translation window (or direct addressing window, when enabled) and the OTQ is not full. For Dual Address Cycles (DACs), ATU internal bus target interface latches the address from the internal bus into the OTQ irrespective of the address. All read transactions are handled as split transactions. When the OTQ is full (previous outbound transactions in progress), the internal bus interface signals a Retry to the transaction initiator.
- When during the completion cycle on the PCI interface, a master abort is encountered, a flag is set and the ATU notifies the internal bus requester about the aborted transaction by generating a Split Completion Error Message (with message class=1h PCI bus error and message index=00h master abort). The OTQ is cleared of the transaction.
- When during the completion cycle on the PCI interface, a target abort is encountered, a flag is set and the ATU notifies the internal bus requester about the aborted transaction by generating a Split Completion Error Message (with message class=1h PCI bus error and message index=01h target abort). The OTQ is cleared of the transaction.
- Once the transaction completes on the PCI bus, the ATU generates a split completion transaction to return data to the internal bus requester.
- When operating in the PCI-X mode, ATU may receive a split completion error message when attempting to read data on the PCI bus. In this case, ATU notifies the internal bus requester about the error by forwarding the Split Completion Error Message from the PCI bus back to the Internal Requester. The OTQ is cleared of the transaction.

Data flow for an outbound read transaction on the PCI bus is summarized as follows:

- The ATU PCI interface requests the PCI bus when the head of the OTQ has at least one entry and the ordering rules are satisfied. Once the bus is granted, the PCI interface transfers the PCI translated address from the OTQ to the PCI bus and wait for the transaction to be claimed.
- When no **DEVSEL**# is asserted, a master abort is signaled. This is passed through to the internal bus target interface.
- When a target abort is signaled from the PCI target, the target abort is returned to the internal bus and the PCI interface returns to idle.
- When operating in the PCI-X mode the read transaction may terminate as a split response termination. Then the ATU receives data during the corresponding split completion transaction. When an error occurs, the ATU may receive a split completion error message.

<sup>1.</sup> For best performance, the user should designate the two Outbound Memory Windows (8000.0000H to 83FF.FFFFH, 8400.0000H to 87FF.FFFFH) as non-cacheable and bufferable from the Intel XScale<sup>®</sup> core. This assignment enables the Intel XScale<sup>®</sup> core to issue multiple outstanding transactions to the Outbound Memory Windows, thereby, taking full advantage of the ATU outbound queue architecture. However, the user needs to be aware that the Outbound ATU queue architecture does not maintain strict ordering between read and write requests as described in Table 117, "ATU Outbound Data Flow Ordering Rules" on page 194. In the event that the user requires strict ordering to be maintained, the user can either change the designation of this region of memory to be non-cacheable and non-bufferable or enforce the requirement in software.



## 3.2.5 Outbound Configuration Cycle Translation

The outbound ATU provides a port programming model for outbound configuration cycles.

Performing an outbound configuration cycle to the PCI bus involves up to two internal bus cycles:

- 1. Writing Outbound Configuration Cycle Address Register (OCCAR) with PCI address used during configuration cycle. See the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a for information regarding configuration address cycle formats. This IB bus cycle enables the transaction.
- 2. Writing or reading the Outbound Configuration Cycle Data Register (OCCDR). A read causes a configuration cycle read to the PCI bus with the address in the outbound configuration cycle address register. Note that the Internal Bus read is executed as a split transaction. Similarly, a write initiates a configuration cycle write to PCI with the write data from the second processor cycle. Configuration cycles are non-burst and restricted to a single DWORD cycle.
- *Note:* Bits 15:11 of the configuration cycle address for Type 0 configuration cycles are defined differently for Conventional versus PCI-X modes. When 80333 software programs the OCCAR to initiate a Type 0 configuration cycle, the OCCAR should always be loaded based on the PCI-X definition for the Type 0 configuration cycle address. When operating in Conventional mode, the 80333 clears bits 15:11 of the OCCAR prior to initiating an outbound Type 0 configuration cycle.
- *Note:* During the attribute phase of a Type 0 configuration transaction, the Secondary Bus Number field (bits 7:0) is set equal to the Requester Bus Number (bits 15:8 of the "PCI-X Status Register PX\_SR" on page 310).

Master aborts during outbound configuration reads result in ATU generating a Split Completion Error Message (class=1h - bridge error and index=00h - master abort on PCI) on internal bus.

Target aborts during outbound configuration reads result in ATU generating a Split Completion Error Message (class=1h - bridge error and index=01h - target abort on PCI) on the internal bus.

Parity errors during outbound configuration reads result in ATU generating a Split Completion Error Message (class=2h - completer error and index=82h - parity error on PCI) on the internal bus.

Parity errors detected by target of an outbound configuration write may result in the ATU receiving either of the two Split Completion Write Data Parity Error Messages (with message class=2h -completer error and message index=01h - split write data parity error or with message class=1h - bridge error and message index=02h - write data parity error) on the PCI bus. When Parity Checking is enabled, the ATU sets error bits in the ATUSR and the PCIXSR. The Intel XScale<sup>®</sup> core is interrupted when the Split Completion Error and/or Master Data Parity interrupt(s) are unmasked.

When the Configuration Cycle Data Register is written, the data is latched and forwarded to the PCI bus with the internal target issuing a single data phase disconnect with data for 32-bits only. This cycle does not receive an **ACK64**# from the ATU and therefore is defined as 32-bit only.

Note that while the programming model uses the register interface for outbound configuration cycles, from a hardware standpoint, the address is entered into the OTQ (reads) or OWADQ (writes), configuration write data goes through the OWQ and configuration read data is returned in the ORQ.

*Note:* Outbound configuration cycle data registers are not physical registers. They are a 80333 memory mapped addresses used to initiate a transaction with the address in the associated address register.

The user should designate the memory region containing the OCCDR as non-cacheable and non-bufferable from the Intel XScale<sup>®</sup> core. This guarantees that all load/stores to the OCCDR is only of DWORD quantities. In event the user inadvertently issues a read to the OCCDR that crosses a DWORD address boundary, the ATU target aborts the transaction. All writes is terminated with a Single-Phase-Disconnect and only bytes 3:0 is relevant.

## 3.3 Messaging Unit

intel

The Messaging Unit (MU) transfers data between the PCI system and the 80333 and notifies the respective system when new data arrives. The MU is described in Chapter 4, "Messaging Unit".

The PCI window for messaging transactions is always the *first* 4 Kbytes of the inbound translation window defined by the Inbound ATU Base Address Register 0 (IABAR0) and the Inbound ATU Limit Register 0 (IALR0).

All of the Messaging Unit errors are reported in the same manner as ATU errors. Error conditions and status can be found in the ATUSR and the ATUISR, see Section 3.7, "ATU Error Conditions" on page 200.



## 3.4 Expansion ROM Translation Unit

The inbound ATU supports one address range (defined by a base/limit register pair) used for the Expansion ROM. Refer to the *PCI Local Bus Specification*, Revision 2.3 for details on Expansion ROM format and usage.

During a powerup sequence, initialization code from Expansion ROM is executed once by the host processor to initialize the associated device. The code can be discarded once executed. Expansion ROM registers are described in Section , and Section 3.10.29.

The inbound ATU supports an inbound Expansion ROM window which works like the inbound translation window. A read from the expansion ROM windows is forwarded to the internal bus. The address translation algorithm is the same as the inbound translation; see Section 3.2.1.1, "Inbound Address Translation" on page 165. As a PCI target, the Expansion ROM interface behaves as a standard ATU interface and is capable of returning a 64-bit access by the assertion of **ACK64#** in response to a 64-bit request.

The Expansion ROM unit uses the ATU inbound transaction queue and the inbound read data queue.

When operating in the conventional PCI mode, the address of the inbound delayed read cycle is entered into the ITQ queue and the delayed read completion data is returned in the IRQ. That is, inbound reads to the Expansion ROM window are handled as delayed transactions on the PCI bus.

When operating in the PCI-X mode, the address of the inbound read cycle is entered into the ITQ queue and the read completion data is returned in the IRQ. That is, inbound reads to the Expansion ROM window are handled as split transactions on the PCI bus. The internal bus initiator interface fills the IRQ read queue with the full byte count before generating the split completion transaction on the PCI bus. That is, the ATU generates a Memory Read Block transaction on the internal bus with byte count set to the byte count specified in the PCI read. The PBI (Flash Interface) terminates the read as a split read request and return data in

- either one or more 1024 byte split completion transactions when byte count is greater than or equal to 1024 bytes or
- one split completion with the full byte count when byte count is less than 1024 bytes.

Expansion ROM writes are not supported and result in a Target Abort.



## 3.5 ATU Queue Architecture

ATU operation and performance depends on the queueing mechanism implemented between the internal bus interface and PCI bus interface. As indicated in Figure 11, the ATU queue architecture consists of separate inbound and outbound queues. The function of each queue is described in the following sections.

## 3.5.1 Inbound Queues

The inbound data queues of the ATU support transactions initiated on a PCI bus and targeted at either 80333 local memory or a 80333 memory mapped register. Table 112 details the name and sizes of the ATU inbound data queues.

#### Table 112.Inbound Queues

Queue Mnemonic	Queue Name	Queue Size (Bytes)
IWQ	Inbound Write Data Queue	4 KBytes (4*1KB)
IWADQ	Inbound Write Address Queue	4 Transaction Addresses
IRQ	Inbound Read Data Queue	4 KBytes (4*1KB)
IDWQ Inbound Delayed Write address/data Queue		1 Transaction
ITQ	Inbound Transaction Queue	8 Addresses/Commands

#### 3.5.1.1 Inbound Write Queue Structure

The ATU Inbound Write Queues consist of the inbound write data queue and the inbound write address queue. The inbound write data queue holds the data for memory write transactions moving from a PCI Bus to the internal bus and the address queue holds the corresponding address of the transactions in the data queues. The inbound write queue, IWQ, has a queue depth of 4 KBytes and moves write transactions from the PCI bus to the internal bus. The corresponding address queue, IWADQ, is capable of holding 4 address entries. The queue pair is capable of holding up to 4 memory write (or MWI when operating in the conventional PCI mode) transactions.

The following rules apply to the PCI bus interface and govern the acceptance of data into IWQ and address into the tail of the IWADQ:

- A memory write operation claimed by the target PCI interface on the PCI bus is accepted into the address and data queues when the IWADQ has at least one address entry available.
- When operating in the conventional PCI mode, when the IWQ reaches a full state while filling, a disconnect with data is signaled to the master of the transaction.
- When operating in the PCI-X mode, when the IWQ reaches a full state while filling, a disconnect at next ADB is signaled to the master of that transaction.

Memory write transactions are drained from the head of the queue when the master interface has acquired bus ownership and transaction ordering and priority have been satisfied (see Section 3.5.3). A memory write transaction is considered drained from the queue when the entire amount of data entered on the PCI bus has been accepted by the internal bus target. Error conditions resulting in the cancellation of a write transaction only flush the transaction at the head of the data and address queue. All other transactions within the queues are considered still valid.



#### 3.5.1.2 Inbound Read Queue Structure

The inbound read queues are responsible for retrieving data from local memory and returning it to the PCI bus in response to a read transaction initiated from a PCI master. When operating in the conventional PCI mode, reads are handled as delayed transactions. When operating in the PCI-X mode reads are handled as split transactions. The address of the transactions are held in the ITQ. Up to 8 read requests can be stored in the ITQ. The read data is returned through IRQ.

When operating in the conventional PCI mode, the IRQ holds data from four PCI bus read transactions. The read request cycle on PCI latches the read command and the address into the ITQ when the cycle is first initiated by the PCI master. The ATU internal bus initiator interface takes the translated address and the command and performs a read on the internal bus. Reads can be any of the PCI memory read command types using the ATU inbound translation or an inbound configuration read using the specific configuration cycle translation. The data from the read on the internal bus is stored in the IRQ until the PCI master initiates a read cycle that matches the initial request cycle in both command and address. Any data left in an IRQ after the delivery of a completion cycle on PCI is flushed. This is possible since all internal bus memory is considered prefetchable with no read side effects.

When operating in the PCI-X mode, the IRQ may hold data from up to four PCI bus read transactions. The read request cycle on PCI latches the read command and the address into the ITQ when the cycle is first initiated by the PCI master. The ATU internal bus initiator interface takes the translated address and the command and performs a read on the internal bus. Reads can be any of the PCI memory read command types using the ATU inbound translation or an inbound configuration read using the specific configuration cycle translation. Once read data is available in the IRQ, the ATU generates one or more split completions to return read data to the PCI requester.

When operating in the conventional PCI mode, the exact amount of data (byte count) read by the master state machine on the internal bus interface depends upon the read command used and how much data the Internal Bus target device delivers. Table 113 shows the amount of data attempted to be read for the different memory read commands for the ATU, when operating in the conventional PCI mode.

Peer Memory Block Read transactions will be issued as Memory Read transactions on the peer segment when the peer segment is operating in conventional PCI mode. Peer Memory Block Read transactions will be issued as 128Byte (max) Memory Block Read transactions on the peer segment when the peer segment is operating in PCI-X mode.

Internal bus error conditions override all prefetch amounts. i.e. a master-abort and target-abort conditions.

#### Table 113. Inbound Read Prefetch Data Sizes

PCI Read Command	Prefetch Size (Bytes)
Memory Read	4 to 32
Memory Read Line	4 to 128
Memory Read Multiple	4 to 1024

## 3.5.1.3 Inbound Delayed Write Queue

The IDWQ is used specifically for inbound configuration write cycles to the ATU. I/O Write transactions are not accepted by the ATU and result in a Master Abort.

The IDWQ contains both the address and data of a configuration write cycle. When operating in the conventional PCI mode, the configuration writes are handled as delayed writes. When operating in the PCI-X mode, the configuration writes are handled as split writes. When the write cycle is initiated on the PCI bus, the address and data are entered into the 8 byte queue, and forwarded to the internal bus. The transaction is forwarded to the internal bus once transaction ordering has been satisfied. The status of the transaction (normal completion) is maintained in the IDWQ for return to the PCI master on the initiating bus. When operating in the PCI-X mode, a write completion message is generated by the ATU to indicate the successful execution of the configuration write transaction.

The IDWQ can only hold 32-bits of data and should never be accessed with **REQ64#** asserted per the *PCI Local Bus Specification*, Revision 2.3 which states that "only memory transactions support 64-bit data transfers". In addition, the cycle should always return only 32-bits of data on the internal bus.

### 3.5.1.4 Inbound Transaction Queues Command Translation Summary

PCI Command	Conventional Mode	PCI-X Mode	Internal Bus Command
Memory Write	<b>~</b>	~	Memory Write
Memory Write and Invalidate	<b>~</b>		Memory Write Block
Memory Write Block		~	Memory Write Block
Alias to Memory Write Block		~	Memory Write Block
Memory Read	~		Memory Read Block
Memory Read Line	~		Memory Read Block
Memory Read Multiple	~		Memory Read Block
Memory Read Block		~	Memory Read Block
Alias to Memory Read Block		~	Memory Read Block
Memory Read DWORD		~	Memory Read DWORD
Configuration Read	~	~	Configuration Read
Configuration Write	<b>~</b>	~	Configuration Write
Split Read Completion		~	Split Read Completion
Split Write Completion		~	None
Split Completion Message		~	Split Completion Message
All Other Commands Not Claimed by the ATU	~	~	N/A

#### Table 114. PCI to Internal Bus Command Translation for All Inbound Transactions

## 3.5.2 Outbound Queues

The outbound queues of the ATU are used to hold read and write transactions from the core processor directed at the PCI bus. Each ATU outbound queue structure has a separate read queue, write queue, and address queue. Table 115 contains information about ATU outbound queues.

#### Table 115. Outbound Queues

Queue Mnemonic	Queue Name	Queue Size (Bytes)
OWQ	Outbound Write Data Queue	4 KBytes (4*1024B)
OWADQ	Outbound Write Address Queue	4 Transaction Addresses
ORQ	Outbound Read Data Queue	2 or 4 KBytes (4* 512B or 4*1024B) <sup>a</sup>
OTQ	Outbound Transaction Queue	8 Addresses/Commands

a. The ORQ can be throttled between 2 Kbytes or 4 Kbytes depending on the setting of the Maximum Memory Read Byte Count (MMRBC) field of the PCI-X Command register (see Section 3.10.64, "PCI-X Command Register -PX\_CMD" on page 309). When the MMRBC is set to 512 bytes (default value), the ORQ is only capable of handling 2 Kbytes of SRC data, otherwise, the ORQ handles 4 Kbytes of SRC data.

The outbound queues are capable of holding outbound memory read, memory write, I/O read, and I/O write transactions. The type of transaction used is defined by the internal bus address and the command used on the internal bus. See Section 3.2.2 and Section 3.2.3 for details on outbound address translation.

When an internal bus agent initiates an outbound write transaction, the address is entered into the OWADQ (when not full). The data from the internal bus write is then entered into the OWQ and the transaction is forwarded to the PCI bus. When the write completes (or an error occurs), the address is flushed from the OWADQ. Data is flushed only for the master abort or target abort cases.

For outbound reads, the address is entered into the OTQ (when not full) and a split response termination is signaled to the requester on the internal bus. Read data is fetched and returned to the requester on the internal bus.

## 3.5.3 Transaction Ordering

Because the ATU can process multiple transactions, they must maintain proper ordering to avoid deadlock conditions and improve throughput. The ATU transaction ordering rules used by the 80333 are listed in Table 116 for the inbound direction and Table 117 on page 194 for the outbound direction. The tables are based on the direction the transaction is moving, i.e. the data for outbound delayed read moves in the same direction as the data for an inbound write or the address/command for an inbound read. When operating in the PCI-X mode, the ATU ignores the Relaxed Ordering Attribute.

- *Note:* In PCI mode, the PCI passing rule enforcement logic within the ATU, allows a read completion to pass write data under certain circumstances. For more information see *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #9.
- *Note:* Outbound Non-Posted Writes are the result of Internal Bus Memory writes that are claimed by either the I/O translation window or the Outbound Configuration Cycle Data Register OCCDR. Though these write requests arrive on the PCI bus as non-posted write requests, it is important to note that from the Intel XScale<sup>®</sup> core point of view, these internal bus memory write requests are posted into the Outbound ATU transaction queue. Furthermore, in PCI-X mode non-posted write requests have the potential to be split. Thus, even though a split write completion may be returned to the ATU on the PCI bus for a given outbound non-posted write request, the split write completion is not passed back through to the internal bus. Additionally, strong ordering between outbound memory (posted) write requests and outbound non-posted write requests is **not** maintained as indicated in Table 117 on page 194.

For best performance, the user should designate the two Outbound Memory Windows (8000.0000H to 83FF.FFFFH, 8400.0000H to 87FF.FFFFH) as non-cacheable and bufferable from the Intel XScale<sup>®</sup> core. This assignment enables the Intel XScale<sup>®</sup> core to issue multiple outstanding transactions to the Outbound Memory Windows, thereby, taking full advantage of the ATU outbound queue architecture. However, the user needs to be aware that the Outbound ATU queue architecture does not maintain strict ordering between read and write requests as described in Table 117, "ATU Outbound Data Flow Ordering Rules" on page 194. In the event that the user requires strict ordering to be maintained, the user can either change the designation of this region of memory to be non-cacheable and non-bufferable or enforce the requirement in software.

Row Pass Column? <sup>a</sup>	ATU Inbound Writes	Inbound Delayed Read Request (PCI mode)	Inbound Split Read Request (PCI-X mode)	Inbound Config- uration Write Request	Outbound Split Read Completion
ATU Inbound Writes	No	Yes	Yes	Yes	Yes
Inbound Delayed Read Request (PCI mode)	No	No	NA	No	Yes
Inbound Split Read Request (PCI-X mode)	No	NA	No	No	Yes
Inbound Configuration Write Request	No	No	No	NA	Yes
Outbound Split Read Completion	No	Yes	Yes	Yes	Yes

#### Table 116.ATU Inbound Data Flow Ordering Rules

a. Outbound Non-Posted Write Completions are not included in this table since these transactions are never passed back to the Internal Bus Requester (Intel XScale<sup>®</sup> core). The reason is that from the Intel XScale<sup>®</sup> core point of view, these write requests are posted into the Outbound ATU transaction queue.

Row Pass Column?	Outbound Write		Outbound Read	Inbound Read Completion		Inbound Write Completion	
	Posted	Non- Posted	Request	DRC	SRC	DWC <sup>a</sup>	SWC <sup>b</sup>
Outbound Posted Write Request	No	Yes	Yes	Yes	Yes	Yes	Yes
Outbound Non-posted <sup>c</sup> Write Request	No	No	No	Yes	Yes	Yes	Yes
Outbound Read Request	No	No	No	Yes	Yes	Yes	Yes
Inbound Delayed Read Completion (DRC)	No	Yes	Yes	Yes	NA	Yes	NA
Inbound Split Read Completion (SRC)	No	Yes	Yes	NA	Yes	NA	No
Inbound Delayed Write Completion (DWC)	No	Yes	Yes	Yes	NA	NA	NA
Inbound Split Write Completion (SWC)	No	Yes	Yes	NA	Yes	NA	NA

#### Table 117.ATU Outbound Data Flow Ordering Rules

a. Since the Inbound DWR transaction queue is one-deep, the passing rule associated with DWC vs. DWC is moot (i.e., NA).

b. Since the Inbound SWR transaction queue is one-deep, the passing rule associated with SWC vs. SWC is moot (i.e., NA).

c. Outbound Non-posted writes include I/O writes and Configuration writes.

Definitions of the terms used in Table 116 and Table 117 are as follows. PCI terms are noted in parenthesis:

- Inbound Write (PMW) Data from a write cycle initiated on PCI and targeted at the internal bus. Note that the address is in a separate transaction queue and is not referenced. Inbound writes can also come in through the Messaging Unit which is part of the ATU.
- Inbound Read Request (DRR-PCI mode and SRR-PCI-X mode) address information from a read transactions retried or split on the PCI bus. is mastered on the internal bus to retrieve data for the Inbound Read Completion.
- Inbound Configuration Write Request (DWR- PCI mode and SWR PCI-X mode) The address and data associated with a configuration write transaction from PCI and targeted at the ATU PCI configuration address space. Once completed on the internal bus, creates an Inbound Configuration Write Completion.
- Outbound Read Completion (SRC) The data read on PCI in the process of being returned to the internal bus. This data is the completion cycle that results from an Outbound Read Request.
- Outbound Write (PMW) The address and data from a write initiated on the internal bus and eventually completing on the PCI bus.
- Outbound Read Request (SRR) The address/command of a split read cycle initiated on the internal bus. The read data is returned in the Outbound Read Completion cycle.
- Inbound Read Completion (DRC-PCI mode and SRC-PCI-X mode) The data read on the internal bus in the process of being returned to the PCI bus. This data is the completion cycle for an Inbound Read Request.
- Inbound Configuration Write Completion (DWC-PCI mode and SWC-PCI-X mode) The status of an inbound write configuration cycle traveling from the internal bus back towards the PCI bus.

## intel

These transaction ordering rules define the way in which data moves in both directions through the ATU. In Table 116 and Table 117 a **NO** response in a box means that based on ordering rules, the current transaction (the row) can not pass the previous transaction (the column) under any circumstance. A **Yes** response in the box means that the current transaction is *allowed* to pass the previous transaction but is not required to, based on whether a consistent view of data or prevention of deadlocks is needed.

In the case of inbound write operations, multiple transactions may exist within the IWQ and the corresponding IWADQ at any point in time. The ordering of these transactions is based on a time stamp basis. Transactions entering the queue are stamped with a relative time in relation to all other transactions moving in a similar direction.



#### Example 1. Inbound Queue Completion

In Example 1 on page 195, the inbound write and outbound read queues of the ATU are shown. In this example, transaction A entered the write queue at **Time 0**. Next, the ATU entered read data into the outbound read queue at **Time 1** (Transaction B). Finally, before the previous transactions could be cleared, another inbound write, Transaction C, was entered into the IWQ. The ordering in Table 116 states that nothing can pass an inbound write and therefore Transaction A must complete on the internal bus before Transaction B since an outbound read completion can not pass an inbound write. Also, Transaction A must complete before Transaction C since an inbound write can not pass another inbound write. Once Transaction A completes, Transaction C moves to the head of the IWQ. The two transactions at the head of the queues moving data in an inbound direction are now Transaction C, an inbound write, and Transaction B, an outbound read completion. Ordering states that an inbound write may pass an outbound read completion. This means that the arbitration mechanism now takes over to decide which completes. Note that ordering enforced the completion of Transaction A but arbitration dictated the completion of Transactions B and C.

The first action performed to determine which transaction is allowed to proceed (either inbound or outbound) is to apply the rules of ordering as defined in Table 116 and Table 117. Any box marked **No** must be satisfied first. For example, when an inbound read request is in ITQ and it was latched *after* the data in the IDWQ arrived (this is a configuration write), then ordering states that an Inbound Read Request may not pass an Inbound Configuration Write Request. Therefore, the Inbound Configuration Write Request must be cleared out of IDWQ before the Inbound Read Request is attempted on the internal bus. Once transaction ordering is satisfied, the boxes marked **Yes** are now resolved.

### 3.5.3.1 Transaction Ordering Summary

Table 118 and Table 119, define transaction ordering in relation to token assignment of the priority mechanism (this is discussed in Section 3.5.3). These tables are read as follows:

- 1. As the transaction enters the head of the respective queue, the question in column 2 is asked.
- 2. When all the answers in column 3 for a given transaction type assigns a token to the transaction at the head of the queue, a token is assigned. Otherwise, no token is assigned signifying that transaction ordering must first be satisfied. Any transaction with a token may be initiated on the bus.

#### Table 118.Inbound Transaction Ordering Summary

Transaction at Head of Queue	Question	Answer	Action
		Yes	Do Not Assign Token
Inbound Write in	Is there an Inbound Write Request with an earlier time stamp?		Allow previous Transaction to Complete
		No	Assign Token
	Is there an Inbound Write with an	Voc	Do Not Assign Token
		res	Allow previous Transaction to Complete
		No	Assign Token
	Is there an Inbound Read Request with	Voo	Do Not Assign Token
Inbound Read		res	Allow previous Transaction to Complete
Requestioning		No	Assign Token
	Is there an Inbound Configuration Write Request with an earlier time stamp?	Yes	Do Not Assign Token
			Allow previous Transaction to Complete
		No	Assign Token
	Is there an Inbound Write with an earlier time stamp?	Yes	Do Not Assign Token
Inbound Configuration Write Request in IDWQ			Allow previous Transaction to Complete
		No	Assign Token
	Is there an Inbound Read Request with	Yes	Do Not Assign Token
			Allow previous Transaction to Complete
		No	Assign Token
Outbound Read Completion in ORQ		Yes	Do Not Assign Token
	Is there an Inbound Write with an earlier time stamp?		Allow previous Transaction to Complete
		No	Assign Token



### Table 119. Outbound Transaction Ordering Summary

Transaction at Head of Queue	Question	Answer	Action
Outbound Write in OWQ	Is there an Outbound Write Request with an earlier time stamp?	Yes	Do Not Assign Token Allow previous Transaction to Complete
		No	Assign Token
	Is there an Outbound Write with an earlier time stamp?	Yes	Do Not Assign Token Allow previous Transaction to Complete
Outbound Read		No	Assign Token
Request in OTQ	Is there an Outbound Read Request with an earlier time stamp?	Yes	Do Not Assign Token Allow previous Transaction to Complete
		No	Assign Token
	Is there an Outbound Posted Write with an earlier time stamp?	Yes	Do Not Assign Token Allow previous Transaction to Complete
Inbound		No	Assign Token
Write Completion in IRQ	Is there an Inbound Read Completion with an earlier time stamp?	Yes	Do Not Assign Token and allow previous Transaction to Complete when in PCI-X Mode.
			Assign Token when in Conventional Mode
		No	Assign Token
	Is there an Outbound Posted Write with an earlier time stamp?	Yes	Do Not Assign Token Allow previous Transaction to Complete
Inbound Read Completion in IRQ		No	Assign Token
	Is there an Inbound Read Completion with an earlier time stamp?	Yes	Do Not Assign Token and allow previous Transaction to Complete when in PCI-X Mode. Assign Token when in Conventional
		No	Assign Token
	Is there a Configuration Write Completion with an earlier time stamp?	Yes	Do Not Assign Token and allow previous Transaction to Complete when in PCI-X Mode. Assign Token when in Conventional Mode
		No	Assign Token



## 3.6 **Private Device Control**

Private devices are hidden from PCI configuration software but are accessible from the Address Translation Unit. To include private devices on the secondary PCI interface, the 80333 bridge unit and the ATU must be configured properly prior to system initialization following **PWRGD** deassertion.

To configure the 80333 to control private devices, the PCI Express-to-PCI Bridge must be configured to:

- Enable private Type 0 commands on the A-Segment Interface
- Enable the private memory space on the A-Segment Interface

Refer to Section 2.3.2, "Addressing" on page 59 for details on how to properly configure this capability.

The ATU is internally mapped to PCI Device Number 14 (0EH) with AD30 used as the ATU IDSEL input. This mapping also has the ATU interrupt wired to the **IRQ14#** IOAPIC pin.

## 3.6.1 Private Type 0 Commands on the Secondary Interface

Type 0 configuration reads and write commands can be generated by the Address Translation Unit of the80333. These Type 0 configuration commands are required to configure private PCI devices on the A-segment PCI bus which are in private PCI address space. These commands are initiated by the Address Translation Unit and not by Type 1 commands on the PCI Express bus. Any device mapped into this private address space *is not* part of the standard A-segment PCI address space and therefore is not configured by the system host processor. These devices are hidden from PCI configuration software but are accessible from the Address Translation Unit.

The bridge can be configured to block 10 address lines during Configuration transactions from the PCI Express interface. A-segment PCI address bits **A\_AD[25:16]** are the address lines that can be masked by the bridge private device control.



Figure 17 shows an example of connecting PCI Address/Data lines to **IDSEL** inputs of PCI devices and private PCI devices.





#### NOTES:

1. A\_AD[27:24] are reserved, see table 9

2. AD30 is dedicated to the Address Translation Unit IDSEL input, and must not be used for an external PCI device.

## 3.6.2 Private Memory Space

When Private Addressing Enable in the BridgeBINIT (see Section 2.9.5.56, "Bridge Initialization Register - BINIT (Offset FC)" on page 137 is set, the bridge overrides its secondary inverse decode logic and not forward upstream any secondary bus initiated DAC Memory transactions with address in the upper half of 64-bit address space (i.e.,AD(63:62)=10b). This establishes a private memory space for A-Segment bus usage, independent of bridge downstream forwarding configuration.

## 3.7 ATU Error Conditions

PCI and internal bus error conditions cause ATU state machines to exit normal operation and return to idle states. In addition, status bits are set to inform error handling code of exact cause of error condition. All of the Messaging Unit errors are reported in the same manner as ATU errors. Error conditions and status can be found in the ATUSR. The basic flow for a PCI error is as follows:

- Set the bit in the ATU Status Register which corresponds to the error condition (master abort, target abort, etc.)
- Set the bit in the ATU Interrupt Status Register which corresponds to the error condition (master abort, target abort, etc.). This function is maskable for all PCI error conditions.
- The setting of the bit in the ATU Interrupt Status Register results in an interrupt being driven to the Intel XScale<sup>®</sup> core.

Error conditions on one side of the ATU are generally propagated to the other side of the ATU and have different effects depending on the error. Error conditions and their effects are described in the following sections.

PCI bus error conditions and the action taken on the bus are defined within the *PCI Local Bus Specification*, Revision 2.3, and the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a. The ATU adheres to the error conditions defined within the PCI specification for both requester and target operation. Error conditions on the internal bus are caused by an ECC error from the Memory Controller (see Section 8.5, "Interrupts/Error Conditions" on page 542 for details on memory controller error conditions) or by incorrect addressing resulting in an internal master abort. All actions on the PCI Bus for error situations are dependent on the error control bits found in the ATU Command Register (see Section 3.10.3, "ATU Command Register - ATUCMD" on page 240) for both Conventional and PCI-X modes. For PCI-X mode, the error response is also dependent on an error control bit in the PCI-X Command Register (see Section 3.10.64, "PCI-X Command Register - PX\_CMD" on page 309).

The following sections detail all ATU error conditions on the PCI and 80333 internal bus, action taken on these conditions, and status and control bits associated with error handling.

## 3.7.1 Address and Attribute Parity Errors on the PCI Interface

The ATUs must detect and report address and attribute (PCI-X mode only) parity errors for transactions on the PCI bus. When an address or attribute parity error occurs on the PCI interface of the ATU, the 80333 performs the following actions based on the constraints specified:

- In Conventional mode, when the Parity Error Response bit in ATUCMD is set, the ATU ignores (Master-Abort) the transaction by not asserting **DEVSEL#**. When clear, the transaction proceeds normally.
- In PCI-X mode, when the Parity Error Response bit in ATUCMD is set, the ATU completes the transaction on the PCI bus as though no error had occurred, but the request or completion is not forwarded to the internal bus. When clear, the transaction proceeds normally.
- Assert **SERR**# when the **SERR**# Enable bit and the Parity Error Response bit in the ATUCMD are set. When the ATU asserts **SERR**#, additional actions is taken:
  - Set the **SERR**# Asserted bit in the ATUSR
  - When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.
  - When the ATU SERR# Detected Interrupt Enable Bit in the ATUCR (see Section 3.10.43, "ATU Configuration Register - ATUCR" on page 282) is set, set the SERR# Detected bit in the ATUISR. When clear, no action.
- Set the Detected Address or Attribute Parity Error bit in PCSR (PCI Configuration and Status Register).
- Set the Detected Parity Error bit in the ATUSR. When the ATU sets the Detected Parity Error bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.
- *Note:* The Detected Parity Error bit with its associated interrupt along with the Detected Address or Attribute Parity Error bit provides software with the ability to distinguish between an Address or Attribute Parity error versus a Data Parity Error during a Detected Parity error interrupt.



## 3.7.2 Data Parity Errors on the PCI Interface

Two kinds of data parity errors can occur on the PCI interface: errors as an initiator and errors as a target.

Errors encountered as an initiator:

- Outbound Read Request
- Outbound Write Request
- Inbound Read Completions
- Inbound Configuration Write Completion Messages

As an initiator, the ATU provides an error response for data parity errors on outbound reads, and data parity errors occurring at the target for outbound writes. However, there is no error response for data parity errors on inbound configuration write completion messages and inbound read completions.

Errors encountered as a target:

- Inbound Read Request (Immediate Data Transfer)
- Inbound Write Request
- Outbound Read Completions
- Outbound Split (I/O or Configuration) Write Data Parity Error Messages
- Inbound Configuration Write
- Split Completion Messages

As a target, the ATU provides an error response for data parity errors on inbound writes, outbound read completions, outbound split write data parity error messages, inbound configuration writes, and split completion messages. However, there is no error response for data parity errors on inbound reads.



## 3.7.2.1 Outbound Read Request Data Parity Errors

#### 3.7.2.1.1 Immediate Data Transfer

As an initiator, the ATU may encounter this error condition in Conventional or PCI-X mode when the target transfers data immediately rather than signalling a Retry<sup>1</sup> (Conventional Delayed Read Request) or a Split Response Termination (PCI-X Split Read Request).

Data parity errors occurring during read operations initiated by the ATU are recorded, **PERR**# is asserted (when enabled) and **SERR**# is asserted (when enabled). Specifically, the following actions with the given constraints are taken by the ATU:

- **PERR**# is asserted two clocks cycles (three clock cycles when operating in the PCI-X mode) following the data phase in which the data parity error is detected on the bus. This is only done when the Parity Error Response bit in the ATUCMD is set. When the ATU asserts **PERR**#, additional actions is taken:
  - The Master Parity Error bit in the ATUSR is set.
  - When the ATU PCI Master Parity Error Interrupt Mask Bit in the ATUIMR is clear, set the PCI Master Parity Error bit in the ATUISR. When set, no action.
  - When the ATU is operating in the PCI-X mode, the SERR# Enable bit in the ATUCMD is set, and the Data Parity Error Recover Enable bit in the PCIXCMD register is clear, assert SERR#, otherwise no action. When the ATU asserts SERR#, additional actions is taken: Set the SERR# Asserted bit in the ATUSR.

When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.

When the ATU **SERR**# Detected Interrupt Enable Bit in the ATUCR is set, set the **SERR**# Detected bit in the ATUISR. When clear, no action.

- A completion error message (with message class=2h completer error and message index=82h - PCI bus read parity error) is generated on the internal bus of the 80333.
- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.

<sup>1.</sup> Retry terminations may also be signaled in PCI-X mode when the target is too busy to handle the current request. However, this is not the same as a Delayed Read Request in Conventional PCI mode since the requester is not required or expected by the target to return with the same read request.



#### 3.7.2.1.2 Split Response Termination

As an initiator, the ATU may encounter this error condition in PCI-X mode when the target signals a Split Response Termination.

Parity errors occurring during Split Response Terminations of Read Requests by the ATU are recorded, **PERR**# is asserted (when enabled) and **SERR**# is asserted (when enabled). Specifically, the following actions with the given constraints are taken by the ATU:

- **PERR**# is asserted two clocks cycles (three clock cycles when operating in the PCI-X mode) following the Split Response Termination in which the parity error is detected on the bus. This is only done when the Parity Error Response bit in the ATUCMD is set. When the ATU asserts **PERR**#, additional actions is taken:
  - The Master Parity Error bit in the ATUSR is set.
  - When the ATU PCI Master Parity Error Interrupt Mask Bit in the ATUIMR is clear, set the PCI Master Parity Error bit in the ATUISR. When set, no action.
  - When the SERR# Enable bit in the ATUCMD is set, and the Data Parity Error Recover Enable bit in the PCIXCMD register is clear, assert SERR#, otherwise no action. When the ATU asserts SERR#, additional actions is taken:
    - Set the SERR# Asserted bit in the ATUSR.
    - When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.
    - When the ATU **SERR**# Detected Interrupt Enable Bit in the ATUCR is set, set the **SERR**# Detected bit in the ATUISR. When clear, no action
- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.

The Outbound Read Request remains enqueued in the ATU as the completer is initiating completion transactions that are associated with this request.



## 3.7.2.2 Outbound Write Request Data Parity Errors

## 3.7.2.2.1 Outbound Writes that are not MSI (Message Signaled Interrupts)

As an initiator, the ATU may encounter this error condition when operating in either the Conventional or PCI-X modes.

Data parity errors occurring during write operations initiated by the ATU may record the assertion of **PERR**# from the target on the PCI Bus. In PCI-X mode, this includes the assertion of **PERR**# from the target on the PCI Bus following the split response termination of a non-posted write request. When an error occurs, the ATUs continues writing data to the target to clear the OWQ of the current outbound write transaction. Specifically, the following actions with the given constraints are taken by the ATU:

- When **PERR**# is sampled active and the Parity Error Response bit in the ATUCMD is set, set the Master Parity Error bit in the ATUSR. When the Parity Error Response bit in the ATUCMD is clear, no action is taken. When the Master Parity Error bit in the ATUSR is set, additional actions is taken:
  - When the ATU PCI Master Parity Error Interrupt Mask Bit in the ATUIMR is clear, set the PCI Master Parity Error bit in the ATUISR. When set, no action.
  - When the ATU is operating in the PCI-X mode, the SERR# Enable bit in the ATUCMD is set, and the Data Parity Error Recover Enable bit in the PCIXCMD register is clear, assert SERR#, otherwise no action. When the ATU asserts SERR#, additional actions is taken:

Set the SERR# Asserted bit in the ATUSR

When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.

When the ATU **SERR**# Detected Interrupt Enable Bit in the ATUCR is set, set the **SERR**# Detected bit in the ATUISR. When clear, no action

Outbound write parity errors, does not result in a master completion. In addition, when the target terminates the transaction (disconnect), the ATU master must reinitiate the transaction to clear the data from the OWQ.



#### 3.7.2.2.2 MSI Outbound Writes

As an initiator, the ATU may encounter this error condition when operating in either the Conventional or PCI-X modes.

Data parity errors occurring during MSI write operations initiated by the ATU may record the assertion of **PERR#** from the target on the PCI Bus. When an error occurs, the ATU completes the transaction normally. Then, the following actions with the given constraints are taken by the ATU:

- When **PERR**# is sampled active and the Parity Error Response bit in the ATUCMD is set, set the Master Parity Error bit in the ATUSR. When the Parity Error Response bit in the ATUCMD is clear, no action is taken. When the Master Parity Error bit in the ATUSR is set, additional actions is taken:
  - When the ATU PCI Master Parity Error Interrupt Mask Bit in the ATUIMR is clear, set the PCI Master Parity Error bit in the ATUISR. When set, no action.
  - When the SERR# Enable bit in the ATUCMD is set, assert SERR#, otherwise no action.
     When the ATU asserts SERR#, additional actions is taken:

Set the **SERR**# Asserted bit in the ATUSR.

When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.

When the ATU **SERR**# Detected Interrupt Enable Bit in the ATUCR is set, set the **SERR**# Detected bit in the ATUISR. When clear, no action.



## 3.7.2.3 Inbound Read Completions Data Parity Errors

As an initiator, the ATU may encounter this error condition when operating in the PCI-X mode.

When as the completer of a Split Read Request the ATU observes **PERR**# assertion during the split completion transaction, the ATU attempts to complete the transaction normally and no further action is taken.

#### 3.7.2.4 Inbound Configuration Write Completion Message Data Parity Errors

As an initiator, the ATU may encounter this error condition when operating in the PCI-X mode.

When as the completer of a Configuration (Split) Write Request the ATU observes **PERR#** assertion during the split completion transaction, the ATU attempts to complete the transaction normally and no further action is taken.

#### 3.7.2.5 Inbound Read Request Data Parity Errors

#### 3.7.2.5.1 Immediate Data Transfer

As a target, the ATU may encounter this error when operating in the Conventional PCI or PCI-X modes.

Inbound read data parity errors occur when read data delivered from the IRQ is detected as having bad parity by the initiator of the transaction who is receiving the data. The initiator may optionally report the error to the system by asserting **PERR#**. As a target device in this scenario, no action is required and no error bits are set.

#### 3.7.2.5.2 Split Response Termination

As a target, the ATU may encounter this error when operating in the PCI-X mode.

Inbound read data parity errors occur during the Split Response Termination. The initiator may optionally report the error to the system by asserting **PERR#**. As a target device in this scenario, no action is required and no error bits are set.

#### 3.7.2.6 Inbound Write Request Data Parity Errors

As a target, the ATU may encounter this error when operating in the Conventional or PCI-X modes.

Data parity errors occurring during write operations received by the ATU may assert **PERR#** on the PCI Bus. When an error occurs, the ATU continues accepting data until the initiator of the write transaction completes or a queue fill condition is reached. Specifically, the following actions with the given constraints are taken by the ATU:

- **PERR**# is asserted two clocks cycles (three clock cycles when operating in the PCI-X mode) following the data phase in which the data parity error is detected on the bus. This is only done when the Parity Error Response bit in the ATUCMD is set.
- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.



## 3.7.2.7 Outbound Read Completion Data Parity Errors

As a target, the ATU may encounter this error when operating in the PCI-X mode.

Data parity errors occurring during read completion transactions that are claimed by the ATU are recorded, **PERR**# is asserted (when enabled) and **SERR**# is asserted (when enabled). Specifically, the following actions with the given constraints are taken by the ATU:

- **PERR**# is asserted two clocks cycles (three clock cycles when operating in the PCI-X mode) following the data phase in which the data parity error is detected on the bus. This is only done when the Parity Error Response bit in the ATUCMD is set. When the ATU asserts **PERR**#, additional actions is taken:
  - The Master Parity Error bit in the ATUSR is set.
  - When the ATU PCI Master Parity Error Interrupt Mask Bit in the ATUIMR is clear, set the PCI Master Parity Error bit in the ATUISR. When set, no action.
  - When the SERR# Enable bit in the ATUCMD is set, and the Data Parity Error Recover Enable bit in the PCIXCMD register is clear, assert SERR#, otherwise no action. When the ATU asserts SERR#, additional actions is taken:

Set the SERR# Asserted bit in the ATUSR.

When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.

When the ATU **SERR**# Detected Interrupt Enable Bit in the ATUCR is set, set the **SERR**# Detected bit in the ATUISR. When clear, no action.

- A completion error message (with message class=2h completer error and message index=82h - PCI bus read parity error) is generated on the internal bus of the 80333.
- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.

## intel

## 3.7.2.8 Outbound Split Write Data Parity Error Message

The ATU claims a Split Completion Error Message that indicates a data parity error has occurred on one of the ATUs non-posted (I/O or Configuration) write requests (Message Class = 2h, Message Index = 01h -- Split Write Data Parity Error or Message Class = 1h, Message Index = 02h -- Write Data Parity Error).

When the ATU receives a Split Completion Error Message indicating a data parity error for an outstanding Configuration or I/O (split) write request, the error is recorded, and **SERR**# is asserted (when enabled). Specifically, the following actions with the given constraints are taken by the ATU:

- When the Parity Error Response bit in the ATUCMD is set, these actions is taken:
  - The Master Parity Error bit in the ATUSR is set.
  - When the ATU PCI Master Parity Error Interrupt Mask Bit in the ATUIMR is clear, set the PCI Master Parity Error bit in the ATUISR. When set, no action.
  - When the SERR# Enable bit in the ATUCMD is set, and the Data Parity Error Recover Enable bit in the PCIXCMD register is clear, assert SERR#, otherwise no action. When the ATU asserts SERR#, additional actions is taken:

Set the SERR# Asserted bit in the ATUSR

When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.

When the ATU **SERR**# Detected Interrupt Enable Bit in the ATUCR is set, set the **SERR**# Detected bit in the ATUISR. When clear, no action.

- The Received Split Completion Error Message bit in the PCIXSR is set (based on bit 30 of the completer attributes being set). When the ATU sets this bit, additional actions is taken:
  - When the ATU Received Split Completion Error Message Interrupt Mask bit in the ATUIMR is clear, set the Received Split Completion Error Message bit in the ATUISR. When set, no action.
- The transaction associated with the Split Completion Error Message is discarded.



### 3.7.2.9 Inbound Configuration Write Request

As a target, the ATU may encounter this error when operating in the Conventional or PCI-X modes.

#### 3.7.2.9.1 Conventional PCI Mode

To allow for correct data parity calculations for delayed write transactions, the ATU delays the assertion of STOP# (signalling a Retry) until **PAR** is driven by the master. A parity error during a delayed write transaction (inbound configuration write cycle) can occur in any of the following parts of the transactions:

- During the initial Delayed Write Request cycle on the PCI bus when the ATU latches the address/command and data for delayed delivery to the internal configuration register.
- During the Delayed Write Completion cycle on the PCI bus when the ATU delivers the status of the operation back to the original master.

The 80333 ATU PCI interface has the following responses to a delayed write parity error for inbound transactions during Delayed Write Request cycles with the given constraints:

• When the Parity Error Response bit in the ATUCMD is set, the ATU asserts TRDY# (disconnects with data) and two clock cycles later asserts **PERR**# notifying the initiator of the parity error. The delayed write cycle is not enqueued and forwarded to the internal bus.

When the Parity Error Response bit in the ATUCMD is cleared, the ATU retries the transaction by asserting STOP# and enqueues the Delayed Write Request cycle to be forwarded to the internal bus. **PERR**# is not asserted.

- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.

For the original write transaction to be completed, the initiator retries the transaction on the PCI bus and the ATU returns the status from the internal bus, completing the transaction.

For the Delayed Write Completion transaction on the PCI bus where a data parity error occurs and therefore does not agree with the status being returned from the internal bus (i.e. status being returned is normal completion) the ATU performs the following actions with the given constraints:

- When the Parity Error Response Bit is set in the ATUCMD, the ATU asserts TRDY# (disconnects with data) and two clocks later asserts **PERR**#. The Delayed Completion cycle in the IDWQ remains since the data of retried command did not match the data within the queue.
- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.

#### 3.7.2.9.2 PCI-X Mode

Data parity errors occurring during configuration write operations received by the ATU may cause **PERR**# assertion and delivery of a Split Completion Error Message on the PCI Bus. When an error occurs, the ATU accepts the write data and complete with a Split Response Termination. Specifically, the following actions with the given constraints are then taken by the ATU:

- When the Parity Error Response bit in the ATUCMD is set, **PERR**# is asserted three clocks cycles following the Split Response Termination in which the data parity error is detected on the bus. When the ATU asserts **PERR**#, additional actions is taken:
  - A Split Write Data Parity Error message (with message class=2h completer error and message index=01h - Split Write Data Parity Error) is initiated by the ATU on the PCI bus that addresses the requester of the configuration write.
  - When the Initiated Split Completion Error Message Interrupt Mask in the ATUIMR is clear, set the Initiated Split Completion Error Message bit in the ATUISR. When set, no action.
  - The Split Write Request is not enqueued and forwarded to the internal bus.
- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.



### 3.7.2.10 Split Completion Messages

As a target, the ATU may encounter this error when operating in the PCI-X mode.

Data parity errors occurring during Split Completion Messages claimed by the ATU may assert **PERR#** (when enabled) or **SERR#** (when enabled) on the PCI Bus. When an error occurs, the ATU accepts the data and complete normally. Specifically, the following actions with the given constraints are taken by the ATU:

- **PERR**# is asserted three clocks cycles following the data phase in which the data parity error is detected on the bus. This is only done when the Parity Error Response bit in the ATUCMD is set. When the ATU asserts **PERR**#, additional actions is taken:
  - The Master Parity Error bit in the ATUSR is set.
  - When the ATU PCI Master Parity Error Interrupt Mask Bit in the ATUIMR is clear, set the PCI Master Parity Error bit in the ATUISR. When set, no action.
  - When the SERR# Enable bit in the ATUCMD is set, and the Data Parity Error Recover Enable bit in the PCIXCMD register is clear, assert SERR#; otherwise no action is taken. When the ATU asserts SERR#, additional actions is taken:
    - Set the SERR# Asserted bit in the ATUSR.

When the ATU **SERR**# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.

When the ATU **SERR**# Detected Interrupt Enable Bit in the ATUCR is set, set the **SERR**# Detected bit in the ATUISR. When clear, no action.

- When the SCE bit (Split Completion Error -- bit 30 of the Completer Attributes) is set during the Attribute phase, the Received Split Completion Error Message bit in the PCIXSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Received Split Completion Error Message Interrupt Mask bit in the ATUIMR is clear, set the Received Split Completion Error Message bit in the ATUISR. When set, no action.
- The Detected Parity Error bit in the ATUSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Detected Parity Error Interrupt Mask bit in the ATUIMR is clear, set the Detected Parity Error bit in the ATUISR. When set, no action.
- The transaction associated with the Split Completion Message is discarded.
- When the discarded transaction was a read, a completion error message (with message class=2h completer error and message index=82h PCI bus read parity error) is generated on the internal bus of the 80333.

# intel

## 3.7.3 Master Aborts on the PCI Interface

As an initiator on the PCI bus, the ATU can encounter master abort conditions during:

- Outbound Read Request
- Outbound Write Request
- Inbound Read Completion
- Inbound Configuration Write Completion Message

As a target, the ATU PCI interface is capable of signaling a master abort case during:

- Address Parity Error (Conventional Mode)
- Inbound Read Request (PCI-X Mode)

#### 3.7.3.1 Master Aborts for Outbound Read or Write Request

This error may be encountered in both the Conventional and the PCI-X modes. For an Outbound transaction, there are two ways in which a Master-Abort may be signaled to the ATU:

- 1. In the Conventional or PCI-X modes, a master abort is signaled when the target of the transaction does not assert **DEVSEL**# within 5 clocks (7 clocks when operating in the PCI-X mode) of the assertion of **FRAME**#.
- 2. In PCI-X mode, ATU may enqueue a Split request (Read or Write) on target-side interface of a PCI-to-PCI Bridge. When PCI-to-PCI Bridge detects a Master Abort on requester-side interface for that Split Request, master abort is signaled to ATU through a Master-Abort Split Completion Error Message (class=1h bridge error and index=00h Master Abort). The following actions with given constraints are performed by ATU when a master abort is detected by the PCI initiator interface or the PCI target interface receives a Master-Abort Split Completion error message:
- Set the Master Abort bit (bit 13) in the ATUSR.
- When the ATU PCI Master Abort Interrupt Mask bit in the ATUIMR is clear, set the PCI Master Abort bit in the ATUISR. When set, no action.
- When an outbound write or inbound completion, flush data and address.
- When the transaction is an MSI outbound write and the **SERR**# Enable bit in the ATUCMD is set, assert **SERR**#, otherwise no action. When the ATU asserts **SERR**#, additional actions is taken:
  - Set the SERR# Asserted bit in the ATUSR
  - When the ATU SERR# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the SERR# Asserted bit in the ATUISR. When set, no action.
  - When the ATU SERR# Detected Interrupt Enable Bit in the ATUCR is set, set the SERR# Detected bit in the ATUISR. When clear, no action
- When operating in PCI-X mode and Master-Abort is signaled via a Split Completion Error Message, the Received Split Completion Error Message bit in PCIXSR is set. When ATU sets this bit, additional actions is taken:
  - When the ATU Received Split Completion Error Message Interrupt Mask bit in the ATUIMR is clear, set the Received Split Completion Error Message bit in the ATUISR. When set, no action.
- For an Outbound Read request, generate a split completion error message (class=1h 80333 Outbound Request error and index=00h master abort) on the internal bus.
- Flush the address from the OTQ.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Address Translation Unit



#### 3.7.3.2 Inbound Read Completion or Inbound Configuration Write Completion Message

The ATU encounters this error only in the PCI-X mode.

A master abort is signaled when the target of the transaction does not assert **DEVSEL#** within 7 clocks of the assertion of **FRAME#**.

When the ATU is signaled a Master-Abort while initiating either a Split Read Completion Transaction or a Split Write Completion Message, the ATU discards the Split Completion and take no further action.

### 3.7.3.3 Master-Aborts Signaled by the ATU as a Target

#### 3.7.3.3.1 Address Parity Errors

The ATU can only signal this error during an Address Parity Error in the Conventional mode.

Please see Section 3.7.1, "Address and Attribute Parity Errors on the PCI Interface" on page 201 for details on the ATU response to an Address Parity Error in the Conventional mode.

#### 3.7.3.3.2 Internal Bus Master-Abort

The ATU can only signal this error during an internal bus master abort in the PCI-X mode.

Please see Section 3.7.7.1, "Master Abort on the Internal Bus" on page 219 for details on the ATU response to an Internal Bus Master Abort in the PCI-X mode.

# intel

## 3.7.4 Target Aborts on the PCI Interface

As an initiator on the PCI bus, the ATU can encounter target abort conditions during:

- Outbound Read Request
- Outbound Write Request
- Inbound Read Completion
- Inbound Configuration Write Completion Message

As a target, the ATU PCI interface is capable of signaling a target abort case during:

- Inbound Read Request (PCI-X and Conventional Modes)
- Inbound Write Request to EROM memory space (PCI-X and Conventional Modes)

#### 3.7.4.1 Target Aborts for Outbound Read Request or Outbound Write Request

This error can be encountered by the ATU in both the Conventional and PCI-X modes. For an Outbound transaction, there are two ways in which a Target-Abort may be signaled to the ATU:

- 1. In the Conventional or PCI-X modes, a target abort is signaled when the target of the transaction simultaneously deasserts **DEVSEL**#, deasserts **TRDY**#, and asserts **STOP**#.
- 2. In PCI-X mode, ATU may enqueue a Split request (Read or Write) on target-side interface of a PCI-to-PCI Bridge. When PCI-to-PCI Bridge detects a Target Abort on requester-side interface for that Split Request, target abort is signaled to ATU through a Target-Abort Split Completion Error Message (class=1h bridge error and index=01h Target Abort). The following actions with the given constraints are performed by the ATU when a target abort is detected by the PCI initiator interface or the PCI target interface receives a Target-Abort Split Completion error message:
- Set the Target Abort (master) bit (bit 12) in the ATUSR.
- When the ATU PCI Target Abort (master) Interrupt Mask bit in the ATUIMR is clear, set the PCI Target Abort (master) bit in the ATUISR. When set, no action.
- When the transaction is an MSI outbound write and the **SERR**# Enable bit in the ATUCMD is set, assert **SERR**#; otherwise, no action is taken. When the ATU asserts **SERR**#, additional actions is taken:
  - Set the **SERR**# Asserted bit in the ATUSR.
  - When the ATU SERR# Asserted Interrupt Mask Bit in the ATUIMR is clear, set the SERR# Asserted bit in the ATUISR. When set, no action.
  - When the ATU SERR# Detected Interrupt Enable Bit in the ATUCR is set, set the SERR# Detected bit in the ATUISR. When clear, no action.
- When operating in the PCI-X mode and the Target-Abort is signaled via a Split Completion Error Message, the Received Split Completion Error Message bit in the PCIXSR is set. When the ATU sets this bit, additional actions is taken:
  - When the ATU Received Split Completion Error Message Interrupt Mask bit in the ATUIMR is clear, set the Received Split Completion Error Message bit in the ATUISR. When set, no action.
- For an Outbound Read request, generate a split completion error message (message class=1h 80333 Outbound Request error and message index=01h target abort) on internal bus.
- Flush the address from the OTQ.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Address Translation Unit



#### 3.7.4.2 Inbound Read Completion or Inbound Configuration Write Completion Message

The ATU encounters this error only in the PCI-X mode.

A target abort is signaled when the target of the transaction simultaneously deasserts **DEVSEL**#, deasserts **TRDY**#, and asserts **STOP**#.

When the ATU is signaled a Target-Abort while initiating either a Split Read Completion Transaction or a Split Write Completion Message, the ATU discards the Split Completion and take no further action

#### 3.7.4.3 Target-Aborts Signaled by the ATU as a Target

#### 3.7.4.3.1 Internal Bus Master Abort

A target abort can be signaled by the ATU during an inbound read request where the internal bus cycle resulted in a master abort on the Internal Bus.

Please see Section 3.7.7.1, "Master Abort on the Internal Bus" on page 219 for details on the ATU response to an Internal Bus Master Abort.

#### 3.7.4.3.2 Internal Bus Target Abort

A target abort can be signaled by the ATU during an inbound read request where the internal bus cycle resulted in a Target Abort from the memory controller due to a non-recoverable multi-bit ECC error.

Please see Section 3.7.7.2, "Target Abort on the Internal Bus" on page 221 for details on the ATU response to an Internal Bus Target Abort.

#### 3.7.4.3.3 Inbound EROM Memory Write

Since the EROM memory window is defined to be read-only by the *PCI Local Bus Specification*, Revision 2.3, the ATU target-aborts when an inbound write transaction is claimed by the EROM memory window.

The following additional actions with the given constraints are performed by the ATU when a target abort is signaled by the PCI target interface during an inbound EROM Memory write transaction:

- Set the Target Abort (target) bit (bit 11) in the ATUSR.
  - When the ATU PCI Target Abort (target) Interrupt Mask bit in the ATUIMR is clear, set the PCI Target Abort (target) bit in the ATUISR. When set, no action.


## 3.7.5 Corrupted or Unexpected Split Completions

*Warning:* When any of the errors discussed in this section actually occur, a catastrophic system failure is likely to result from which the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a provides no recovery mechanism. In these cases, the ATU may be communicating with a non-compliant target device or the system may not be configured properly.

## 3.7.5.1 Completer Address

The ATU only asserts **DEVSEL**# for split completion transactions where the Sequence ID (Requester ID and Tag) matches that of a currently outstanding split request in the OTQ.

Conversely, the ATU does not assert **DEVSEL#** for any split completion transaction where either the Requester ID does not match that of the ATU or the Tag does not match that of any currently outstanding split request. No further action is taken.

When the Sequence ID of a split completion transaction matches that of an outstanding request, but the Lower Address field is not valid, the ATU accepts the split completion transaction in its' entirety according to the invalid Lower Address field and set the Unexpected Split Completion bit in the PCIXSR. No further action is taken.

## 3.7.5.2 Completer Attributes

When the Sequence ID of a split completion transaction matches that of an outstanding request, but the Byte Count is not valid, the ATU accepts the split completion transaction in its' entirety according to the invalid byte count field and set the Unexpected Split Completion bit in the PCIXSR. In this case, the ATU discards all the data. No further action is taken.



## 3.7.6 SERR# Assertion and Detection

The ATU is capable of reporting error conditions through the use of the SERR# output.

The following conditions may result in the assertion of **SERR**# by the ATU:

- An address parity error (or an attribute parity error when operating in the PCI-X mode) is detected by the ATU PCI interface (see Section 3.7.1, "Address and Attribute Parity Errors on the PCI Interface" on page 201 for details).
- A Master Data Parity Error is recorded in the ATUSR while operating in the PCI-X mode (see Section 3.7.2, "Data Parity Errors on the PCI Interface" on page 202 for details).
- An outbound MSI write transaction is either signaled a Master-Abort or a Target-Abort by the target.
- An inbound write transaction is master aborted on the internal bus (see Section 3.7.7.1, "Master Abort on the Internal Bus" on page 219 for details).
- The **SERR**# Manual Assertion bit in the ATUCR has been set and the **SERR**# Enable bit is set in the ATUCMD.

Note that the **SERR**# manual assertion bits must be cleared manually before they can be set again resulting in **SERR**# asserted. Refer to Section 3.10.43, "ATU Configuration Register - ATUCR" on page 282 for details.

The following actions with the given constraints are performed by the ATU when **SERR#** is asserted by the PCI interface:

- Set the **SERR**# Asserted bit in the ATUSR.
- When the ATU **SERR**# Asserted Interrupt Mask bit in the ATUIMR is clear, set the **SERR**# Asserted bit in the ATUISR. When set, no action.
- When **SERR**# is asserted and the ATU **SERR**# Detected interrupt enable is set in the ATUCR, set the **SERR**# Detected bit in the ATUISR. When clear, no action.

The following actions with the given constraints are performed by the ATU when **SERR#** is detected by the PCI interface:

- When **SERR**# is detected and the ATU **SERR**# Detected interrupt enable is set in the ATUCR, set the **SERR**# Detected bit in the ATUISR. When clear, no action.
- *Note:* Whenever the ATU asserts **SERR**#, both the asserted and detected status bits may be set in the corresponding ISR. To mask an interrupt to the core when the ATU asserts **SERR**#, the **SERR**# asserted mask bit must be set and the **SERR**# detected interrupt enable bit must be clear.

# intel

# 3.7.7 Internal Bus Error Conditions

The 80333 internal bus uses a protocol similar to the *PCI-X Addendum to the PCI Local Bus Specification,* Revision 1.0a. As such, master abort and target abort conditions are valid error states on the bus. The error handling protocol for internal bus conditions is similar to the PCI bus error protocol. An internal bus error results in a bit being set in the Interrupt Status Registers at which time an interrupt is driven to the Intel XScale<sup>®</sup> core. Unlike PCI errors, internal bus error conditions are not maskable.

The following sections detail internal bus error conditions for the ATU.

## 3.7.7.1 Master Abort on the Internal Bus

A master abort on the internal bus is seen by the ATU when the inbound translated address presented on the internal bus is not claimed by the assertion of **I\_DEVSEL#**.

### 3.7.7.1.1 Inbound Write Request

The following action with the given constraints are performed by the ATU when a master abort is detected by the internal master interface during an inbound write request transaction:

- Set the Internal Bus Master Abort bit (bit 7) in the ATUISR.
- When the Inbound Error **SERR**# Enable bit is set in the ATUIMR and the **SERR**# Enable bit is set in the ATUCMD, assert **SERR**# on the PCI interface. When both bits are not set, take no action. When **SERR**# is asserted, additional actions are taken:
  - Set the **SERR**# Asserted bit in the ATUSR
  - When the ATU SERR# Asserted Interrupt Mask bit in the ATUIMR is clear, set the SERR# Asserted bit in the ATUISR. When set, no action
  - When the ATU SERR# Detected interrupt enable is set in the ATUCR, set the SERR# Detected bit in the ATUISR. When clear, no action
- Flush the transaction that was master aborted from the IWQ.

The Internal Bus Master Abort bit is non-maskable and always results in an interrupt being driven to the core processor.

#### 3.7.7.1.2 Inbound Read Request

When operating in the Conventional mode, the following actions with the given constraints are performed by the ATU when a master abort is detected by the internal initiator interface during an inbound read transaction:

- Set the Internal Bus Master Abort bit (bit 7) in the ATUISR
- Return a target abort condition to the initiating master during the delayed completion cycle on the PCI bus. No data is ever read from the internal bus and returned to the PCI bus.

The following additional actions with the given constraints are performed by the ATU when a target abort is signaled by the PCI interface during an inbound delayed read completion cycle:

- Set the Target Abort (target) bit (bit 11) in the ATUSR.
- When the ATU PCI Target Abort (target) Interrupt Mask bit in the ATUIMR is clear, set the PCI Target Abort (target) bit in the ATUISR. When set, no action.
- Flush the transaction that was master aborted from the ITQ after the target abort is delivered on the PCI interface.

When operating in the PCI-X mode, the following actions with the given constraints are performed by the ATU when a master abort is detected by the internal master interface during an inbound split read transaction:

- Set the Internal Bus Master Abort bit (bit 7) in the ATUISR.
- Generate a Split Completion Error Message (with message class=2h completer error and message index=80h 80333 internal bus master abort) on the PCI bus.
- When the Initiated Split Completion Error Message Interrupt Mask in the ATUIMR is clear, set the Initiated Split Completion Error Message bit in the ATUISR. When set, no action.
- Flush the transaction that was master aborted.
- *Note:* This split completion error message includes a device specific message index. The error handler would need to have knowledge of the device specific error messages of the 80333 in order to fully diagnose the problem.

The Internal Bus Master Abort bit is non-maskable and always results in an interrupt being driven to the core processor.



## 3.7.7.2 Target Abort on the Internal Bus

Target Aborts can be seen by the internal bus requester interface during inbound read operations to the memory controller. During inbound read operations, the memory controller is capable of signalling a target abort when a multi-bit, unrecoverable ECC error is encountered. This can occur during any read operation.

Note target aborts are signalled on a Qword basis. When either Dword of a Qword target aborts, both is considered to have target aborted.

The Memory Controller is responsible for creating an interrupt to the Intel XScale<sup>®</sup> core for any multi-bit ECC errors.

#### 3.7.7.2.1 Conventional Mode

When operating in the Conventional PCI mode, when the data word which was target aborted on the internal bus is actually requested and delivered on the PCI Bus, and the ATU ECC Target Abort Enable bit is set in the ATUIMR, a target abort is returned to the PCI initiator on that data word. When the ATU ECC Target Abort Enable bit is clear in the ATUIMR, a disconnect with data is returned to the PCI initiator during the data word that was target aborted on the internal bus. In both cases, the IRQ is flushed after the completion cycle is performed on the PCI bus

The following additional actions with the given constraints are performed by the ATU when a target abort is signaled by the PCI target interface during an inbound read transaction:

- Set the Target Abort (target) bit (bit 11) in the ATUSR.
- When the ATU PCI Target Abort (target) Interrupt Mask bit in the ATUIMR is clear, set the PCI Target Abort (target) bit in the ATUISR. When set, no action.

#### 3.7.7.2.2 PCI-X Mode

When operating in the PCI-X mode, a Target-Abort of an inbound read transaction (split read request) on the Internal Bus results in the following actions.

- The ATU initiates a Split Completion Error Message (with message class=2h completer error and message index=81h 80333 internal bus target abort) on the PCI bus.
- When the Initiated Split Completion Error Message Interrupt Mask in the ATUIMR is clear, set the Initiated Split Completion Error Message bit in the ATUISR. When set, no action.
- *Note:* This split completion error message includes a device specific message index. The error handler would need to have knowledge of the device specific error messages of the 80333 in order to fully diagnose the problem.

## 3.7.8 ATU Error Summary

Table 120 summarizes the ATU error reporting for PCI bus errors and Table 121 summarizes the ATU error reporting for internal bus errors. The tables assume that all error reporting is enabled through the appropriate command registers (unless otherwise noted). The ATU Status Register records PCI bus errors. Note that the **SERR**# Asserted bit in the Status Register is set only when the **SERR**# Enable bit in the Command Register is set. The ATU Interrupt Status Registers record Intel XScale<sup>®</sup> core interrupt status information.

Table 120. ATU Error Reporting Summary - PCI Interface (Sheet 1 of 3)

Error Condition <sup>a</sup> (Bus Mode <sup>b</sup> )	Bits Set in ATU Status Register (ATUSR <sup>c</sup> ) or PCI-X Status Register (PCIXSR <sup>d</sup> )	Bits Set in ATU Interrupt Status Register (ATUISR)	Interrupt Mask Bit in ATUIMR or ATUCR
	PCI Bus Error Response (	i.e., signal Target-Abort, signal	Master-Abort etc.)
Address or Attribute Parity Error (Both)	In Conventional Mode, ignore (Ma PCI-X Mode, claim the transactio signal <b>SERR#</b> upon parity error d	aster-Abort) the transaction, and the n and complete as though no erro letection.	nen signal <b>SERR#</b> . In r had occurred; and,
(Both)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6
(Both)	N/A	SERR# Detected - bit 4	ATUCR bit 9
(Both)	Detected Address or Attribute Parity Error - bit 18 of the PCI Configuration and Status Register	N/A	N/A
(Both)	Detected Parity Error - bit 15	Detected Parity Error - bit 9	ATUIMR bit 7
Outbound Read Request Parity Error (Both)	Signal PERR# and SERR# (PCI-	-X Mode Only).	
(Both)	Master Parity Error - bit 8	Master Parity Error - bit 0	ATUIMR bit 2
(PCI-X)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6
(PCI-X)	N/A	SERR# Detected - bit 4	ATUCR bit 9
(Both)	Detected Parity Error - bit 15	Detected Parity Error - bit 9	ATUIMR bit 7
Outbound Write Request Parity Error (Both)	Signal SERR# (only for PCI-X or	MSI Writes).	
(Both)	Master Parity Error - bit 8	Master Parity Error - bit 0	ATUIMR bit 2
(PCI-X or MSI)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6
(PCI-X or MSI)	N/A	SERR# Detected - bit 4	ATUCR bit 9
Inbound Read Completions Parity Error (PCI-X)	None.		
Inbound Configuration Write Completion Message Parity Error (PCI-X)	None.		
Inbound Read Request Parity Error (Both)	None.		

# intel®

Table 120.	ATU Error Reporting		- PCI Interface	(Sheet 2 of 3)	
	Alo Enor Reporting	Journar			,

Error Condition <sup>a</sup> (Bus Mode <sup>b</sup> )	Bits Set in ATU Status Register (ATUSR <sup>c</sup> ) or PCI-X Status Register (PCIXSR <sup>d</sup> )	Bits Set in ATU Interrupt Status Register (ATUISR)	Interrupt Mask Bit in ATUIMR or ATUCR
	PCI Bus Error Response (	i.e., signal Target-Abort, signal	Master-Abort etc.)
Inbound Write Request Parity Error (Both)	Signal <b>PERR#</b> .		
(Both)	Detected Parity Error - bit 15	Detected Parity Error - bit 9	ATUIMR bit 7
Outbound Split Write Data Parity Error Message (PCI-X)	Signal SERR#.		
(PCI-X)	Master Parity Error - bit 8	Master Parity Error - bit 0	ATUIMR bit 2
(PCI-X)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6
(PCI-X)	N/A	SERR# Detected - bit 4	ATUCR bit 9
(PCI-X)	Received Split Completion Error Message - bit 29	Received Split Completion Error Message - bit 12	ATUIMR bit 9
Inbound Configuration Write Request Parity Error (Both)	Signal <b>PERR#</b> . Initiate a Split Wri (PCI-X Mode Only).	ite Data Parity Error Message add	ressed to the Requester
(PCI-X)	N/A	Initiated Split Completion Error Message - bit 13	ATUIMR bit 10
(Both)	Detected Parity Error - bit 15	Detected Parity Error - bit 9	ATUIMR bit 7
Split Completion Message Parity Error (PCI-X)	Signal <b>PERR#</b> and <b>SERR#</b> .		
(PCI-X)	Master Parity Error - bit 8	Master Parity Error - bit 0	ATUIMR bit 2
(PCI-X)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6
(PCI-X)	N/A	SERR# Detected - bit 4	ATUCR bit 9
(PCI-X and SCE <sup>e</sup> )	Received Split Completion Error Message - bit 29	Received Split Completion Error Message - bit 12	ATUIMR bit 9
(PCI-X)	Detected Parity Error - bit 15	Detected Parity Error - bit 9	ATUIMR bit 7
Outbound Read Request Master-Abort (Both)	None.		
(Both)	Master Abort - bit 13	PCI Master Abort - bit 3	ATUIMR bit 5
(PCI-X and SCE)	Received Split Completion Error Message - bit 29	Received Split Completion Error Message - bit 12	ATUIMR bit 9
Outbound Write Request Master-Abort (Both)	None.		
(Both)	Master Abort - bit 13	PCI Master Abort - bit 3	ATUIMR bit 5
(MSI)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6
(MSI)	N/A	SERR# Detected - bit 4	ATUCR bit 9
(PCI-X and SCE)	Received Split Completion Error Message - bit 29	Received Split Completion Error Message - bit 12	ATUIMR bit 9



#### Table 120. ATU Error Reporting Summary - PCI Interface (Sheet 3 of 3)

Error Condition <sup>a</sup> (Bus Mode <sup>b</sup> )	Bits Set in ATU Status Register (ATUSR <sup>c</sup> ) or PCI-X Status Register (PCIXSR <sup>d</sup> )	Bits Set in ATU Interrupt Status Register (ATUISR)	Interrupt Mask Bit in ATUIMR or ATUCR	
	PCI Bus Error Response	i.e., signal Target-Abort, signal	Master-Abort etc.)	
Inbound Read Completions Master-Abort (PCI-X)	None.			
Inbound Configuration Write Completion Message Master-Abort (PCI-X)	None.			
Outbound Read Request Target-Abort (Both)	None.			
(Both)	Target Abort (master) - bit 12	PCI Target Abort (master) - bit 2	ATUIMR bit 4	
(PCI-X and SCE)	Received Split Completion Error Message - bit 29	Received Split Completion Error Message - bit 12	ATUIMR bit 9	
Outbound Write Request Target-Abort (Both)	None.			
(Both)	Target Abort (master) - bit 12	PCI Target Abort (master) - bit 2	ATUIMR bit 4	
(MSI)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6	
(MSI)	N/A	SERR# Detected - bit 4	ATUCR bit 9	
(PCI-X and SCE)	Received Split Completion Error Message - bit 29	ATUIMR bit 9		
Inbound EROM Write Request Target-Abort (Both)	Signal Target-Abort.			
(Both)	Target Abort (target) - bit 11	PCI Target Abort (target) - bit 1	ATUIMR bit 3	
Unexpected Split Completion (PCI-X)	In the PCI-X mode, the transaction completes normally according to the invalid lower address field or invalid byte count.			
(PCI-X)	Unexpected Split Completion - bit 19	N/A	N/A	

a. All parity errors refer to data parity errors except where otherwise noted.

b. Codes for bus mode in which this error response applies: PCI-X means PCI-X Mode Only, Conventional means Conventional PCI Mode Only, and Both means that the error response applies both in the Conventional and PCI-X mode of operation. MSI stands for Message-Signaled Interrupts and refers to an Outbound Write transaction that is actually an MSI write transaction.

d. Table assumes that Data Parity Recovery Enable - bit 0 of the PCIXCMD is clear.

c. Table assumes that Parity Error Response - bit 6 of the ATUCMD register is set.

e. When the SCE bit (bit 30 of the Completer Attributes) and the SCM bit (bit 29 of the Completer Attributes) are set during the Attribute phase of a Split Completion Transaction, the transaction is a Split Completion Message that is an Error Message. In this case, the Received Split Completion Error Message - bit 29 of the PCIXSR is set.

# intel

### Table 121. ATU Error Reporting Summary - Internal Bus Interface

Error Condition <sup>a</sup> (Bus Mode <sup>b</sup> )	Bits Set in ATU Status Register (ATUSR <sup>c</sup> )	Bits Set in ATU Interrupt Status Register (ATUISR)	Interrupt Mask Bit in ATUIMR or ATUCR		
	PCI Bus Error Response	(i.e., signal Target-Abort, signal	Master-Abort etc.)		
Inbound Write Request Master-Abort (Both)	Assert <b>SERR#</b> .				
(Both)	N/A	Internal Bus Master Abort - bit 7	N/A		
(Both)	SERR# Asserted - bit 14	SERR# Asserted - bit 10	ATUIMR bit 6		
(Both)	N/A	SERR# Detected - bit 4	ATUCR bit 9		
Inbound Read Request Master-Abort (Both)	In the Conventional Mode signal Split Completion Error Message	Target-Abort. In the PCI-X Mode s to the Requester.	end a device specific		
(Both)	N/A	Internal Bus Master Abort - bit 7	N/A		
(Conventional)	Target Abort (target) - bit 11	PCI Target Abort (target) - bit 1	ATUIMR bit 3		
(PCI-X)	N/A	Initiated Split Completion Error Message - bit 13	ATUIMR bit 10		
Inbound Read Request Target-Abort (Both)	In the Conventional Mode signal Split Completion Error Message	Target-Abort. In the PCI-X Mode s to the Requester.	end a device specific		
(Conventional)	Target Abort (target) - bit 11	PCI Target Abort (target) - bit 1 ATUIMR bit 3			
(PCI-X)	N/A	Initiated Split Completion Error Message - bit 13 ATUIMR bit 10			

a. There are no Inbound Write Request Target-Abort Error Conditions.

b. Codes for bus mode in which this error response applies: PCI-X means PCI-X Mode Only, Conventional means Conventional PCI Mode Only, and Both means that the error response applies both in the Conventional and PCI-X mode of operation. MSI stands for Message-Signaled Interrupts and refers to an Outbound Write transaction that is actually an MSI write transaction.

c. Table assumes that the ATU Inbound SERR# Enable bit (bit 1 of the ATUIMR), the ATU ECC Target Abort Enable (bit 0 of the ATUIMR), and the SERR# Enable bit (bit 8 of the ATUCMD) are set.



# 3.8 Message-Signaled Interrupts

The Messaging Unit is responsible for the generation of all of the Outbound Interrupts from the 80333. These interrupts can be delivered to the Host Processor via the **IRQ14**# IOAPIC pin or the Message Signaled Interrupt (MSI) mechanism.

When a host processor enables Message-Signaled Interrupts (MSI) on the 80333, an outbound interrupt is signaled to the host via a PCI write instead of the assertion of the **IRQ14**# IOAPIC pin.

In support of MSI, the 80333 implements the MSI capability structure. The capability structure includes the Section 4.9.20, "MSI Capability Identifier Register - MSI\_CAPID" on page 355, the Section 4.9.21, "MSI Next Item Pointer Register - MSI\_NXTP" on page 356, the Section 4.9.23, "MSI Message Address Register - MSI\_MAR" on page 358, the Section 4.9.24, "MSI Message Upper Address Register - MSI\_MUAR" on page 359 and the Section 4.9.25, "MSI Message Data Register- MSI\_MDR" on page 360.

The Message Unit generates MSIs by writing to the MSI port via the internal bus. The ATU generates a write transaction whenever the Message Unit writes to the MSI port, using the address specified in the Section 4.9.23, "MSI Message Address Register - MSI\_MAR" on page 358, the Section 4.9.24, "MSI Message Upper Address Register - MSI\_MUAR" on page 359 and the Section 4.9.25, "MSI Message Data Register - MSI\_MDR" on page 360.

# 3.9 Vital Product Data

Vital Product Data (VPD) provides detailed information to the system regarding the hardware, software and microcode elements of a device. This information may include Part Number, Serial Number or other detailed information. This information resides on a non-volatile storage device (i.e., Flash Memory) attached to the 80333. In addition VPD also provides a mechanism for storing information such as performance or failure data on the device being monitored.

Support of VPD involves the implementation of the VPD Extended Capabilities List Item in the Primary ATU. The VPD Extended capabilities header consists of five registers, the "VPD Capability Identifier Register - VPD\_CAPID" on page 298, the "VPD Next Item Pointer Register - VPD\_NXTP" on page 299, the "VPD Address Register - VPD\_AR" on page 300, and the "VPD Data Register - VPD\_DR" on page 301.

Scheduled by Intel XScale<sup>®</sup> core interrupts, the 80333 may be used to retrieve or store VPD information through the VPD extended capabilities list item.

Please consult Appendix I of the *PCI Local Bus Specification*, Revision 2.3 for the definitions of compliant VPD format.

## 3.9.1 Configuring Vital Product Data Operation

By default, the 80333 VPD functionality is not configured for operation. Specifically, the VPD Extended Capabilities List Item is not discovered during a PCI bus scan and the ATUs VPD interrupt status bit in the "ATU Interrupt Status Register - ATUISR" on page 287 is masked by the "ATU Interrupt Mask Register - ATUIMR" on page 289. The following steps should be followed to properly configure the 80333 support for VPD:

- 1. The 80333 must be strapped to Retry Type 0 Configuration cycles following the deassertion of **PWRGD**. Enabling this configuration cycle retry mechanism guarantees that the Intel XScale<sup>®</sup> core can make the VPD Extended Capabilities List Item visible before the system configures the 80333. The configuration retry mechanism is controlled through bit 2 of the Table 168, "PCI Configuration and Status Register PCSR" on page 283.
- 2. When the configuration retry mechanism is strapped enabled as described in step 1, typically, the 80333 would also be strapped such that the Intel XScale<sup>®</sup> core would immediately boot following the deassertion of **PWRGD** (bit 1 of the PCSR), though this is not required.
- 3. The Intel XScale<sup>®</sup> core writes E8H to the "PCI-X Next Item Pointer Register PX\_NXTP" on page 308. This links the PCI-X Capabilities List Item to the VPD Capabilities List Item.
- 4. The Intel XScale<sup>®</sup> core clears bit 12 of the ATUIMR to enable the ATUs VPD interrupt status bit.



## 3.9.2 Accessing Vital Product Data

The VPD Capabilities List Item provides three fields which the system uses to access the Vital Product Data:

- VPD AddressDWORD Aligned Byte address of the VPD to be accessed which is represented by VPDAR[14:0]. Note that this means that the maximum size of the VPD is 128 Kbytes. The user may pick any 128 Kbyte block of memory in the storage component for the VPD.
- Flag The flag register is used to indicate when the transfer between the VPD Data Register and the storage component is completed. The flag is in VPDAR[15] which means that the Flag is written at the same time that VPD address is written.
- VPD Data Four bytes of VPD Data can be read or written through this field which is represented by VPDDR[31:0]. The least significant byte of this register represents the byte at the VPD Address (VPDAR[14:0]). Four bytes are always transferred between this register and the VPD storage component.
- *Note:* Bit 19 of the VPD Data Register has a bit type of RC (read clear). When the VPD feature is used, care must be taken to mask bit 19. See *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #11.

## 3.9.2.1 Reading Vital Product Data

Using the fields defined in the VPD Capabilities List Item, the 80333 reads Vital Product Data using the following sequence of events:

- 1. Host processor executes a configuration write of the VPD address to the VPDAR with the Flag cleared.
- 2. An interrupt to the Intel XScale<sup>®</sup> core is triggered and bit 17 of the ATUISR is set. Meanwhile, the host processor polls the VPDAR register waiting for the Flag to be set.
- *Warning:* When any configuration writes to either the VPDAR or the VPDDR occur prior to the Flag being set, the results of the original read operation are unpredictable.
  - 3. Using the VPD Address, the Intel XScale<sup>®</sup> core retrieves the Vital Product Data from the VPD storage component (i.e., Flash Memory).
  - 4. The Intel XScale<sup>®</sup> core then writes this data to VPD Data Register (VPDDR).
  - 5. The Intel XScale<sup>®</sup> core clears the VPD interrupt status bit in the ATUISR.
  - 6. The Intel XScale<sup>®</sup> core then sets the Flag in the VPDAR register.
  - 7. When the host processor detects that the Flag has been set, the host processor then reads the retrieved VPD from the VPDDR.



## 3.9.2.2 Writing Vital Product Data

Using the fields defined in the VPD Capabilities List Item, the 80333writes Vital Product Data using the following sequence of events:

- 1. Host processor executes a configuration write of the VPD data to be written to the VPDDR.
- 2. Host processor executes a configuration write of the VPD address to the VPDAR with the Flag set.
- 3. An interrupt to the Intel XScale<sup>®</sup> core is triggered and bit 17 of the ATUISR is set. Meanwhile, the host processor polls the VPDAR register waiting for the Flag to be cleared.
- *Warning:* When any configuration writes to either the VPDAR or the VPDDR occur prior to the Flag being cleared, the results of the original write operation are unpredictable.
  - 4. Using the VPD Address, the Intel XScale<sup>®</sup> core writes the Vital Product Data from the VPDDR to the VPD storage component (i.e., Flash Memory).
  - 5. The Intel XScale<sup>®</sup> core clears the VPD interrupt status bit in the ATUISR.
  - 6. The Intel XScale<sup>®</sup> core then clears the Flag in the VPDAR register.
  - 7. When the host processor detects that the Flag has been cleared, the host processor has been informed that the VPD write operation is complete.

# 3.10 Register Definitions

Every PCI device implements its own separate configuration address space and configuration registers. The *PCI Local Bus Specification*, Revision 2.3 requires that configuration space be 256 bytes, and the first 64 bytes must adhere to a predefined header format.

Figure 18 defines the header format. Table 123 shows the PCI configuration registers, listed by internal bus address offset. Table 123 shows the entire ATU configuration space (including header and extended registers) and the corresponding section that describes each register. Note that all configuration read and write transactions is accepted on the internal bus as 32-bit transactions. Refer to Chapter 19, "Peripheral Registers".

#### Figure 18. ATU Interface Configuration Header Format

ATU De	evice ID	Vendor ID		00
Sta	atus	Com	mand	04
	ATU Class Code		Revision ID	08
ATUBISTR	Header Type	Latency Timer	Cacheline Size	00
	Inbound ATU E	Base Address 0		10
	Inbound ATU Upp	er Base Address 0		1
	Inbound ATU E	Base Address 1		1
Inbound ATU Upper Base Address 1				
	Inbound ATU E	Base Address 2		2
	Inbound ATU Upp	er Base Address 2		2
	Rese	erved		2
ATU Subs	system ID	ATU Subsyste	em Vendor ID	2
Expansion ROM Base Address				3
	Reserved		Capabilities Pointer	3
Reserved				3
Maximum Latency	Minimum Grant	Interrupt Pin	Interrupt Line	3

The ATU is programmed via a Type 0 configuration command on the PCI interface. See Section 3.2.1.4, "Inbound Configuration Cycle Translation" on page 173. ATU configuration space is function number zero of the 80333 single-function PCI device.

Beyond the required 64 byte header format, ATU configuration space implements extended register space in support of the units functionality. Refer to the *PCI Local Bus Specification*, Revision 2.3 for details on accessing and programming configuration register space.

The ATU unit includes two 8 byte and one 16 byte extended capability configuration spaces beginning at configuration offset C0H, D0H and E0H. The extended configuration spaces can be accessed by a device on the PCI interface through a mechanism defined in the *PCI Local Bus Specification*, Revision 2.3.

In the ATU Status Register (Section 3.10.4) the appropriate bit is set indicating that the Extended Capability Configuration space is supported. When this bit is read, the device can then read the Capabilities Pointer register (Section 3.10.20) to determine the configuration offset of the Extended Capabilities Configuration Header. The format of these headers are depicted in Figure 19, Figure 20, Figure 21 and Figure 22.

#### Figure 19. ATU Interface Extended Configuration Header Format (Power Management)

Power Management Capabilities	Next Item Pointer	Capability Identifier	C0H
Reserved	Power Managemer	nt Control/Status	C4H

The first byte at the Extended Configuration Offset COH is the ATU Capability Identifier Register (Section). This identifies this Extended Configuration Header space as the type defined by the *PCI* Bus Power Management Interface Specification, Revision 1.1.

Following the Capability Identifier Register is the single byte Next Item Pointer Register (Section 3.10.58) which indicates the configuration offset of an additional Extended Capabilities Header, when supported. In the ATU, the Next Item Pointer Register is set to DOH indicating that there is an additional Extended Capabilities Headers supported in the ATUs configuration space.

To enable the *PCI Bus Power Management Interface Specification*, Revision 1.1 compliance support, the Power State Transition interrupt mask in bit 8 of the ATUIMR needs to be cleared. It is the configuration software's responsibility to properly enable and initialize the ATUs Power Management Interface before the Configuration Cycle Retry Bit in the Table 168, "PCI Configuration and Status Register - PCSR" on page 283 is cleared in order for the ATU to be *Advanced Configuration and Power Interface Specification*, Revision 2.0 compliant.

#### Figure 20. ATU Interface Extended Configuration Header Format (MSI Capability)

MSI Message Control	MSI Next Item Point	ter MSI Capability ID	D0H
MSI Message Address			D4H
MSI Message Upper Address			D8H
Reserved MSI Message Data			DCH

The first byte at the Extended Configuration Offset D0H is the "MSI Capability Identifier Register - MSI\_CAPID" on page 355. This identifies this Extended Configuration Header space as the type defined by the *PCI Local Bus Specification*, Revision 2.3.

Following the Capability Identifier Register is the single byte "MSI Next Item Pointer Register - MSI\_NXTP" on page 356, which indicates the configuration offset of an additional Extended Capabilities Header, when supported. In the ATU, the Next Item Pointer Register is set to EOH indicating that there is an additional Extended Capabilities Headers supported in the ATUs configuration space.

#### Figure 21. ATU Interface Extended Configuration Header Format (PCI-X Capability)

PCI-X Command		Next Item Pointer	PCI-X Capability ID	E0H
PCI-X Status				



The first byte at the Extended Configuration Offset E0H is the PCI-X Capability Identifier Register (Section 3.10.62). This identifies this Extended Configuration Header space as the type defined by the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a.

Following the Capability Identifier Register is the single byte Next Item Pointer Register (Section 3.10.63) which indicates the configuration offset of an additional Extended Capabilities Header, when supported. In the ATU, the Next Item Pointer Register is set to 00H by default to indicate there are no additional Extended Capabilities Headers in the ATUs configuration space. Software can set this pointer to B8H indicating that there is an additional Extended Capabilities Headers supported in the ATUs configuration space.

#### Figure 22. ATU Interface Extended Configuration Header Format (VPD Capability)

VPD Address	Next Item Pointer	VPD Capability ID	B8H
VF	D Data		BCH

The first byte at the Extended Configuration Offset B8H is the VPD Capability Identifier Register (Section 3.10.53). This identifies this Extended Configuration Header space as the type defined by the *PCI Local Bus Specification*, Revision 2.3.

Following the Capability Identifier Register is the single byte Next Item Pointer Register (Section 3.10.54) which indicates the configuration offset of an additional Extended Capabilities Header, when supported. In the ATU, the Next Item Pointer Register is set to 00H indicating that there are no additional Extended Capabilities Headers supported in the ATUs configuration space.

The following sections describe the ATU and Expansion ROM configuration registers. Configuration space consists of 8, 16, 24, and 32-bit registers arranged in a predefined format. Each register is described in functionality, access type (read/write, read/clear, read only) and reset default condition.

See Section 1.4, "Terminology and Conventions" on page 47 for a description of *reserved*, *read* only, and *read/clear*. All registers adhere to the definitions found in the *PCI Local Bus* Specification, Revision 2.3 unless otherwise noted.

The PCI register number for each register is given in Table 123. As stated, a Type 0 configuration command on the bus with an active **IDSEL** or a memory-mapped internal bus access is required to read or write these registers.

*Note:* Each configuration register's access type is individually defined for PCI configuration accesses. Some PCI read-only configuration registers have read/write capability from the 80333 core CPU. See also Chapter 19, "Peripheral Registers".



## Table 122. Address Translation Unit Registers (Sheet 1 of 3)

Register Name	Bits	PCI Configu- ration Cycle Register Number	Internal Bus Address
ATU Vendor ID Register - ATUVID	16	0	FFFF.E100H
ATU Device ID Register - ATUDID	16	0	FFFF.E102H
ATU Command Register - ATUCMD	16	1	FFFF.E104H
ATU Status Register - ATUSR	16	1	FFFF.E106H
ATU Revision ID Register - ATURID	8	2	FFFF.E108H
ATU Class Code Register - ATUCCR	24	2	FFFF.E109H
ATU Cacheline Size Register - ATUCLSR	8	3	FFFF.E10CH
ATU Latency Timer Register - ATULT	8	3	FFFF.E10DH
ATU Header Type Register - ATUHTR	8	3	FFFF.E10EH
ATU BIST Register - ATUBISTR	8	3	FFFF.E10FH
Inbound ATU Base Address Register 0 - IABAR0	32	4	FFFF.E110H
Inbound ATU Upper Base Address Register 0 - IAUBAR0	32	5	FFFF.E114H
Inbound ATU Base Address Register 1 - IABAR1	32	6	FFFF.E118H
Inbound ATU Upper Base Address Register 1 - IAUBAR1	32	7	FFFF.E11CH
Inbound ATU Base Address Register 2 - IABAR2	32	8	FFFF.E120H
Inbound ATU Upper Base Address Register 2 - IAUBAR2	32	9	FFFF.E124H
Reserved	32	10	FFFF.E128H
ATU Subsystem Vendor ID Register - ASVIR	16	11	FFFF.E12CH
ATU Subsystem ID Register - ASIR	16	11	FFFF.E12EH
Expansion ROM Base Address Register -ERBAR	32	12	FFFF.E130H
ATU Capabilities Pointer Register - ATU_Cap_Ptr	8	13	FFFF.E134H
Reserved	8	13	FFFF.E135H
Reserved	16	13	FFFF.E136H
Reserved	32	14	FFFF.E138H
ATU Interrupt Line Register - ATUILR	8	15	FFFF.E13CH
ATU Interrupt Pin Register - ATUIPR	8	15	FFFF.E13DH
ATU Minimum Grant Register - ATUMGNT	8	15	FFFF.E13EH
ATU Maximum Latency Register - ATUMLAT	8	15	FFFF.E13FH
Inbound ATU Limit Register 0 - IALR0	32	16	FFFF.E140H
Inbound ATU Translate Value Register 0 - IATVR0	32	17	FFFF.E144H
Expansion ROM Limit Register - ERLR	32	18	FFFF.E148H
Expansion ROM Translate Value Register - ERTVR	32	19	FFFF.E14CH
Inbound ATU Limit Register 1 - IALR1	32	20	FFFF.E150H
Inbound ATU Limit Register 2 - IALR2	32	21	FFFF.E154H
Inbound ATU Translate Value Register 2 - IATVR2	32	22	FFFF.E158H
Outbound I/O Window Translate Value Register - OIOWTVR	32	23	FFFF.E15CH



## Table 122. Address Translation Unit Registers (Sheet 2 of 3)

Register Name	Bits	PCI Configu- ration Cycle Register Number	Internal Bus Address
Outbound Memory Window Translate Value Register 0- OMWTVR0	32	24	FFFF.E160H
Outbound Upper 32-bit Memory Window Translate Value Register 0- OUMWTVR0	32	25	FFFF.E164H
Outbound Memory Window Translate Value Register 1- OMWTVR1	32	26	FFFF.E168H
Outbound Upper 32-bit Memory Window Translate Value Register 1- OUMWTVR1	32	27	FFFF.E16CH
Reserved	32	28	FFFF.E170H
Reserved	32	29	FFFF.E174H
Outbound Upper 32-bit Direct Window Translate Value Register - OUDWTVR	32	30	FFFF.E178H
PCI Express-to-PCI Bridge Secondary A-Segment Bus Number Register - PEBSABNR	8	31	FFFF.E17CH
PCI Express-to-PCI Bridge Secondary B-Segment Bus Number Register - PEBSBBNR	8	31	FFFF.E17DH
PCI Express-to-PCI Bridge Primary Bus Number Register - PEBPBNR	8	31	FFFF.E17EH
PCI Express-to-PCI Bridge Device Number Register - PEBDNUM	8	31	FFFF.E17FH
ATU Configuration Register - ATUCR	32	32	FFFF.E180H
PCI Configuration and Status Register - PCSR	32	33	FFFF.E184H
ATU Interrupt Status Register - ATUISR	32	34	FFFF.E188H
ATU Interrupt Mask Register - ATUIMR	32	35	FFFF.E18CH
Inbound ATU Base Address Register 3 - IABAR3	32	36	FFFF.E190H
Inbound ATU Upper Base Address Register 3 - IAUBAR3	32	37	FFFF.E194H
Inbound ATU Limit Register 3 - IALR3	32	38	FFFF.E198H
Inbound ATU Translate Value Register 3 - IATVR3	32	39	FFFF.E19CH
Reserved	32	40	FFFF.E1A0H
Outbound Configuration Cycle Address Register - OCCAR	32	41	FFFF.E1A4H
Reserved	32	42	FFFF.E1A8H
Outbound Configuration Cycle Data Register - OCCDR	32	Not Available in PCI Configu- ration Space	FFFF.E1ACH
Reserved	32	44	FFFF.E1B0H
Reserved	32	45	FFFF.E1B4H
VPD Capability Identifier Register - VPD_CAPID	8	46	FFFF.E1B8H
VPD Next Item Pointer Register - VPD_NXTP	8	46	FFFF.E1B9H
VPD Address Register - VPD_AR	16	47	FFFF.E1BAH
VPD Data Register - VPD_DR	32	47	FFFF.E1BCH
Power Management Capability Identifier Register - PM_CAPID	8	48	FFFF.E1C0H
Power Management Next Item Pointer Register - PM_NXTP	8	48	FFFF.E1C1H
Power Management Capabilities Register - PM_CAP	16	48	FFFF.E1C2H



## Table 122. Address Translation Unit Registers (Sheet 3 of 3)

Register Name	Bits	PCI Configu- ration Cycle Register Number	Internal Bus Address
Power Management Control/Status Register - PM_CSR	16	49	FFFF.E1C4H
Reserved	16	49	FFFF.E1C6H
Reserved	32	50	FFFF.E1C8H
Reserved	32	51	FFFF.E1CCH
MSI Capability Identifier Register - MSI_CAPID	8	52	FFFF.E1D0H
MSI Next Item Pointer Register - MSI_NXTP	8	52	FFFF.E1D1H
MSI Message Control Register - MSI_MCR	16	52	FFFF.E1D2H
MSI Message Address Register - MSI_MAR	32	53	FFFF.E1D4H
MSI Message Upper Address Register - MSI_MUAR	32	54	FFFF.E1D8H
MSI Message Data Register- MSI_MDR	16	55	FFFF.E1DCH
Reserved	16	55	FFFF.E1DEH
PCI-X_Capability Identifier Register - PX_CAPID	8	56	FFFF.E1E0H
PCI-X Next Item Pointer Register - PX_NXTP	8	56	FFFF.E1E1H
PCI-X Command Register - PX_CMD	16	56	FFFF.E1E2H
PCI-X Status Register - PX_SR	32	57	FFFF.E1E4H
"PCI Interrupt Routing Select Register - PIRSR" on page 801	32	59	FFFF.E1ECH
PCI-A Drive Strength Control Register - PADSCR	32	n/a	FFFF.F5C0H
PCI-A Drive Strength Value Register - PADSVR	32	n/a	FFFF.F5C8H
PCI-B Drive Strength Control Register - PBDSCR	32	n/a	FFFF.F5D0H
PCI-B Drive Strength Value Register - PBDSVR	32	n/a	FFFF.F5D8H



## Table 123. ATU PCI Configuration Register Space (Sheet 1 of 2)

Address Offset	ATU PCI Configuration Register Section, Name, Page		
00H	Section 3.10.1, "ATU Vendor ID Register - ATUVID" on page 238		
02H	Section 3.10.2, "ATU Device ID Register - ATUDID" on page 239		
04H	Section 3.10.3, "ATU Command Register - ATUCMD" on page 240		
06H	Section 3.10.4, "ATU Status Register - ATUSR" on page 241		
08H	Section 3.10.5, "ATU Revision ID Register - ATURID" on page 243		
09H	Section 3.10.6, "ATU Class Code Register - ATUCCR" on page 244		
0CH	Section 3.10.7, "ATU Cacheline Size Register - ATUCLSR" on page 245		
0DH	Section 3.10.8, "ATU Latency Timer Register - ATULT" on page 246		
0EH	Section 3.10.9, "ATU Header Type Register - ATUHTR" on page 247		
0FH	Section 3.10.10, "ATU BIST Register - ATUBISTR" on page 248		
10H	Section 3.10.11, "Inbound ATU Base Address Register 0 - IABAR0" on page 249		
14H	Section 3.10.12, "Inbound ATU Upper Base Address Register 0 - IAUBAR0" on page 250		
18H	Section 3.10.13, "Inbound ATU Base Address Register 1 - IABAR1" on page 251		
1CH	Section 3.10.14, "Inbound ATU Upper Base Address Register 1 - IAUBAR1" on page 252		
20H	Section 3.10.15, "Inbound ATU Base Address Register 2 - IABAR2" on page 253		
24H	Section 3.10.16, "Inbound ATU Upper Base Address Register 2 - IAUBAR2" on page 254		
2CH	Section 3.10.17, "ATU Subsystem Vendor ID Register - ASVIR" on page 255		
2EH	Section 3.10.18, "ATU Subsystem ID Register - ASIR" on page 256		
30H	Section 3.10.19, "Expansion ROM Base Address Register - ERBAR" on page 257		
34H	Section 3.10.20, "ATU Capabilities Pointer Register - ATU_Cap_Ptr" on page 258		
3CH	Section 3.10.22, "ATU Interrupt Line Register - ATUILR" on page 261		
3DH	Section 3.10.23, "ATU Interrupt Pin Register - ATUIPR" on page 262		
3EH	Section 3.10.24, "ATU Minimum Grant Register - ATUMGNT" on page 263		
3FH	Section 3.10.25, "ATU Maximum Latency Register - ATUMLAT" on page 264		
40H	Section 3.10.26, "Inbound ATU Limit Register 0 - IALR0" on page 265		
44H	Section 3.10.27, "Inbound ATU Translate Value Register 0 - IATVR0" on page 266		
48H	Section 3.10.28, "Expansion ROM Limit Register - ERLR" on page 267		
4CH	Section 3.10.29, "Expansion ROM Translate Value Register - ERTVR" on page 268		
50H	Section 3.10.30, "Inbound ATU Limit Register 1 - IALR1" on page 269		
54H	Section 3.10.31, "Inbound ATU Limit Register 2 - IALR2" on page 270		
58H	Section 3.10.32, "Inbound ATU Translate Value Register 2 - IATVR2" on page 271		
5CH	Section 3.10.33, "Outbound I/O Window Translate Value Register - OIOWTVR" on page 272		
60H	Section 3.10.34, "Outbound Memory Window Translate Value Register 0 - OMWTVR0" on page 273		
64H	Section 3.10.35, "Outbound Upper 32-bit Memory Window Translate Value Register 0 - OUMWTVR0" on page 274		
68H	Section 3.10.36, "Outbound Memory Window Translate Value Register 1 - OMWTVR1" on page 275		
6CH	Section 3.10.37, "Outbound Upper 32-bit Memory Window Translate Value Register 1 - OUMWTVR1" on page 276		
7CH	Section 3.10.39, "PCI Express-to-PCI Bridge Secondary A-Segment Bus Number Register - PEBSABNR" on page 278		
7DH	Section 3.10.40, "PCI Express-to-PCI Bridge Secondary B-Segment Bus Number Register - PEBSBBNR" on page 279		
7EH	Section 3.10.41, "PCI Express-to-PCI Bridge Primary Bus Number Register - PEBPBNR" on page 280		

# intel

## Table 123. ATU PCI Configuration Register Space (Sheet 2 of 2)

Address Offset	ATU PCI Configuration Register Section, Name, Page		
7FH	Section 3.10.42, "PCI Express-to-PCI Bridge Device Number Register - PEBDNUM" on page 281		
80H	Section 3.10.43, "ATU Configuration Register - ATUCR" on page 282		
84H	Section 3.10.44, "PCI Configuration and Status Register - PCSR" on page 283		
88H	Section 3.10.45, "ATU Interrupt Status Register - ATUISR" on page 287		
8CH	Section 3.10.46, "ATU Interrupt Mask Register - ATUIMR" on page 289		
90H	Section 3.10.47, "Inbound ATU Base Address Register 3 - IABAR3" on page 292		
94H	Section 3.10.48, "Inbound ATU Upper Base Address Register 3 - IAUBAR3" on page 293		
98H	Section 3.10.49, "Inbound ATU Limit Register 3 - IALR3" on page 294		
9CH	Section 3.10.50, "Inbound ATU Translate Value Register 3 - IATVR3" on page 295		
A4H	Section 3.10.51, "Outbound Configuration Cycle Address Register - OCCAR" on page 296		
ACH	Section 3.10.52, "Outbound Configuration Cycle Data Register - OCCDR" on page 297		
B8H	Section 3.10.53, "VPD Capability Identifier Register - VPD_CAPID" on page 298		
B9H	Section 3.10.54, "VPD Next Item Pointer Register - VPD_NXTP" on page 299		
BAH	Section 3.10.55, "VPD Address Register - VPD_AR" on page 300		
BCH	Section 3.10.56, "VPD Data Register - VPD_DR" on page 301		
C0H	Section 3.10.57, "Power Management Capability Identifier Register - PM_CAPID" on page 302		
C1H	Section 3.10.58, "Power Management Next Item Pointer Register - PM_NXTP" on page 303		
C2H	Section 3.10.59, "Power Management Capabilities Register - PM_CAP" on page 304		
C4H	Section 3.10.60, "Power Management Control/Status Register - PM_CSR" on page 305		
D0H	Section 4.9.20, "MSI Capability Identifier Register - MSI_CAPID" on page 355		
D1H	Section 4.9.21, "MSI Next Item Pointer Register - MSI_NXTP" on page 356		
D2H	Section 4.9.22, "MSI Message Control Register - MSI_MCR" on page 357		
D4H	Section 4.9.23, "MSI Message Address Register - MSI_MAR" on page 358		
D8H	Section 4.9.24, "MSI Message Upper Address Register - MSI_MUAR" on page 359		
DCH	Section 4.9.25, "MSI Message Data Register- MSI_MDR" on page 360		
E0H	Section 3.10.62, "PCI-X Capability Identifier Register - PX_CAPID" on page 307		
E1H	Section 3.10.63, "PCI-X Next Item Pointer Register - PX_NXTP" on page 308		
E2H	Section 3.10.64, "PCI-X Command Register - PX_CMD" on page 309		
E4H	Section 3.10.65, "PCI-X Status Register - PX_SR" on page 310		
ECH	Section 17.8.17, "PCI Interrupt Routing Select Register - PIRSR" on page 801		
n/a	Section 3.10.67, "PCI-A Drive Strength Control Register - PADSCR" on page 313		
n/a	Section 3.10.68, "PCI-A Drive Strength Value Register - PADSVR" on page 314		
n/a	Section 3.10.69, "PCI-B Drive Strength Control Register - PBDSCR" on page 315		
n/a	Section 3.10.70, "PCI-B Drive Strength Value Register - PBDSVR" on page 316		



# 3.10.1 ATU Vendor ID Register - ATUVID

ATU Vendor ID Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3.







# 3.10.2 ATU Device ID Register - ATUDID

ATU Device ID Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3.

### Table 125. ATU Device ID Register - ATUDID

		IOP       15       12       8       4       0         Attributes       rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/r
Internal FFFF.E	Bus Addre 102H	ss PCI Configuration Address Offset 02H - 03H Attribute Legend: RW = Read/Write RV = Reserved RC = Read Clear PR = Preserved RO = Read Only RS = Read/Set NA = Not Accessible
Bit	Default	Description
15:00	0374H	ATU Device ID - This is a 16-bit value assigned to the ATU. This ID, combined with the VID, uniquely identify any PCI device.



# 3.10.3 ATU Command Register - ATUCMD

ATU Command Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3 and in most cases, affect the behavior of the PCI ATU and devices on the PCI bus.



E.

		IOP     15     12     8       Attributes     rv     rv     rv     rv     rv     rv     rv       PCI     rv     rv     rv     rv     rv     rv     rv     rv		
FFFF.E	al Bus Address 104H	PCI Configuration Address Offset 04H - 05H	Attribute Legend: RV = Reserved PR = Preserved RS = Read/Set	RW = Read/Write RC = Read Clear RO = Read Only NA = Not Accessible
Bit	Default	Description		
15:11	000000 <sub>2</sub>	Reserved		
10	0	Interrupt Disable - This bit disables 80333 from asser 0 = enables the assertion of interrupt signal. 1 = disables the assertion of its interrupt signal. Refer to Section 3.10.23, "ATU Interrupt Pin Register interrupt signal.	ing the ATU interrupt sign - ATUIPR" on page 262 fo	nal. or details on the ATU
09	02	Fast Back to Back Enable - When cleared, the ATU interface is not allowed to generate fast back-to-back cycles on its bus. Ignored when operating in the PCI-X mode.		
08	02	SERR# Enable - When cleared, the ATU interface is not allowed to assert SERR# on the PCI interface.		
07	1 <sub>2</sub>	Address/Data Stepping Control - Address stepping is ATU inserts 2 clock cycles of address stepping for Co stepping for PCI-X mode.	implemented for configur nventional Mode and 4 cl	ation transactions. The ock cycles of address
06	02	Parity Error Response - When set, the ATU takes nor cleared, parity checking is disabled.	mal action when a parity of	error is detected. When
05	02	VGA Palette Snoop Enable - The ATU interface does perform VGA palette snooping.	not support I/O writes an	d therefore, does not
04	02	Memory Write and Invalidate Enable - When set, ATU use Memory Write commands instead of MWI. Ignore	may generate MWI com d when operating in the F	mands. When clear, ATU PCI-X mode.
03	02	Special Cycle Enable - The ATU interface does not re implemented and a reserved bit field.	spond to special cycle co	mmands in any way. Not
02	02	Bus Master Enable - The ATU interface can act as a master on the PCI bus. When cleared, disables the device from generating PCI accesses. When set, allows the device to behave as a PCI bus master. When operating in the PCI-X mode, ATU initiates a split completion transaction regardless of the state of this bit.		
01	02	Memory Enable - Controls the ATU interface's response to PCI memory addresses. When cleared, the ATU interface does not respond to any memory access on the PCI bus.		
00	02	I/O Space Enable - Controls the ATU interface response to I/O transactions. Not implemented and a reserved bit field.		



# 3.10.4 ATU Status Register - ATUSR

The ATU Status Register bits adhere to the *PCI Local Bus Specification*, Revision 2.3 definitions. The *read/clear* bits can only be set by internal hardware and cleared by either a reset condition or by writing a  $1_2$  to the register.



		IOP       15       12       8       4       0         Attributes       rc       rc			
Internal Bus Address FFFF.E106H		PCI Configuration Address OffsetAttribute Legend:RW = Read/Write06H - 07HRV = ReservedRC = Read ClearPR = PreservedRO = Read OnlyRS = Read/SetNA = Not Accessible			
Bit	Default	Description			
15	02	<ul> <li>Detected Parity Error - set when a parity error is detected in data received by the ATU on the PCI bus even when the ATUCMD register's Parity Error Response bit is cleared. Set under the following conditions:</li> <li>Write Data Parity Error when the ATU is a target (inbound write).</li> <li>Read Data Parity Error when the ATU is a requester (outbound read).</li> <li>Any Address or Attribute (PCI-X Only) Parity Error on the Bus (including one generated by the ATU).</li> </ul>			
14	02	SERR# Asserted - set when SERR# is asserted on the PCI bus by the ATU.			
13	02	Master Abort - set when a transaction initiated by the ATU PCI master interface, ends in a Master-Abort or when the ATU receives a Master Abort Split Completion Error Message in PCI-X mode.			
12	02	Target Abort (master) - set when a transaction initiated by the ATU PCI master interface, ends in a target abort or when the ATU receives a Target Abort Split Completion Error Message in PCI-X mode.			
11	02	Target Abort (target) - set when the ATU interface, acting as a target, terminates the transaction on the PCI bus with a target abort.			
10:09	012	DEVSEL# Timing - These bits are read-only and define the slowest DEVSEL# timing for a target device in Conventional PCI Mode regardless of the operating mode (except configuration accesses). $00_2 = Fast$ $01_2 = Medium$ $10_2 = Slow$ $11_2 = Reserved$ The ATU interface uses Medium timing.			
08	02	<ul> <li>Master Parity Error - The ATU interface sets this bit under the following conditions:</li> <li>The ATU asserted <b>PERR#</b> itself or the ATU observed <b>PERR#</b> asserted.</li> <li>And the ATU acted as the requester for the operation in which the error occurred.</li> <li>And the ATUCMD register's Parity Error Response bit is set</li> <li>Or (PCI-X Mode Only) the ATU received a Write Data Parity Error Message</li> <li>And the ATUCMD register's Parity Error Response bit is set</li> </ul>			
07	1 <sub>2</sub> (Conventional mode) 0 <sub>2</sub> (PCI-X mode)	Fast Back-to-Back - The ATU/Messaging Unit interface is capable of accepting fast back-to-back transactions in Conventional PCI mode when the transactions are not to the same target. Since fast back-to-back transactions do not exist in PCI-X mode, this bit is forced to 0 in the PCI-X mode.			
06	02	UDF Supported - User Definable Features are not supported			
05	1 <sub>2</sub>	66 MHz. Capable - 66 MHz operation is supported.			
04	12	Capabilities - When set, this function implements extended capabilities.			



#### Table 127. ATU Status Register - ATUSR (Sheet 2 of 2)





# 3.10.5 ATU Revision ID Register - ATURID

Revision ID Register bit definitions adhere to PCI Local Bus Specification, Revision 2.3.

Table 128. ATU Revision ID Register - ATURID





# 3.10.6 ATU Class Code Register - ATUCCR

Class Code Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3. Auto configuration software reads this register to determine the PCI device function.







# 3.10.7 ATU Cacheline Size Register - ATUCLSR

Cacheline Size Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3. This register is programmed with the system cacheline size in DWORDs (32-bit words). Cacheline Size is restricted to either 0, 8 or 16 DWORDs; the ATU interprets any other value as "0".

 Table 130.
 ATU Cacheline Size Register - ATUCLSR





## 3.10.8 ATU Latency Timer Register - ATULT

ATU Latency Timer Register bit definitions apply to the PCI interface.

#### Table 131. ATU Latency Timer Register - ATULT





## 3.10.9 ATU Header Type Register - ATUHTR

Header Type Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3. This register indicates the layout of ATU configuration space bytes 10H to 3FH. The MSB indicates whether or not the device is multi-function.

 Table 132.
 ATU Header Type Register - ATUHTR





## 3.10.10 ATU BIST Register - ATUBISTR

The ATU BIST Register controls the functions the Intel XScale<sup>®</sup> core performs when BIST is initiated. This register is the interface between the host processor requesting BIST functions and the 80333 replying with the results from the software implementation of the BIST functionality.

 Table 133.
 ATU BIST Register - ATUBISTR



## 3.10.11 Inbound ATU Base Address Register 0 - IABAR0

The Inbound ATU Base Address Register 0 (IABAR0) together with the Inbound ATU Upper Base Address Register 0 (IAUBAR0) defines the block of memory addresses where the inbound translation window 0 begins. The inbound ATU decodes and forwards the bus request to the 80333 internal bus with a translated address to map into 80333 local memory. The IABAR0 and IAUBAR0 define the base address and describes the required memory block size; see Section 3.10.21, "Determining Block Sizes for Base Address Registers" on page 259. Bits 31 through 12 of the IABAR0 is either read/write bits or read only with a value of 0 depending on the value located within the IALR0. This configuration allows the IABAR0 to be programmed per *PCI Local Bus Specification*, Revision 2.3.

The first 4 Kbytes of memory defined by the IABAR0, IAUBAR0 and the IALR0 is reserved for the Messaging Unit.

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

- *Warning:* When IALR0 is cleared prior to host configuration, the user should also clear the Prefetchable Indicator and the Type Indicator. Assuming IALR0 is not cleared:
  - a. Since non prefetchable memory windows can never be placed above the 4 Gbyte address boundary, when the Prefetchable Indicator is cleared prior to host configuration, the user should also set the Type Indicator for 32 bit addressability.
  - b. For compliance to the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a, when the Prefetchable Indicator is set prior to host configuration, the user should also set the Type Indicator for 64 bit addressability. This is the default for IABAR0.



#### Table 134. Inbound ATU Base Address Register 0 - IABAR0



## 3.10.12 Inbound ATU Upper Base Address Register 0 - IAUBAR0

This register contains the upper base address when decoding PCI addresses beyond 4 GBytes. Together with the Translation Base Address this register defines the actual location the translation function is to respond to when addressed from the PCI bus for addresses > 4GBytes (for DACs).

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

*Note:* When the Type indicator of IABAR0 is set to indicate 32 bit addressability, the IAUBAR0 register attributes are read-only.



#### Table 135. Inbound ATU Upper Base Address Register 0 - IAUBAR0

#### 3.10.13 Inbound ATU Base Address Register 1 - IABAR1

The Inbound ATU Base Address Register (IABAR1) together with the Inbound ATU Upper Base Address Register 1 (IAUBAR1) defines the block of memory addresses where the inbound translation window 1 begins. This window is used merely to allocate memory on the PCI bus and, the ATU does not process any PCI bus transactions to this memory range.

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the PCI Local Bus Specification, Revision 2.3 for additional information on programming base address registers.

When enabled, the ATU interrupts the Intel XScale<sup>®</sup> core when the IABAR1 register is written from the PCI bus. Please see Section 3.10.45, "ATU Interrupt Status Register - ATUISR" on page 287 for more details.

- Warning: When a non-zero value is not written to IALR1 prior to host configuration, the user should not set either the Prefetchable Indicator or the Type Indicator for 64 bit addressability. This is the default for IABAR1. Assuming a non-zero value is written to IALR1, the user may set the Prefetchable Indicator or the Type Indicator
  - c. Since non prefetchable memory windows can never be placed above the 4 Gbyte address boundary, when the Prefetchable Indicator is not set prior to host configuration, the user should also leave the Type Indicator set for 32 bit addressability. This is the default for IABAR1.
  - d. For compliance to the PCI-X Addendum to the PCI Local Bus Specification, Revision 1.0a, when the Prefetchable Indicator is set prior to host configuration, the user should also set the Type Indicator for 64 bit addressability.





March 2005

occupy I/O space, thus this bit must be zero.



## 3.10.14 Inbound ATU Upper Base Address Register 1 - IAUBAR1

This register contains the upper base address when locating this window for PCI addresses beyond 4 GBytes. Together with the IABAR1 this register defines the actual location for this memory window for addresses > 4GBytes (for DACs). This window is used merely to allocate memory on the PCI bus and, the ATU does not process any PCI bus transactions to this memory range.

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

When enabled, the ATU interrupts the Intel XScale<sup>®</sup> core when the IAUBAR1 register is written from the PCI bus. Please see Section 3.10.45, "ATU Interrupt Status Register - ATUISR" on page 287 for more details.

*Note:* When the Type indicator of IABAR1 is set to indicate 32 bit addressability, the IAUBAR1 register attributes are read-only. This is the default for IABAR1.



Table 137. Inbound ATU Upper Base Address Register 1 - IAUBAR1
## 3.10.15 Inbound ATU Base Address Register 2 - IABAR2

The Inbound ATU Base Address Register 2 (IABAR2) together with the Inbound ATU Upper Base Address Register 2 (IAUBAR2) defines the block of memory addresses where the inbound translation window 2 begins. The inbound ATU decodes and forwards the bus request to the 80333 internal bus with a translated address to map into 80333 local memory. The IABAR2 and IAUBAR2 define the base address and describes the required memory block size; see Section 3.10.21, "Determining Block Sizes for Base Address Registers" on page 259. Bits 31 through 12 of the IABAR2 is either read/write bits or read only with a value of 0 depending on the value located within the IALR2. This configuration allows the IABAR2 to be programmed per *PCI Local Bus Specification*, Revision 2.3.

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

- *Warning:* When a non-zero value is not written to IALR2 prior to host configuration, the user should not set either the Prefetchable Indicator or the Type Indicator for 64 bit addressability. This is the default for IABAR2. Assuming a non-zero value is written to IALR2, the user may set the Prefetchable Indicator or the Type Indicator:
  - 1. Since non prefetchable memory windows can never be placed above the 4 Gbyte address boundary, when the Prefetchable Indicator is not set prior to host configuration, the user should also leave the Type Indicator set for 32 bit addressability. This is the default for IABAR2.
  - 2. For compliance to the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a, when the Prefetchable Indicator is set prior to host configuration, the user should also set the Type Indicator for 64 bit addressability.



#### Table 138. Inbound ATU Base Address Register 2 - IABAR2



## 3.10.16 Inbound ATU Upper Base Address Register 2 - IAUBAR2

This register contains the upper base address when decoding PCI addresses beyond 4 GBytes. Together with the Translation Base Address this register defines the actual location the translation function is to respond to when addressed from the PCI bus for addresses > 4GBytes (for DACs).

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

*Note:* When the Type indicator of IABAR2 is set to indicate 32 bit addressability, the IAUBAR2 register attributes are read-only. This is the default for IABAR2.



#### Table 139. Inbound ATU Upper Base Address Register 2 - IAUBAR2



## 3.10.17 ATU Subsystem Vendor ID Register - ASVIR

ATU Subsystem Vendor ID Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3.

#### Table 140. ATU Subsystem Vendor ID Register - ASVIR





## 3.10.18 ATU Subsystem ID Register - ASIR

ATU Subsystem ID Register bit definitions adhere to PCI Local Bus Specification, Revision 2.3.

 Table 141.
 ATU Subsystem ID Register - ASIR



## 3.10.19 Expansion ROM Base Address Register - ERBAR

The Expansion ROM Base Address Register defines the block of memory addresses used for containing the Expansion ROM. It permits the inclusion of multiple code images, allowing the device to be initialized. The code image supplied consists of either executable code or an interpreted code. Each code image must start on a 512 byte boundary and each must contain the PCI Expansion ROM header. Image placement in ROM space depends on the length of code images which precede it within ROM. ERBAR defines the base address and describes the required memory block size; see Section 3.10.21. Expansion ROM address space (limit size) can be a maximum of 16 MBytes. Bits 31 through 12 of the ERBAR is either read/write bits or read only with a value of 0 depending on the value located within the ERLR. This configuration allows the ERBAR to be programmed per *PCI Local Bus Specification*, Revision 2.3.

The Expansion ROM Base Address Register's programmed value must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming Expansion ROM base address registers.



#### Table 142. Expansion ROM Base Address Register - ERBAR



## 3.10.20 ATU Capabilities Pointer Register - ATU\_Cap\_Ptr

The Capabilities Pointer Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3. This register provides an offset in this function's PCI Configuration Space for the location of the first item in the first Capability list. In the case of the 80333, this is the PCI Bus Power Management extended capability as defined by the *PCI Bus Power Management Interface Specification*, Revision 1.1.

#### Table 143. ATU Capabilities Pointer Register - ATU\_Cap\_Ptr



## 3.10.21 Determining Block Sizes for Base Address Registers

The required address size and type can be determined by writing ones to a base address register and reading from the registers. By scanning the returned value from the least-significant bit of the base address registers upwards, the programmer can determine the required address space size. The binary-weighted value of the first non-zero bit found indicates the required amount of space. Table 144 describes the relationship between the values read back and the byte sizes the base address register requires.

Response After Writing all 1s to the Base Address Register	Size (in Bytes)	Response After Writing all 1s to the Base Address Register	Size (in Bytes)
FFFFFF0H	16	FFF00000H	1 M
FFFFFE0H	32	FFE00000H	2 M
FFFFFC0H	64	FFC00000H	4 M
FFFFF80H	128	FF800000H	8 M
FFFFF00H	256	FF000000H	16 M
FFFFE00H	512	FE000000H	32 M
FFFFC00H	1K	FC000000H	64 M
FFFF800H	2K	F8000000H	128 M
FFFF000H	4K	F000000H	256 M
FFFE000H	8K	E000000H	512 M
FFFFC000H	16K	C000000H	1 G
FFFF8000H	32K	8000000H	2 G
FFFF0000H	64K		Register not
FFFE0000H	128K	0000000	mented, no
FFFC0000H	256K		address
FFF80000H	512K		required.

#### Table 144. Memory Block Size Read Response

As an example, assume that FFFF.FFFFH is written to the ATU Inbound Base Address Register 0 (IABAR0) and the value read back is FFF0.0008H. Bit zero is a zero, so the device requires memory address space. Bit three is one, so the memory does supports prefetching. Scanning upwards starting at bit four, bit twenty is the first one bit found. The binary-weighted value of this bit is 1,048,576, indicated that the device requires 1 Mbyte of memory space.

The ATU Base Address Registers and the Expansion ROM Base Address Register use their associated limit registers to enable which bits within the base address register are read/write and which bits are read only (0). This allows the programming of these registers in a manner similar to other PCI devices even though the limit is variable.



#### Table 145. ATU Base Registers and Associated Limit Registers

Base Address Register	Limit Register	Description	
Inbound ATU Base Address Register 0	Inbound ATU Limit Register 0	Defines the inbound translation window 0 from the PCI bus.	
Inbound ATU Upper Base Address Register 0	N/A	Together with ATU Base Address Register 0 defines the inbound translation window 0 from the PCI bus for DACs.	
Inbound ATU Base Address Register 1 <sup>a</sup>	Inbound ATU Limit Register 1 Defines inbound window 1 from t PCI bus.		
Inbound ATU Upper Base Address Register 1	N/A	Together with ATU Base Address Register 1 defines inbound window 1 from the PCI bus for DACs.	
Inbound ATU Base Address Register 2	Inbound ATU Limit Register 2	Defines the inbound translation window 2 from the PCI bus.	
Inbound ATU Upper Base Address Register 2	N/A	Together with ATU Base Address Register 2 defines the inbound translation window 2 from the PC bus for DACs.	
Inbound ATU Base Address Register 3	Inbound ATU Limit Register 3	Defines the inbound translation window 3 from the PCI bus.	
Inbound ATU Upper Base Address Register 3	N/A	Together with ATU Base Address Register 3 defines the inbound translation window 3 from the PCI bus for DACs. <b>NOTE:</b> This is a private BAR that resides outside of the standard PCI configuration header space (offsets 00H-3FH).	
Expansion ROM Base Address Register	Expansion ROM Limit Register	Defines the window of addresses used by a bus master for reading from an Expansion ROM.	

a. ATU Inbound Window 1 is **not** a translate window. The ATU does not claim any PCI accesses that fall within this range. This window is used to allocate host memory for use by Private Devices. When enabled, the ATU interrupts the Intel XScale® core when either the IABAR1 register or the IAUBAR1 register is written from the PCI bus. Please see Section 3.10.45, "ATU Interrupt Status Register - ATUISR" on page 287 for more details.

## 3.10.22 ATU Interrupt Line Register - ATUILR

ATU Interrupt Line Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3. This register identifies the system interrupt controller's interrupt request lines which connect to the device's PCI interrupt request lines (as specified in the interrupt pin register).

In a PC environment, for example, the register values and corresponding connections are:

- 0 (00H) through 15 (0FH) correspond to IRQ0 through IRQ15
- 16 (10H) through 254 (FEH) are reserved
- 255 (FFH) indicates "unknown" or "no connection"

The operating system or device driver can examine each device's interrupt pin and interrupt line register to determine which system interrupt request line the device uses to issue requests for service.

#### Table 146. ATU Interrupt Line Register - ATUILR





## 3.10.23 ATU Interrupt Pin Register - ATUIPR

ATU Interrupt Pin Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3. This register identifies the interrupt pin the ATU and Messaging Unit interface uses. The 80333 is, a PCI single-function device and, as such, generates only one interrupt output. The interrupt output is for the Messaging Unit on **INTA**#.

The ATU interrupt output is wired to the A-segment IOAPIC **IRQ14**# pin. This correlates to an Assert\_INTC and Deassert\_INTC messages on PCI Express for legacy interrupt mode of the IOAPIC. This interrupt wiring coincides with the ATU device number 14 (0EH) based on A\_AD30 used as the IDSEL input to the ATU device.

#### Table 147. ATU Interrupt Pin Register - ATUIPR





## 3.10.24 ATU Minimum Grant Register - ATUMGNT

ATU Minimum Grant Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3. This register specifies the burst period the device requires in increments of 8 PCI clocks.

This register and the ATU Maximum Latency register are information-only registers which the configuration uses to determine how often a bus master typically requires access to the PCI bus and the duration of a typical transfer when it does acquire the bus. This information is useful in determining the values to be programmed into the bus master latency timers and in programming the algorithm to be used by the PCI bus arbiter.

#### Table 148. ATU Minimum Grant Register - ATUMGNT





## 3.10.25 ATU Maximum Latency Register - ATUMLAT

ATU Maximum Latency Register bit definitions adhere to *PCI Local Bus Specification*, Revision 2.3. This register specifies how often the device needs to access the PCI bus in increments of 8 PCI clocks.

This register and the Minimum Grant Register are information-only registers which the configuration uses to determine how often a bus master typically requires access to the PCI bus and the duration of a typical transfer when it does acquire the bus. This information is useful in determining the values to be programmed into the bus master latency timers and in programming the algorithm to be used by the PCI bus arbiter.

#### Table 149. ATU Maximum Latency Register - ATUMLAT



## 3.10.26 Inbound ATU Limit Register 0 - IALR0

Inbound address translation for memory window 0 occurs for data transfers occurring from the PCI bus (originated from the PCI bus) to the 80333 internal bus. The address translation block converts PCI addresses to internal bus addresses.

The inbound translation base address for inbound window 0 is specified in Section 3.10.11. When determining block size requirements — as described in Section 3.10.21 — the translation limit register provides the block size requirements for the base address register. The remaining registers used for performing address translation are discussed in Section 3.2.1.1.

The 80333 translate value register's programmed value must be naturally aligned with the base address register's programmed value. The limit register is used as a mask; thus, the lower address bits programmed into the 80333 translate value register are invalid. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

Bits 31 to 12 within the IALR0 have a direct effect on the IABAR0 register, bits 31 to 12, with a one to one correspondence. A value of 0 in a bit within the IALR0 makes the corresponding bit within the IABAR0 a read only bit which always returns 0. A value of 1 in a bit within the IALR0 makes the corresponding bit within the IABAR0 read/write from PCI. Note that a consequence of this programming scheme is that unless a valid value exists within the IALR0, all writes to the IABAR0 has no effect since a value of all zeros within the IALR0 makes the IABAR0 a read only register.



#### Table 150. Inbound ATU Limit Register 0 - IALR0



## 3.10.27 Inbound ATU Translate Value Register 0 - IATVR0

The Inbound ATU Translate Value Register 0 (IATVR0) contains the internal bus address used to convert PCI bus addresses. The converted address is driven on the internal bus as a result of the inbound ATU address translation.



#### Table 151. Inbound ATU Translate Value Register 0 - IATVR0

# intel

## 3.10.28 Expansion ROM Limit Register - ERLR

The Expansion ROM Limit Register (ERLR) defines the block size of addresses the ATU defines as Expansion ROM address space. The block size is programmed by writing a value into the ERLR.

Bits 31 to 12 within the ERLR have a direct effect on the ERBAR register, bits 31 to 12, with a one to one correspondence. A value of 0 in a bit within the ERLR makes the corresponding bit within the ERBAR a read only bit which always returns 0. A value of 1 in a bit within the ERLR makes the corresponding bit within the ERBAR read/write from PCI.



#### Table 152. Expansion ROM Limit Register - ERLR



## 3.10.29 Expansion ROM Translate Value Register - ERTVR

The Expansion ROM Translate Value Register contains the 80333 internal bus address which the ATU converts the PCI bus access. This address is driven on the internal bus as a result of the Expansion ROM address translation.





# intel

## 3.10.30 Inbound ATU Limit Register 1 - IALR1

Bits 31 to 12 within the IALR1 have a direct effect on the IABAR1 register, bits 31 to 12, with a one to one correspondence. A value of 0 in a bit within the IALR1 makes the corresponding bit within the IABAR1 a read only bit which always returns 0. A value of 1 in a bit within the IALR1 makes the corresponding bit within the IABAR1 read/write from PCI. Note that a consequence of this programming scheme is that unless a valid value exists within the IALR1, all writes to the IABAR1 has no effect since a value of all zeros within the IALR1 makes the IABAR1 a read only register.

The inbound memory window 1 is used merely to allocate memory on the PCI bus. The ATU does not process any PCI bus transactions to this memory range.

*Warning:* The ATU does not claim any PCI accesses that fall within the range defined by IABAR1, IAUBAR1, and IALR1.



Table 154. Inbound ATU Limit Register 1 - IALR1



### 3.10.31 Inbound ATU Limit Register 2 - IALR2

Inbound address translation for memory window 2 occurs for data transfers occurring from the PCI bus (originated from the PCI bus) to the 80333 internal bus. The address translation block converts PCI addresses to internal bus addresses.

The inbound translation base address for inbound window 2 is specified in Section 3.10.15. When determining block size requirements — as described in Section 3.10.21 — the translation limit register provides the block size requirements for the base address register. The remaining registers used for performing address translation are discussed in Section 3.2.1.1.

The 80333 translate value register's programmed value must be naturally aligned with the base address register's programmed value. The limit register is used as a mask; thus, the lower address bits programmed into the 80333 translate value register are invalid. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

Bits 31 to 12 within the IALR2 have a direct effect on the IABAR2 register, bits 31 to 12, with a one to one correspondence. A value of 0 in a bit within the IALR2 makes the corresponding bit within the IABAR2 a read only bit which always returns 0. A value of 1 in a bit within the IALR2 makes the corresponding bit within the IABAR2 read/write from PCI. Note that a consequence of this programming scheme is that unless a valid value exists within the IALR2, all writes to the IABAR2 has no effect since a value of all zeros within the IALR2 makes the IABAR2 a read only register.



#### Table 155. Inbound ATU Limit Register 2 - IALR2



## 3.10.32 Inbound ATU Translate Value Register 2 - IATVR2

The Inbound ATU Translate Value Register 2 (IATVR2) contains the internal bus address used to convert PCI bus addresses. The converted address is driven on the internal bus as a result of the inbound ATU address translation.



#### Table 156. Inbound ATU Translate Value Register 2 - IATVR2



### 3.10.33 Outbound I/O Window Translate Value Register - OIOWTVR

The Outbound I/O Window Translate Value Register (OIOWTVR) contains the PCI I/O address used to convert the internal bus access to a PCI address. This address is driven on the PCI bus as a result of the outbound ATU address translation. See Section 3.2.2.1, "Outbound Address Translation - Single Address Cycle (SAC) Internal Bus Transactions" on page 175 for details on outbound address translation.

The I/O window is from 80333 internal bus address 9000 000H to 9000 FFFFH with the fixed length of 64 Kbytes.



#### Table 157. Outbound I/O Window Translate Value Register - OIOWTVR



## 3.10.34 Outbound Memory Window Translate Value Register 0 - OMWTVR0

The Outbound Memory Window Translate Value Register 0 (OMWTVR0) contains the PCI address used to convert 80333 internal bus addresses for outbound transactions. This address is driven on the PCI bus as a result of the outbound ATU address translation. See Section 3.2.2.1, "Outbound Address Translation - Single Address Cycle (SAC) Internal Bus Transactions" on page 175 for details on outbound address translation.

The memory window is from internal bus address 8000 000H to 83FF FFFFH with the fixed length of 64 Mbytes.



#### Table 158. Outbound Memory Window Translate Value Register 0- OMWTVR0



## 3.10.35 Outbound Upper 32-bit Memory Window Translate Value Register 0 - OUMWTVR0

The Outbound Upper 32-bit Memory Window Translate Value Register 0 (OUMWTVR0) defines the upper 32-bits of address used during a dual address cycle. This enables the outbound ATU to directly address anywhere within the 64-bit host address space. When this register is all-zero, then a SAC is generated on the PCI bus.

The memory window is from internal bus address 8000 000H to 83FF FFFFH with the fixed length of 64 Mbytes.



#### Table 159. Outbound Upper 32-bit Memory Window Translate Value Register 0- OUMWTVR0



## 3.10.36 Outbound Memory Window Translate Value Register 1 - OMWTVR1

The Outbound Memory Window Translate Value Register 1 (OMWTVR1) contains the PCI address used to convert 80333 internal bus addresses for outbound transactions. This address is driven on the PCI bus as a result of the outbound ATU address translation. See Section 3.2.2.1, "Outbound Address Translation - Single Address Cycle (SAC) Internal Bus Transactions" on page 175 for details on outbound address translation.

The memory window is from internal bus address 8400 000H to 87FF FFFFH with the fixed length of 64 Mbytes.



#### Table 160. Outbound Memory Window Translate Value Register 1- OMWTVR1



#### 3.10.37 **Outbound Upper 32-bit Memory Window** Translate Value Register 1 - OUMWTVR1

The Outbound Upper 32-bit Memory Window Translate Value Register 1 (OUMWTVR1) defines the upper 32-bits of address used during a dual address cycle. This enables the outbound ATU to directly address anywhere within the 64-bit host address space. When this register is all-zero, then a SAC is generated on the PCI bus.

The memory window is from internal bus address 8400 000H to 87FF FFFFH with the fixed length of 64 Mbytes.



#### Table 161. Outbound Upper 32-bit Memory Window Translate Value Register 1- OUMWTVR1



## 3.10.38 Outbound Upper 32-bit Direct Window Translate Value Register - OUDWTVR

The Outbound Upper 32-bit Direct Window Translate Value Register (OUDWTVR) defines the upper 32-bits of address used during a dual address cycle for the transactions via Direct Addressing Window. This enables the outbound ATU to directly address anywhere within the 64-bit host address space. When this register is all-zero, then a SAC is generated on the PCI bus.

Table 162. Outbound Upper 32-bit Direct Window Translate Value Register - OUDWTVR





### 3.10.39 PCI Express-to-PCI Bridge Secondary A-Segment Bus Number Register - PEBSABNR

The Secondary Bus Number register of the A-Segment PCI Express-to-PCI Bridge is mirrored in this register for direct access by software. This register can be used to determine when the PCI bus number changed of the A-Segment PCI bus (and ATU) is operating in PCI mode.

#### Table 163. PCI Express-to-PCI Bridge Secondary A\_Segment Bus Number Register -PEBSABNR



## intel®

## 3.10.40 PCI Express-to-PCI Bridge Secondary B-Segment Bus Number Register - PEBSBBNR

The Secondary Bus Number register of the B-Segment PCI Express-to-PCI Bridge is mirrored in this register for direct access by software. This register can be used to determine when the PCI bus number changed of the B-Segment PCI bus.

*Note:* The PCI Express-to-PCI Bridge Secondary B-Segment Bus Number Register (PEBSBBNR) does not provide the B-segment bus number. PEBSBBNR is suppose to mirror the Secondary Bus Number(SCBN) field of the BNUM register (bits 15:8 at offset 18h) in the B-segment configuration header space. Instead, PEBSBBNR provides the Subordinate Bus Number (SBBN) of the BNUM register (bits 23:16 at offset 18h) from the A-segment configuration header space. For more details, see *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #4.

## Table 164.PCI Express-to-PCI Bridge Secondary B-Segment<br/>Bus Number Register - PEBSBBNR





#### 3.10.41 PCI Express-to-PCI Bridge Primary Bus Number Register - PEBPBNR

The Primary Bus Number register of the A-Segment PCI Express-to-PCI Bridge is mirrored in this register for direct access by software.

#### Table 165. PCI Express-to-PCI Bridge Primary Bus Number Register - PEBPBNR





### 3.10.42 PCI Express-to-PCI Bridge Device Number Register - PEBDNUM

The Device Number field of the A-Segment PCI Express-to-PCI Bridge is mirrored in this register for direct access by software.

Table 166. PCI Express-to-PCI Bridge Device Number Register - PEBDNUM





## 3.10.43 ATU Configuration Register - ATUCR

The ATU Configuration Register controls the outbound address translation for address translation unit. It also contains bits for Conventional PCI Delayed Read Command (DRC) aliasing, discard timer status, **SERR#** manual assertion, **SERR#** detection interrupt masking, and ATU BIST interrupt enabling.



Table 167. ATU Configuration Register - ATUCR



## 3.10.44 PCI Configuration and Status Register - PCSR

The PCI Configuration and Status Register has additional bits for controlling and monitoring various features of the PCI bus interface.



Attributes       31       28       24       20       16       12       8       4       0         IOP Attributes			
FFFF.E18	Bus Address 34H	PCI Configuration Address Offset       Attribute Legend:       RW = Read/Write         84H - 87H       RV = Reserved       RC = Read Clear         CO = Clear Only       RO = Read Only         RS = Read/Set       NA = Not Accessible	
Bit	Default	Description	
31:19	0000H	Reserved	
18	02	<ul> <li>Detected Address or Attribute Parity Error - set when a parity error is detected during either the address or attribute phase of a transaction on the PCI bus even when the ATUCMD register's Parity Error Response bit is cleared. Set under the following conditions:</li> <li>Any Address or Attribute (PCI-X Only) Parity Error on the Bus (including one generated by the ATU).</li> </ul>	
17:16	Varies with external state of DEVSEL#, STOP#, and TRDY#, during <b>PWRGD</b>	PCI-X capability - These two bits define the mode of the PCI bus (conventional or PCI-X) as well as the operating frequency in the case of PCI-X mode.         00 - Conventional PCI mode         01 - PCI-X 66         10 - PCI-X 100         11 - PCI-X 133         As defined by the PCI-X Addendum to the PCI Local Bus Specification, Revision 1.0a, the operating mode is determined by an initialization pattern on the PCI bus during PWRGD assertion:         DEVSEL#       STOP#         TRDY#       Mode         Deasserted       Deasserted         Deasserted       Deasserted         Deasserted       Asserted         PCI-X 100         Deasserted       Asserted         PCI-ST Addendum to the PCI Local Bus Specification, Revision 1.0a, the operating mode is determined by an initialization pattern on the PCI bus during PWRGD assertion:         DEVSEL#       STOP#         TRDY#       Mode         Deasserted       Deasserted         Deasserted       Asserted         PCI-X 100       Deasserted         Asserted       Asserted         PCI-X 133         All other patterns are reserved.	
15	02	Outbound Transaction Queue Busy: 0 = Outbound Transaction Queue Empty 1 = Outbound Transaction Queue Busy	
14	02	Inbound Transaction Queue Busy: 0 = Inbound Transaction Queue Empty 1 = Inbound Transaction Queue Busy	
13	02	Reserved	
12	02	Discard Timer Value - This bit controls the time-out value for the four discard timers attached to the queues holding read data. A value of 0 indicates the time-out value is $2^{15}$ clocks. A value of 1 indicates the time-out value is $2^{10}$ clocks.	
11	02	Reserved	





#### Table 168. PCI Configuration and Status Register - PCSR (Sheet 2 of 4)





Att At Internal E FFFF.E18	IOP     31       ributes     rv rv       PCI     rv rv       tributes     rv rv       Bus Address     444	28       24       20       16         1       1       1       1       1       1         2       1       1       1       1       1       1         2       1       1       1       1       1       1       1         2       1       1       1       1       1       1       1       1       1         2       1	12 8 ro rv rw rv ro rv ro rv rv rv rw ro rv rw rv ro rv ro rv rv rv rw Attribute Legend: RV = Reserved CO = Clear Only RS = Read/Set	4 0 pr rv rw co rv pr rv ro co rv RW = Read/Write RC = Read Clear RO = Read Only NA = Not Accessible
Bit	Default	De	escription	
05	02	<ul> <li>Reset Internal Bus - This bit controls the reset of the Intel XScale<sup>®</sup> core and all units on the internal bus. In addition to the internal bus initialization, this bit triggers the assertion of the M_RST# pin for initialization of registered DIMMs. When set:</li> <li>When operating in the conventional PCI mode:</li> <li>All current PCI transactions being mastered by the ATU will complete, and the ATU master interfaces will proceed to an idle state. No additional transactions will be mastered by these units until the internal bus reset is complete.</li> <li>All current transactions being slaved by the ATU on either the PCI bus or the internal bus will complete, and the ATU target interfaces will proceed to an idle state. All future slave transactions will master abort, with the exception of the completion cycle for the transaction that set the Reset Internal Bus bit in the PCSR.</li> <li>When the value of the Core Processor Reset bit in the PCSR (upon PWRGD assertion) is set, the Intel XScale<sup>®</sup> core will be held in reset when the internal bus reset is complete.</li> <li>The ATU will ignore configuration cycles, and they will appear as master aborts for: 32 Internal Bus clocks.</li> <li>The A0333 hardware will clear this bit after the reset operation completes.</li> <li>When operating in the PCI-X mode:</li> <li>The ATU hardware will respond the same as in Conventional PCI-X mode. However, this may create a problem in PCI-X mode for split requests in that there may still be an outstanding split completion that the ATU CMD. This will ensure that no new transactions, either outbound or inbound are enqueud.</li> <li>Wait for both the Outbound (bit 15 of the ATUCMD) and the Memory Enable (bit 1 of the ATUCMD) bits in the ATUCMD. This will ensure that no new transactions, either outbound or inbound are enqueud.</li> <li>Wait for both the Outbound (bit 15 of the PCSR) and Inbound Read (bit 14 of the PCSR) Transaction queue busy bits to be clear.</li> <li>Set the Reset Internal Bus bit</li> <li>As a result, t</li></ul>		





#### Table 168. PCI Configuration and Status Register - PCSR (Sheet 4 of 4)

# intel

## 3.10.45 ATU Interrupt Status Register - ATUISR

The ATU Interrupt Status Register is used to notify the core processor of the source of an ATU interrupt. In addition, this register is written to clear the source of the interrupt to the interrupt unit of the 80333. All bits in this register are Read/Clear.

Bits 4:0 are a direct reflection of bits 14:11 and bit 8 (respectively) of the ATU Status Register (these bits are set at the same time by hardware but need to be cleared independently). Bit 7 is set by an error associated with the internal bus of the 80333. Bit 8 is for software BIST. The conditions that result in an ATU interrupt are cleared by writing a 1 to the appropriate bits in this register.

*Note:* Bits 4:0, and bits 15 and 13:7 can result in an interrupt being driven to the Intel XScale<sup>®</sup> core.

Table 169. ATU Interrupt Status Register - ATUISR (Sheet 1 of 2)





Att Att Internal E FFFF.E18	IOP ributes rv rv PCI rv rv Bus Address 38H	28       24       20       16       12       8       4       0         rv       r	
Bit	Default	Description	
09	02	<ul> <li>Detected Parity Error - set when a parity error is detected on the PCI bus even when the ATUCMD register's Parity Error Response bit is cleared. Set under the following conditions:</li> <li>Write Data Parity Error when the ATU is a target (inbound write).</li> <li>Read Data Parity Error when the ATU is an initiator (outbound read).</li> <li>Any Address or Attribute (PCI-X Only) Parity Error on the Bus.</li> </ul>	
08	02	ATU BIST Interrupt - When set, generates the ATU BIST Start Interrupt and indicates the host processor has set the start BIST, ATUBISTR register bit 6, and the ATU BIST interrupt enable (ATUCR register bit 3) is enabled. The Intel XScale <sup>®</sup> core can initiate the software BIST and store the result in ATUBISTR register bits 3:0. Configuration register writes to the ATUBISTR will NOT result in bit 15 also being set or the assertion of the ATU Configure Register Write Interrupt.	
07	02	Internal Bus Master Abort - set when a transaction initiated by the ATU internal bus initiator interface ends in a Master-abort.	
06:05	00 <sub>2</sub>	Reserved.	
04	02	SERR# Detected - set when SERR# is detected on the PCI bus by the ATU.	
03	02	PCI Master Abort - set when a transaction initiated by the ATU PCI initiator interface ends in a Master-abort.	
02	02	PCI Target Abort (master) - set when a transaction initiated by the ATU PCI master interface ends in a Target-abort.	
01	02	PCI Target Abort (target) - set when the ATU interface, acting as a target, terminates the transaction on the PCI bus with a target abort.	
00	02	<ul> <li>PCI Master Parity Error - Master Parity Error - The ATU interface sets this bit under the following conditions:</li> <li>The ATU asserted PERR# itself or the ATU observed PERR# asserted.</li> <li>And the ATU acted as the requester for the operation in which the error occurred.</li> <li>And the ATUCMD register's Parity Error Response bit is set</li> <li>Or (PCI-X Mode Only) the ATU received a Write Data Parity Error Message</li> <li>And the ATUCMD register's Parity Error Response bit is set</li> </ul>	

#### Table 169. ATU Interrupt Status Register - ATUISR (Sheet 2 of 2)


## 3.10.46 ATU Interrupt Mask Register - ATUIMR

The ATU Interrupt Mask Register contains the control bit to enable and disable interrupts generated by the ATU.



#### Table 170. ATU Interrupt Mask Register - ATUIMR (Sheet 1 of 3)



At At Internal FFFF.E18	IOP tributes / tv/ PCI / tr/ Bus Address BCH	28       24       20       16       12       8       4       0         rv       r		
Bit	Default	Description		
08	1 <sub>2</sub>	Power State Transition Interrupt Mask - Controls the setting of bit 12 of the ATUISR and generation of the ATU Error interrupt when ATU Power Management Control/Status Register is written to transition the ATU Function Power State from D0 to D3, D0 to D1, D1 to D3 or D3 to D0. 0 = Not Masked 1 = Masked		
07	02	ATU Detected Parity Error Interrupt Mask - Controls the setting of bit 9 of the ATUISR and generation of the ATU Error interrupt when a parity error detected on the PCI bus that sets bit 15 of the ATUSR. 0 = Not Masked 1 = Masked		
06	02	ATU SERR# Asserted Interrupt Mask - Controls the setting of bit 10 of the ATUISR and generation of the ATU Error interrupt when SERR# is asserted on the PCI interface resulting in bit 14 of the ATUSR being set. 0 = Not Masked 1 = Masked NOTE: This bit is specific to the ATU asserting SERR# and not detecting SERR# from another master.		
05	02	ATU PCI Master Abort Interrupt Mask - Controls the setting of bit 3 of the ATUISR and generation of the ATU Error interrupt when a master abort error resulting in bit 13 of the ATUSR being set. 0 = Not Masked 1 = Masked		
04	02	ATU PCI Target Abort (Master) Interrupt Mask- Controls the setting of bit 12 of the ATUISR and ATU Error generation of the interrupt when a target abort error resulting in bit 12 of the ATUSR being set 0 = Not Masked 1 = Masked		
03	02	ATU PCI Target Abort (Target) Interrupt Mask- Controls the setting of bit 1 of the ATUISR and generation of the ATU Error interrupt when a target abort error resulting in bit 11 of the ATUSR being set. 0 = Not Masked 1 = Masked		
02	02	02 ATU PCI Master Parity Error Interrupt Mask - Controls the setting of bit 0 of the ATUISR and generation of the ATU Error interrupt when a parity error resulting in bit 8 of the ATUSR being set. 0 = Not Masked 1 = Masked		
01	02	ATU Inbound Error <b>SERR#</b> Enable - Controls when the ATU will assert (when enabled through the ATUCMD) <b>SERR#</b> on the PCI interface in response to a master abort on the internal bus during an inbound write transaction. 0 = <b>SERR#</b> Not Asserted due to error 1 = <b>SERR#</b> Asserted due to error		

### Table 170. ATU Interrupt Mask Register - ATUIMR (Sheet 2 of 3)









## 3.10.47 Inbound ATU Base Address Register 3 - IABAR3

The Inbound ATU Base Address Register 3 (IABAR3) together with the Inbound ATU Upper Base Address Register 3 (IAUBAR3) defines the block of memory addresses where the inbound translation window 3 begins. The inbound ATU decodes and forwards the bus request to the 80333 internal bus with a translated address to map into 80333 local memory. The IABAR3 and IAUBAR3 define the base address and describes the required memory block size; see Section 3.10.21, "Determining Block Sizes for Base Address Registers" on page 259. Bits 31 through 12 of the IABAR3 is either read/write bits or read only with a value of 0 depending on the value located within the IALR3. This configuration allows the IABAR3 to be programmed per *PCI Local Bus Specification*, Revision 2.3.

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

- *Note:* Since IABAR3 does not appear in the standard PCI configuration header space (offsets 00H 3CH), IABAR3 is not configured by the host during normal system initialization.
- *Warning:* When a non-zero value is not written to IALR3, the user should not set either the Prefetchable Indicator or the Type Indicator for 64 bit addressability. This is the default for IABAR3. Assuming a non-zero value is written to IALR3, the user may set the Prefetchable Indicator or the Type Indicator:
  - 1. Since non prefetchable memory windows can never be placed above the 4 Gbyte address boundary, when the Prefetchable Indicator is not set, the user should also leave the Type Indicator set for 32 bit addressability. This is the default for IABAR3.
  - 2. For compliance to the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a, when the Prefetchable Indicator is set, the user should also set the Type Indicator for 64 bit addressability.



#### Table 171. Inbound ATU Base Address Register 3 - IABAR3



## 3.10.48 Inbound ATU Upper Base Address Register 3 - IAUBAR3

This register contains the upper base address when decoding PCI addresses beyond 4 GBytes. Together with the Translation Base Address this register defines the actual location the translation function is to respond to when addressed from the PCI bus for addresses > 4GBytes (for DACs).

The programmed value within the base address register must comply with the PCI programming requirements for address alignment. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

*Note:* When the Type indicator of IABAR3 is set to indicate 32 bit addressability, the IAUBAR3 register attributes are read-only. This is the default for IABAR3.



#### Table 172. Inbound ATU Upper Base Address Register 3 - IAUBAR3



## 3.10.49 Inbound ATU Limit Register 3 - IALR3

Inbound address translation for memory window 3 occurs for data transfers occurring from the PCI bus (originated from the PCI bus) to the 80333 internal bus. The address translation block converts PCI addresses to internal bus addresses.

The inbound translation base address for inbound window 3 is specified in Section 3.10.15. When determining block size requirements — as described in Section 3.10.21 — the translation limit register provides the block size requirements for the base address register. The remaining registers used for performing address translation are discussed in Section 3.2.1.1.

The 80333 translate value register's programmed value must be naturally aligned with the base address register's programmed value. The limit register is used as a mask; thus, the lower address bits programmed into the 80333 translate value register are invalid. Refer to the *PCI Local Bus Specification*, Revision 2.3 for additional information on programming base address registers.

Bits 31 to 12 within the IALR3 have a direct effect on the IABAR3 register, bits 31 to 12, with a one to one correspondence. A value of 0 in a bit within the IALR3 makes the corresponding bit within the IABAR3 a read only bit which always returns 0. A value of 1 in a bit within the IALR3 makes the corresponding bit within the IABAR3 read/write from PCI. Note that a consequence of this programming scheme is that unless a valid value exists within the IALR3, all writes to the IABAR3 has no effect since a value of all zeros within the IALR3 makes the IABAR3 a read only register.



#### Table 173. Inbound ATU Limit Register 3 - IALR3



## 3.10.50 Inbound ATU Translate Value Register 3 - IATVR3

The Inbound ATU Translate Value Register 3 (IATVR3) contains the internal bus address used to convert PCI bus addresses. The converted address is driven on the internal bus as a result of the inbound ATU address translation.



Table 174. Inbound ATU Translate Value Register 3 - IATVR3



## 3.10.51 Outbound Configuration Cycle Address Register - OCCAR

The Outbound Configuration Cycle Address Register is used to hold the 32-bit PCI configuration cycle address. The Intel XScale<sup>®</sup> core writes the PCI configuration cycles address which then enables the outbound configuration read or write. The Intel XScale<sup>®</sup> core then performs a read or write to the Outbound Configuration Cycle Data Register to initiate the configuration cycle on the PCI bus.

*Note:* Bits 15:11 of the configuration cycle address for Type 0 configuration cycles are defined differently for Conventional versus PCI-X modes. When 80333 software programs the OCCAR to initiate a Type 0 configuration cycle, the OCCAR should always be loaded based on the PCI-X definition for the Type 0 configuration cycle address. When operating in Conventional mode, the 80333 clears bits 15:11 of the OCCAR prior to initiating an outbound Type 0 configuration cycle. See the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a for details on the two formats.



#### Table 175. Outbound Configuration Cycle Address Register - OCCAR

## 3.10.52 Outbound Configuration Cycle Data Register - OCCDR

The Outbound Configuration Cycle Data Register is used to initiate a configuration read or write on the PCI bus. The register is logical rather than physical meaning that it is an address not a register. The Intel XScale<sup>®</sup> core reads or writes the data registers memory-mapped address to initiate the configuration cycle on the PCI bus with the address found in the OCCAR. For a configuration write, the data is latched from the internal bus and forwarded directly to the OWQ. For a read, the data is returned directly from the ORQ to the Intel XScale<sup>®</sup> core and is never actually entered into the data register (which does not physically exist).

The OCCDR is only visible from 80333 internal bus address space and appears as a reserved value within the ATU configuration space.



#### Table 176. Outbound Configuration Cycle Data Register - OCCDR



## 3.10.53 VPD Capability Identifier Register - VPD\_CAPID

The Capability Identifier Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3. This register in the PCI Extended Capability header identifies the type of Extended Capability contained in that header. In the case of the 80333, this is the VPD extended capability with an ID of 03H as defined by the *PCI Local Bus Specification*, Revision 2.3.





## intel

## 3.10.54 VPD Next Item Pointer Register - VPD\_NXTP

The Next Item Pointer Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3. This register describes the location of the next item in the function's capability list. For the 80333, this the final capability list, and hence, this register is set to 00H.

#### Table 178. VPD Next Item Pointer Register - VPD\_NXTP





## 3.10.55 VPD Address Register - VPD\_AR

The VPD Address register (VPDAR) contains the DWORD-aligned byte address of the VPD to be accessed. The register is read/write and the initial value at power-up is indeterminate.

A PCI Configuration Write to the VPDAR interrupts the Intel XScale<sup>®</sup> core. Software can use the Flag setting to determine whether the configuration write was intended to initiate a read or write of the VPD through the VPD Data Register.







## 3.10.56 VPD Data Register - VPD\_DR

This register is used to transfer data between the 80333 and the VPD storage component.

*Note:* Bit 19 of the VPD Data Register has a bit type of RC (read clear). When the VPD feature is used, care must be taken to mask bit 19. See *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #11.



#### Table 180. VPD Data Register - VPD\_DR



## 3.10.57 Power Management Capability Identifier Register - PM\_CAPID

The Capability Identifier Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3. This register in the PCI Extended Capability header identifies the type of Extended Capability contained in that header. In the case of the 80333, this is the PCI Bus Power Management extended capability with an ID of 01H as defined by the *PCI Bus Power Management Interface Specification*, Revision 1.1.

#### Table 181. Power Management Capability Identifier Register - PM\_CAPID





### 3.10.58 Power Management Next Item Pointer Register - PM\_NXTP

The Next Item Pointer Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3. This register describes the location of the next item in the function's capability list. For the 80333, the next capability (MSI capability list) is located at off-set D0H.

#### Table 182. Power Management Next Item Pointer Register - PM\_NXTP





## 3.10.59 Power Management Capabilities Register - PM\_CAP

Power Management Capabilities bits adhere to the definitions in the *PCI Bus Power Management Interface Specification*, Revision 1.1. This register is a 16-bit read-only register which provides information on the capabilities of the ATU function related to power management.

#### Table 183. Power Management Capabilities Register - PM\_CAP

Int	ernal Bus Addr FF.E1C2H	Attributes	15 12 8 ro ro r	4 0 ro ro rv ro ro ro ro ro ro rv ro ro ro ro Attribute Legend: RV = Reserved PR = Preserved	RW = Read/Write RC = Read Clear RO = Read Only
				RS = Read/Set	NA = Not Accessible
Bit	Default		Description		
15:11	00000 <sub>2</sub>	PME_Support - This function is not capable of asserting the <b>PME#</b> signal in any state, since <b>PME#</b> is not supported by the 80333.			
10	02	D2_Support - This bit is set to $0_2$ indicating that the 80333 does not support the D2 Power Management State			
9	1 <sub>2</sub>	D1_Support - This I	D1_Support - This bit is set to 1 <sub>2</sub> indicating that the 80333 supports the D1 Power Management State		
8:6	000 <sub>2</sub>	Aux_Current - This field is set to 000 <sub>2</sub> indicating that the 80333 has no current requirements for the 3.3Vaux signal as defined in the <i>PCI Bus Power Management Interface Specification</i> , Revision 1.1			
5	02	DSI - This field is set to $0_2$ meaning that this function does not require a device specific initialization sequence following the transition to the D0 uninitialized state.			
4	02	Reserved.			
3	02	PME Clock - Since	PME Clock - Since the 80333 does not support <b>PME#</b> signal generation this bit is cleared to $0_2$ .		
2:0	0102	Version - Setting these bits to 010 <sub>2</sub> means that this function complies with <i>PCI Bus Power Management</i> Interface Specification, Revision 1.1			



## 3.10.60 Power Management Control/Status Register - PM\_CSR

Power Management Control/Status bits adhere to the definitions in the *PCI Bus Power Management Interface Specification*, Revision 1.1. This 16-bit register is the control and status interface for the power management extended capability.

#### Table 184. Power Management Control/Status Register - PM\_CSR

	IOP       15       12       8       4       0         Attributes       ro       ro       rv       rv				
Internal Bus Address FFFF.E1C4H			PCI Configuration Offset C4- C5H	Attribute Legend: RV = Reserved PR = Preserved RS = Read/Set	RW = Read/Write RC = Read Clear RO = Read Only NA = Not Accessible
Bit	Default		Desc	cription	
15	02	PME_Status - This function is not capable of asserting the PME# signal in any state, since <b>PME#</b> # is not supported by the 80333.			
14:9	00H	Reserved			
8	02	PME_En - This bit is hardwired to read-only 0 <sub>2</sub> since this function does not support <b>PME#</b> generation from any power state.			
7:2	000000 <sub>2</sub>	Reserved			
1:0	00 <sub>2</sub>	Power State - This 2-bit field is used both to determine the current power state of a function and to set the function into a new power state. The definition of the values is: 00 <sub>2</sub> - D0 01 <sub>2</sub> - D1 10 <sub>2</sub> - D2 (Unsupported) 11 <sub>2</sub> - D3 <sub>hot</sub> The 80333 supports only the D0 and D3 <sub>hot</sub> states.			



## 3.10.61 MSI Capability Registers

The MSI capability registers are defined in the Messaging Unit. See "MSI Capability Identifier Register - MSI\_CAPID" on page 355, "MSI Next Item Pointer Register - MSI\_NXTP" on page 356, "MSI Message Control Register - MSI\_MCR" on page 357, "MSI Message Address Register - MSI\_MAR" on page 358, "MSI Message Upper Address Register - MSI\_MUAR" on page 359 and "MSI Message Data Register- MSI\_MDR" on page 360.

## intel

## 3.10.62 PCI-X Capability Identifier Register - PX\_CAPID

The Capability Identifier Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3. This register in the PCI Extended Capability header identifies the type of Extended Capability contained in that header. In the case of the 80333, this is the PCI-X extended capability with an ID of 07H as defined by the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a.

#### Table 185. PCI-X\_Capability Identifier Register - PX\_CAPID





## 3.10.63 PCI-X Next Item Pointer Register - PX\_NXTP

The Next Item Pointer Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.3. This register describes the location of the next item in the function's capability list.

By default, the PCI-X capability is the last capabilities list for the 80333, thus this register defaults to 00H.

However, this register may be written to B8H prior to host configuration to include the VPD capability located at off-set B8H.

*Warning:* Writing this register to any value other than 00H (default) or B8H is not supported and may produce unpredictable system behavior.

In order to guarantee that this register is written prior to host configuration, the 80333 must be initialized at **PWRGD** assertion to Retry Type 0 configuration cycles (bit 2 of PCSR). Typically, the Intel XScale<sup>®</sup> core would be enabled to boot immediately following **PWRGD** assertion in this case (bit 1 of PCSR), as well. Please see Table 168, "PCI Configuration and Status Register - PCSR" on page 283 for more details on the 80333 initialization modes.

#### Table 186. PCI-X Next Item Pointer Register - PX\_NXTP





## 3.10.64 PCI-X Command Register - PX\_CMD

This register controls various modes and features of ATU and Message Unit when operating in the PCI-X mode.

#### Table 187. PCI-X Command Register - PX\_CMD

		IOP     15     12       Attributes     rv     rv     rv       PCI     rv     rv     rv	8 4 0 v rv rw rw rw rw ro rw v rv rw rw rw rw rw ro rw
Internal Bus Address FFFF.E1E2H		PCI Configuration Offset E2- E3H	Attribute Legend:RW = Read/WriteRV = ReservedRC = Read ClearPR = PreservedRO = Read OnlyRS = Read/SetNA = Not Accessible
Bit	Default		Description
15:7	000000000 <sub>2</sub>	Reserved.	
6:4	011 <sub>2</sub>	Maximum Outstanding Split Transactions - Th the device is permitted to have outstanding atRegisterMaximum Outstanding0112233448512616732	is register sets the maximum number of Split Transactions one time.
3:2	00 <sub>2</sub>	Maximum Memory Read Byte Count - This reinitiating a Sequence with one of the burst me         Register       Maximum Byte Count         0       512         1       1024         2       2048         3       4096	gister sets the maximum byte count the device uses when mory read commands.
1	02	Enable Relaxed Ordering - The 80333 does r of Transactions.	not set the relaxed ordering bit in the Requester Attributes
0	02	Data Parity Error Recovery Enable - The devi recover from data parity errors. When this bit SERR# (when enabled) whenever the Master	ce driver sets this bit to enable the device to attempt to is 0 and the device is in PCI-X mode, the device asserts Data Parity Error bit (Status register, bit 8) is set.



## 3.10.65 PCI-X Status Register - PX\_SR

This register identifies the capabilities and current operating mode of ATU, DMAs and Message Unit when operating in the PCI-X mode.

*Note:* The PCI-X Status Register has meaning only in PCI-X mode. The device number and bus number fields are always updated when a configuration write is detected. The ATU always grabs AD[15:11] for configuration writes, whether it is in PCI or PCI-X mode. When a PCI configuration transaction occurs, the Device Number bits (7:3) are updated to a value of "00000" from the default value of "11111". The bus number bits (15:8) are only grabbed during the attribute phase (which does not exist for PCI).



Table 188. PCI-X Status Register - PX\_SR (Sheet 1 of 2)



Table 188.       PCI-X Status Register - PX_SR (Sheet 2 of 2)					
Attr Att Internal FFFF.E1 Bit	31       28       24       20       16       12       8       4       0         Attributes				
18	02	Split Completion Discarded - This bit is set when the device discards a Split Completion because the requester would not accept it. See Section 5.4.4 of the PCI-X Addendum to the PCI Local Bus Specification, Revision 1.0a for details. Once set, this bit remains set until software writes a 1 to this location.         0 = no Split Completion has been discarded.         1 = a Split Completion has been discarded.         NOTE: The 80333 does not set this bit since there is no Inbound address responding to Inbound Read Requests with Split Responses (Memory or Register) that has "read side effects."			
17	1 <sub>2</sub>	80333 is a 133 MHz capable device.			
16	1 <sub>2</sub>	The ATU of 80333 has a 64-bit interface on the secondary A-segment therefore this bit is always set 0 = The bus is 32 bits wide. 1 = The bus is 64 bits wide.			
15:8	FFH	Bus Number - This register is read for diagnostic purposes only. It indicates the number of the bus segment for the device containing this function. The function uses this number as part of its Requester ID and Completer ID. For all devices other than the source bridge, each time the function is addressed by a Configuration Write transaction, the function must update this register with the contents of AD[7::0] of the attribute phase of the Configuration Write, regardless of which register in the function is addressed by the transaction. The function is addressed by a Configuration Write transaction. The function is addressed by a Configuration Write transaction. The function is addressed by a Configuration Write transaction when all of the following are true: 1. The transaction uses a Configuration Write command. 2. IDSEL is asserted during the address phase. 3. AD[1::0] are 00b (Type 0 configuration transaction). 4. AD[10::08] of the configuration address contain the appropriate function number.			
7:3	<ul> <li>7:3</li> <li>1FH</li> <li>Device Number - This register is read for diagnostic purposes only. It indicates the number of the device containing this function, i.e., the number in the Device Number field (AD[15::11]) of the address of a Type 0 configuration transaction that is assigned to the device containing this function by the connection of the system markased by a the source number of the support of the source bridge). The function uses this number as part of its Requester ID and Completer ID. Each time the function is addressed by a Configuration Write transaction, the device must update this register with the contents of AD[15::11] of the address phase of the Configuration Write, regardless of which register in the function is addressed by the transaction. The function is addressed by a Configuration Write transaction when all of the following are true:         <ol> <li>The transaction uses a Configuration Write command.</li> <li>IDSEL is asserted during the address phase.</li> <li>AD[1::0] are 00b (Type 0 configuration transaction).</li> <li>AD[1::08] of the configuration address contain the appropriate function number.</li> </ol> </li> </ul>				
2:0	000 <sub>2</sub>	Function Number - This register is read for diagnostic purposes only. It indicates the number of this function; i.e., the number in the Function Number field (AD[10::08]) of the address of a Type 0 configuration transaction to which this function responds. The function uses this number as part of its Requester ID and Completer ID.			



## 3.10.66 PCI Interrupt Routing Select Register - PIRSR

The PCI Interrupt Routing Select Register (PIRSR) determines the routing of the external input pins. This register is addressed in the ATU register space, but is defined in Section 17.8.17, "PCI Interrupt Routing Select Register - PIRSR" on page 801. Refer to that page for functional details and programming information.



## 3.10.67 PCI-A Drive Strength Control Register - PADSCR

When necessary, the PCI-A drive strength control register is used to manually control the slew rate and drive strength of the PCI bus interface.

*Note:* By default, the user is **not** required to program this register.

 Table 189.
 PCI-A Drive Strength Control Register - PADSCR

At A Intel XS	10P tributes rw ro ro PCI na na na ccale® Core Local Bu	28 24 20 16 Tro ro r	12 8 ro/rw/rw/rw/rw/rw/ro/ro/rw/ na na na na na na na na na na Attribute Legend: RV = Reserved	4 0 rw rw rw rw rw na na na na na RW = Read/Write RC = Read Clear
	00011	NOT ACCESSIBLE	PR = Preserved I RS = Read/Set I	RO = Read Only NA = Not Accessible
Bit	Default	Description		
31	0	<b>Overdrive Enable</b> : When set, enables the PCI-A Bus interface drive strengths to be programmed with the values in this register.		
30:14	0 0000H	Reserved		
13:8	B 11 1111 Pull-up Drive Strength: When bit 31 is set, this field programs the strength of the n-drivers for the PCI-A Bus interface.			
7:6	00	Reserved		
5:0	11 1111 <b>Pull-Down Drive Strength</b> : When bit 31 is set, this field programs the strength of the n-drivers for the PCI-A Bus interface.			



## 3.10.68 PCI-A Drive Strength Value Register - PADSVR

When necessary, the PCI-A drive strength value register is used to read the drive strength value the RCOMP state machine calculates for the PCI bus interface.



#### Table 190. PCI-A Drive Strength Value Register - PADSVR



## 3.10.69 PCI-B Drive Strength Control Register - PBDSCR

When necessary, the PCI-B drive strength control register is used to manually control the slew rate and drive strength of the PCI bus interface.

*Note:* By default, the user is **not** required to program this register.

 Table 191.
 PCI-B Drive Strength Control Register - PBDSCR

At A Intel XS FFFF F5	10P tributes rw ro ro PCI na na na cale <sup>®</sup> Core Local Bu 5D0H	28 24 20 ro ro r	16 12 8 o ro ro rw rw rw rw rw ro ro a na na Attribute Legend: RV = Reserved PR = Preserved	4 0 rw rw rw rw rw rw na na na na na na RW = Read/Write RC = Read Clear RO = Read Only	
	1		RS = Read/Set	NA = Not Accessible	
Bit	Default		Description		
31	0	<b>Overdrive Enable</b> : When set, enables the PCI-B Bus interface drive strengths to be programmed with the values in this register.			
30:14	000 0000H	Reserved			
13:8	3 11 1111 <b>Pull-up Drive Strength</b> : When bit 31 is set, this field programs the strength of the n-drivers for the PCI-B Bus interface.				
7:6	00	00 Reserved			
5:0	11 1111	Pull-Down Drive Strength: When bit 31 is set, this field programs the strength of the n-drivers for the PCI-B Bus interface.			



## 3.10.70 PCI-B Drive Strength Value Register - PBDSVR

When necessary, the PCI-B drive strength value register is used to read the drive strength value the RCOMP state machine calculates for the PCI bus interface.



#### Table 192. PCI-B Drive Strength Value Register - PBDSVR



## 4

This chapter describes the Messaging Unit (MU) of the Intel<sup>®</sup> 80333 I/O processor (80333). The MU is closely related to the Address Translation Unit (ATU) described in Chapter 3, "Address Translation Unit".

## 4.1 **Overview**

The Messaging Unit (MU) provides a mechanism for data to be transferred between the PCI bus and the Intel XScale<sup>®</sup> core and notifying the respective system of the arrival of new data through an interrupt. The MU can be used to send and receive messages.

The MU has four distinct messaging mechanisms. Each allows a host processor or external PCI agent and the 80333 to communicate through message passing and interrupt generation. The four mechanisms are:

- Message Registers allow the 80333 and external PCI agents to communicate by passing messages in one of four 32-bit Message Registers. In this context, a message is any 32-bit data value. Message registers combine aspects of mailbox registers and doorbell registers. Writes to the message registers may optionally cause interrupts.
- **Doorbell Registers** allow the 80333 to assert the PCI interrupt signal and allow external PCI agents to generate an interrupt to the Intel XScale<sup>®</sup> core.
- Circular Queues support a message passing scheme that uses four circular queues.
- **Index Registers** support a message passing scheme that uses a portion of the 80333 local memory to implement a large set of message registers.

Each of the above are available to the system designer at the same time. No special mode selection is needed.

## 4.2 Theory of Operation

The MU has four independent messaging mechanisms.

There are four Message Registers that are similar to a combination of mailbox and doorbell registers. Each holds a 32-bit value and generates an interrupt when written.

The two Doorbell Registers support software interrupts. When a bit is set in a Doorbell Register, an interrupt is generated.

The Circular Queues support a message passing scheme that uses 4 circular queues. The 4 circular queues are implemented in 80333 local memory. Two queues are used for inbound messages and two are used for outbound messages. Interrupts may be generated when the queue is written.

The Index Registers use a portion of the 80333 local memory to implement a large set of message registers. When one of the Index Registers is written, an interrupt is generated and the address of the register written is captured.

Interrupt status for all interrupts is recorded in the Inbound Interrupt Status Register and the Outbound Interrupt Status Register. Each interrupt generated by the Messaging Unit can be masked.

Multi-DWORD PCI burst accesses are not supported by the Messaging Unit, with the exception of Multi-DWORD reads to the index registers. In Conventional mode, the MU terminates Multi-DWORD PCI transactions (other than index register reads) with a disconnect at the next Qword boundary, with the exception of queue ports. In PCI-X mode, the MU terminates a Multi-DWORD PCI read transaction with a Split Response and the data is returned through split completion transaction(s); however, when the burst request crosses into or through the range of offsets 40h to 4Ch (e.g., this includes the queue ports) the transaction is signaled target-abort immediately on the PCI bus. In PCI-X mode, Multi-DWORD PCI writes is signaled a Single-Data-Phase Disconnect which means that no data beyond the first Qword (Dword when the MU does not assert **P\_ACK64#**) is written.

All registers needed to configure and control the Messaging Unit are memory-mapped registers.

The MU uses the first 4 Kbytes of the inbound translation window in the Address Translation Unit (ATU). This PCI address window is used for PCI transactions that access the 80333 local memory. The PCI address of the inbound translation window is contained in the Inbound ATU Base Address Register. See Chapter 3, "Address Translation Unit" for more details on inbound ATU addressing and the ATU.

From the PCI perspective, the Messaging Unit is part of the Address Translation Unit. The Messaging Unit uses the PCI configuration registers of the ATU for control and status information. The Messaging Unit must observe all PCI control bits in the ATU Command Register and ATU Configuration Register. The Messaging Unit reports all PCI errors in the ATU Status Register.

Parts of the Messaging Unit can be accessed as a 64-bit PCI device. The register interface, message registers, doorbell registers, and index registers returns a **P\_ACK64#** in response to a **P\_REQ64#** on the PCI interface. Up to 1 Qword of data can be read or written per transaction (except Index Register reads, see Section 4.6, "Index Registers" on page 332). The Inbound and Outbound Queue Ports are always 32-bit addresses and the MU does not assert **P\_ACK64#** to offsets 40H and 44H.



#### Figure 23. PCI Memory Map

First 4	Kbytes of the ATU Inbound PCI Add	Ire	ss Space
0000H	Reserved	1	
0004H	Reserved	1	
0008H	Reserved		
000CH	Reserved		
0010H	Inbound Message Register 0	-	
0014H	Inbound Message Register 1		4 Magazara Registera
0018H	Outbound Message Register 0		4 message Registers
001CH	Outbound Message Register 1		
0020H	Inbound Doorbell Register		
0024H	Inbound Interrupt Status Register		
0028H	Inbound Interrupt Mask Register		2 Doorbell Registers and
002CH	Outbound Doorbell Register		4 Interrupt Registers
0030H	Outbound Interrupt Status Register		
0034H	Outbound Interrupt Mask Register	_	
0038H	Reserved		
003CH	Reserved		
0040H	Inbound Queue Port		2 Queue Ports
0044H	Outbound Queue Port		
0048H	Reserved		
004CH	Reserved		
0050H	Intel Xscale <sup>®</sup> Microarchitecture Local Memory		1004 Index Registers
0FFCH		]_	

Note: See Table 199, "Message Unit Register".

FFFF E300H

intel

#### Figure 24. Internal Bus Memory Map

FFFF E310H	Inbound Message Register 0
FFFF E314H	Inbound Message Register 1
FFFF E318H	Outbound Message Register 0
FFFF E31CH	Outbound Message Register 1
FFFF E320H	Inbound Doorbell Register
FFFF E324H	Inbound Interrupt Status Register
FFFF E328H	Inbound Interrupt Mask Register
FFFF E32CH	Outbound Doorbell Register
FFFF E330H	Outbound Interrupt Status Register
FFFF E334H	Outbound Interrupt Mask Register
FFFF E338H	reserved
FFFF E33CH	reserved
FFFF E340H	reserved
FFFF E344H	reserved
FFFF E348H	reserved
FFFF E34CH	reserved
FFFF E350H	MU Configuration Register
FFFF E354H	Queue Base Address Register
FFFF E358H	reserved
FFFF E35CH	reserved
FFFF E360H	Inbound Free Head Pointer Register
FFFF E364H	Inbound Free Tail Pointer Register
FFFF E368H	Inbound Post Head pointer Register
FFFF E36CH	Inbound Post Tail Pointer Register
FFFF E370H	Outbound Free Head Pointer Register
FFFF E374H	Outbound Free Tail Pointer Register
FFFF E378H	Outbound Post Head pointer Register
FFFF E37CH	Outbound Post Tail Pointer Register
FFFF E380H	Index Address Register



Table 193 provides a summary of the four messaging mechanisms used in the Messaging Unit.

#### Table 193. MU Summary

Mechanism	Quantity	Assert PCI Interrupt Signals	Generate I/O Processor Interrupt
Message Registers	2 Inbound 2 Outbound	Optional	Optional
Doorbell Registers	1 Inbound 1 Outbound	Optional	Optional
Circular Queues	4 Circular Queues	Under certain conditions	Under certain conditions
Index Registers	1004 32-bit Memory Locations	No	Optional

## 4.2.1 Transaction Ordering

From a PCI standpoint, the Messaging Unit is a piece of the ATU and therefore must maintain ordering requirements against ATU transactions. Transaction ordering is achieved for the Index Registers, the Doorbell Register, and the Message Registers since these transactions are routed through the standard set of ATU read/write queues.

The Circular Queues (Inbound/Outbound Queue Port) are separate queue structures and therefore require ordering. The Inbound Post Queue (contains PCI writes) must be ordered against the inbound write queue of the ATU to allow the data that is represented by the Inbound Post interrupt to be written to local memory before the interrupt is delivered. See Table 194 for a summary of Messaging Unit transaction ordering.

#### Table 194. Circular Queue Ordering Requirements

Messaging	Unit Feature	Transaction Ordering Mechanism	
Message	Registers	Through ATU Queues	
Doorbell Registers Index Registers		Through ATU Queues	
		Through ATU Queues	
	Inbound Post	Ordered Against ATU Inbound Write Queue (PMW Can Not Pass Another PMW)	
Circular Queues	Inbound Free	No Specific Hardware Ordering	
	Outbound Post	No Specific Hardware Ordering	
	Outbound Free	No Specific Hardware Ordering	



## 4.3 Message Registers

Messages can be sent and received by the 80333 through the use of the Message Registers. When written, the message registers may cause an interrupt to be generated to either the Intel XScale<sup>®</sup> core or the host processor. Inbound messages are sent by the host processor and received by the 80333. Outbound messages are sent by the 80333 and received by the host processor.

The interrupt status for outbound messages is recorded in the Outbound Interrupt Status Register. Interrupt status for inbound messages is recorded in the Inbound Interrupt Status Register.

## 4.3.1 Outbound Messages

When an outbound message register is written by the Intel XScale<sup>®</sup> core, an interrupt may be generated through the IOAPIC **IRQ14**# interrupt pin or a message signaled interrupt is generated when MSI is enabled.

The PCI interrupt is recorded in the Outbound Interrupt Status Register. The interrupt causes the Outbound Message Interrupt bit to be set in the Outbound Interrupt Status Register. This is a Read/Clear bit that is set by the MU hardware and cleared by software.

The interrupt is cleared when an external PCI agent writes a value of 1 to the Outbound Message Interrupt bit in the Outbound Interrupt Status Register to clear the bit.

The interrupt may be masked by the mask bits in the Outbound Interrupt Mask Register.

### 4.3.2 Inbound Messages

When an inbound message register is written by an external PCI agent, an interrupt may be generated to the Intel XScale<sup>®</sup> core. The interrupt may be masked by the mask bits in the Inbound Interrupt Mask Register.

The Intel XScale<sup>®</sup> core interrupt is recorded in the Inbound Interrupt Status Register. The interrupt causes the Inbound Message Interrupt bit to be set in the Inbound Interrupt Status Register. This is a Read/Clear bit that is set by the MU hardware and cleared by software.

The interrupt is cleared when the Intel XScale<sup>®</sup> core writes a value of 1 to the Inbound Message Interrupt bit in the Inbound Interrupt Status Register.

# intel

## 4.4 Doorbell Registers

There are two Doorbell Registers: the Inbound Doorbell Register and the Outbound Doorbell Register. The Inbound Doorbell Register allows external PCI agents to generate interrupts to the Intel XScale<sup>®</sup> core. The Outbound Doorbell Register allows the Intel XScale<sup>®</sup> core to generate a PCI interrupt. Both Doorbell Registers may generate interrupts whenever a bit in the register is set.

## 4.4.1 Outbound Doorbells

When the Outbound Doorbell Register is written by the Intel XScale<sup>®</sup> core, an interrupt may be generated through the IOAPIC **IRQ14**# or a message signaled interrupt is generated when MSI is enabled. An interrupt is generated when any of the bits in the doorbell register is written to a value of 1. Writing a value of 0 to any bit does not change the value of that bit and does not cause an interrupt to be generated. Once a bit is set in the Outbound Doorbell Register, it cannot be cleared by the Intel XScale<sup>®</sup> core.

The interrupt is recorded in the Outbound Interrupt Status Register.

The interrupt may be masked by the mask bits in the Outbound Interrupt Mask Register. When the mask bit is set for a particular bit, no interrupt is generated for that bit. The Outbound Interrupt Mask Register affects only the generation of the interrupt and not the values written to the Outbound Doorbell Register.

The interrupt is cleared when an external PCI agent writes a value of 1 to the bits in the Outbound Doorbell Register that are set. Writing a value of 0 to any bit does not change the value of that bit and does not clear the interrupt.

In summary, the Intel XScale<sup>®</sup> core generates an interrupt and external PCI agents clear the interrupt by setting bits in the Outbound Doorbell Register.

## 4.4.2 Inbound Doorbells

When the Inbound Doorbell Register is written by an external PCI agent, an interrupt may be generated to the Intel XScale<sup>®</sup> core. An interrupt is generated when any of the bits in the doorbell register is written to a value of 1. Writing a value of 0 to any bit does not change the value of that bit and does not cause an interrupt to be generated. Once a bit is set in the Inbound Doorbell Register, it cannot be cleared by any external PCI agent. The interrupt is recorded in the Inbound Interrupt Status Register.

The interrupt may be masked by the Inbound Doorbell Interrupt mask bit in the Inbound Interrupt Mask Register. When the mask bit is set for a particular bit, no interrupt is generated for that bit. The Inbound Interrupt Mask Register affects only the generation of the normal messaging unit interrupt and not the values written to the Inbound Doorbell Register. One bit in the Inbound Doorbell Register is reserved for an Error Doorbell interrupt.

The interrupt is cleared when the Intel XScale<sup>®</sup> core writes a value of 1 to the bits in the Inbound Doorbell Register that are set. Writing a value of 0 to any bit does not change the value of that bit and does not clear the interrupt.

## 4.5 Circular Queues

The MU implements four circular queues. There are 2 inbound queues and 2 outbound queues. In this case, inbound and outbound refer to the direction of the flow of posted messages.

Inbound messages are either:

- posted messages by other processors for the Intel XScale<sup>®</sup> core to process or
- *free* (or empty) messages that can be reused by other processors.

Outbound messages are either:

- *posted* messages by the Intel XScale<sup>®</sup> core for other processors to process or
- *free* (or empty) messages that can be reused by the Intel XScale<sup>®</sup> core.

Therefore, free inbound messages flow away from the and free outbound messages flow toward the 80333.

The four Circular Queues are used to pass messages in the following manner. The two inbound queues are used to handle inbound messages and the two outbound queues are used to handle outbound messages. One of the inbound queues is designated the Free queue and it contains inbound free messages. The other inbound queue is designated the Post queue and it contains inbound posted messages. Similarly, one of the outbound queues is designated the Free queue and the other outbound queue is designated the Post queue. Table 195 contains a summary of the queues.

#### Table 195.Circular Queue Summary

Queue Name	Purpose	Action on PCI Interface
Inbound Post Queue	Queue for inbound messages from other processors waiting to be processed by the 80333	Written
Inbound Free Queue	Queue for empty inbound messages from the 80333 available for use by other processors	Read
Outbound Post Queue	Queue for outbound messages from the 80333 that are being posted to the other processors	Read
Outbound Free Queue	Queue for empty outbound messages from other processors available for use by the 80333	Written

The two outbound queues allow the Intel XScale<sup>®</sup> core to post outbound messages in one queue and to receive free messages returning from the host processor. The Intel XScale<sup>®</sup> core posts outbound messages, the host processor receives the posted message and when it is finished with the message, places it back on the outbound free queue for reuse by the Intel XScale<sup>®</sup> core.

The two inbound queues allow the host processor to post inbound messages for the 80333 in one queue and to receive free messages returning from the 80333. The host processor posts inbound messages, the Intel XScale<sup>®</sup> core receives the posted message and when it is finished with the message, places it back on the inbound free queue for reuse by the host processor.

Figure 25 provides an overview of the Circular Queue operation.

The circular queues are accessed by external PCI agents through two port locations in the PCI address space: Inbound Queue Port and Outbound Queue Port. The Inbound Queue Port is used by external PCI agents to read the Inbound Free Queue and write the Inbound Post Queue. The
Outbound Queue Port is used by external PCI agents to read the Outbound Post Queue and write the Outbound Free Queue. Note that a PCI transaction to the inbound or outbound queue ports with null byte enables ( $P_C/BE[3:0]$ # = 1111<sub>2</sub>) does not cause the MU hardware to increment the queue pointers. This is treated as when the PCI transaction did not occur. The Inbound and Outbound Queue Ports never respond with  $P_ACK64$ # on the PCI interface.

Figure 25. Overview of Circular Queue Operation



The data storage for the circular queues must be provided by the 80333 local memory. The base address of the circular queues is contained in the Queue Base Address Register (Section 4.9.10, "Queue Base Address Register - QBAR" on page 345). Each entry in the queue is a 32-bit data value. Each read from or write to the queue may access only one queue entry. Multi-DWORD accesses to the circular queues are not allowed. Sub-DWORD accesses are promoted to DWORD accesses.

Each circular queue has a head pointer and a tail pointer. The pointers are offsets from the Queue Base Address. Writes to a queue occur at the head of the queue and reads occur from the tail. The head and tail pointers are incremented by either the Intel XScale<sup>®</sup> core or the Messaging Unit hardware. Which unit maintains the pointer is determined by the writer of the queue. More details about the pointers are given in the queue descriptions below. The pointers are incremented after the queue access. Both pointers wrap around to the first address of the circular queue when they reach the circular queue size.



The Messaging Unit generates an interrupt to the Intel XScale<sup>®</sup> core or generate a PCI interrupt under certain conditions. In general, when a Post queue is written, an interrupt is generated to notify the receiver that a message was posted.

The size of each circular queue can range from 4K entries (16 Kbytes) to 64K entries (256 Kbytes). All four queues must be the same size and may be contiguous. Therefore, the total amount of local memory needed by the circular queues ranges from 64 Kbytes to 1 Mbytes. The Queue size is determined by the Queue Size field in the MU Configuration Register.

There is one base address for all four queues. It is stored in the Queue Base Address Register (QBAR). The starting addresses of each queue is based on the Queue Base Address and the Queue Size field. Table 196 shows an example of how the circular queues should be set up based on the *Intelligent I/O* ( $I_2O$ ) Architecture Specification. Other ordering of the circular queues is possible.

#### Table 196. Queue Starting Addresses

Queue	Starting Address	
Inbound Free Queue	QBAR	
Inbound Post Queue	QBAR + Queue Size	
Outbound Post Queue	QBAR + 2 * Queue Size	
Outbound Free Queue	QBAR + 3 * Queue Size	



Figure 26 provides a more detailed diagram of the usage of the Circular Queues.

Figure 26. Circular Queue Operation



March 2005



#### 4.5.1 Inbound Free Queue

The Inbound Free Queue holds free inbound messages placed there by the Intel XScale<sup>®</sup> core for other processors to use. This queue is read from the queue tail by external PCI agents. It is written to the queue head by the Intel XScale<sup>®</sup> core. The tail pointer is maintained by the MU hardware. The head pointer is maintained by the Intel XScale<sup>®</sup> core.

For a PCI read transaction that accesses the Inbound Queue Port, the MU attempts to read the data at the local memory address in the Inbound Free Tail Pointer. When the queue is not empty (head and tail pointers are not equal) or full (head and tail pointers are equal but the head pointer was last written by software), the data is returned. When the queue is empty (head and tail pointers are equal and the tail pointer was last updated by hardware), the value of -1 (FFFF.FFFFH) is returned. When the queue was not empty and the MU succeeded in returning the data at the tail, the MU hardware must increment the value in the Inbound Free Tail Pointer Register.

To reduce latency for the PCI read access, the MU implements a prefetch mechanism to anticipate accesses to the Inbound Free Queue. The MU hardware prefetches the data at the tail of the Inbound Free Queue and load it into an internal prefetch register. When the PCI read access occurs, the data is read directly from the prefetch register.

The prefetch mechanism loads a value of -1 (FFFF.FFFFH) into the prefetch register when the head and tail pointers are equal and the queue is empty. In order to update the prefetch register when messages are added to the queue and it becomes non-empty, the prefetch mechanism automatically starts a prefetch when the prefetch register contains FFFF.FFFFH and the Inbound Free Head Pointer Register is written. The Intel XScale<sup>®</sup> core needs to update the Inbound Free Head Pointer Register when it adds messages to the queue.

A prefetch must appear atomic from the perspective of the external PCI agent. When a prefetch is started, any PCI transaction that attempts to access the Inbound Free Queue is signalled a Retry until the prefetch is completed.

The Intel XScale<sup>®</sup> core may place messages in the Inbound Free Queue by writing the data to the local memory location pointed to by the Inbound Free Head Pointer Register. The processor must then increment the Inbound Free Head Pointer Register.

## 4.5.2 Inbound Post Queue

The Inbound Post Queue holds posted messages placed there by other processors for the Intel XScale<sup>®</sup> core to process. This queue is read from the queue tail by the Intel XScale<sup>®</sup> core. It is written to the queue head by external PCI agents. The tail pointer is maintained by the Intel XScale<sup>®</sup> core. The head pointer is maintained by the MU hardware.

For a PCI write transaction that accesses the Inbound Queue Port, the MU writes the data to the local memory location address in the Inbound Post Head Pointer Register.

When the data written to the Inbound Queue Port is written to local memory, the MU hardware increments the Inbound Post Head Pointer Register.

An Intel XScale<sup>®</sup> core interrupt may be generated when the Inbound Post Queue is written. The Inbound Post Queue Interrupt bit in the Inbound Interrupt Status Register indicates the interrupt status. The interrupt is cleared when the Inbound Post Queue Interrupt bit is cleared. The interrupt can be masked by the Inbound Interrupt Mask Register. Software must be aware of the state of the Inbound Post Queue Interrupt Mask bit to guarantee that the full condition is recognized by the core processor. In addition, to guarantee that the queue does not get overwritten, software must process messages from the tail of the queue before incrementing the tail pointer and clearing this interrupt. Once cleared, an interrupt is NOT generated when the head and tail pointers remain unequal (i.e. queue status is Not Empty). Only a new message posting the in the inbound queue generates a new interrupt. Therefore, when software leaves any unprocessed messages in the post queue status.

From the time that the PCI write transaction is received until the data is written in local memory and the Inbound Post Head Pointer Register is incremented, any PCI transaction that attempts to access the Inbound Post Queue Port is signalled a Retry.

The Intel XScale<sup>®</sup> core may read messages from the Inbound Post Queue by reading the data from the local memory location pointed to by the Inbound Post Tail Pointer Register. The Intel XScale<sup>®</sup> core must then increment the Inbound Post Tail Pointer Register. When the Inbound Post Queue is full (head and tail pointers are equal and the head pointer was last updated by hardware), the hardware retries any PCI writes until a slot in the queue becomes available. A slot in the post queue becomes available by the Intel XScale<sup>®</sup> core incrementing the tail pointer.



#### 4.5.3 Outbound Post Queue

The Outbound Post Queue holds outbound posted messages placed there by the Intel XScale<sup>®</sup> core for other processors to process. This queue is read from the queue tail by external PCI agents. It is written to the queue head by the Intel XScale<sup>®</sup> core. The tail pointer is maintained by the MU hardware. The head pointer is maintained by the Intel XScale<sup>®</sup> core.

For a PCI read transaction that accesses the Outbound Queue Port, the MU attempts to read the data at the local memory address in the Outbound Post Tail Pointer Register. When the queue is not empty (head and tail pointers are not equal) or full (head and tail pointers are equal but the head pointer was last written by software), the data is returned. When the queue is empty (head and tail pointers are equal and the tail pointer was last updated by hardware), the value of -1 (FFFF.FFFFH) is returned. When the queue was not empty and the MU succeeded in returning the data at the tail, the MU hardware must increment the value in the Outbound Post Tail Pointer Register.

To reduce latency for the PCI read access, the MU implements a prefetch mechanism to anticipate accesses to the Outbound Post Queue. The MU hardware prefetches the data at the tail of the Outbound Post Queue and load it into an internal prefetch register. When the PCI read access occurs, the data is read directly from the prefetch register.

The prefetch mechanism loads a value of -1 (FFFF.FFFFH) into the prefetch register when the head and tail pointers are equal and the queue is empty. In order to update the prefetch register when messages are added to the queue and it becomes non-empty, the prefetch mechanism automatically starts a prefetch when the prefetch register contains FFFF.FFFFH and the Outbound Post Head Pointer Register is written. The Intel XScale<sup>®</sup> core needs to update the Outbound Post Head Pointer Register when it adds messages to the queue.

A prefetch must appear atomic from the perspective of the external PCI agent. When a prefetch is started, any PCI transaction that attempts to access the Outbound Post Queue is signalled a Retry until the prefetch is completed.

A PCI interrupt may be generated when data in the prefetch buffer is valid. When the prefetch queue is clear, no interrupt is generated. The Outbound Post Queue Interrupt bit in the Outbound Interrupt Status Register shall indicate the status of the prefetch buffer data and therefore the interrupt status. The interrupt is cleared when any prefetched data has been read from the Outbound Queue Port. The interrupt can be masked by the Outbound Interrupt Mask Register.

The Intel XScale<sup>®</sup> core may place messages in the Outbound Post Queue by writing the data to the local memory address in the Outbound Post Head Pointer Register. The processor must then increment the Outbound Post Head Pointer Register.

## 4.5.4 Outbound Free Queue

The Outbound Free Queue holds free messages placed there by other processors for the Intel XScale<sup>®</sup> core to use. This queue is read from the queue tail by the Intel XScale<sup>®</sup> core. It is written to the queue head by external PCI agents. The tail pointer is maintained by the Intel XScale<sup>®</sup> core. The head pointer is maintained by the MU hardware.

For a PCI write transaction that accesses the Outbound Queue Port, the MU writes the data to the local memory address in the Outbound Free Head Pointer Register. When the data written to the Outbound Queue Port is written to local memory, the MU hardware increments the Outbound Free Head Pointer Register.

When the head pointer and the tail pointer become equal and the queue is full, the MU may signal an interrupt to the Intel XScale<sup>®</sup> core to register the queue full condition. This interrupt is recorded in the Inbound Interrupt Status Register. The interrupt is cleared when the Outbound Free Queue Full Interrupt bit is cleared and not by writing to the head or tail pointers. The interrupt can be masked by the Inbound Interrupt Mask Register. Software must be aware of the state of the Outbound Free Queue Interrupt Mask bit to guarantee that the full condition is recognized by the core processor.

From the time that a PCI write transaction is received until the data is written in local memory and the Outbound Free Head Pointer Register is incremented, any PCI transaction that attempts to access the Outbound Free Queue Port is signalled a retry.

The Intel XScale<sup>®</sup> core may read messages from the Outbound Free Queue by reading the data from the local memory address in the Outbound Free Tail Pointer Register. The processor must then increment the Outbound Free Tail Pointer Register. When the Outbound Free Queue is full, the hardware must retry any PCI writes until a slot in the queue becomes available.

Queue Name	PCI Port	Generate PCI Interrupt	Generate Intel Xscale <sup>®</sup> Core Interrupt	Head Pointer maintained by	Tail Pointer maintained by
Inbound Post Queue	Inbound Queue	No	Yes, when queue is written	MU hardware	Intel XScale <sup>®</sup> microarchitecture
Inbound Free Queue	Port	No	No	Intel XScale <sup>®</sup> microarchitecture	MU hardware
Outbound Post Queue	Outbound Queue	Yes, when prefetch buffer data is valid	No	Intel XScale <sup>®</sup> microarchitecture	MU hardware
Outbound Free Queue	Port	No	Yes, when queue is full	MU hardware	Intel XScale <sup>®</sup> microarchitecture

#### Table 197. Circular Queue Summary

#### Table 198. Circular Queue Status Summary

Queue Name	Queue Status	Head & Tail Pointer	Last Pointer Update
Inbound Post Queue	Empty	Equal	Tail pointer last updated by software
&	Not Empty	Not Equal	N/A
Outbound Free Queue	Full	Equal	Head pointer last updated by hardware
Inbound Free Queue	Empty	Equal	Tail pointer last updated by hardware
&	Not Empty	Not Equal	
Outbound Post Queue	Full	Equal	Head pointer last updated by software

## 4.6 Index Registers

The Index Registers are a set of 1004 registers that when written by an external PCI agent can generate an interrupt to the Intel XScale<sup>®</sup> core. These registers are for inbound messages only. The interrupt is recorded in the Inbound Interrupt Status Register.

The storage for the Index Registers is allocated from the 80333 local memory. PCI write accesses to the Index Registers write the data to local memory. PCI read accesses to the Index Registers read the data from local memory. The local memory used for the Index Registers ranges from Inbound ATU Translate Value Register + 050H to Inbound ATU Translate Value Register + FFFH. Chapter 3, "Address Translation Unit" describes how PCI addresses are translated to local memory addresses.

The address of the first write access is stored in the Index Address Register. This register is written during the earliest write access and provides a means to determine which Index Register was written. Once updated by the MU, the Index Address Register is not updated until the Index Register Interrupt bit in the Inbound Interrupt Status Register is cleared. When the interrupt is cleared, the Index Address Register is re-enabled and stores the address of the next Index Register write access.

Writes by the Intel XScale<sup>®</sup> core to the local memory used by the Index Registers does not cause an interrupt and does not update the Index Address Register.

The index registers can be accessed with Multi-DWORD reads and single QWORD aligned writes.

## 4.7 Messaging Unit Error Conditions

The Messaging Unit, like the ATU, encounters error conditions on the PCI interface as well as the internal bus interface. As a PCI target, all PCI errors (parity and aborts) are captured and recorded in the ATU Status Register and can be masked using the ATU mechanisms. Refer to Chapter 3, "Address Translation Unit" for further details.

## 4.8 Message-Signaled Interrupts

The Messaging Unit is responsible for the generation of all of the Outbound Interrupts from the 80333. These interrupts can be delivered to the Host Processor via the **IRQ14**# IOAPIC pin or the Message Signaled Interrupt (MSI) mechanism.

When a host processor enables Message-Signaled Interrupts (MSI) on the 80333, an outbound interrupt is signaled to the host via a PCI write instead of the assertion of the **IRQ14**# IOAPIC pin.

The *PCI-X* Addendum to the *PCI* Local Bus Specification, Revision 1.0a states that "PCI-X devices that generate interrupts are required to support message-signaled interrupts, as defined by the *PCI* Local Bus Specification, Revision 2.2 and must support a 64-bit message address." "Devices that require interrupts in systems that do not support message-signaled interrupts, must implement interrupt pins." Thus, the 80333 needs to implement both wired and message-signaled interrupt delivery mechanisms.

In support of MSI, the 80333 implements the MSI capability structure. The capability structure includes the "MSI Message Control Register - MSI\_MCR" on page 357, the "MSI Message Address Register - MSI\_MAR" on page 358, the "MSI Message Upper Address Register - MSI\_MUAR" on page 359 and the "MSI Message Data Register - MSI\_MDR" on page 360.

During system initialization, the configuration software for an MSI system reads the Message Control Register to determine that the 80333 supports a 64-bit Message Address, and that it is capable of generating two unique interrupt messages.

After gathering this data from all of the MSI capable devices in the system, the configuration software decides where to initialize the Message Address and how many unique messages each MSI capable device is allowed. Then, software writes the Message Address Registers (and the Message Upper Address Registers when Message Address is above the 4G address boundary<sup>1</sup>), and the Message Data Register. This system specified data is used to route the interrupt request message to the appropriate entry in a host processor Local APIC table.

Configuration of MSI completes with a write to the Message Control Register which includes an update to the Multiple Message Enable field and the MSI enable bit of each device. This informs the device how many unique messages (Local APIC table entries) have been allocated for exclusive use by that device and enable that device for MSI. Device hardware is required to handle allocation of fewer unique interrupt messages than requested by the Multiple Message Capable field.

80333 is able to handle generating only one message, even though the device is capable of generating two unique messages. When two unique messages are enabled, one message is reserved for Outbound Post Queue Interrupt and the other message represents all of Outbound Doorbell and Outbound Message Interrupts. When only one message is enabled, all interrupts are represented by a single message. Interrupt handler software needs to read 80333 Outbound Interrupt Status Register to determine the interrupt cause when more than one source is represented by a single message.

To signal an Outbound Interrupt with MSI enabled, the 80333 creates an outbound write transaction using the Message Address and the Message Data. When two unique messages are enabled, the lowest order bit of the Message Data is modified by hardware so that the host processor can distinguish between them.

*Note:* When host software enables MSI, a Messaging Unit Interrupt does not result in the assertion of the **IRQ14#** IOAPIC pin. However, all the **XINT[7:0]#** pins is functional for steering of interrupts from other PCI devices that may not be MSI capable.

<sup>1.</sup> When host software writes the Message Upper address register to a non-zero value, device hardware uses a write transaction with a Dual Address Cycle (DAC) to present the full 64-bit address to the bus.



### 4.8.1 Level-Triggered Versus Edge-Triggered Interrupts

When MSI is disabled, the **IRQ14**# IOAPIC pin remains asserted and pended to the host when **any** of the MU interrupt sources requires service (Outbound Post Queue, Outbound Doorbell and Outbound Message). Since the PCI pin signaled interrupt is **level-triggered**, the interrupt service routine does not drop out of the service routine until the interrupt signal is deasserted. This guarantees that an interrupt is not missed.

MSI interrupts are inherently edge-triggered, in that an interrupt is only pended to the host as a write event when any of the MU interrupt sources requires service. The MU generates a MSI message when a interrupt status bit in the OISR is set. Additional MSI messages is generated by the MU when other bits are set in the OISR, that are previously clear.

*Note:* The MU does NOT generate a MSI message when one pending interrupt is cleared in the OISR and other pending interrupt status bits remain set. Interrupt service routines must service and clear all pending interrupts from the OISR before exiting to insure no interrupts are missed.

## 4.9 **Register Definitions**

The following registers are located in the PCI address space and in the Peripheral Memory-Mapped Register (PMMR) address space. They are accessible through PCI bus transactions and through Intel XScale<sup>®</sup> core internal bus accesses. In the PCI address space, they are mapped into the first 80 bytes of the inbound address window of the ATU.

- Inbound Message 0 Register
- Inbound Message 1 Register
- Outbound Message 0 Register
- Outbound Message 1 Register
- Inbound Doorbell Register

- Inbound Interrupt Status Register
- Inbound Interrupt Mask Register
- Outbound Doorbell Register
- Outbound Interrupt Status Register
- Outbound Interrupt Mask Register

The following registers are located in the Peripheral Memory-Mapped Register (PMMR) address space as described in Chapter 19, "Peripheral Registers".

- MU Configuration Register
- Queue Base Address Register
- Inbound Free Head Pointer Register
- Inbound Free Tail Pointer Register
- Inbound Post Head Pointer Register
- Inbound Post Tail Pointer Register
- Outbound Free Head Pointer Register
- Outbound Free Tail Pointer Register
- Outbound Post Head Pointer Register
- Outbound Post Tail Pointer Register
- Index Address Register

Reading or writing a register that is reserved is undefined.

#### Table 199.Message Unit Register

Internal Bus Address	Section, Register Name - Acronym (Page)		
FFFF.E310H	Section 4.9.1, "Inbound Message Register - IMRx" on page 336		
FFFF.E314H	Section 4.9.1, "Inbound Message Register - IMRx" on page 336		
FFFF.E318H	Section 4.9.2, "Outbound Message Register - OMRx" on page 337		
FFFF.E31CH	Section 4.9.2, "Outbound Message Register - OMRx" on page 337		
FFFF.E320H	Section 4.9.3, "Inbound Doorbell Register - IDR" on page 338		
FFFF.E324H	Section 4.9.4, "Inbound Interrupt Status Register - IISR" on page 339		
FFFF.E328H	Section 4.9.5, "Inbound Interrupt Mask Register - IIMR" on page 340		
FFFF.E32CH	Section 4.9.6, "Outbound Doorbell Register - ODR" on page 341		
FFFF.E330H	Section 4.9.7, "Outbound Interrupt Status Register - OISR" on page 342		
FFFF.E334H	Section 4.9.8, "Outbound Interrupt Mask Register - OIMR" on page 343		
FFFF.E350H	Section 4.9.9, "MU Configuration Register - MUCR" on page 344		
FFFF.E354H	Section 4.9.10, "Queue Base Address Register - QBAR" on page 345		
FFFF.E360H	Section 4.9.11, "Inbound Free Head Pointer Register - IFHPR" on page 346		
FFFF.E364H	Section 4.9.12, "Inbound Free Tail Pointer Register - IFTPR" on page 347		
FFFF.E368H	Section 4.9.13, "Inbound Post Head Pointer Register - IPHPR" on page 348		
FFFF.E36CH	Section 4.9.14, "Inbound Post Tail Pointer Register - IPTPR" on page 349		
FFFF.E370H	Section 4.9.15, "Outbound Free Head Pointer Register - OFHPR" on page 350		
FFFF.E374H	Section 4.9.16, "Outbound Free Tail Pointer Register - OFTPR" on page 351		
FFFF.E378H	Section 4.9.17, "Outbound Post Head Pointer Register - OPHPR" on page 352		
FFFF.E37CH	Section 4.9.18, "Outbound Post Tail Pointer Register - OPTPR" on page 353		
FFFF.E380H	Section 4.9.19, "Index Address Register - IAR" on page 354		

March 2005



#### 4.9.1 Inbound Message Register - IMRx

There are two Inbound Message Registers: IMR0 and IMR1. When the IMR register is written, an interrupt to the Intel XScale<sup>®</sup> core may be generated. The interrupt is recorded in the Inbound Interrupt Status Register and may be masked by the Inbound Message Interrupt Mask bit in the Inbound Interrupt Mask Register.



#### Table 200. Inbound Message Register - IMRx

### 4.9.2 Outbound Message Register - OMRx

There are two Outbound Message Registers: OMR0 and OMR1. When the OMR register is written, a PCI interrupt may be generated. The interrupt is recorded in the Outbound Interrupt Status Register and may be masked by the Outbound Message Interrupt Mask bit in the Outbound Interrupt Mask Register.



#### Table 201. Outbound Message Register - OMRx



### 4.9.3 Inbound Doorbell Register - IDR

The Inbound Doorbell Register (IDR) is used to generate interrupts to the Intel XScale<sup>®</sup> core. Bit 31 is reserved for generating an Error Doorbell interrupt. When bit 31 is set, an Error interrupt may be generated to the Intel XScale<sup>®</sup> core. All other bits, when set, cause the Normal Messaging Unit interrupt line of the Intel XScale<sup>®</sup> core to be asserted, when the interrupt is not masked by the Inbound Doorbell Interrupt Mask bit in the Inbound Interrupt Mask Register. The bits in the IDR register can only be set by an external PCI agent and can only be cleared by the Intel XScale<sup>®</sup> core.



#### Table 202. Inbound Doorbell Register - IDR

## 4.9.4 Inbound Interrupt Status Register - IISR

The Inbound Interrupt Status Register (IISR) contains hardware interrupt status. It records the status of Intel XScale<sup>®</sup> core interrupts generated by the Message Registers, Doorbell Registers, and the Circular Queues. All interrupts are routed to the Normal Messaging Unit interrupt input of the Intel XScale<sup>®</sup> core, except for the Error Doorbell Interrupt and the Outbound Free Queue Full interrupt; these two are routed to the Messaging Unit Error interrupt input. The generation of interrupts recorded in the Inbound Interrupt Status Register may be masked by setting the corresponding bit in the Inbound Interrupt Mask Register. Some of the bits in this register are Read Only. For those bits, the interrupt must be cleared through another register.



#### Table 203. Inbound Interrupt Status Register - IISR



### 4.9.5 Inbound Interrupt Mask Register - IIMR

The Inbound Interrupt Mask Register (IIMR) provides the ability to mask Intel XScale<sup>®</sup> core interrupts generated by the Messaging Unit. Each bit in the Mask register corresponds to an interrupt bit in the Inbound Interrupt Status Register.

Setting or clearing bits in this register does not affect the Inbound Interrupt Status Register. They only affect the generation of the Intel XScale<sup>®</sup> core interrupt.



Table 204. Inbound Interrupt Mask Register - IIMR

## 4.9.6 Outbound Doorbell Register - ODR

The Outbound Doorbell Register (ODR) allows software interrupt generation. It allows the Intel XScale<sup>®</sup> core to generate PCI interrupts to the host processor by writing to this register. The generation of PCI interrupts through the Outbound Doorbell Register may be masked by setting the Outbound Doorbell Interrupt Mask bit in the Outbound Interrupt Mask Register.

The Software Interrupt bits in this register can only be set by the Intel XScale<sup>®</sup> core and can only be cleared by an external PCI agent.



#### Table 205. Outbound Doorbell Register - ODR



### 4.9.7 Outbound Interrupt Status Register - OISR

The Outbound Interrupt Status Register (OISR) contains hardware interrupt status. It records the status of PCI interrupts generated by the Message Registers, Doorbell Registers, and the Circular Queues. The generation of PCI interrupts recorded in the Outbound Interrupt Status Register may be masked by setting the corresponding bit in the Outbound Interrupt Mask Register. Some of the bits in this register are Read Only. For those bits, the interrupt must be cleared through another register.



#### Table 206. Outbound Interrupt Status Register - OISR

## 4.9.8 Outbound Interrupt Mask Register - OIMR

The Outbound Interrupt Mask Register (OIMR) provides the ability to mask outbound PCI interrupts generated by the Messaging Unit. Each bit in the mask register corresponds to a hardware interrupt bit in the Outbound Interrupt Status Register. When the bit is set, the PCI interrupt is not generated. When the bit is clear, the interrupt is allowed to be generated.

Setting or clearing bits in this register does not affect the Outbound Interrupt Status Register. They only affect the generation of the PCI interrupt.



Table 207. Outbound Interrupt Mask Register - OIMR



#### 4.9.9 MU Configuration Register - MUCR

The MU Configuration Register (MUCR) contains the Circular Queue Enable bit and the size of one Circular Queue. The Circular Queue Enable bit enables or disables the Circular Queues. The Circular Queues are disabled at reset to allow the software to initialize the head and tail pointer registers before any PCI accesses to the Queue Ports. Each Circular Queue may range from 4 K entries (16 Kbytes) to 64 K entries (256 Kbytes) and there are four Circular Queues.



#### Table 208. MU Configuration Register - MUCR



### 4.9.10 Queue Base Address Register - QBAR

The Queue Base Address Register (QBAR) contains the local memory address of the Circular Queues. The base address is required to be located on a 1 Mbyte address boundary.

All Circular Queue head and tail pointers are based on the QBAR. When the head and tail pointer registers are read, the Queue Base Address is returned in the upper 12 bits. Writing to the upper 12 bits of the head and tail pointer registers does not affect the Queue Base Address or Queue Base Address Register.

*Warning:* The QBAR must designate a range allocated to the 80333 DDR SDRAM interface (Chapter 8, "Memory Controller").



#### Table 209. Queue Base Address Register - QBAR



### 4.9.11 Inbound Free Head Pointer Register - IFHPR

The Inbound Free Head Pointer Register (IFHPR) contains the local memory offset from the Queue Base Address of the head pointer for the Inbound Free Queue. The Head Pointer must be aligned on a DWORD address boundary. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored. This register is maintained by software.



#### Table 210. Inbound Free Head Pointer Register - IFHPR

## 4.9.12 Inbound Free Tail Pointer Register - IFTPR

The Inbound Free Tail Pointer Register (IFTPR) contains the local memory offset from the Queue Base Address of the tail pointer for the Inbound Free Queue. The Tail Pointer must be aligned on a DWORD address boundary. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored.



#### Table 211. Inbound Free Tail Pointer Register - IFTPR



### 4.9.13 Inbound Post Head Pointer Register - IPHPR

The Inbound Post Head Pointer Register (IPHPR) contains the local memory offset from the Queue Base Address of the head pointer for the Inbound Post Queue. The Head Pointer must be aligned on a DWORD address boundary. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored.





## 4.9.14 Inbound Post Tail Pointer Register - IPTPR

The Inbound Post Tail Pointer Register (IPTPR) contains the local memory offset from the Queue Base Address of the tail pointer for the Inbound Post Queue. The Tail Pointer must be aligned on a DWORD address boundary. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored.



#### Table 213. Inbound Post Tail Pointer Register - IPTPR



#### 4.9.15 Outbound Free Head Pointer Register - OFHPR

The Outbound Free Head Pointer Register (OFHPR) contains the local memory offset from the Queue Base Address of the head pointer for the Outbound Free Queue. The Head Pointer must be aligned on a DWORD address boundary. This register is maintained by software. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored.



Table 214. Outbound Free Head Pointer Register - OFHPR

## 4.9.16 Outbound Free Tail Pointer Register - OFTPR

The Outbound Free Tail Pointer Register (OFTPR) contains the local memory offset from the Queue Base Address of the tail pointer for the Outbound Free Queue. The Tail Pointer must be aligned on a DWORD address boundary. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored.







### 4.9.17 Outbound Post Head Pointer Register - OPHPR

The Outbound Post Head Pointer Register (OPHPR) contains the local memory offset from the Queue Base Address of the head pointer for the Outbound Post Queue. The Head Pointer must be aligned on a DWORD address boundary. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored.





## 4.9.18 Outbound Post Tail Pointer Register - OPTPR

The Outbound Post Tail Pointer Register (OPTPR) contains the local memory offset from the Queue Base Address of the tail pointer for the Outbound Post Queue. The Tail Pointer must be aligned on a DWORD address boundary. When read, the Queue Base Address is provided in the upper 12 bits of the register. Writes to the upper 12 bits of the register are ignored.



#### Table 217. Outbound Post Tail Pointer Register - OPTPR

### 4.9.19 Index Address Register - IAR

The Index Address Register (IAR) contains the offset of the least recently accessed Index Register. It is written by the MU when the Index Registers are written by a PCI agent. The register is not updated until the Index Interrupt bit in the Inbound Interrupt Status Register is cleared.

The local memory address of the Index Register least recently accessed is computed by adding the Index Address Register to the Inbound ATU Translate Value Register.



Table 218. Index Address Register - IAR



## 4.9.20 MSI Capability Identifier Register - MSI\_CAPID

The MSI Capability Identifier Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.2. This register in the PCI Extended Capability header identifies the type of Extended Capability contained in that header. The value of 05H in this field identifies the function as message signaled interrupt capable.

Table 219. MSI Capability Identifier Register - MSI\_CAPID





### 4.9.21 MSI Next Item Pointer Register - MSI\_NXTP

The Next Item Pointer Register bits adhere to the definitions in the *PCI Local Bus Specification*, Revision 2.2. This register describes the location of the next item in the function capability list. For the 80333 that is the PCI-X capability header at offset E0H.

#### Table 220. MSI Next Item Pointer Register - MSI\_NXTP





## 4.9.22 MSI Message Control Register - MSI\_MCR

The Message Control Register provides system software control over MSI. After reset, MSI is disabled. System software is permitted to modify the Message Control register's read/write bits and fields while a device driver is not permitted to modify them.



IOP 15 12 8 4 0   Attributes rv					
PCI Con	figuration Offse	t Intel Xscale <sup>™</sup> Core Local Bus Address	Attribute Legend: RV = Reserved	RW = Read/Write RC = Read Clear	
D2H - D.	38	FFFF EIDZH	PR = Preserved	RO = Read Only	
			RS = Read/Set	NA = Not Accessible	
Bit	Default	Description			
15:8	00H	Reserved			
7	1 <sub>2</sub>	64-bit Address Support - This field is set to $1_2$ indicating that the 80333 is capable of generating a 64-bit message address.			
6:4	6:4 000 <sub>2</sub> Multiple Message Enable - System software writes to this field to indicate the number of messages allocated to the 80333. While, the 80333 requests two messages, it is possible that system software only allocates one message. The device hardware is designed to handle both cases.				
3:1	001 <sub>2</sub>	Multiple Message Capable - This field is set to $001_2$ indicating that the 80333 can issue up to two unique interrupt messages.			
		inten apt incoordgoon			



### 4.9.23 MSI Message Address Register - MSI\_MAR

The Message address register specifies the DWORD aligned address for the MSI memory write transaction. The value is set by system software during initialization.



#### Table 222. MSI Message Address Register - MSI\_MAR



#### 4.9.24 MSI Message Upper Address Register - MSI\_MUAR

The Message Upper Address register is set during system initialization when system software wishes to place the MSI address location above the 4G address boundary. When this register is set to a non-zero value, the 80333 generates a dual address cycle for the MSI write command and uses the contents of this register as the upper 32-bits of that address.



#### Table 223. MSI Message Upper Address Register - MSI\_MUAR



#### 4.9.25 MSI Message Data Register- MSI\_MDR

The value in the Message Data Register contains the data used during an MSI write transaction. When two unique messages are enabled, one message is reserved for the Outbound Post Queue Interrupt and the other message represents all of the Outbound Doorbell and Outbound Message Interrupts. When only one message is enabled, all of these interrupts is represented by a single message. Interrupt handler software needs to read the 80333 Outbound Interrupt Status Register to determine the cause of the interrupt when more than one source is represented by a single message.

During an MSI write data phase, the value in the Message Data Register is driven on to **AD**[15:0] while **AD**[31:16] is driven to zero. **C**/**BE**[3:0]# are asserted during the data phase of the memory write transaction.



#### Table 224. MSI Message Data Register - MSI\_MDR

## 4.10 Power/Default Status

The software is responsible for initializing the Circular Queue Size in the MU Configuration Register and all head and tail pointer registers before setting the Circular Queue Enable bit.


# Intel XScale<sup>®</sup> Core Bus Interface Unit 5

This chapter describes the Intel XScale<sup>®</sup> core (ARM\* architecture compliant) Bus Interface Unit (BIU) for the Intel<sup>®</sup> 80333 I/O processor (80333).

## 5.1 **Overview**

The BIU provides the interfaces between the Intel XScale<sup>®</sup> core and the Internal Bus and Memory Controller outbound ports. The BIU decodes transactions from the core and sends them to the appropriate outbound port. Transactions which address the DDR SDRAM are sent to the MCU port, and all other transactions are sent to the 80333 64-bit internal bus.

The high performance Internal Bus is used in the 80333 to interconnect major peripheral blocks as shown in Figure 27. A detailed block diagram of the BIU and Memory Controller Unit (MCU) is shown in Figure 27.

The BIU acts as a Master for all the read and write requests issued by the Intel XScale<sup>®</sup> core targeting the Internal Bus, and acts as a Slave for split completions of BIU split requests. The BIU has software accessible state information in the form of registers which reside in Peripheral Memory Mapped Register space.



## 5.2 Theory of Operation

The BIU accepts Intel XScale<sup>®</sup> core bus accesses. It decodes core transactions and forwards them to the memory controller or internal bus outbound ports accordingly. The BIU translates the Intel XScale<sup>®</sup> core bus transactions to the protocol for the internal bus port, and the format for the Core Memory Transaction Queue of the MCU for the Core MCU port.

## 5.2.1 Functional Blocks

The memory controller and Intel XScale<sup>®</sup> core processor bus interface unit (BIU) logically comprises the blocks illustrated in Figure 27. The memory controller is a multiported unit, supporting an inbound path for both the BIU and Internal Bus (IB) to the DDR SDRAM.

#### Figure 27. Intel<sup>®</sup> 80333 I/O Processor BIU and MCU Architecture



#### 5.2.1.1 Core Processor Port Address Decode

The address decode block for the Core Processor transactions resides in the BIU. The Core MCU port therefore does not require any decode, and any transactions received from the BIU via the Core MCU Port are decoded to determine when they address the DDR SDRAM Memory Space. When a transaction addresses the DDR SDRAM Memory Space, it is sent to the Core MCU port of the MCU which is a direct connection to the Core Memory Transaction Queue. When the transaction does not address that space, it is sent to the Core Internal Bus Transaction Queue to be issued on the Internal Bus. Core processor transactions addressing the MCU MMR Space, are directed to the Internal Bus, and subsequently claimed by the MCU Internal Bus Port interface.

#### 5.2.1.2 Transaction Queues

The BIU has one transaction queue for transactions directed to the Internal Bus. The Core IB Transaction Queue (CIBTQ) has 8 read entries, 4 write entries, and 4 return data entries. The CIBTQ holds 4 outstanding read transactions up to 32-Bytes each, and 4 posted write transactions up to 16-Bytes each. The BIU acts as a Master for all the read and write requests issued by the Intel XScale<sup>®</sup> core targeting the Internal Bus, and acts as a Slave for split completions of BIU split requests. The transaction queue for MCU directed transactions exists in the MCU, and is defined in Section 8.3.1.3, "Memory Transaction Queues" on page 494. The BIU has software accessible state information in the form of memory mapped registers. All BIU registers reside in the Peripheral Memory Mapped Register space.

The Core Internal Bus port propagates read accesses to the Internal Bus and returns the read data to the Intel XScale<sup>®</sup> core. It completes write accesses for the Intel XScale<sup>®</sup> core. The BIU follows the ordering rules below for Internal Bus targeted transactions:

- Writes can pass Read requests
- Read requests does not pass read requests
- Writes does not pass Writes

The BIU may address any target on the Internal Bus. The BIU divides the core processor clock domain from the Internal Bus clock domain. The BIU has several address/data buffers:

- 4-Deep Write Transaction Queue
- 8-Deep Read Transaction Queue
- Write Data Buffer
- Read Data Buffer

In the Read transaction queue, each entry can be associated with an instruction or data cache line fill, a non-cacheable DWORD read. In the Write transaction queue, each entry can be associated with a write varying from 1 to 16 bytes of data.

The Write Data Buffer associated with each transaction queue entry can temporarily store write accesses that are destined for the Internal Bus. The BIU is responsible for forwarding all write accesses to the Internal Bus and ensuring their completion.

The Read Data Buffer associated with each transaction queue entry can temporarily store return data from the Internal Bus. The BIU is responsible for returning all read accesses to the Intel XScale<sup>®</sup> core.

A given transaction queue entry remains busy until the last byte of write data is transmitted or the last byte of read data is received across the Internal Bus.

## 5.3 Addressing

The BIU byte enables need to be adjusted depending on whether the 32-bit address resides in the even (A2='0') or odd (A2='1') DWORD.

See Table 225 for the Byte Enable encodings.

The BIU only reads and writes data on the Internal Bus as indicated by the Byte address generated by the Intel XScale<sup>®</sup> core. The BIU assumes that all accesses are to non-prefetchable memory. It does not promote Intel XScale<sup>®</sup> microarchitecture Byte or Intel XScale<sup>®</sup> microarchitecture Short accesses to DWORD accesses.

See Table 225 for the Byte Enable encodings.

#### Table 225.Contiguous Byte Enable Encodings<sup>a</sup>

Access		Intel XScale <sup>®</sup> Core Address	Internal Bus							
Type	Audress	A[2:0]	BE7	BE6	BE5	BE4	BE3	BE2	BE1	BE0
Mard	Even	0	1	1	1	1	0	0	0	0
word	Odd	4	0	0	0	0	1	1	BE1 0 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1	1
	Even	0	1	1	1	1	1	1	0	0
Short	Lven	2	1	1	1	1	0	0	BE2     BE1       0     0       1     1       1     0       0     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1       1     1	1
Onon	Odd	4	1	1	0	0	1	1	1	1
	Ouu	6	0	0	1	1	1	1	1 1 1	1
		0	1	1	1	1	1	1	1	0
	Evon	1	1	1	1	1	1	1	0	1
	Lven	2	1	1	1	1	1	0	1	1
Bute		3	1	1	1	1	0	1	1	1
Dyle		4	1	1	1	0	1	1	1	1
	Odd	5	1	1	0	1	1	1	1	1
	Cuu	6	1	0	1	1	1	1	1	1
		7	0	1	1	1	1	1	1	1

a. Table assumes that write coalescing is turned off. When write coalescing is enabled, many other non-contiguous byte enable patterns are possible including all disabled.

b. For the purposes of this column, definitions are: Word equals 32-bits, Short = 16-bits, Byte = 8-bits.

c. For the purposes of this column, Even and Odd is based on Address bit A2.

## 5.3.1 Bus Width

The BIU supports a bus width of 32-bit for the core interface, 64-bits for the Internal Bus data bus and a width of 128 bits for the Core MCU port to the Core Memory Transaction Queue in the MCU.

intel

## 5.4 Internal Bus Commands

For data read accesses, the BIU uses the Memory Read Block command for Intel XScale<sup>®</sup> microarchitecture cacheline (32 bytes) fills, and Memory Read DWORD command for read accesses include Byte, Short, and Word accesses (1-, 2-, or 4-byte). The BIU supports Critical Word First (CWF) address ordering for Intel XScale<sup>®</sup> microarchitecture cacheline (32 bytes) fills. For Intel XScale<sup>®</sup> core transactions targeting the PCI bus through the Address Translation Unit, refer to Chapter 3, "Address Translation Unit" for command translation.

For all write accesses including cache line flushes, the BIU uses the Memory Write command.

*Note:* The Intel XScale<sup>®</sup> core promotes all Intel XScale<sup>®</sup> microarchitecture Byte and Short read accesses to Intel XScale<sup>®</sup> microarchitecture Word read accesses when the data cache is enabled.

Table 226 contains a summary of the Internal Bus commands used by the BIU.

#### Table 226.Internal Bus Command Summary

Internal Bus Command	Intel XScale <sup>®</sup> Microarchitecture Word Size	Usage		
Memory Read DWORD	Intel XScale <sup>®</sup> microarchitecture Byte, Short, or Word	Read 1, 2, 4 bytes		
Memory Read Block	Intel XScale <sup>®</sup> microarchitecture core cacheline fills	Read 32 bytes (CWF)		
Split Completion (Target Only)	All	Read data may be returned to the BIU using a Split completion		
Memory Write	All	For all writes		

## 5.5 Features

The BIU supports the following additional features.

## 5.5.1 Multi-Transaction Timer

The BIU has an associated Multi-Transaction Timer (MTT1) in the Internal Bus Arbiter. When programmed properly, the MTT allows for a guaranteed quantum of time for the BIU. See Chapter 13, "Arbitration Unit" for more details.

## 5.5.2 ATU Accesses

All read accesses from the BIU to the Address Translation Units (ATU) are processed as Split Transactions.

All writes from the BIU to the Address Translation Unit (ATU) are posted.

Instruction fetches from the BIU to the ATU are supported.

Atomic accesses from the BIU to the ATU are not supported.



## 5.6 Interrupts and Error Conditions

The BIU records error conditions that result from accesses initiated by the BIU on behalf of the Intel XScale<sup>®</sup> core on the Internal Bus or MCU port. The errors are recorded in the BIU Status Register (BIUSR).

## 5.6.1 Internal Bus Errors

There are two ways that the BIU can receive an error from transactions issued on the Internal Bus:

- IB Master-Abort: No target on receives a Target-Abort when accessing the Configuration Outbound Data Register of the ATU when the ATU configuration space is incorrectly configured as cacheable. In this case the BIU would issue a burst read request (cacheline fill), and the ATU responds with a Target-Abort.
- PCI Master-Abort: The ATU, as a PCI master on behalf of the BIU, receives a Master-Abort on the PCI bus and returns the Master-Abort to the BIU with a Split Completion Error Message (Message Class = 1H, Message Index = 00H).
- PCI Target-Abort: The ATU, as a PCI master on behalf of the BIU, receives a Target-Abort on the PCI bus and returns the Target-Abort to the BIU via a Split Completion Error Message (Message Class = 1H, Message Index = 01H).

There are also two ways that the BIU can detect an error from transactions issued on the Internal Bus:

- Split completion Target Abort: On a read completion, the remaining byte count is out of range from the original request.
- Split completion Target Abort: The lower 7-bits of the address from the original request are corrupted on the first Split Completion Transaction.

#### 5.6.1.1 Internal Bus Master-Abort

When an Internal Bus access initiated by the BIU receives a Master-Abort, the BIU records the Master-Abort condition in the BIU Status Register (BIUSR), the address in BIU Error Address Register (BEAR) and generates an interrupt to the Interrupt Controller Unit. On read transactions which receive a Master-Abort, the BIU will also signal a Data Abort to the Intel XScale<sup>®</sup> core.

*Note:* Software must comprehend the two instances of Master-Abort error recording on BIU reads; one in the Interrupt Controller Unit, and the second as a Data Abort to the Intel XScale<sup>®</sup> core.

When a Master Abort to a Read request is received on the Internal Bus, the BIU signals an Imprecise Data Abort to the Intel XScale<sup>®</sup> core. This results in a Data Access Memory Abort exception by the Intel XScale<sup>®</sup> core; for more details on exceptions, please refer to Section 17.3, "The Intel XScale<sup>®</sup> Core Exceptions Architecture" on page 752.

#### 5.6.1.2 Internal Bus Target-Abort

When an Internal Bus access initiated by the BIU receives a Target-Abort, the BIU records the Target-Abort condition in the BIUSR, the address in BEAR and signals a Data Abort to the Intel XScale<sup>®</sup> core.

### 5.6.1.3 PCI Master-Abort

The PCI Master Abort is recorded in the ATUISR. The ATU generates an interrupt to the Intel XScale<sup>®</sup> core.

For PCI read accesses, the ATU returns a Split Completion Error Message to the BIU. This causes the BIU to signal an imprecise data abort to the Intel XScale<sup>®</sup> core.

The software is able to diagnose that the error occurred via an outbound BIU transaction by examining the BIU Error Address Register (BEAR) and the BIU Control Register.

Because writes are posted, the ATU is responsible for signalling an interrupt to the Intel XScale<sup>®</sup> core when a PCI Master-Abort is received during a write access to the ATU.

#### 5.6.1.4 PCI Target-Abort

The PCI Target Abort is recorded in the ATUISR. The ATU generates an interrupt to the Intel XScale<sup>®</sup> core. When the ATU detects a Target-Abort on the PCI bus for a read access, the ATU returns a Split Completion Error Message to the BIU indicating the Target-Abort.

For PCI read accesses, the ATU returns a Split Completion Error Message to the BIU. This causes the BIU to signal an imprecise data abort to the Intel XScale<sup>®</sup> core.

The software is able to diagnose that the error occurred via an outbound BIU transaction by examining the BIU Error Address Register (BEAR) and the BIU Control Register.

Because writes are posted, the ATU is responsible for signalling an interrupt to the Intel XScale<sup>®</sup> core when a PCI Target-Abort is received during a write access to the ATU.

#### 5.6.1.5 Split Transaction Errors

As a requester on the internal bus, the BIU may detect the following error conditions due to a corrupted split completion:

- On a read completion, the remaining byte count is out of range from the original request.
- The lower 7-bits of the address from the original request are corrupted on the first Split Completion Transaction.

Upon detecting either of these errors, the BIU signals an Imprecise Data Abort to the Intel XScale<sup>®</sup> core and records the error in the BIUSR.

The BIU never acts as a Completer on internal bus for split transactions, and therefore no errors as a completer are possible.



## 5.6.2 Core MCU Errors

There are two ways that the BIU can receive an error from transactions issued to the MCU port:

- DDR SDRAM Multi-bit ECC Error
- Invalid Address Region Error

#### 5.6.2.1 Multi-bit ECC Error

When a multi-bit ECC error is detected by the MCU, the error is also reported by the BIU. A multi-bit ECC error reported by the BIU can only occur on a read transaction by the Intel XScale<sup>®</sup> core.

For write accesses, the core processor transaction is posted to the MCU. Therefore to the BIU, the transaction appears as when it has completed before the transaction completes to the DDR SDRAM. In this case, the MCU is the only unit that notifies the core processor of any error condition.

For read accesses, the BIU signals an Imprecise Data Abort to the Intel XScale<sup>®</sup> core. The MCU also records the error in the ELOGx registers and generates an interrupt to the Intel XScale<sup>®</sup> core.

#### 5.6.2.2 Invalid Address Region Error

When a 32-bit region is defined with a 64-bit wide memory, an invalid address region exists. The region is illustrated in Figure 28 and is equal size to the 32-bit region as defined by Section 8.7.7, "SDRAM 32-bit Region Size Register - S32SR" on page 554, and at a contiguous address to the 32-bit region.

Intel XScale<sup>®</sup> core transactions which address the invalid address region are sent to the MCU via the Core MCU port. The MCU detects and reports this error in the Section 8.7.12, "Memory Controller Interrupt Status Register - MCISR" on page 559.

#### Figure 28. DDR SDRAM 64-bit Memory Address Map



# intel

## 5.7 Reset and Initialization Conditions

## 5.7.1 Reset

The BIU is reset:

- when **PWRGD** is deasserted.
- When RSTIN# is asserted.
- when the Intel XScale<sup>®</sup> core is reset.

When reset, all Core IB transaction buffers are marked as invalid, and BIU registers return to the reset values.

## 5.7.2 Initialization

For proper operation of 80333, the configuration of the MCU must be complete prior to enabling the MCU port of the BIU. Once the MCU is fully configured (specifically the SDCR[1:0], SDIR, SDBR, SBR[1:0], and S32SR), the BIU control register can be programmed to enable the MCU port of the BIU.



## 5.8 Register Definitions

There are two Registers for the BIU. These registers are mapped into Intel XScale<sup>®</sup> core memory space.

Table 227 lists the PMMR registers for the BIU.

#### Table 227.Bus Interface Unit Register Tables

Section, Register Name - Acronym (Page)
Section 5.8.1, "BIU Status Register - BIUSR" on page 371
Section 5.8.2, "BIU Error Address Register - BEAR" on page 373
Section 5.8.3, "BIU Control Register - BIUCR" on page 374



## 5.8.1 BIU Status Register - BIUSR

The BIU Status Register (BIUSR) allows software to determine the cause of any BIU error. When an error is logged and this register is updated, an imprecise data abort is signaled to the Intel XScale<sup>®</sup> core.







#### Table 228. BIU Status Register - BIUSR (Sheet 2 of 2)



## 5.8.2 BIU Error Address Register - BEAR

This register is responsible for logging the addresses where the errors were detected on the Internal Bus. One error can be detected and logged. For error details, see Section 5.6, "Interrupts and Error Conditions" on page 366).



#### Table 229. BIU Error Address Register - BEAR

## 5.8.3 BIU Control Register - BIUCR

The BIU Control Register (BIUCR) provides enable control of the Core MCU port interface of the BIU. Monitoring the Read and Write transactions queue depth status bits of the "BIU Status Register - BIUSR" on page 371 to determine when the pending Intel XScale<sup>®</sup> core processor transactions have completed and drained through the Core Memory Transaction Queue of the Memory Controller Unit.



#### Table 230. BIU Control Register - BIUCR

#### Intel<sup>®</sup> 80333 I/O Processor Developer's Manual DMA Controller Unit

# **DMA Controller Unit**

# 6

This chapter describes the integrated Direct Memory Access (DMA) Controller Unit. DMA Controller operation modes, setup, external interface, and implementation are detailed in this chapter.

## 6.1 **Overview**

**int**el<sup>®</sup>

The DMA Controller provides low-latency, high-throughput data transfer capability. The DMA Controller optimizes block transfers of data between the Segment-A PCI bus and the local processor memory. The DMA is an initiator on the Internal bus with burst capabilities providing a maximum throughput of 1064 Mbytes/sec. when the Segment-A PCI bus is operating in 64-bit/133 MHz PCI-X mode. When addressing host memory space, the PCI Express-to-PCI Bridge forwards the transaction to the upstream PCI Express port. In addition, the DMA provides a hardware assist to the iSCSI\* application by calculating CRC-32C on the data during the block transfer.

The DMA Controller hardware is responsible for executing data transfers and for providing the programming interface. The DMA Controller features:

- Two Independent Channels
- Three 1-KByte queues in both Ch-0 and Ch-1
- Utilization of the 80333 Memory Controller Interface
- 2<sup>32</sup> addressing range on the Internal Bus interface
- 2<sup>64</sup> addressing range on the PCI interface by using PCI Dual Address Cycle (DAC)
- Hardware support for unaligned data transfers for both the PCI bus and the Internal Bus
- Up to 1064 Mbytes/sec burst support for the PCI bus in the PCI-X mode
- · Direct addressing to and from the PCI bus
- Memory to Memory DMA transfer mode
- Calculation of CRC-32C on the transferred data stream
- Fully programmable directly from the internal bus
- Support for automatic data chaining for gathering and scattering of data blocks
- 64-bit/333 MHz Intel<sup>®</sup> 80333 I/O processor (80333) internal bus interface.



Figure 29 shows the connections of the DMA channels to the Internal Bus.

#### Figure 29. DMA Controller



## 6.2 Theory of Operation

The DMA Controller provides two channels of high throughput PCI-to-Memory or Memory-to-Memory transfers. Channels 0 and 1 transfer blocks of data between the PCI bus and I/O processor local memory.

The two DMA channels operate identically. Each channel has an internal bus interface. Figure 30 shows the block diagram for one channel of the DMA Controller.



Figure 30. DMA Channel Block Diagram

Each DMA channel uses direct addressing for both the PCI bus and the internal bus. It supports data transfers to and from the full 64-bit address range of the PCI bus. This includes 64-bit addressing using the PCI DAC command. The channel provides a special register which contains the upper 32 address bits for the 64-bit address. The DMA channels do not support data transfers that cross a 32-bit address (4 GByte) boundary.

When Memory-to-Memory transfer mode is enabled, the DMA can be programmed to transfer data across the entire 4 Gbyte memory space on the Internal Bus. This includes but is not limited to transferring data to and from the 80333 Memory Controller (MCU), and from the Peripheral Bus Interface (PBI) to the MCU.

Both the PCI interface and the internal bus interface support large burst lengths up to 4 KBytes.

The channel programming interface is accessible from the internal bus through a memory-mapped register interface. Each channel is programmed independently and has its own set of registers. A normal DMA transfer is configured by writing the source address, destination address, number of bytes to transfer, and various control information into a chain descriptor in 80333 local memory. Chain descriptors are described in detail in Section 6.3.



Each channel supports chaining. Chain descriptors that describe one DMA transfer each can be linked together in 80333 local memory to form a linked list. Each chain descriptor contains all the necessary information for transferring a block of data in addition to a pointer to the next chain descriptor. The end of the chain is indicated when the pointer is zero.

Each channel contains a hardware data alignment unit. This unit enables data transfers from or to unaligned addresses in either the PCI address space or the I/O processor local address space. All combinations of unaligned data are supported with the data alignment unit.

Each channel includes a CRC Engine to calculate the CRC-32C algorithm on the data stream being transferred. The CRC engine also initiates the read of a CRC seed and the subsequent write back of the 32-bit result to a CRC Address specified in the descriptor. In addition, a Transfer Complete status bit is updated to the descriptor in memory.

## 6.3 DMA Transfer

intel

A DMA transfer is a block move of data from one memory address space to another. DMA transfers are configured and initiated through a set of memory-mapped registers and one or more chain descriptors located in local memory. A DMA transfer is defined by the source address, destination address, number of bytes to transfer, and control values. These values are loaded into the chain descriptor before a DMA transfer begins. On the 80333 internal bus, the DMA controller attempts all transactions as 64-bit transfers.

#### Table 231. DMA Registers

Register	Abbreviation	Description			
Channel Control Register x	CCRx	Channel Control Word			
Channel Status Register x	CSRx	Channel Status Word			
Descriptor Address Register x	DARx	Address of Current Chain Descriptor			
Next Descriptor Address Register x	NDARx	Address of Next Chain Descriptor			
PCI Address Register x	PADRx	Lower 32-bit PCI Address of Source/Destination			
PCI Upper Address Register x	PUADRx	Upper 32-bit PCI Address of Source/Destination			
Local Address Register x	LADRx	Local Address of Source/Destination			
Byte Count Register x	BCRx	Number of Bytes to transfer			
Descriptor Control Register x	DCRx	Chain Descriptor Control Word			

## 6.3.1 Chain Descriptors

All DMA transfers are controlled by chain descriptors located in local memory. A chain descriptor contains necessary information to complete one data transfer. A single DMA transfer has only one chain descriptor in memory. Chain descriptors can be linked together to form more complex DMA operations. To perform a DMA transfer, one or more chain descriptors must first be written to local memory. Every descriptor requires seven contiguous words in local memory and is required to be aligned on an 8-word boundary. When CRC is not enabled, this word is ignored. It is still read, but does not have any effect. Each word in the chain descriptor is analogous to control register values. Bit definitions for words in the chain descriptor are the same as for the DMA control registers.

*Warning:* Chain descriptors can only be located in DDR SDRAM memory in order for the DMA to function properly. Location of the chain descriptors anywhere else (e.g., either on the Peripheral Bus or on PCI) is not supported and the results would be unpredictable.

Figure 31 shows the format of an individual chain descriptor.

- First word is 80333 memory address of the next chain descriptor. A value of "0" specifies the end of chain. This value is loaded into the Next Descriptor Address Register. Because chain descriptors are aligned on an 8-word boundary, the channel may ignore bits 04:00 of this address.
- Second word is lower 32-bit PCI source/destination address. This address is generated on the PCI bus. This value is loaded into the PCI Address Register.
- Third word is upper 32-bit PCI source/destination address, when needed. When non-zero, this address is used during Dual Address Cycles for driving 64-bit PCI addresses. When zero, Single Address Cycles are used for driving 32-bit PCI addresses. This value is loaded into the PCI Upper Address Register.
- Fourth word is internal bus source/destination address, driven on the internal bus. This value is loaded into Intel XScale<sup>®</sup> core (ARM\* architecture compliant) Local Address Register.
- Fifth word is Byte Count value. This value determines the number of bytes to transfer. This value is loaded into the Byte Count Register.
- Sixth word is Descriptor Control word. This word configures the DMA channel for one DMA transfer. It contains the PCI transaction type, which determines the direction of the data transfer. This value is loaded into the Descriptor Control Register.
- Seventh word is CRC Address word. This address is used to load a 32-bit value to seed the CRC calculation (when enabled). Also, this address is used as the destination for the write back of the CRC result. This is a descriptor field, not a memory mapped register.

There are no data alignment requirements for either the PCI address or the internal bus address.

Refer to Section 6.14 for additional descriptions about the DMA Controller registers

#### Figure 31. DMA Chain Descriptor

Next Descriptor Address [NDA]	Address of Next Chain Descriptor
PCI Address (31:0) [PAD]	Lower 32-bit PCI Source/Destination Address
PCI Upper Address (63:32) [PUAD]	Upper 32-bit PCI Source/Destination Address
Local Address [LAD]	Local Source/Destination Address
Byte Count [BC]	Number of Bytes to Transfer
Descriptor Control [DC]	Descriptor Control
CRC Address [CAD]	CRC Address

A9082-01

# intel

## 6.3.2 Chaining DMA Descriptors

A series of chain descriptors can be built in local memory to transfer data between the PCI bus and the internal bus. For example, the application can build multiple chain descriptors to transfer many blocks of data which have different source addresses within the local memory. When multiple chain descriptors are built in 80333 local memory, the application can link each of these chain descriptors using the Next Descriptor Address in the chain descriptor. This address logically links chain descriptors together. This allows the application to build a list of DMA transfers which may not require the Intel XScale<sup>®</sup> core until all DMA transfers are complete. Figure 32 shows a list of DMA transfers built in local memory and how they are linked.



#### Figure 32. DMA Chaining Operation



## 6.3.3 Initiating DMA Transfers

A DMA transfer is started by building one or more chain descriptors in 80333 local memory. Each chain descriptor takes the form shown in Figure 31. The chain descriptors are required to be aligned on an 8-word boundary in the 80333 local memory.

The following describes the steps for initiating a new DMA transfer:

- 1. In order to begin the processing of a new chain of DMA transfers, the channel must be inactive. This can be checked by software by reading the *Channel Active* bit in the Channel Status Register (CSR). When this bit is clear, the channel is inactive. When this bit is set, the channel is currently active with a DMA transfer.
- 2. The CSR must be cleared of all error conditions.
- 3. The software writes the address of the first chain descriptor to the Next Descriptor Address Register.
- 4. The software sets the *Channel Enable* bit in the Channel Control Register (CCR). Since this is the start of a new DMA transfer and not the resumption of a previous transfer, the *Chain Resume* bit in the CCR should be clear.
- 5. The channel starts the DMA transfer by reading the chain descriptor at the address contained in the Next Descriptor Address Register. The channel loads the chain descriptor values into the channel control registers and begins data transfer. The Descriptor Address Register now contains the address of the chain descriptor just read and the Next Descriptor Address Register now contains the Next Descriptor Address from the chain descriptor just read.

The last descriptor in the DMA chain list has zero in the next descriptor address field specifying the last chain descriptor. The NULL value notifies the DMA channel not to read additional chain descriptors from memory.

Once a DMA transfer is active, it may be temporarily suspended by clearing the *Channel Enable* bit in the CCR. Note that this does not abort the DMA transfer. The channel resumes the DMA transfer when the *Channel Enable* bit is set.

When descriptors are read from local memory, bus latency and memory speed affect chaining latency. Chaining latency is defined as the time required for the channel to access the next chain descriptor plus the time required to set up for the next DMA transfer.

See Section 6.9 for a state diagram of the channel programming model.



## 6.3.4 Scatter Gather DMA Transfers

The DMA Controller can be used to perform typical scatter gather data transfers. This consists of programming the chain descriptors to gather the data which may be located in non-contiguous blocks of memory. The chain descriptor specifies the destination location such that once all data has been transferred, the data is contiguous in memory. Figure 33 shows how the destination pointers can gather data.

#### Figure 33. Example of Gather Chaining





## 6.3.5 Synchronizing a Program to Chained Transfers

Chained DMA transfers can be synchronized to a program executing on the Intel XScale<sup>®</sup> core through the use of processor interrupts. The channel generates an interrupt to the Intel XScale<sup>®</sup> core under certain conditions. They are:

- [Interrupt and Continue] The channel completes the data transfer for a chain descriptor and the Next Descriptor Address Register is non-zero. When the *Interrupt Enable* bit within the Descriptor Control Register is set, an interrupt is generated to the Intel XScale<sup>®</sup> core. This interrupt is for synchronization purposes only. The channel sets the *End Of Transfer Interrupt* flag in the Channel Status Register. Since it is not the last chain descriptor in the list, the DMA channel starts to process the next chain descriptor without requiring any processor interaction.
- 2. [End of Chain] The DMA channel completes the data transfer for a DMA chain descriptor and the Next Descriptor Address Register is zero, and the chain resume bit is clear, specifying the end of the chain. When the *Interrupt Enable* bit within the Descriptor Control Register is set, an interrupt is generated to the Intel XScale<sup>®</sup> core. The channel sets the *End Of Chain Interrupt* flag in the Channel Status Register.
- 3. [Error] An error condition occurs (refer to Section 6.12 for DMA error conditions) during a DMA transfer. The DMA channel halts operation on the current chain descriptor and not proceed to the next chain descriptor.

Each chain descriptor can independently set the *Interrupt Enable* bit in the Descriptor Control word. This bit enables an independent channel interrupt upon completion of the data transfer for the chain descriptor. This bit can be set or clear within each chain descriptor. Control of interrupt generation within each descriptor aids in synchronization of the executing software with DMA transfers.

Figure 34 shows two examples of program synchronization. The left shows program synchronization based on individual chain descriptors. Descriptor 1A generated an interrupt to the processor, while descriptor 2A did not because the *Interrupt Enable* bit was clear. The last descriptor nA, generated an interrupt to signify the end of the chain has been reached. The right shows an example where the interrupt was generated only on the last descriptor signifying the end of chain.

#### Figure 34. Synchronizing to Chained Transfers



March 2005

## 6.3.6 Appending to The End of a Chain

Once a channel has started processing a chain of DMA descriptors, the application software may need to append a chain descriptor to the current chain without interrupting the transfer in progress. This action is controlled by the Channel Control Register *Chain Resume* bit.

The channel reads the subsequent chain descriptor each time the channel completes the current chain descriptor and the Next Descriptor Address Register is non-zero. The Next Descriptor Address Register always contains the address of the next chain descriptor to be read and the Descriptor Address Register always contains the address of the current chain descriptor.

The procedure for appending chains requires the software to find the last chain descriptor in the current chain and change the Next Descriptor Address in that descriptor to the address of the new chain to be appended. The software then sets the *Chain Resume* bit in the Channel Control Register for the channel. It does not matter when the channel is active or not.

The channel examines the *Chain Resume* bit of the CCR when the channel is idle or upon completion of a chain of DMA transfers. When this bit is set, the channel re-reads the Next Descriptor Address of the current chain descriptor and load it into the Next Descriptor Address Register. The address of the current chain descriptor is contained in the Descriptor Address Register. The channel clears the *Chain Resume* bit and then examine the Next Descriptor Address Register. When the Next Descriptor Address Register is not zero, the channel reads the chain descriptor using this new address and begin a new DMA transfer. When the Next Descriptor Address Register is zero, the channel remains or return to idle and set the end of chain and generates an interrupt when enabled.

There are three cases to consider:

- 1. The channel completes a DMA transfer and it is not the last descriptor in the chain. In this case, the channel clears the *Chain Resume* bit and reads the next chain descriptor. The appended descriptor is read when the channel reaches the end of the original chain.
- 2. The channel completes a DMA transfer and it is the last descriptor in the chain. In this case, the channel examines the state of the *Chain Resume* bit. When the bit is set, the channel re-reads the current descriptor to get the address of the appended chain descriptor. When the bit is clear, the channel returns to idle.
- 3. The channel is idle. In this case, the channel examines the state of the *Chain Resume* bit when the CCR is written. When the bit is set, the channel re-reads the last descriptor from the most-recent chain to get the appended chain descriptor.



## 6.4 Data Transfers

The 80333 DMA controller is optimized to perform data transfers between the PCI bus and local memory, and between the local memory and the local memory. These transfers are summarized in the following sections.

## 6.4.1 PCI-to-Local Memory Transfers

PCI-to-Local memory transfers perform read cycles on the PCI bus and place the data into the DMA channel queues. Once data is placed into the queue, the internal bus interface of the DMA channel requests the internal bus and drain the queue by writing the data to the local memory.

DMA Internal Bus transactions that are destined for the external PCI Bus are automatically forwarded through the Outbound ATU to the PCI bus. In this way, the PCI direct addressing programming model defined for the DMA descriptors is preserved.

The DMA controller attempts to perform a memory write on the internal bus once the DMA queue has 1K Bytes of data or the Internal Bus transaction has disconnected. The byte count is set to the full byte count of the DMA and the byte count modified bit is set to 0 for memory write transactions on the internal bus. The write transaction continues on the internal bus until the byte count is satisfied, latency timer expires and **GNT#** is deasserted, or the target initiates a disconnect at an ADB. When the transaction is not complete the DMA restarts a memory write using the modified byte count, the starting address, and the same sequence ID to complete the write sequence.

*Note:* The Maximum Read Burst size of the DMA on the PCI bus can be throttled in the ATU by bits 3:2 of the PCIXCMD register (see Section 3.10.64, "PCI-X Command Register - PX\_CMD" on page 309). The maximum burst length supported by the DMA is 1 Kbytes.

## 6.4.2 Local Memory to PCI Transfers

Local memory to PCI transfers perform read cycles on the internal bus and place the data into the DMA channel queues. The DMA queue size is 1 Kbyte, thus when the byte count of the descriptor is greater than 1 Kbyte then the DMA is completed in bursts with maximum burst size of 1 Kbyte. Otherwise, the full byte count is used for the byte count. Once data is placed into the queue, the DMA channel requests the internal bus and then drain the queue by writing the data to the PCI bus through the ATU.

DMA Internal Bus transactions that are destined for the external PCI Bus are automatically forwarded through the Outbound ATU to the PCI bus. In this way, the PCI direct addressing programming model defined for the DMA descriptors is preserved.

When operating in PCI mode depending on the Internal Bus transaction's byte count and starting address, the ATU selects between two PCI write command types: Memory Write and Memory Write and Invalidate.

In PCI-X mode, the ATU uses the Memory Write Block command, exclusively.

## 6.4.3 Local Memory to Local Memory DMA

When bit 6 of the DCR is set, the DMA functions as a Local Memory to Local Memory DMA engine. The transfer takes place from the range of 32 bit addresses defined by the PADR (PCI Address Register -- Lower 32 bits) and the byte count to the range of addresses defined by the LADR (Local Address Register) and the byte count.

Local Memory to Local Memory transfers perform read cycles on the internal bus using PADR based addresses and place the data into the DMA channel queues. The maximum byte count for a DMA transaction is 1 Kbyte, thus when the byte count of the descriptor is greater than 1 Kbyte then the byte count is set to 1 Kbyte. Otherwise, the full byte count is used for the byte count. Once data is placed into the queue, the DMA channel requests the internal bus and then drain the queue by writing the data to the Internal Bus using LADR based addresses.

Local Memory to Local Memory transfers utilizes the Memory Read Block (MRB) and Memory Write Block (MWB) commands on the Internal Bus.

*Note:* Since the Internal Bus only supports 32 bit addressing, the PCI Upper Address Register (PUADR) is ignored. Because data transfers only occur in one direction in this mode, the command field (bits 3 to 0) of the DCR are also ignored.

## 6.4.4 Exclusive Access

The DMA Controller does not support exclusive access through the PCI LOCK# signal.

## 6.5 Data Queues

DMA Ch-0 and Ch-1 each contain three 1 Kbyte, data buffers. These queues temporarily hold data to increase performance of data transfers in both directions.

## 6.6 Data Alignment

Each channel contains a hardware data alignment unit to support unaligned data transfers between the source and destination busses. The data alignment unit optimizes data transfers to and from 32 and 64-bit memory. The channel reformats data words for the correct bus data width.

Aligned data transfers involve data accesses that fall on natural boundaries. For example; double words are aligned on 8-byte boundaries and words are aligned on 4-byte boundaries. DMA transfers can occur with both the source and destination addresses unaligned.

## 6.6.1 64-bit Unaligned Data Transfers

Figure 35 illustrates a DMA transfer between unaligned 64-bit source and destination addresses.

#### Figure 35. Optimization of an Unaligned DMA





## 6.6.2 64/32-bit Unaligned Data Transfers

Figure 36 illustrates a DMA transfer between an unaligned 32-bit source address and an unaligned 64-bit destination address.





## 6.7 CRC Generation

When enabled, the DMA generates a 32-bit CRC based on the programmed data stream and a 32-bit seed. The CRC engine uses the CRC-32C algorithm required by the iSCSI\* Specification.

- *Note:* The DMA CRC result value is byte reversed. See *Intel<sup>®</sup> 80333 I/O Processor Specification Update*, Erratum #13.
- *Note:* Use local address sources only when calculating CRC. See *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #14.

## 6.7.1 CRC Mode Configuration and Operation

In addition to the normal DMA Descriptor configuration, the following additional steps are required to configure the CRC engine:

- 1. When writing out the descriptor, bit 8 of the DC is set to enable CRC generation.
- 2. Word 7 of the descriptor is written with the CRC Address.

When a descriptor is configured to generate CRC, the DMA performs the following steps in addition to the Data Transfer:

- 1. Prior to start of Data Transfer, the DMA loads the 32-bit seed value from the CRC Address.
- 2. The DMA accumulates the 32-bit CRC as the Data Transfer is occurring. Specifically, the CRC Value is recirculated through the CRC-32C algorithm using the previous CRC Value and the most recent data transferring through the DMA.
  - a. CRC Value (n) = CRC-32C (CRC Value (n-1), Data(n)).
- 3. Following completion of Data Transfer, the DMA writes the CRC Value to the CRC Address.
- 4. Finally, the DMA writes the Descriptor Control Register back to the descriptor with the Transfer Complete bit set.
- *Note:* DCR bit 7 may be used to disable the Data Transfer during a CRC calculation. This provides the ability to generate CRC on a source data block without writing a destination data block.
- *Note:* DCR bit 9 is used to enable chaining of descriptors to calculate CRC across fragmented data source. Bit is clear for first descriptor of chain, then set for subsequent descriptors in chain for complete data.

## 6.7.2 CRC-32C Algorithm

The CRC Engine generates a 32-bit CRC of the programmed data stream using the CRC-32C generator polynomial required by the iSCSI\* Specification, (11EDC6F41).

Figure 37. CRC-32C Generator Polynomial



*Note:* The seed value located at the CRC Address (CAD) must be all 1s (FFFF FFFFh) to generate an iSCSI\* compatible CRC-32C.

## 6.8 Channel Priority

The Intel XScale<sup>®</sup> core internal bus arbitration logic determines which internal bus master has access to the internal bus. Each DMA channel has an independent bus Request/Grant signal pair to the internal bus arbitration. Chapter 13, "Arbitration Unit" further describes the priority scheme between all the bus masters on the internal bus. Also described is the priority mechanism used between the two DMA channels.

In addition, the internal bus arbitration unit has a Multi-Transaction timer (Section 13.4.3, "Multi-Transaction Timer Register 2 - MTTR2" on page 692) that affects the throughput of a DMA channel. The default value for MTT2 of 152 clocks was chosen to ensure that once an internal bus agent (in this case a DMA channel) is granted the internal bus that it is guaranteed an opportunity to burst data into DDR SDRAM memory. However, when the bus is busy the DMA channel loses grant before the burst is completed. This means that the DMA channel is able to complete only one burst for each arbitration cycle.

Alternatively, the user may wish to increase the value of MTT2 to guarantee that two or more bursts is able to complete within an arbitration cycle.

For example, assuming 1 Kbyte bursts and a 64-bit memory subsystem, an MTT2 setting of 192 clocks would be sufficient to support two 1 Kbyte bursts for a given DMA channel in a single arbitration cycle.

*Warning:* Increasing the MTT2 value may also increase the latency to memory for the Intel XScale<sup>®</sup> core on the average. Before changing the MTT2 value, it's imperative that the overall impact to the performance of the application is considered.



## 6.9 **Programming Model State Diagram**

The channel programming model diagram is shown in Figure 38. Error condition states are not shown.

Figure 38. DMA Programming Model State Diagram





## 6.10 DMA Channel Programming Examples

The software for the DMA channels falls into the following categories:

- Channel initialization
- Start DMA transfer
- Suspend channel

Examples for each of the software is shown in the following sections as pseudo code flow.

## 6.10.1 Software DMA Controller Initialization

The DMA Controller is designed to have independent control of the interrupts, enables, and control. The initialization consists of virtually no overhead as shown in Figure 39.

#### Figure 39. Software Example for Channel Initialization

CCR0 = 0x0000 0000 ; Disable channel Call setup\_channel

## 6.10.2 Software Start DMA Transfer

The DMA channel control register provides independent control per channel. This provides the greatest flexibility to the applications programmer. The example shown in Figure 41 describes the pseudo code for starting a DMA transfer on channel 0.

#### Figure 40. Software Example for PCI-to-Local DMA Transfer

; Set up descriptor in local memory at address d
d.nad = 0 ; No chaining
d.pad = 0x0000F000 ; Source address of 0000F000H
d.puad = 0 ; DAC is not used
$d.lad = 0 \times B00000000$ ; Destination address B0000000H
d.bc = 64 ; byte count of 64
d.dc = 0x0000001E ; PCI Memory Read Block command
; Interrupt processor after transfer
; Check for inactive channel & no interrupts pending
if (CSR0 != 0) exit; If channel is not ready, exit
; Start transfer
NDAR0 = &d ; Set up descriptor address
CCR0 = 0x00000001 ; Set Channel Enable bit to start

#### Figure 41. Software Example for Local Memory-to-Local Memory DMA Transfer

March 2005



## 6.10.3 Software Suspend Channel

The channel may need to be suspended for various reasons. The channel provides the ability to suspend the state of the channel without losing the current status. The channel resumes DMA operation without requiring the software to save the channel configuration. The example shown in Figure 42 describes the pseudo code for suspending channel 0.

#### Figure 42. Software Example for Channel Suspend

CCR0	=	0x0000	0000	;	Suspend	Channel		0
Cł	ha	nnel su	Ispend	lec	1			
CCR0	=	0x0000	0001	;	Resume	Channel	0	



## 6.11 Interrupts

Each channel can generate an interrupt to the Intel XScale<sup>®</sup> core. The *Interrupt Enable* bit in the Descriptor Control Register (DCR) determines whether the channel generates an interrupt upon successful error-free completion of a DMA transfer. Error conditions described in Section 6.12 also generate an interrupt. Each channel has one interrupt output connected to the PCI and Peripheral Interrupt Controller described in Chapter 17, "Interrupt Controller Unit and IOAPIC" summarizes the status flags and conditions when interrupts are generated in the Channel Status Register (CSRx).

Table 232.DMA Interrupt Summary

	Channel Status Register (CSR) Flags								Interrupt Generated?		
Interrupt Condition	Active	End of Transfer	End of Chain	PCI Master Abort	PCI Target Abort	Unknown Split Transaction	Internal Bus Error	Interrupt Enabled	Interrupt Disabled		
Byte count == 0 && (NDARx != NULL    Resume == 1) (End of Transfer)	1	1	0	0	0	0	0	Y	Ν		
Byte Count == 0 && NDARx == NULL (End of Chain) & resume == 0	0	0	1	0	0	0	0	Y	Ν		
PCI Master-Abort	0	0	0	1	0	0	0	Y	Y		
PCI Target-Abort	0	0	0	0	1	0	0	Y	Y		
Unknown Split Transaction Error	0	0	0	0	0	1	0	Y	Y		
Internal Bus Error	0	0	0	0	0	0	1	Y	Y		

*Note:* End-of-Transfer and End-of-Chain flags is set only when interrupt is enabled. Otherwise, the above flags are always set to 0. End-of-Transfer Interrupt and End-of-Chain Interrupt can only be reported in the CSR when the DMA transfer completed without any reportable errors. The channel shall never report an End-of-Transfer interrupt or End-of-Chain interrupt along with any PCI error conditions. Multiple error conditions may occur and be reported together. Also, because the channel does not stop after reporting the End-of- Transfer Interrupt, internal bus errors may occur before the End-of-Transfer interrupt is acknowledged and cleared.



## 6.12 Error Conditions

There are four error conditions that may occur during a DMA transfer that are recorded by the channel. All error conditions are reported by setting the appropriate bit in the Channel Status Register (CSR).

The possible error conditions are:

- PCI Master-Abort
- PCI Target-Abort
- Unknown PCI-X Split Transaction Error
- Internal Bus Errors

#### 6.12.1 PCI Errors

- When a PCI Master-Abort occurs during a PCI-to-Local DMA transfer, the channel sets bit 3 in the CSR. Refer to Section 3.7, "ATU Error Conditions" on page 200 for complete details on the ATU error response.
- When a PCI Target-Abort occurs during a PCI-to-Local DMA transfer, the channel sets bit 2 in the CSR. The ATU also records this PCI Bus error condition by setting the appropriate bit in its' status register. Refer to Section 3.7, "ATU Error Conditions" on page 200 for complete details on the ATU error response.
- When an Unknown PCI-X Split transaction error occurs during a PCI-to-Local DMA transfer, the channel sets bit 1 in the CSR. For PCI parity errors, data with incorrect parity is never transferred to the DMA. Refer to Section 3.7, "ATU Error Conditions" on page 200 for complete details on the ATU error response.
- *Note:* Errors (including address/data parity errors) that occur on the PCI bus during Local-to-PCI or PCI-to-Local transfers is logged by the ATU. Refer to Chapter 3, "Address Translation Unit" for complete details.
# 6.12.2 Internal Bus Errors

- When a Master-Abort occurs during a read of the Chain Descriptor or Next Descriptor Address, the channel sets the Internal Bus Master-abort error flag which is bit 5 in the CSR. Then, the channel loads the registers (when possible) and stop.
- When a Master-Abort occurs during a data transfer, the channel sets the Internal Bus Master-abort error flag which is bit 5 in the CSR. Then, the channel stops operation.
- When a Target-Abort occurs, the DMA stops operation, but the error is recorded by the MCU.

When an error condition occurs, the actions taken are detailed below:

- The channel shall cease data transfers for the current chain descriptor and clear the *Channel Active* flag in the CSR.
- The channel invalidates any data held in the queue and not read any new chain descriptors.
- The channel sets the appropriate error flag in the Channel Status Register. For example; when a PCI Master-Abort occurred during a DMA transfer, the channel sets bit 3 in the CSR.
- The channel also signals an interrupt to the Intel XScale<sup>®</sup> core.
- The channel does not restart a DMA transfer after any error condition. It is the responsibility of the application software to configure the channel to complete any remaining transfers.
- *Note:* Internal Bus errors (**Target-abort only**) that occur while a DMA channel is the master on the Internal Bus are recorded by the MCU and interrupt the core. For correct operation of the DMA channel, user software has to disable the channel before clearing the error condition. Further, the channel needs to be re-enabled by writing a 1 to the CCR channel enable before initiating a new operation.



# 6.13 **Power-up/Default Status**

Upon power-up, an external hardware reset, the DMA Registers is initialized to their default values.

# 6.14 Register Definitions

The DMA controller contains registers for controlling each channel. Each channel has nine memory-mapped control registers for independent operation. In register titles, x is 0 or 1 for channel 0 or 1, respectively.

There is read/write access only to the Channel Control Register, Channel Status Register, and the Next Descriptor Address Register. The remaining registers are read-only and are loaded with new values from the chain descriptor whenever the channel reads a chain descriptor from memory.

#### Table 233. DMA Controller Unit Registers

Section, Register Name, Acronym, Page
Section 6.14.1, "Channel Control Register x - CCRx" on page 399
Section 6.14.2, "Channel Status Register x - CSRx" on page 400
Section 6.14.4, "Next Descriptor Address Register x - NDARs" on page 402
Section 6.14.3, "Descriptor Address Register x - DARx" on page 401
Section 6.14.8, "Byte Count Register x - BCRx" on page 406
Section 6.14.5, "PCI Address Register x - PADRx" on page 403
Section 6.14.6, "PCI Upper Address Register x - PUADRx" on page 404
Section 6.14.7, "Local Address Register x - LADRx" on page 405
Section 6.14.9, "Descriptor Control Register x - DCRx" on page 407



### 6.14.1 Channel Control Register x - CCRx

The Channel Control Register (CCRx) specifies parameters that dictate the overall channel operating environment. The CCRx should be initialized prior to any other DMA register following a system reset. Figure 32 shows the register format for the CCRx. This register can be read or written while the DMA channel is active.



#### Table 234. Channel Control Register x - CCRx



### 6.14.2 Channel Status Register x - CSRx

The Channel Status Register (CSRx) contains status flags that indicate the channel status. This register is typically read by software to examine the source of an interrupt. See Section 6.12 for a description of the error conditions that are reported in the CSRx. See Section 6.11 for a description of interrupts caused by the DMA channel.

When a DMA error occurs, application software must check the status of the Channel Active flag before processing the interrupt. It is possible that the channel may still be active completing outstanding PCI transactions.



#### Table 235. Channel Status Register x - CSRx



### 6.14.3 Descriptor Address Register x - DARx

The Descriptor Address Register (DARx) contains the address of the current chain descriptor in 80333 local memory for a DMA transfer. This register is read-only and is loaded when a new chain descriptor is read. Table 236 depicts the Descriptor Address Register.

All chain descriptors are aligned on an eight DWORD boundary.







# 6.14.4 Next Descriptor Address Register x - NDARs

The Next Descriptor Address Register (NDARx) contains the address of the next chain descriptor in 80333 local memory for a DMA transfer. When starting a DMA transfer, this register contains the address of the first chain descriptor. Table 237 depicts the Next Descriptor Address Register.

All chain descriptors are required to be aligned on an eight DWORD boundary. The channel may set bits 04:00 to zero when loading this register.

*Note:* The *Channel Enable* bit in the CCRx and the *Channel Active* bit in the CSRx must both be clear prior to writing the Next Descriptor Address Register. Writing a value to this register while the channel is active may result in undefined behavior.



#### Table 237. Next Descriptor Address Register x - NDARs

# 6.14.5 PCI Address Register x - PADRx

The PCI Address Register (PADRx) contains the 32-bit PCI address for SAC cycles or the lower 32-bit PCI address of a 64-bit PCI address for DAC cycles. This address represents the source or destination of the DMA transfer.

Table 238 shows the PCI Address Register.

The channel uses the full address bus to indicate the starting address. Refer to the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a for additional information.

*Note:* The application programmer must not program the channel to transfer data across a 4 Gbyte boundary (i.e., the lower 32-bit address must not increment past the maximum address of FFFF.FFFFH). The channel does not notify the application of this condition.



#### Table 238. PCI Address Register x - PADRx



# 6.14.6 PCI Upper Address Register x - PUADRx

The PCI Upper Address Register (PUADRx) contains the upper 32-bits of a 64-bit PCI address. Table 239 shows the register. This register is read-only and is loaded when a chain descriptor is read from memory







### 6.14.7 Local Address Register x - LADRx

The 80333 Local Address Register (LADRx) contains the 32-bit 80333 local address. The 80333 address space is a 32-bit, byte addressable address space. Table 240 shows the Local Address Register. This read-only register is loaded when a chain descriptor is read from memory.







# 6.14.8 Byte Count Register x - BCRx

The Byte Count Register (BCRx) contains the number of bytes to transfer for a DMA transfer. This is a read-only register that is loaded from the Byte Count word in a chain descriptor. It allows for a maximum DMA transfer of 16 Mbytes. A value of zero is a valid byte count and results in no data words being transferred and no cycles generated on either the PCI bus or the internal bus. This register specifies the amount of data that is moved from the source to the destination.

Table 241 shows the Byte Count Register

*Note:* The byte count value is not required to be aligned to a DWORD boundary.



#### Table 241.Byte Count Register x - BCRx



### 6.14.9 Descriptor Control Register x - DCRx

The Descriptor Control Register (DCRx) contains control values for the DMA transfer on a per-chain descriptor basis. This read-only register is loaded whenever a chain descriptor is read from memory. These values may vary from chain descriptor to chain descriptor. Table 242 shows the definition of the Descriptor Control Register.



#### Table 242. Descriptor Control Register x - DCRx

### 6.14.9.1 PCI Transactions Support

The Memory Write Block Command on the bus in PCI-X mode performs similarly to the PCI Memory Write and Invalidate command and is fully supported by both channels of the DMA controller. Refer to Section 6.4.2, "Local Memory to PCI Transfers" on page 386 for a complete description of the behavior of the DMA channel during this PCI bus cycle.

DMA Internal Bus transactions that are destined for the external PCI Bus are automatically forwarded through the Outbound ATU to the PCI bus. In this way, the backward compatibility with the previous generation programming model defined for the DMA descriptors is preserved.

#### Table 243.

DCR[3:0]	PCI Conventional Mode	PCI-X Mode
6H, CH, or EH	MR, MRL, and MRM as appropriate	Memory Read Block is used
7H or FH	Memory Write and Memory Write Invalidate Command as appropriate.	Memory Write Block is use.
all other values	Reserved	

# **Application Accelerator Unit**

This chapter describes the integrated Application Accelerator (AA) Unit. The operation modes, setup, external interface, and implementation of the AA unit are detailed in this chapter.

# 7.1 Overview

The Application Accelerator provides low-latency, high-throughput data transfer capability between the AA unit and 80333 local memory. It executes data transfers to and from 80333 local memory, checks for all-zero result across local memory blocks, performs memory block fills, and provides the necessary programming interface. The Application Accelerator performs the following functions:

- Transfers data (read) from memory controller.
- Performs an optional boolean operation (XOR) on read data.
- Transfers data (write) to memory controller or PCI.
- Checks for All-zero result across local memory blocks.
- Performs memory block fills.
- Optional Dual-XOR for XOR-based RAID-6 application single strip write.
- Optional Galois Field (GF) Multiply calculation for P+Q RAID-6 in conjunction with XOR operations.

The AA unit features:

- 1Kbyte/512-byte store queue.
- Utilization of the 80333 memory controller Interface.
- $2^{32}$  addressing range on the 80333 local memory interface.
- Hardware support for unaligned data transfers for the internal bus.
- Fully programmable from the Intel XScale<sup>®</sup> core.
- Support for automatic data chaining for gathering and scattering of data blocks.
- Support for writing a constant value to a memory block (block fill).
- Support for writing descriptor status to local memory.
- Hardware to perform Galois Field (GF) Multiply function on the Source Data Streams during an XOR operation, when enabled.

# 7.2 Theory of Operation

The Application Accelerator is a master on the internal bus and performs data transfers to and from local memory. It does not interface to the PCI bus. AA uses direct addressing for memory controller.

The AA implements XOR algorithm in hardware. It performs XOR operation on multiple blocks of source (incoming) data and stores result back in 80333 local memory. The source and destination addresses are specified through chain descriptors resident in 80333 local memory. A Dual-XOR operation is also supported for optimized processing of two different, but related XOR operations in a single operation. The AA can also check for all-zero result across local memory blocks or fill a memory block with arbitrary data. Figure 43 shows a block diagram of the AA unit. The AA can also perform memory-to-memory transfers of data blocks controlled by 80333 memory controller unit.

#### Figure 43. Application Accelerator Block Diagram



AA programming interface is accessible from the internal bus through a memory-mapped register interface. Data for XOR operation is configured by writing source addresses, destination address, number of bytes to transfer, and various control information into a local memory chain descriptor. Chain descriptors are described in detail in Section 7.3.2, "Chain Descriptors" on page 413.

The AA unit contains a hardware data packing and unpacking unit. This unit enables data transfers from and to unaligned addresses in 80333 local memory. All combinations of unaligned data are supported with the packing and unpacking unit. Data is held internally in the AA until ready to be stored back to local memory. This is done using a 1KByte/512Byte holding queue. Data to be written back to 80333 local memory can either be aligned or unaligned.

Each chain descriptor contains necessary information for initiating an XOR operation on blocks of data specified by the source addresses. The AA unit supports chaining. Chain descriptors that specify the source data to be XORed can be linked together in 80333 local memory to form a linked list.

Similar to XOR operations, AA can be programed to compute parity across multiple memory blocks specified by chain descriptors. In addition, AA is also used for memory block fills. A Dual-XOR operation is available for use when calculating two parity blocks for a XOR-based RAID-6 single strip write.

In conjunction with the XOR, a GF Multiply calculation can be applied to source data in support of P+Q RAID-6. The AA will perform a GF Multiply between source data and a control byte for each source before the XOR operation when enabled. P+Q RAID-6 is enabled through an enable bit in the "Accelerator Control Register - ACR" (bit 3).



# 7.3 Hardware-Assist XOR Unit

The AAU implements the XOR algorithm in hardware. It performs the XOR operation on multiple blocks of source (incoming) data and stores the result back in 80333 local memory.

- The process of reading source data, executing the XOR algorithm, and storing the XOR data will hereafter is referred to as *XOR-transfer*.
- The operation of two *XOR-transfers* defined in a single descriptor will hereafter be referred to as *Dual-XOR transfer*.
- The process of reading or writing data will hereafter is referred to as *data transfer*.

Source and destination addresses specified through chain descriptors resident in 80333 local memory.

### 7.3.1 Data Transfer

All transfers are configured and initiated through a set of memory-mapped registers and one or more chain descriptors located in local memory. A transfer is defined by the source address, destination address, number of bytes to transfer, and control values. These values are loaded in the chain descriptor before a transfer begins. Table 244 describes the registers that need to be configured for any operation.

When P+Q RAID-6 is enabled, the GF Multiplier bytes also act as control values in the data transfer.

Register	Abbreviation	Description
Accelerator Control Register	ACR	Application Accelerator Control Word
Accelerator Status Register	ASR	Application Accelerator Status Word
Accelerator Descriptor Address Register	ADAR	Address of Current Chain Descriptor
Accelerator Next Descriptor Address Register	ANDAR	Address of Next Chain Descriptor
Data / Source Address Register1	D/SAR1	Data to be written or Local memory addresses of source data
Source Address Register 232	SAR2 SAR32	Local memory addresses of source data
P+Q RAID-6 Source Address Register 216	PQSAR2 PQSAR16	Local memory addresses of source data when operating in P+Q RAID-6 mode
P+Q RAID-6 GF Multiplier Register 15,D	GFMR1 GFMR5, GFMRD	P+Q RAID-6 GF Multiplier bytes when operating in P+Q RAID-6 mode
Destination Address Register	DAR	Address of result data
Accelerator Byte Count Register	ABCR	Number of Bytes to transfer
Data Register	DR	Data to be written to the destination
Accelerator Descriptor Control Register	ADCR	Chain Descriptor Control Word

#### Table 244. Register Description



### 7.3.2 Chain Descriptors

All transfers are controlled by chain descriptors located in local memory. A chain descriptor contains the necessary information to complete one transfer. A single transfer has only one chain descriptor in memory. Chain descriptors can be linked together to form more complex operations.

*Warning:* Chain descriptors can only be located in DDR SDRAM memory in order for the AAU to function properly. Location of the chain descriptors anywhere else (e.g., either on the Peripheral Bus or on PCI) is not supported and the results would be unpredictable.

To perform a transfer, one or more chain descriptors must first be written to 80333 local memory. The words of a descriptor are contiguous in local memory. Descriptors can be different sizes, each dependent on the number of sources being addressed by the operation. The sizes supported by the Application Accelerator are four, eight, sixteen and thirty-two sources. The alignment of the descriptor in local memory is dependent on the descriptor size and is defined for each in the following sections. Not all sources must be used in a given descriptor.

Descriptor formats for P+Q RAID-6 enabled operation are defined separately. When P+Q RAID-6 is enabled, only the P+Q RAID-6 descriptor formats are valid.



#### 7.3.2.1 Principle / Four-Source Descriptor Format

Figure 44 shows the format of an individual chain descriptor. This four-source descriptor is the smallest supported descriptor. The four-source descriptor requires eight contiguous words in 80333 local memory and is required to be aligned on an 8-word boundary. All eight words are required.

#### Figure 44. Principle / Four Source Descriptor Format

Chain Descriptor in Local Memory	Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (D/SAR1)	Immediate Data or Source Address for first block of data
Source Address (SAR2)	Source Address for second block of data
Source Address (SAR3)	Source Address for third block of data
Source Address (SAR4)	Source Address for fourth block of data
Destination Address (DAR)	Destination Address
Byte Count (BC)	Number of bytes
Descriptor Control (DC)	Descriptor Control

Each word in the chain descriptor is analogous to control register values. Bit definitions for the words in the chain descriptor are the same as for the control registers.

- First word is local memory address of next chain descriptor. A value of zero specifies the end of the chain. This value is loaded into the Accelerator Next Descriptor Address Register. Because chain descriptors must be aligned on a minimum 8-word boundary, the unit may ignore bits 04:00 of this address.
- Second word is address of the first block of data resident in local memory, or immediate data when performing a Memory Block Fill. This value is loaded into the Data / Source Address Register 1.
- Third word is the address of the second block of data resident in local memory. This address is driven on the internal bus. This value is loaded into the Source Address Register 2.
- Fourth word is the address of the third block of data resident in local memory. This address is driven on the internal bus. This value is loaded into the Source Address Register 3.
- Fifth word is the address of the fourth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into the Source Address Register 4.
- Sixth word is the destination address where data is stored in local memory or PCI. This address is driven on the internal bus. This value is loaded into the Destination Address Register.
- Seventh word is the Byte Count value. This value specifies the number of bytes of data in the current chain descriptor. This value is loaded into the Accelerator Byte Count Register.
- Eighth word is the Descriptor Control Word. This word configures the Application Accelerator for one operation. This value is loaded into the Accelerator Descriptor Control Register.

There are no data alignment requirements for any source addresses or destination address. However, maximum performance is obtained from aligned transfers, especially small transfers. See Section 7.4.

Refer to Section 7.13 for additional description on the control registers.

### 7.3.2.2 Eight-Source Descriptor Format

To perform an *XOR-transfer* with up to eight source blocks of data, a special chain descriptor needs to be configured:

- The first part is the four-source descriptor (referred to as the *principal-descriptor*) containing the address of the first 4 source data blocks along with other information.
- The second part (*mini-descriptor*) contains 4, DWORDs containing the address of the additional four (SAR5 SAR8) source data blocks. The mini-descriptor is written to a contiguous address immediately following the principal descriptor.

To perform a transfer, both parts (principal and mini-descriptor) must be written to local memory. Figure 45 shows the format of this eight-source descriptor. The eight-source descriptor requires twelve contiguous words in local memory and is required to be aligned on an 16-word boundary. All twelve words are required.

#### Figure 45. Chain Descriptor Format for Eight Source Addresses (XOR Function)

Chain Descriptor in I/O Processor Memory	Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (D/SAR1)	Source Address for first block of data
Source Address (SAR2)	Source Address for second block of data
Source Address (SAR3)	Source Address for third block of data
Source Address (SAR4)	Source Address for fourth block of data
Destination Address (DAR)	Destination Address of XOR-ed data
Byte Count (BC)	Number of bytes to XOR
Descriptor Control (DC)	Descriptor Control
Source Address (SAR5)	Source Address for fifth data block
Source Address (SAR6)	Source Address for sixth data block
Source Address (SAR7)	Source Address for seventh data block
Source Address (SAR8)	Source Address for eighth data block

- The first eight words are defined in the four-source descriptor definition. See Section 7.3.2.1, "Principle / Four-Source Descriptor Format" for the definition of these words.
- The ninth word (1st word of mini-descriptor) is the address of the fifth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR5.
- The tenth word (2nd word of mini-descriptor) is the address of the sixth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR6.
- The eleventh word (3rd word of mini-descriptor) is the address of the seventh block of data resident in local memory. This address is driven on the internal bus. This value is loaded SAR7.
- The twelfth word (4th word of mini-descriptor) is the address of the eighth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 8.



#### 7.3.2.3 Sixteen-Source Descriptor Format

To perform an *XOR-transfer* with up to sixteen source blocks of data, a special chain descriptor needs to be configured:

- The first part (*principal-descriptor*) contains the address of the first 4 source data blocks along with other information.
- The second part (*mini-descriptor*) contains four, DWORDs containing the address of the additional four (SAR5 SAR8) source data blocks. The mini-descriptor is written to a contiguous address immediately following the principal descriptor.
- The third part (*extended-descriptor 0*) contains nine, DWORDs containing the address of the additional eight (SAR9 SAR16) source data blocks and the command/control for these data blocks. The extended-descriptor 0 is written to a contiguous address immediately following the mini descriptor.

To perform a transfer, all three parts (principal descriptor, mini-descriptor and extended-descriptor 0) must be written to local memory. Figure 46 shows the format of this configuration. Every descriptor requires twenty one contiguous words in local memory and is required to be aligned on an 32-word boundary. All twenty one words are required.

#### Figure 46. Chain Descriptor Format for Sixteen Source Addresses (XOR Function)

Chain Descriptor in Intel XScale® Core Memo	ory Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (D/SAR1)	Source Address for first block of data
Source Address (SAR2)	Source Address for second block of data
Source Address (SAR3)	Source Address for third block of data
Source Address (SAR4)	Source Address for fourth block of data
Destination Address (DAR)	Destination Address of XOR-ed data
Byte Count (BC)	Number of bytes to XOR
Descriptor Control (DC)	Descriptor Control
Source Address (SAR5)	Source Address for fifth data block
Source Address (SAR6)	Source Address for sixth data block
Source Address (SAR7)	Source Address for seventh data block
Source Address (SAR8)	Source Address for eighth data block
Extended Descriptor Control 0 (EDC0)	Extended Descriptor 0 control
Source Address (SAR9)	Source Address for ninth block of data
Source Address (SAR10)	Source Address for tenth block of data
Source Address (SAR11)	Source Address for eleventh block of data
Source Address (SAR12)	Source Address for twelfth block of data
Source Address (SAR13)	Source Address for thirteenth block of data
Source Address (SAR14)	Source Address for fourteenth block of data
Source Address (SAR15)	Source Address for fifteenth block of data
Source Address (SAR16)	Source Address for sixteenth block of data

- The first eight words are defined in the four-source descriptor definition. See Section 7.3.2.1, "Principle / Four-Source Descriptor Format" for the definition of these words.
- Words nine through twelve are defined in the eight-source descriptor definition. See Section 7.3.2.2, "Eight-Source Descriptor Format" for the definition of these words.
- The thirteenth word (1st word of extended-descriptor 0) is the Extended Descriptor Control Word 0. This word configures the Application Accelerator for one operation. The value is loaded into the Extended Descriptor Control Register 0.
- The fourteenth word (2nd word of extended-descriptor 0) is the address of the ninth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 9.
- The fifteenth word (3rd word of extended-descriptor 0) is the address of the tenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 10.
- The sixteenth word (4th word of extended-descriptor 0) is the address of the eleventh block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 11.
- The seventeenth word (5th word of extended-descriptor 0) is the address of the twelfth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 12.
- The eighteenth word (6th word of extended-descriptor 0) is the address of the thirteenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 13.
- The nineteenth word (7th word of extended-descriptor 0) is the address of the fourteenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 14.
- The twentieth word (8th word of extended-descriptor 0) is the address of the fifteenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 15.
- The twenty first word (9th word of extended-descriptor 0) is the address of the sixteenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 16.



#### 7.3.2.4 Thirty-two-Source Descriptor Format

To perform an *XOR-transfer* with up to thirty two source blocks of data, a special chain descriptor needs to be configured:

- The first part (*principal-descriptor*) contains the address of the first 4 source data blocks along with other information.
- The second part (*mini-descriptor*) contains four, DWORDs containing the address of the additional four (SAR5 SAR8) source data blocks. The mini-descriptor is written to a contiguous address immediately following the principal descriptor.
- The third part (*extended-descriptor 0*) contains nine, DWORDs containing the address of the additional eight (SAR9 SAR16) source data blocks and the command/control for these data blocks. The extended-descriptor 0 is written to a contiguous address immediately following the mini descriptor.
- The fourth part (*extended-descriptor 1*) contains nine, DWORDs containing the address of the additional eight (SAR17 SAR24) source data blocks and the command/control for these data blocks. The extended-descriptor 1 is written to a contiguous address immediately following extended-descriptor 0.
- The fifth part (*extended-descriptor 2*) contains nine, DWORDs containing the address of the additional eight (SAR25 SAR32) source data blocks and the command/control for these data blocks. The extended-descriptor 2 is written to a contiguous address immediately following extended-descriptor 1.

To perform a transfer, all five parts (principal descriptor, mini-descriptor, extended-descriptor 0, extended-descriptor 1, and extended-descriptor 2) must be written to local memory. Figure 47 shows the format of this configuration. The full descriptor requires thirty nine contiguous words in local memory and is required to be aligned on an 64-word boundary. All thirty nine words are required.

#### Figure 47. Chain Descriptor Format for Thirty Two Source Addresses (XOR Function)

Chain Descriptor in I/O Processor Memory	Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (D/SAR1)	Source Address for first block of data
Source Address (SAR2)	Source Address for second block of data
Source Address (SAR3)	Source Address for third block of data
Source Address (SAR4)	Source Address for fourth block of data
Destination Address (DAR)	Destination Address of XOR-ed data
Byte Count (BC)	Number of bytes to XOR
Descriptor Control (DC)	Descriptor Control
Source Address (SAR5)	Source Address for fifth data block
Source Address (SAR6)	Source Address for sixth data block
Source Address (SAR7)	Source Address for seventh data block
Source Address (SAR8)	Source Address for eighth data block
Extended Descriptor Control 0 (EDC0)	Extended Descriptor 0 control
Source Address (SAR9)	Source Address for ninth block of data
Source Address (SAR10)	Source Address for tenth block of data
Source Address (SAR11)	Source Address for eleventh block of data
Source Address (SAR12)	Source Address for twelfth block of data
Source Address (SAR13)	Source Address for thirteenth block of data
Source Address (SAR14)	Source Address for fourteenth block of data
Source Address (SAR15)	Source Address for fifteenth block of data
Source Address (SAR16)	Source Address for sixteenth block of data
Extended Descriptor Control 1 (EDC1)	Extended Descriptor 1 control
Source Address (SAR17)	Source Address for seventeenth block of data
Source Address (SAR18)	Source Address for eighteenth block of data
Source Address (SAR19)	Source Address for nineteenth block of data
Source Address (SAR20)	Source Address for twentieth block of data
Source Address (SAR21)	Source Address for twenty first block of data
Source Address (SAR22)	Source Address for twenty second block of data
Source Address (SAR23)	Source Address for twenty third block of data
Source Address (SAR24)	Source Address for twenty fourth block of data
Extended Descriptor Control 2 (EDC2)	Extended Descriptor 2 control
Source Address (SAR25)	Source Address for twenty fifth block of data
Source Address (SAR26)	Source Address for twenty sixth block of data
Source Address (SAR27)	Source Address for twenty seventh block of data
Source Address (SAR28)	Source Address for twenty eighth block of data
Source Address (SAR29)	Source Address for twenty ninth block of data
Source Address (SAR30) Source Address (SAR31)	Source Address for thirty first block of data
Source Address (SAR32)	Source Address for thirty second block of data
· · ·	

March 2005



- The first eight words are defined in the Four-source descriptor definition. See Section 7.3.2.1, "Principle / Four-Source Descriptor Format" for the definition of these words.
- Words nine through twelve are defined in the Eight-source descriptor definition. See Section 7.3.2.2, "Eight-Source Descriptor Format" for the definition of these words.
- Words thirteen through twenty-one are defined in the Sixteen-source descriptor definition. See Section 7.3.2.3, "Sixteen-Source Descriptor Format" for the definition of these words.
- The twenty second word (1st word of extended-descriptor 1) is the Extended Descriptor Control Word 1. This word configures the Application Accelerator for one operation. The value is loaded into the Extended Descriptor Control Register 1.
- The twenty third word (2nd word of extended-descriptor 1) is the address of the seventeenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR17.
- The twenty fourth word (3rd word of extended-descriptor 1) is the address of the eighteenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 18.
- The twenty fifth word (4th word of extended-descriptor 1) is the address of the nineteenth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 19.
- The twenty sixth word (5th word of extended-descriptor 1) is the address of the twentieth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 20.
- The twenty seventh word (6th word of extended-descriptor 1) is the address of the twenty first block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 21.
- The twenty eighth word (7th word of extended-descriptor 1) is the address of the twenty second block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 22.
- The twenty ninth word (8th word of extended-descriptor 1) is the address of the twenty third block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 23.
- The thirtieth word (9th word of extended-descriptor 1) is the address of the twenty fourth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 24.
- The thirty first word (1st word of extended-descriptor 2) is the Extended Descriptor Control Word 2. This word configures the Application Accelerator for one operation. The value is loaded into the Extended Descriptor Control Register 2.
- The thirty second word (2nd word of extended-descriptor 2) is the address of the twenty fifth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 25.
- The thirty third word (3rd word of extended-descriptor 2) is the address of the twenty sixth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 26.
- The thirty fourth word (4th word of extended-descriptor 2) is the address of the twenty seventh block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 27.

- The thirty fifth word (5th word of extended-descriptor 2) is the address of the twenty eighth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 28.
- The thirty sixth word (6th word of extended-descriptor 2) is the address of the twenty ninth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 29.
- The thirty seventh word (7th word of extended-descriptor 2) is the address of the thirtieth block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 30.
- The thirty eighth word (8th word of extended-descriptor 2) is the address of the thirty first block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 31.
- The thirty ninth word (9th word of extended-descriptor 2) is the address of the thirty second block of data resident in local memory. This address is driven on the internal bus. This value is loaded into SAR 32.



#### 7.3.2.5 Dual-XOR-Transfer Descriptor Format

To perform a *Dual-XOR-transfer*, a special chain descriptor needs to be configured:

- The descriptor contains addresses for 4 source data blocks and 2 destination data buffers along with other information.
- This descriptor format is only valid when the *Dual-XOR-transfer* Enable bit (bit 27) in the Descriptor Control word is set.
- The format is based on the Eight Source Descriptor for *XOR-transfers*, and the control registers of the corresponding words take on a different meaning when processing this descriptor.

Figure 48 shows the format of this *Dual-XOR-transfer* descriptor. The *Dual-XOR-transfer* descriptor requires nine contiguous words in local memory and is required to be aligned on an 16-word boundary. All nine words are required.

#### Figure 48. Chain Descriptor Format for Dual-XOR-transfer

Chain Descriptor in I/O Processor Memory	Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (SAR1)	Source Address for first block of data
Source Address (SAR2)	Source Address for second block of data
Horizontal Source Address (SAR_H)	Source Address for Horizontal data
Diagonal Source Address (SAR_D)	Source Address for Diagonal data
Horizontal Destination Address (DAR_H)	Destination Address of Horizontal XOR-ed data
Byte Count (BC)	Number of bytes to XOR
Descriptor Control (DC)	Descriptor Control
Diagonal Destination Address (DAR_D)	Destination Address for Diagonal XOR-ed data

Each word in the chain descriptor is analogous to control register values. Bit definitions for the words in the chain descriptor are the same as for the control registers.

- First word is local memory address of next chain descriptor. A value of zero specifies the end of the chain. This value is loaded into the Accelerator Next Descriptor Address Register. Because chain descriptors must be aligned on a minimum 8-word boundary, the unit may ignore bits 04:00 of this address.
- Second word is the address of the first block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into the Data / Source Address Register 1 (SAR1).
- Third word is the address of the second block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into the Source Address Register 2 (SAR2)
- Fourth word is the address of the third block of data resident in local memory. This address will be driven on the internal bus. This source is referred to as the Horizontal source, and is associated with the Horizontal Destination for the *Dual-XOR-transfer*. This value is loaded into the Horizontal Source Address Register (SAR3/SAR\_H).
- Fifth word is the address of the fourth block of data resident in local memory. This address will be driven on the internal bus. This source is referred to as the Diagonal source, and is associated with the Diagonal Destination for the *Dual-XOR-transfer*. This value is loaded into the Diagonal Source Address Register 4 (SAR4/SAR\_D).

- Sixth word is the destination address where the first XOR result will be stored in local memory. This address will be driven on the internal bus. This destination is referred to as the Horizontal Destination, and is associated with the Horizontal Source for the *Dual-XOR-transfer*. This value is loaded into the Destination Address Register (DAR/DAR\_H).
- Seventh word is the Byte Count value. This value specifies the number of bytes of data in the current chain descriptor. This value is loaded into the Accelerator Byte Count Register.
- Eighth word is the Descriptor Control Word. This word configures the Application Accelerator for one operation. This value is loaded into the Accelerator Descriptor Control Register.
- The ninth word is the destination address where the second XOR result will be stored in local memory. This address will be driven on the internal bus. This destination is referred to as the Diagonal Destination, and is associated with the Diagonal Source for the *Dual-XOR-transfer*. This value is loaded into the Diagonal Destination Address Register (SAR5/DAR\_D).

There are no data alignment requirements for any source addresses. While the destinations addresses (Horizontal and Diagonal) also have no data alignment requirements relative to memory, the alignment of the Horizontal and Diagonal destination addresses must match (relative to 16 Byte address).

Refer to Section 7.13 for additional description on the control registers.

#### 7.3.2.6 P+Q Three-Source Descriptor Format

Figure 49 shows the format of an individual chain descriptor when P+Q RAID-6 is enabled. This three-source descriptor is the smallest supported descriptor for P+Q RAID-6 operations. The three-source descriptor requires eight contiguous words in Intel<sup>®</sup> 80333 I/O Processor local memory and is required to be aligned on an 8-word boundary. All eight words are required.

#### Figure 49. P+Q Base Chain Descriptor Format

Chain Descriptor in Local Memory	Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (D/PQSAR1)	Immediate Data or Source Address for first block of data
Source Address (PQSAR2)	Source Address for second block of data
Source Address (PQSAR3)	Source Address for third block of data
Data Multiplier Values (PQMR1)	Data Multiplier Values for Sources 1 through 3
Destination Address (DAR)	Destination Address
Byte Count (BC)	Number of bytes
Descriptor Control (DC)	Descriptor Control

Each word in the chain descriptor is analogous to control register values. Bit definitions for the words in the chain descriptor are the same as for the control registers.

- First word is local memory address of next chain descriptor. A value of zero specifies end of chain. Value is loaded into the Accelerator Next Descriptor Address Register. Because chain descriptors must be aligned on a minimum 8-word boundary, unit may ignore bits 04:00 of this address.
- Second word is address of the first block of data resident in local memory, or immediate data when performing a Memory Block Fill. This value is loaded into the Data / P+Q RAID-6 Source Address Register 1 (D/PQSAR1).
- Third word is address of second block of data resident in local memory. This address is driven on the internal bus. This value is loaded into the P+Q RAID-6 Source Address Register 2 (PQSAR2).
- Fourth word is address of third block of data resident in local memory and is driven on the internal bus. This value is loaded into P+Q RAID-6 Source Address Register 3 (PQSAR3).
- Fifth word contains Data Multiplier Values (DMLTx) for source addresses 1 through 3. These bytes are used as control input for GF Multiply of corresponding source. The respective byte will be driven to the GF Multiply when source data is being fetched. The lowest order byte of this word contains the data multiplier for source address 1, the second byte for source address 2, the third byte for source address 3, and the highest order byte is not used and is reserved. This value is loaded into the P+Q RAID-6 GF Multiply Multiplier Register 1 (GFMR1).
- Sixth word is the destination address where data will be stored in local memory. This address will be driven on the internal bus. This value is loaded into the Destination Address Register.
- Seventh word is the Byte Count value. This value specifies the number of bytes of data in the current chain descriptor. This value is loaded into the Accelerator Byte Count Register.
- Eighth word is the Descriptor Control Word. This word configures the Application Accelerator for one operation. This value is loaded into the Accelerator Descriptor Control Register.

There are no data alignment requirements for any source addresses or destination address. However, maximum performance is obtained from aligned transfers, especially small transfers. See Section 7.4.

Refer to Section 7.13 for additional description on the control registers.



### 7.3.2.7 P+Q Six-Source Descriptor Format

To perform an P+Q RAID-6 *XOR-transfer* with up to six source blocks of data, a special chain descriptor needs to be configured:

- First part: three-source descriptor (referred to as *principal-descriptor*) containing source address and data multiplier values of first 3 source data blocks along with other information.
- Second part: (*mini-descriptor*) contains 4 DWORDs containing the address of the additional three source data blocks and Data Multiplier values. The mini-descriptor is written to a contiguous address immediately following the principal descriptor.

To perform a transfer, both parts (principal and mini-descriptor) must be written to local memory. Figure 50 shows the format of this eight-source descriptor. The six-source descriptor requires twelve contiguous words in local memory and is required to be aligned on an 16-word boundary. All twelve words are required.

#### Figure 50. P+Q Chain Descriptor Format for Six Source Addresses (XOR Function)

Chain Descriptor in I/O Processor Memory	Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (PQSAR1)	Source Address for first block of data
Source Address (PQSAR2)	Source Address for second block of data
Source Address (PQSAR3)	Source Address for third block of data
Data Multiplier Values (GFMR1)	Data Multiplier Values for Sources 1 through 3
Destination Address (DAR)	Destination Address of XOR-ed data
Byte Count (BC)	Number of bytes to XOR
Descriptor Control (DC)	Descriptor Control
Source Address (PQSAR4)	Source Address for fourth data block
Source Address (PQSAR5)	Source Address for fifth data block
Source Address (PQSAR6)	Source Address for sixth data block
Data Multiplier Values (GFMR2)	Data Multiplier Values for Sources 4 through 6

- The first eight words are defined in the three-source descriptor definition. See Section 7.3.2.6, "P+Q Three-Source Descriptor Format" for the definition of these words.
- The ninth word (1st word of mini-descriptor) is the address of the fourth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR4.
- The tenth word (2nd word of mini-descriptor) is the address of the fifth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR5.
- The eleventh word (3rd word of mini-descriptor) is the address of the sixth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded PQSAR6
- The twelfth word contains the Data Multiplier Values (DMLT) for source addresses 4 through 6. These bytes are used as the control input for the GF Multiply of the corresponding source. The respective byte will be driven to the GF Multiply when source data is being fetched. The lowest order byte of this word contains the data multiplier for source address 4, the second byte for source address 5, the third byte for source address 6, and the highest order byte is not used and is reserved. This value is loaded into the P+Q RAID-6 GF Multiply Multiplier Register 2 (GFMR2).



#### 7.3.2.8 P+Q Twelve-Source Descriptor Format

To perform an *XOR-transfer* with a GF Multiply data multiplier on up to twelve source blocks of data, a special chain descriptor needs to be configured:

- The first part (*principal-descriptor*) contains the address of the first 3 source data blocks and data multiplier values along with other information.
- The second part (*mini-descriptor*) contains four, DWORDs containing the address of the additional three source data blocks and data multiplier values. The mini-descriptor is written to a contiguous address immediately following the principal descriptor.
- The third part (*extended-descriptor 0*) contains nine, DWORDs containing the address of the additional six source data blocks and data multiplier values. The extended-descriptor 0 is written to a contiguous address immediately following the mini descriptor.

To perform a transfer, all three parts (principal descriptor, mini-descriptor and extended-descriptor 0) must be written to local memory. Figure 51 shows the format of this configuration. Every descriptor requires twenty one contiguous words in local memory and is required to be aligned on an 32-word boundary. All twenty one words are required.

#### Figure 51. P+Q Chain Descriptor Format for Twelve Source Addresses (XOR Function)

Chain Descriptor in Intel XScale <sup>®</sup> Core Memo	ory Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (D/PQSAR1)	Source Address for first block of data
Source Address (PQSAR2)	Source Address for second block of data
Source Address (PQSAR3)	Source Address for third block of data
Data Multiplier Values (GFMR1)	Data Multiplier Values for Sources 1 through 3
Destination Address (DAR)	Destination Address of XOR-ed data
Byte Count (BC)	Number of bytes to XOR
Descriptor Control (DC)	Descriptor Control
Source Address (PQSAR4)	Source Address for fourth data block
Source Address (PQSAR5)	Source Address for fifth data block
Source Address (PQSAR6)	Source Address for sixth data block
Data Multiplier Values (GFMR2	Data Multiplier Values for Sources 4 through 6
Extended Descriptor Control 0 (EDC0)	Extended Descriptor 0 control
Source Address (PQSAR7)	Source Address for seventh block of data
Source Address (PQSA8)R	Source Address for eighth block of data
Source Address (PQSAR9)	Source Address for ninth block of data
Data Multiplier Values (GFMR3)	Data Multiplier Values for Sources 7 through 9
Source Address (PQSAR10)	Source Address for tenth block of data
Source Address (PQSAR11)	Source Address for eleventh block of data
Source Address (PQSAR12)	Source Address for twelfth block of data
Data Multiplier Values (GFMR4)	Data Multiplier Values for Sources 10 through 12

- The first eight words are defined in the three-source descriptor definition. See Section 7.3.2.6, "P+Q Three-Source Descriptor Format" for the definition of these words.
- Words nine through twelve are defined in the six-source descriptor definition. See Section 7.3.2.7, "P+Q Six-Source Descriptor Format" for the definition of these words.
- The thirteenth word (1st word of extended-descriptor 0) is the Extended Descriptor Control Word 0. This word configures the Application Accelerator for one operation. The value is loaded into the Extended Descriptor Control Register 0.
- The fourteenth word (2nd word of extended-descriptor 0) is the address of the seventh block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR7.
- The fifteenth word (3rd word of extended-descriptor 0) is the address of the eighth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR8.
- The sixteenth word (4th word of extended-descriptor 0) is the address of the ninth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR9.
- The seventeenth word (5th word of extended-descriptor 0) contains the Data Multiplier Values (DMLTx) for source addresses 7 through 9. These bytes are used as the control input for the TDIfn of the corresponding source. The respective byte will be driven to the GF Multiply when source data is being fetched. The lowest order byte of this word contains the data multiplier for source address 7, the second byte for source address 8, the third byte for source address 9, and the highest order byte is not used and is reserved. This value is loaded into the P+Q RAID-6 GF Multiply Multiplier Register 3 (GFMR3).
- The eighteenth word (6th word of extended-descriptor 0) is the address of the tenth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR10.
- The nineteenth word (7th word of extended-descriptor 0) is the address of the eleventh block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR11.
- The twentieth word (8th word of extended-descriptor 0) is the address of the twelfth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR12.
- The twenty first word (9th word of extended-descriptor 0) contains the Data Multiplier Values (DMLTx) for source addresses 10 through 12. These bytes are used as the control input for the GF Multiply of the corresponding source. The respective byte will be driven to the GF Multiply when source data is being fetched. The lowest order byte of this word contains the data multiplier for source address 10, the second byte for source address 11, the third byte for source address 12, and the highest order byte is not used and is reserved. This value is loaded into the P+Q RAID-6 GF Multiply Multiplier Register 4 (GFMR4).



#### 7.3.2.9 P+Q Sixteen-Source Descriptor Format

To perform an P+Q RAID-6 *XOR-transfer* with up to sixteen source blocks of data, a special chain descriptor needs to be configured:

- The first part (*principal-descriptor*) contains the address of the first 3 source data blocks (PQSAR1 PQSAR3) and data multiplier values along with other information.
- The second part (*mini-descriptor*) contains four, DWORDs containing the address of the additional three (PQSAR4 PQSAR6) source data blocks and data multiplier values. The mini-descriptor is written to a contiguous address immediately following the principal descriptor.
- The third part (*extended-descriptor 0*) contains nine, DWORDs containing the address of the additional six (PQSAR7 PQSAR12) source data blocks and data multiplier values. The extended-descriptor 0 is written to a contiguous address immediately following the mini descriptor.
- The fourth part (*extended-descriptor 1*) contains nine, DWORDs containing the address of the additional four (PQSAR13 PQSAR16) source data blocks and data multiplier values along with the command/control for these data blocks. The extended-descriptor 1 is written to a contiguous address immediately following extended-descriptor 0.

To perform a transfer, all four parts (principal descriptor, mini-descriptor, extended-descriptor 0, and extended-descriptor 1) must be written to local memory. Figure 52 shows the format of this configuration. The full descriptor requires twenty-seven contiguous words in local memory and is required to be aligned on an 32-word boundary. All twenty-seven words are required.

• The first eight words are defined in the three-source descriptor definition. See Section 7.3.2.6,

#### Figure 52. P+Q Chain Descriptor Format for Sixteen Source Addresses (XOR Function)

Chain Descriptor in I/O Processor Memory	Description
Next Descriptor Address (NDA)	Address of Next Chain Descriptor
Source Address (D/PQSAR1)	Source Address for first block of data
Source Address (PQSAR2)	Source Address for second block of data
Source Address (PQSAR3)	Source Address for third block of data
Data Multiplier Values (GFMR1)	Data Multiplier Values for Sources 1 through 3
Destination Address (DAR)	Destination Address of XOR-ed data
Byte Count (BC)	Number of bytes to XOR
Descriptor Control (DC)	Descriptor Control
Source Address (PQSAR4)	Source Address for fourth data block
Source Address (PQSAR5)	Source Address for fifth data block
Source Address (PQSAR6)	Source Address for sixth data block
Data Multiplier Values (GFMR2)	Data Multiplier Values for Sources 4 through 6
Extended Descriptor Control 0 (EDC0)	Extended Descriptor 0 control
Source Address (PQSAR7)	Source Address for seventh block of data
Source Address (PQSAR8)	Source Address for eighth block of data
Source Address (PQSAR9)	Source Address for ninth block of data
Data Multiplier Values (GFMR3)	Data Multiplier Values for Sources 4 through 9
Source Address (PQSAR10)	Source Address for tenth block of data
Source Address (PQSAR11)	Source Address for eleventh block of data
Source Address (PQSAR12)	Source Address for twelfth block of data
Data Multiplier Values (GFMR4)	Data Multiplier Values for Sources 10 through 12
Reserved	Reserved - not used
Source Address (PQSAR13)	Source Address for thirteenth block of data
Source Address (PQSAR14)	Source Address for fourteenth block of data
Source Address (PQSAR15)	Source Address for fifteenth block of data
Source Address (PQSAR16)	Source Address for sixteenth block of data
Data Multiplier Values (GFMR5)	Data Multiplier Values for Sources 13 through 16

"P+Q Three-Source Descriptor Format" for the definition of these words.

- Words nine through twelve are defined in the six-source descriptor definition. See Section 7.3.2.7, "P+Q Six-Source Descriptor Format" for the definition of these words.
- Words thirteen through twenty-one are defined in the Twelve-source descriptor definition. See Section 7.3.2.8, "P+Q Twelve-Source Descriptor Format" for the definition of these words.
- The twenty second word (1st word of extended-descriptor 1) is reserved and not used by the AA in the processing of this descriptor.
- The twenty third word (2nd word of extended-descriptor 1) is the address of the thirteenth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR12.



- The twenty fourth word (3rd word of extended-descriptor 1) is the address of the fourteenth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR12.
- The twenty fifth word (4th word of extended-descriptor 1) is the address of the fifteenth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR12.
- The twenty sixth word (5th word of extended-descriptor 1) is the address of the sixteenth block of data resident in local memory. This address will be driven on the internal bus. This value is loaded into PQSAR12.
- The twenty seventh word (6th word of extended-descriptor 1) contains the Data Multiplier Values (DMLTx) for source addresses 13 through 16. These bytes are used as the control input for the GF Multiply of the corresponding source. The respective byte will be driven to the GF Multiply when source data is being fetched. The lowest order byte of this word contains the data multiplier for source address 13, the second byte for source address 14, the third byte for source address 15, and the highest order byte for source address 16. This value is loaded into the P+Q RAID-6 GF Multiply Multiplier Register 5 (GFMR5).
- *Note:* The highest order byte (bits[31:24]) is unused in other Data Multiplier words GFMR[4:1], but is defined in GFMR5.



- 1

# 7.3.3 Descriptor Summary

The following table summarizes the content of the descriptors defined in previous sections.

Т

Register Address	4-Source XOR	8-Source XOR	16-Source XOR	32-Source XOR	Dual XOR	P+Q 3-Source	P+Q 6-Source	P+Q 12-Source	P+Q 16-Source
FFFF E80Ch	NDA	NDA	NDA	NDA	NDA	NDA	NDA	NDA	NDA
FFFF E810h	SAR1	SAR1	SAR1	SAR1	SAR1	SAR1	SAR1	SAR1	SAR1
FFFF E814h	SAR2	SAR2	SAR2	SAR2	SAR2	SAR2	SAR2	SAR2	SAR2
FFFF E818h	SAR3	SAR3	SAR3	SAR3	SAR_H	SAR3	SAR3	SAR3	SAR3
FFFF E81Ch	SAR4	SAR4	SAR4	SAR4	SAR_D	GFMR1	GFMR1	GFMR1	GFMR1
FFFF E820h	DAR	DAR	DAR	DAR	DAR_H	DAR	DAR	DAR	DAR
FFFF E824h	BC	BC	BC	BC	BC	BC	BC	BC	BC
FFFF E828h	DC	DC	DC	DC	DC	DC	DC	DC	DC
FFFF E82Ch		SAR5	SAR5	SAR5	DAR_D		SAR4	SAR4	SAR4
FFFF E830h		SAR6	SAR6	SAR6			SAR5	SAR5	SAR5
FFFF E834h		SAR7	SAR7	SAR7			SAR6	SAR6	SAR6
FFFF E838h		SAR8	SAR8	SAR8			GFMR2	GFMR2	GFMR2
FFFF E83Ch			EDCR0	EDCR0				EDCR0	EDCR0
FFFF E840h			SAR9	SAR9				SAR7	SAR7
FFFF E844h			SAR10	SAR10				SAR8	SAR8
FFFF E848h			SAR11	SAR11				SAR9	SAR9
FFFF E84Ch			SAR12	SAR12				GFMR3	GFMR3
FFFF E850h			SAR13	SAR13				SAR10	SAR10
FFFF E854h			SAR14	SAR14				SAR11	SAR11
FFFF E858h			SAR15	SAR15				SAR12	SAR12
FFFF E85Ch			SAR16	SAR16				GFMR4	GFMR4

Table 245.Descriptor Summary (Sheet 1 of 2)

FFFF E860h	EDCR1			rsvd
FFFF E864h	SAR17			SAR13
FFFF E868h	SAR18			SAR14
FFFF E86Ch	SAR19			SAR15
FFFF E870h	SAR20			SAR16
FFFF E874h	SAR21			GFMR5
FFFF E878h	SAR22			
FFFF E87Ch	SAR23			
FFFF E880h	SAR24			
FFFF E884h	EDCR2			
FFFF E888h	SAR25			
FFFF E88Ch	SAR26			
FFFF E890h	SAR27			
FFFF E894h	SAR28			
FFFF E898h	SAR29			
FFFF E89Ch	SAR30			
FFFF E8A0h	SAR31			
FFFF E8A4h	SAR32			

#### Table 245. Descriptor Summary (Sheet 2 of 2)

intel


## 7.3.4 Descriptor Chaining

To perform an AA operation, a series of chain descriptors can be built in local memory to operate on multiple blocks of source data resident in local memory. The result can then be stored back in local memory. An application can build multiple chain descriptors to operate on many blocks of data which have different source addresses within the local memory.

When multiple chain descriptors are built in local memory, the application can link each of these chain descriptors using the Next Descriptor Address in the chain descriptor. This address logically links the chain descriptors together. This allows the application to build a list of transfers which may not require the processor until all transfers are complete. Figure 53 shows an example of a linked-list of transfers using only four-source descriptors specified in external memory.



#### Figure 53. XOR Chaining Operation



# 7.4 AA Descriptor Processing

An AA operation is initiated by building one or more chain descriptors in Intel XScale<sup>®</sup> core local memory (ARM\* architecture compliant). Figure 54 shows the format of a principal descriptor.

#### Figure 54. Example of Gather Chaining for Four Source Blocks



# intel

The following describes the steps for initiating a new AA operation:

- 1. The AA must be inactive prior to starting an AA operation. This can be checked by software by reading the *Accelerator Active* bit in the Accelerator Status Register. When this bit is clear, the unit is inactive. When this bit is set, the unit is currently active.
- 2. The ASR must be cleared of all error conditions.
- 3. The software writes the address of the first chain descriptor to the Accelerator Next Descriptor Address Register (ANDAR).
- 4. The software sets the *Accelerator Enable* bit in the Accelerator Control Register (ACR). Because this is the start of a new AA operation and not the resumption of a previous operation, the *Chain Resume* bit in the ACR should be clear.
- 5. The AA starts the AA operation by reading the chain descriptor at the address contained in ANDAR. The AA loads the chain descriptor values into the ADAR and begins data transfer. The Accelerator Descriptor Address Register (ADAR) contains the address of the chain descriptor just read and ANDAR now contains the Next Descriptor Address from the chain descriptor just read.

The last descriptor in the AA chain list has zero in the next descriptor address field specifying the last chain descriptor. A NULL value notifies the AA not to read additional chain descriptors from memory.

Once an AA operation is active, it can be temporarily suspended by clearing the *Accelerator Enable* bit in the ACR. Note that this does not abort the AA operation. The unit resumes the process when the *Accelerator Enable* bit is set.

When descriptors are read from external memory, bus latency and memory speed affect chaining latency. Chaining latency is defined as the time required for the AA to access the next chain descriptor plus the time required to set up the next AA operation.



## 7.4.1 Scatter Gather Transfers

The Application Accelerator can be used to perform typical scatter gather transfers. This consists of programming the chain descriptors to gather data which may be located in non-contiguous blocks of memory. The chain descriptor specifies the destination location such that once all data has been processed, the data is contiguous in memory. Figure 54 shows how the destination pointers can gather data.

## 7.4.2 Synchronizing a Program to Chained Operation

Any operation involving the AA can be synchronized to a program executing on the Intel XScale<sup>®</sup> core through the use of processor interrupts. The AA generates an interrupt to the Intel XScale<sup>®</sup> core under certain conditions. They are:

- 1. [Interrupt and Continue] The AA completes processing a chain descriptor and the Accelerator Next Descriptor Address Register (ANDAR) is non-zero. When the *Interrupt Enable* bit within the Accelerator Descriptor Control Register (ADCR) is set, an interrupt is generated to the Intel XScale<sup>®</sup> core. This interrupt is for synchronization purposes. The AA sets the *End Of Transfer Interrupt* flag in the Accelerator Status Register (ASR). Since it is not the last chain descriptor in the list, the AA starts to process the next chain descriptor without requiring any processor interaction.
- 2. [End of Chain] The AA completes processing a chain descriptor and the Accelerator Next Descriptor Address Register is zero specifying the end of the chain. When the *Interrupt Enable* bit within the ADCR is set, an interrupt is generated to the Intel XScale<sup>®</sup> core. The AA sets the *End Of Chain Interrupt* flag in the ASR.
- 3. [Error] An error condition occurs (refer to Section 7.11, "Error Conditions" on page 464 for Application Accelerator error conditions) during a transfer. The AA halts operation on the current chain descriptor and not proceed to the next chain descriptor.

Each chain descriptor can independently set the *Interrupt Enable* bit in the Descriptor Control word. This bit enables an independent interrupt once a chain descriptor is processed. This bit can be set or clear within each chain descriptor. Control of interrupt generation within each descriptor aids in synchronization of the executing software with AA operation.

Figure 55 shows two examples of program synchronization. The left column shows program synchronization based on individual chain descriptors. Descriptor 1A generated an interrupt to the processor, while descriptor 2A did not because the *Interrupt Enable* bit was clear. The last



descriptor nA, generated an interrupt to signify the end of the chain has been reached. The right column in Figure 55 shows an example where the interrupt was generated only on the last descriptor signifying the end of chain.







## 7.4.3 Appending to The End of a Chain

Once the AA has started processing a chain of descriptors, application software may need to append a chain descriptor to the current chain without interrupting the transfer in progress. The mechanism used for performing this action is controlled by the *Chain Resume* bit in the Accelerator Control Register (ACR).

The AA reads the subsequent chain descriptor each time it completes the current chain descriptor and the Accelerator Next Descriptor Address Register (ANDAR) is non-zero. ANDAR always contains the address of the next chain descriptor to be read and the Accelerator Descriptor Address Register (ADAR) always contains the address of the current chain descriptor.

The procedure for appending chains requires the software to find the last chain descriptor in the current chain and change the Next Descriptor Address in that descriptor to the address of the new chain to be appended. The software then sets the *Chain Resume* bit in the ACR. It does not matter when the unit is active or not.

The AA examines the *Chain Resume* bit of the ACR when the unit is idle or upon completion of a chain of transfers. When this bit is set, the AA re-reads the Next Descriptor Address of the current chain descriptor and loads it into ANDAR. The address of the current chain descriptor is contained in ADAR. The AA clears the *Chain Resume* bit and then examines ANDAR. When ANDAR is not zero, the AA reads the chain descriptor using this new address and begins a new operation. When ANDAR is zero, the AA remains or returns to idle.

There are three cases to consider:

- 1. The AA completes an AA operation and it is not the last descriptor in the chain. In this case, the AA clears the *Chain Resume* bit and reads the next chain descriptor. The appended descriptor is read when the AA reaches the end of the original chain.
- 2. The channel completes an AA transfer and it is the last descriptor in the chain. In this case, the AA examines the state of the *Chain Resume* bit. When the bit is set, the AA re-reads the current descriptor to get the address of the appended chain descriptor. When the bit is clear, the AA returns to idle.
- 3. The AA is idle. In this case, the AA examines the state of the *Chain Resume* bit when the ACR is written. When the bit is set, the AA re-reads the last descriptor from the most-recent chain to get the next descriptor address of the appended chain descriptor.

# 7.5 AA Operations

The AA can be configured on a per descriptor basis through the Descriptor Control Word to perform three distinct operations:

- 1. In an **XOR operation**, the AA generates a parity data stream in local memory that is comprised of the XOR of up to 32 distinct data streams (i.e., SAR1..32) on a per byte basis. All of the source data streams and the parity data stream can be up to 16 MB long.
- 2. In a Dual-XOR operation, the AA will generate two parity data streams in local memory.. The two parity data streams are the Horizontal and Diagonal parities for XOR based RAID-6. Each parity stream is comprised of the XOR of 2 common data streams (SAR1 and SAR2) and 1 distinct data streams (SAR\_H, and SAR\_D) respectively on a per byte basis. All of the source data streams and the parity data stream can be up to 16 MB long.
- 3. Perform a **Memory Block Fil**l of up to 16 MB of local memory with a 32-bit constant (DATA/SAR1).
- 4. With the **Zero Result Buffer Check**, the AA confirms that the XOR of all the bytes of source data results in 00H for the entire byte count. All the source data streams can be up to 16 MB long. The results of the check is written back to the Descriptor Control Word in local memory.
- *Note:* P+Q RAID-6 operation is controlled for the entire AA, and is applicable to all descriptors processed, not on a per-descriptor basis.

Table 246 documents the combination of AA operations, modes and Descriptor Control features which are valid. The typical application usage of each combination is provided. Combinations of descriptor control features not listed are not valid.

Descri	ptor Control F	eature	AA	
Zero-Result Check	Zero-Result Check Dual-XOR Destination Write Enable Comn		Operation (Source Command)	Application Usage Description
0	0	1	XOR	Typical usage for RAID Applications, up to 32 sources for RAID-3, RAID-5 and 2D-XOR RAID-6. (up to16 sources for P+Q RAID-6)
			Block Fill <sup>a</sup>	Memory Block fill with constant data
0	1	1	XOR	Two parity calculation for single strip write I/O in 2D-XOR RAID-6
1	0	0	XOR	Parity Scrub for RAID array w/o saving check buffer (typical use). Can be used in conjunction with P+Q RAID-6 mode.
1	0	1	XOR	Parity Scrub for RAID array with check buffer saved to memory (not typical use). Can be used in conjunction with P+Q RAID-6 mode.
x	х	x	Direct Fill <sup>a</sup>	First Source moved into result buffer (not XORed with current contents) Normal use for RAID applications

### Table 246. AA Operation and Command Combination Summary

a. Specified only in Block 1 Command of Accelerator Descriptor Control.

## 7.5.1 AA Addressing

All source address operated on by the AA must be local memory addresses. The destination address may be either a local memory address, or a PCI address mapped through the ATU outbound memory windows. Table 247 summarizes the application usage of the AA operations for the valid destination addressing options.

### Table 247. Typical AA Operation and Addressing Summary

Descripto Feat	or Control ure <sup>a</sup>	AA Operation	<b>Destination</b> <sup>b</sup>	Application Usage Description				
Zero-Result Check	Dual-XOR	(Source Command)	Address					
0	0	XOR	PCI	RAID Application Degraded Read				
0	0	XOR	Local	Typical usage for RAID Applications				
0	0	Block Fill <sup>c</sup>	Local	Memory Block fill with constant data				
0	1	XOR	Local	Two parity calculation for single strip write I/O in 2D-XOR RAID-6				
1	0	XOR	Local	Parity Scrub for RAID array with check buffer saved to memory. Can be used in conjunction with P+Q RAID-6 mode.				

a. Destination Write Enable set for all cases listed. Cases with DWE clear are not listed.

b. All AA sources must be local memory addresses.

c. Specified only in Block 1 Command of Accelerator Descriptor Control.

The following sections describes the AA operations in detail.



## 7.5.2 XOR Operation

Figure 56 describes the XOR algorithm implementation. In this illustrative example, there are four blocks of source data to be XOR-ed. The intermediate result is kept by the store queue in the AA before being written back to local memory. The source data is located at addresses A000 0400H, A000 0800H, A000 0C00H and A000 1000H respectively.

All data transfers needed for this operation are controlled by chain descriptors located in local memory. The Application Accelerator as a master on the internal bus initiates data transfer. The algorithm is implemented such that as data is read from local memory, the boolean unit executes the XOR operation on incoming data.





March 2005







The XOR algorithm and methodology followed once a chain descriptor has been configured is detailed below:

- 1. The Application Accelerator as a master on the bus initiates data transfer from the address pointed at by the First Source Address Register (SAR1). The total number of bytes to *XOR-transfer* is specified by the Byte Count (BC) field in the chain descriptor.
  - a. When the Direct Fill command is selected for SAR1, this is designated as the first block of data in the current XOR operation, and the data is transferred directly to the store queue. The number of bytes transferred to the store queue is 1KByte/512Bytes (based on bit 2 of the Accelerator Control Register).
  - b. When the XOR command is selected for SAR1, the boolean unit performs the XOR operation on the data currently existing in the store queue with the data being transferred from memory (see steps 3-7 for SAR2). This may be done to XOR more than 32 blocks of data together with a byte count of 1KByte or less.
- *Note:* When the Byte Count Register contains a value greater than the buffer size, the AA completes the *XOR-transfer* operation on the first buffer of data obtained from each Source Register (D/SAR1, SAR2- SAR4), then proceeds with the next buffer of data. This process is repeated until the BCR contains a zero value.
  - 2. The Application Accelerator transfers the first eight bytes of data from the address pointed at by the Second Source Address Register (SAR2).
  - 3. The boolean unit performs the bit-wise XOR algorithm on the input operands. The input operands are the first eight bytes of data read from D/SAR1 (bytes 1-8) which are stored in the queue and the first eight bytes of data just read from SAR2 (bytes 1-8).

# intel

- 4. The XOR-ed result is transferred to the store queue and stored in the first eight bytes (bytes 1-8) overwriting previously stored data.
- 5. The Application Accelerator transfers the next eight bytes of data (bytes 9-16) from address pointed at by the Second Source Address Register (SAR2).
- 6. The boolean unit performs the bit-wise XOR algorithm on the input operands. The input operands are the next eight bytes of data read from D/SAR1 (bytes 9-16 stored in the queue) and the eight bytes of data read from SAR2 in Step-5.
- 7. Step-5 and Step-6 (Data transfer and XOR) are repeated until all data pointed at by SAR1 is XOR-ed with the corresponding data pointed at by SAR2. The store queue now contains a buffer full of XOR-ed data, the source addresses for which were specified in SAR1 and SAR2.
- 8. Steps 1-7 are repeated once again. The first input to the XOR unit is the data held in the store queue and the second input is the data pointed at by SAR3.
- 9. The above steps are repeated once more. The first input to the XOR unit is the data held in the store queue and the second input is the data pointed at by SAR4.
- 10. Once Steps 1-9 are completed, the XOR operation is complete for the first full buffer of the current chain descriptor. When the Destination Write Enable Bit in the Accelerator Descriptor Control Register (ADCR) is set, the data in the store queue is written to local memory at the address pointed to by the Destination Address Register (DAR). When the Destination Write Enable Bit in the ADCR is not set, the data is not written to local memory and is held in the queue. Steps 1-9 are repeated until all the bytes of data have undergone the *XOR-transfer* operation.
- *Note:* The Destination Write Enable bit should be SET when Descriptor Byte Count is larger than the AA buffer size. When the ABCR register contains a value greater than the buffer size and the ADCR.dwe bit is cleared, the AAU only reads the first buffer of data and performs the specified function. It does not read the remaining bytes specified in the ABCR. Furthermore, the AAU proceeds to process the next chain descriptor when it is specified.



## 7.5.3 XOR Operation with P+Q RAID-6 Mode

Figure 58 describes the XOR with P+Q RAID-6 mode implementation. In this illustrative example, there are three blocks of source data to have a P+Q RAID-6 mode function performed on them followed by the an XOR function. The intermediate result is kept by the XOR store queue in the AA before being written back to local memory. The source data is located at addresses A000 0400H, A000 0800H, A000 0C00H and A000 1000H respectively.

All data transfers needed for this operation are controlled by chain descriptors located in local memory. The Application DMA as a master on the internal bus initiates a data transfer. The algorithm is implemented such that as data is read from local memory, the boolean unit executes the XOR operation on incoming data.

*Note:* Two descriptors are required for P+Q RAID-6 modes, one for each check value. Each descriptor would be processed as illustrated in Figure 58.



#### Figure 58. The Bit-wise XOR Algorithm including the P+Q RAID-6 Mode



The GF Multiply function is shown in Figure 59 for 8-bits of data. This function is replicated across each byte lane of the data path, where the G(j) input is the same for each byte lane for a given source.

#### Figure 59. **GF Multiply Function**



The blocks for gflog and gfilog are the Galois Field Logarithm and Inverse Logarithm transformations respectively. These transformations for 8-bit words are provided in Figure 60 and Figure 61, with the upper nibble used to index the rows, and the lower nibble used to index the columns. These are based on the primitive polynomial listed in Equation 5.

#### Figure 60. Galois Field Logarithm Transformation Table

	Table of gflog(2^8): gflog(xy)																
afloa	(v))	У															
griog(xy)		0	1	2	3	4	5	6	7	8	9	а	b	С	d	e	f
	0		00	01	19	02	32	1a	c6	03	df	33	ee	1b	68	c7	4b
	1	04	64	e0	0e	34	8d	ef	81	1c	c1	69	f8	c8	08	4c	71
	2	05	8a	65	2f	e1	24	Of	21	35	93	8e	da	f0	12	82	45
	3	1d	b5	c2	7d	6a	27	f9	b9	c9	9a	09	78	4d	e4	72	a6
	4	06	bf	8b	62	66	dd	30	fd	e2	98	25	b3	10	91	22	88
	5	36	d0	94	се	8f	96	db	bd	f1	d2	13	5c	83	38	46	40
	6	1e	42	b6	a3	c3	48	7e	6e	6b	3a	28	54	fa	85	ba	3d
	7	ca	5e	9b	9f	0a	15	79	2b	4e	d4	e5	ac	73	f3	а7	57
<b>^</b>	8	07	70	c0	f7	8c	80	63	0d	67	4a	de	ed	31	c5	fe	18
	9	e3	a5	99	77	26	b8	b4	7c	11	44	92	d9	23	20	89	2e
	а	37	3f	d1	5b	95	bc	cf	cd	90	87	97	b2	dc	fc	be	61
	b	f2	56	d3	ab	14	2a	5d	9e	84	3c	39	53	47	6d	41	a2
	С	1f	2d	43	d8	b7	7b	a4	76	c4	17	49	ec	7f	0c	6f	f6
	d	6c	a1	3b	52	29	9d	55	aa	fb	60	86	b1	bb	CC	3e	5a
	е	cb	59	5f	b0	9c	a9	a0	51	0b	f5	16	eb	7a	75	2c	d7
	f	4f	ae	d5	e9	e6	e7	ad	e8	74	d6	f4	ea	a8	50	58	af



Table of gfilog(2^8): gfilog(xy)																	
gfilog(xy)		<u>у</u>															
		0	1	2	3	4	5	6	7	8	9	а	b	С	d	е	f
	0	01	02	04	08	10	20	40	80	1d	3a	74	e8	cd	87	13	26
	1	4c	98	2d	5a	b4	75	ea	c9	8f	03	06	0c	18	30	60	c0
	2	9d	27	4e	9c	25	4a	94	35	6a	d4	b5	77	ee	c1	9f	23
	3	46	8c	05	0a	14	28	50	a0	5d	ba	69	d2	b9	6f	de	a1
	4	5f	be	61	c2	99	2f	5e	bc	65	са	89	Of	1e	3c	78	f0
	5	fd	e7	d3	bb	6b	d6	b1	7f	fe	e1	df	a3	5b	b6	71	e2
	6	d9	af	43	86	11	22	44	88	0d	1a	34	68	d0	bd	67	ce
<b>v</b>	7	81	1f	3e	7c	f8	ed	c7	93	3b	76	ес	c5	97	33	66	CC
<b>^</b>	8	85	17	2e	5c	b8	6d	da	a9	4f	9e	21	42	84	15	2a	54
	9	a8	4d	9a	29	52	a4	55	aa	49	92	39	72	e4	d5	b7	73
	а	e6	d1	bf	63	c6	91	3f	7e	fc	e5	d7	b3	7b	f6	f1	ff
	b	e3	db	ab	4b	96	31	62	c4	95	37	6e	dc	a5	57	ae	41
	С	82	19	32	64	c8	8d	07	0e	1c	38	70	e0	dd	a7	53	a6
	d	51	a2	59	b2	79	f2	f9	ef	c3	9b	2b	56	ac	45	8a	09
	е	12	24	48	90	3d	7a	f4	f5	f7	f3	fb	eb	cb	8b	0b	16
	f	2c	58	b0	7d	fa	e9	cf	83	1b	36	6c	d8	ad	47	8e	

### Figure 61. Galois Field Inverse Logarithm Transformation table

Equation 5. Galois Field Primitive Polynomial (0x11D)

 $X^8 + X^4 + X^3 + X^2 + 1$ 



The P+Q RAID-6 Generation matrix contains the coefficients used in calculating the P and Q check values, based on the Galois Field Matrix multiplication provided in Figure 62. Where P is the simple XOR of the source data (The result of the GF Multiply is equal to the source data with coefficient equal to 1). The Q generation coefficients (G(0), G(1), G(2),...G(N-1)), are computed from the Vandermond matrix based on the array data disk count N.

The P and Q check values are computed by the AA using separate descriptors. One descriptor is required to generate the "P" check value, where the GF Multiplier byte values are equal to 1 (first row of generation matrix). A second descriptor is required to generate the "Q" check value, where the GF Multiplier byte values are G(0), G(1), G(2),...G(N-1) (second row of generation matrix).

#### Figure 62. P+Q RAID-6 Generation Equation



## 7.5.4 Dual-XOR Operation

The Dual-XOR operation can be used in RAID-6 applications when a single strip write requires updating of two parity blocks. For this use, the two parity blocks are based on the same data being updated, and this operation performs the calculation of both updated parity blocks. In the illustrative Dual-XOR example in Figure 63, there are four blocks of source data to be XOR-ed. Two sources are used for both updated parity results, and there is one unique source for each parity result. The intermediate results are kept in store queues in the AA before being written back to local memory.

The two common source data blocks are located at addresses A000 0400H, A000 0800H. The Horizontal data source is located at address A000 0C00H and the Diagonal data source is located at address A000 1000H. The Horizontal and Diagonal destination addresses are located at addresses B000 0400H and B000 0800H respectively.

All data transfers needed for this operation are controlled by chain descriptors located in local memory. The Application Accelerator as a master on the internal bus initiates data transfer. The algorithm is implemented such that as data is read from local memory, the boolean unit executes the XOR operation on incoming data.

*Note:* Dual-XOR operation is intended for use with single strip write I/Os to XOR-based RAID-6 arrays. To generate two check values for a full stripe of data in a RAID-6 array, XOR operations defined by separate descriptors for each check value must be used. See Section 7.5.2, "XOR Operation" on page 441 for details.



Figure 63. The Bit-wise Dual-XOR Algorithm





The XOR algorithm and methodology followed once a chain descriptor has been configured is similar to that described for the basic XOR transfer. The steps followed by the AA when processing a *Dual-XOR-transfer* are detailed below:

- 1. The Application Accelerator as a master on the bus initiates data transfer from the address pointed at by the First Source Address Register (SAR1). The total number of bytes to *XOR-transfer* is specified by the Byte Count (BC) field in the chain descriptor.
- *Note:* The Direct Fill command must be selected for SAR1, designating it as the first block of data in the current Dual-XOR operation, resulting in data being transferred directly to the horizontal and diagonal store queues. The number of bytes transferred to the store queues is 1KByte/512Bytes (based on bit 2 of the Accelerator Control Register).
- *Note:* If the Byte Count Register contains a value greater than the buffer size, the AA completes the *XOR-transfer* operation on the first buffer (store queue size) of data obtained from each Source Register (SAR1, SAR2, SAR\_H, SAR\_D), then proceeds with the next buffer of data. This process is repeated until the BCR contains a zero value.
  - 2. The Application Accelerator transfers the first eight bytes of data from the address pointed at by the Second Source Address Register (SAR2).
    - a. The XOR command must be selected for SAR2, so that the boolean unit performs the XOR operation on the data currently existing in the two store queues (SAR1) with the data being transferred from memory (SAR2).
  - 3. The boolean unit performs the bit-wise XOR algorithm on the input operands. The input operands are the first eight bytes of data read from SAR1 (bytes 1-8) which are stored in the store queues and the first eight bytes of data just read from SAR2 (bytes 1-8).
  - 4. The XOR-ed result is transferred to both store queues and stored in the first eight bytes (bytes 1-8) overwriting previously stored data.
  - 5. The Application Accelerator transfers the next eight bytes of data (bytes 9-16) from address pointed at by the Second Source Address Register (SAR2).
  - 6. The boolean unit performs the bit-wise XOR algorithm on the input operands. The input operands are the next eight bytes of data read from SAR1 (bytes 9-16 stored in the queue) and the eight bytes of data read from SAR2 in Step-5.
  - 7. Step-5 and Step-6 (Data transfer and XOR) are repeated until all data pointed at by SAR1 is XOR-ed with the corresponding data pointed at by SAR2. The two store queues now both contain a buffer full of XOR-ed data, the source addresses for which were specified in SAR1 and SAR2.
  - 8. Steps 2-7 are repeated with the Horizontal Source address used for the next source data with the following exceptions
    - a. Only the Horizontal Store Queue is overwritten with the new XOR-ed result.
    - b. Upon completion, the Horizontal Store Queue holds the bit-wise XOR of Source 1 (SAR1), Source 2 (SAR2) and the Horizontal Source (SAR\_H).
    - c. The Diagonal Store Queue remains unchanged.
  - 9. Once Step 8 is completed, the XOR operation is complete for the first full buffer of the Horizontal XOR operation. The Destination Write Enable Bit in the Accelerator Descriptor Control Register (ADCR) must be set. The data in the horizontal store queue is written to local memory at the address pointed to by the Horizontal Destination Address Register (DAR\_H).

# intel

- 10. Steps 2-7 are then repeated with the Diagonal Source address used for the next source data with the following exceptions
  - a. Only the Diagonal Store Queue is overwritten with the new XOR-ed result.
  - b. Upon completion, the Diagonal Store Queue holds the bit-wise XOR of Source 1 (SAR1), Source 2 (SAR2) and the Diagonal Source (SAR\_D).
  - c. The Horizontal Store Queue remains unchanged.
- 11. Once Step 10 is completed, the XOR operation is complete for the first full buffer of the Diagonal XOR operation. The Destination Write Enable Bit in the Accelerator Descriptor Control Register (ADCR) must be set. The data in the diagonal store queue is written to local memory at the address pointed to by the Diagonal Destination Address Register (DAR\_D).
- 12. Steps 1-11 are repeated until all the bytes of data have undergone the *Dual-XOR-transfer* operation.

## 7.5.5 Zero Result Buffer Check

The AA can be used to verify parity across memory blocks specified by the SARx registers. XOR operation descriptors are used to specify the memory blocks on which the AA performs the Zero Result Buffer Check. Figure 64 illustrates a Zero Result Buffer Check performed by the AA. After processing all source data, the AA updates the Transfer Complete and Result Buffer Not Zero bit of the eighth word of the descriptor (ADCR) pointed to by the ADAR in local memory.







## 7.5.6 Zero Result Buffer Check with P+Q RAID-6

The AA can be used to verify check values for P+Q RAID-6 implementations. P+Q XOR operation descriptors are used to specify the memory blocks on which the AA performs the Zero Result Buffer Check. Figure 65 illustrates a P+Q RAID-6 Zero Result Buffer Check performed by the AA. After processing all source data, the AA updates the Transfer Complete and Result Buffer Not Zero bit of the eighth word of the descriptor (ADCR) pointed to by the ADAR in local memory. If the Destination Write Enable bit is set, the result buffer is also stored to the memory location pointed to by the DAR.

#### Figure 65. An example of Zero Result Buffer Check with P+Q RAID-6





## 7.5.7 Memory Block Fill Operation

The AA can be used to write a constant value to a memory block in the 80333 local memory. As with XOR operations, descriptors are used to specify the memory blocks to which the AA writes the data contained in the Data / Source Address Register1. All memory block fill operations are controlled by chain descriptors located in the Intel XScale<sup>®</sup> core local memory. Figure 66 illustrates a Block Fill Operation to an arbitrary destination address.







# 7.6 **Programming Model State Diagram**

The AA programming model diagram is shown in Figure 67. Error condition states are not shown.

#### Figure 67. Application Accelerator Programming Model State Diagram





# 7.7 Application Accelerator Priority

The internal bus arbitration logic determines which internal bus master has access to the 80333 internal bus. The Application Accelerator has an independent Bus Request/Grant signal pair to the internal bus arbitration logic.

In addition, the internal bus arbitration unit has a Multi-Transaction timer that affects the throughput of the AA. The default value for MTT2 of 152 clocks was chosen to ensure that once an internal bus agent (in this case the AA) is granted the internal bus that it is guaranteed an opportunity to burst data into DDR SDRAM memory. However, when the bus is busy the AA loses grant before the burst is completed. This means that the AA is able to complete only one burst for each arbitration cycle.

Alternatively, the user may wish to increase the value of MTT2 to guarantee that two or more bursts are able to complete within an arbitration cycle.

For example, assuming 1 Kbyte bursts and a 64-bit memory subsystem, an MTT2 setting of 192 clocks would be sufficient to support two 1 Kbyte bursts for the AA in a single arbitration cycle.

*Warning:* Increasing the MTT2 value may also increase the latency to peripheral memory mapped registers or PCI addresses for the Intel XScale<sup>®</sup> core on the average. Before changing the MTT2 value, it's imperative that the overall impact to the performance of the application is considered.

# 7.8 Packing and Unpacking

int<sub>el</sub>.

The Application Accelerator contains a hardware data packing and unpacking unit to support data transfers between unaligned source and destination addresses. Source and destination addresses can either be unaligned or aligned on natural boundaries. The packing unit optimizes data transfers to and from 32 and 64-bit memory. It reformats data words for the correct bus data width. When the read data needs to be packed or unpacked, the data is held internally and does not need to be re-read.

Aligned data transfers fall on natural boundaries. For example; DWORDs are aligned on 8-byte boundaries and words are aligned on 4-byte boundaries. Data transfers take place in two instances:

- The source and destination addresses are both aligned.
- All or some source addresses are unaligned and the destination address is aligned or unaligned.

## 7.8.1 64-bit Unaligned Data Transfers

Figure 68 illustrates a data transfer between unaligned 64-bit, source and destination addresses.

#### Figure 68. Optimization of an Unaligned Data Transfer





## 7.9 Programming the Application Accelerator

The operations for Application Accelerator software falls into the following categories:

- AA initialization
- Suspend AA
- Appending Descriptors
- Resume AA Operation

An example for each category is shown in the following sections as pseudo code flow.

The AA control register provides independent control each time the AA is configured. This provides the greatest flexibility to the applications programmer.

The most efficient method for operating the AA is to use the *Chain Resume* capability described in Section 7.3.4, "Descriptor Chaining". To use of the *Chain Resume* capability for appending descriptors to chains in normal operation, an initial AA descriptor must be executed. This initialization step is described in Section 7.9.1, "Application Accelerator Initialization". The example AA operations provided later in this section use the *Chain Resume* capability as follows:

- Store Descriptor in Local Memory
- Append Descriptor to Chain
- Resume AA Operation



## 7.9.1 Application Accelerator Initialization

The AA is designed to have independent control of the interrupts, enables, and control. The initialization consists of virtually no overhead as shown in Figure 69.

#### Figure 69. Pseudo Code: Application Accelerator Initialization

ACR = 0x0000 0000 ; Disable the application accelerator Call setup accelerator

The following example illustrates how AA initialization S/W prepares the AA for descriptor *Chain Resume* operation. Initializing the AA for chaining requires an initial descriptor be created and executed. This descriptor is then the start of the chain, and future descriptors are appended to this descriptor to create the chain. This descriptor is a NULL descriptor, requiring no source or destination data buffers be allocated. To start an operation, software simply sets the AA Enable bit in the "Accelerator Control Register - ACR" as shown in Figure 70.

#### Figure 70. Pseudo Code: Application Accelerator Chain Resume Initialization

```
; Set up descriptor in Intel XScale core local memory at address d
d.nda = 0
                  /* No chaining */
d.D/SAR1 = 0x0000 0000/* Source address of Data Block 1
                                                            */
                                                          */
d.SAR2 = 0x0000 0000/* Source address of Data Block 2
d.SAR3 = 0x0000 0000/* Source address of Data Block 3
                                                          */
d.SAR4 = 0x0000 0000/* Source address of Data Block 4
                                                          */
d.DAR = 0x0000 0000/* Destination address of XOR-ed data */
                  /* Byte Count of zero */
d.ABCR = 0x0
d.ADCR = 0x000 0000/* Null Descriptor, No Interrupt*/
; Start operation
ANDAR = &d ; Setup descriptor address
ACR = 0x0000 0001 ; Set AA Enable bit
```

## 7.9.2 Suspending and Resuming the Application Accelerator

The Application Accelerator unit provides the ability to suspend the current state without losing status information. The AA resumes without requiring application software to restore the previous configuration. The example shown in Figure 71 describes pseudo-code for suspending the ongoing operation and then restarting.

#### Figure 71. Pseudo Code: Suspend Application Accelerator

```
;Suspend Application Accelerator
ACR = 0x0000 0000 ; Suspend ongoing AA transfer
;Restart Application Accelerator
ACR = 0x0000 0001 ; Restart AA operation
```

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Application Accelerator Unit



## 7.9.3 Appending Descriptor for XOR Operations

The example shown in Figure 72 describes the pseudo code for initiating an XOR operation with the AA. The examples illustrates appending the XOR operation to an existing chain, and taking advantage of the Chain Resume capability as described in Section 7.9.2, "Suspending and Resuming the Application Accelerator".

```
Figure 72. Pseudo Code: XOR Transfer Operation
```

```
; Set up descriptor in Intel XScale® core local memory at address d
d.nda = 0 /* No chaining */
d.D/SAR1 = 0xA000 0400/* Source address of Data Block 1
                                                          */
d.SAR2 = 0xA000 0800/* Source address of Data Block 2
                                                        */
d.SAR3 = 0xA000 0C00/* Source address of Data Block 3
                                                        */
d.SAR4 = 0xA000 1000/* Source address of Data Block 4
                                                        */
d.DAR = 0xB000 0100/* Destination address of XOR-ed data */
d.ABCR = 1024
                /* Byte Count of 1024 */
d.ADCR = 0x8000 049F/* Direct fill data from Block 1
                                                       */
                 /* XOR with data from Block 2, Block 3 and
                    Block 4
                                                       */
                  /* Store the result and interrupt processor */
; Append descriptor to end of last chain at address c
c.nda = d
; Resume AA operation
ACR = 0x00000003 ; Set AA Enable and Resume bits
```



## 7.9.4 Appending Descriptor for Dual-XOR Operations

The example shown in Figure 73 describes the pseudo code for initiating a Dual-XOR operation with the AA. The examples illustrates appending the Dual-XOR operation to an existing chain, and taking advantage of the Chain Resume capability as described Section 7.9.2, "Suspending and Resuming the Application Accelerator".

```
Figure 73. Pseudo Code: Dual-XOR Transfer Operation
```

```
; Set up descriptor in Intel XScale core local memory at address d
                  /* No chaining */
d_n da = 0
d.D/SAR1 = 0xA000 0400/* Source address of Data Block 1
                                                             */
d.SAR2 = 0xA000 0800/* Source address of Data Block 2
                                                           */
d.SAR3 = 0xA000 0C00/* Source address of Horizontal Data Block */
d.SAR4 = 0xA000 1000/* Source address of Diagonal Data Block */
d.DAR H = 0xB000 0100/* Destination address of Horizontal XOR-ed data */
d_{ABCR} = 1024
                  /* Byte Count of 1024 */
d.ADCR = 0x8800 049F/* Dual-XOR Operation */
                  /* Required: Direct fill from Block 1 */
                   /* XOR enabled for Blocks 2, 3 and 4 */
                  /* Store the results */
                  /* Optional: interrupt processor */
d.DAR D = 0xB000 4100/* Destination address of Diagonal XOR-ed data */
; Append descriptor to end of last chain at address c
c.nda = d
; Resume AA operation
ACR = 0 \times 0 0 0 0 0 0 3
                 ; Set AA Enable and Resume bits
```

## 7.9.5 Appending Descriptor for Memory Block Fill Operations

The example shown in Figure 74 describes the pseudo code for initiating a Memory Block Fill operation with the AA.

### Figure 74. Pseudo Code: Memory Block Fill Operation

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Application Accelerator Unit



## 7.9.6 Appending Descriptor for Zero Result Buffer Check

The example shown in Figure 75 describes the pseudo code for initiating an XOR operation with the AA.

#### Figure 75. Pseudo Code: Zero Result Buffer Check Operation<sup>a</sup>

```
; Set up descriptor in Intel XScale® core local memory at address d
            /* No chaining */
d.nda = 0
d.D/SAR1 = 0xA000 0400/* Source address of Data Block 1
                                                          */
d.SAR2 = 0xA000 0800/* Source address of Data Block 2
                                                         */
d.SAR3 = 0xA000 0C00/* Source address of Data Block 3
                                                         */
d.SAR4 = 0xA000 1000/* Source address of Data Block 4
                                                         */
d.ABCR = 1024 /* Byte Count of 1024 */
d.ADCR = 0x4000 049F/* Direct fill data from Block 1
                                                        */
                  /* XOR with data from Block 2, Block 3 and
                     Block 4
                                                       */
                  /* Check Result, Write Status (ADCR) and interrupt processor */
; Append descriptor to end of last chain at address c
c.nda = d
; Resume AA operation
ACR = 0x00000003 ; Set AA Enable and Resume bits
```

a. Notice that ADCR.dwe is cleared and that the DAR is not programmed. The reason is that for Zero Result Buffer Check operations, there is no need to write out a destination parity stripe.

## 7.10 Interrupts

The Application Accelerator can generate an interrupt to the Intel XScale<sup>®</sup> core. The *Interrupt Enable* bit in the Accelerator Descriptor Control Register (ADCR.ie) determines whether the AA generates an interrupt upon successful, error-free completion. Error conditions described in Section 7.11 also generate an interrupt. The AA has one interrupt output connected to the PCI and Peripheral Interrupt Controller.

Once the AA is enabled, the AA loads the chain descriptor fields into the respective registers. A special case exists when data write enable is clear, then an interrupt is generated (when enabled) after the descriptor is fetched and processed as defined by the block control fields in the ADCR. Table 248 summarizes the status flags and conditions when interrupts are generated in the Accelerator Status Register (ASR).

	Accel	erator S (ASR)	Interrupt Generated?			
Interrupt Condition	Active	End of Transfer	End of Chain	IB Master Abort	Interrupt Enabled	Interrupt Disabled
(Data Write Enable == 0    byte count == 0) && (ANDAR != NULL    Resume == 1) (End of Transfer)	1	1	0	0	Y	Ν
(Data Write Enable == 0    byte count == 0) && ANDAR == NULL && Resume == 0 (End of Chain)	0	0	1	0	Y	Ν
IB Master Abort	0	0	0	1	Y	Y
IB Target Abort	0	0	0	0	Ν	Ν

#### Table 248. AA Interrupts

*Note:* End-of-Transfer and End-of-Chain flags is set only when Interrupt Enable is set. When Interrupt Enable is clear, then the above flags are always set to 0. End-of-Transfer Interrupt and End of Chain Interrupt can only be reported in the ASR when the descriptor fetch and processing completed without any reportable errors. However, multiple error conditions may occur and be reported together. Also, because the AA does not stop after reporting the End-of-Transfer interrupt, an IB master-abort error may occur before the End-of-Transfer interrupt is serviced and cleared.



# 7.11 Error Conditions

Master Aborts that occur during a transfer are recorded by the Application Accelerator.

When an error occurs, the actions taken are detailed below:

- The AA ceases the ongoing transfer for the current chain descriptor and clear the *Application Accelerator Active* flag in the ASR.
- The AA does not read any new chain descriptors.
- The AA sets the error flag in the Accelerator Status Register. For example; when an IB master-abort occurred during a transfer, the channel sets bit 5 in the ASR.
- The AA signals an interrupt to the Intel XScale<sup>®</sup> core.
- The Application Accelerator does not restart the transfer after an error condition. It is the responsibility of the application software to reconfigure the AA to complete any remaining transfers.
- *Note:* Target-aborts during AAU reads result from multi-bit ECC errors that are recorded by the MCU. For correct operation of the AAU, user software has to disable the AAU before clearing the error condition. Furthermore, the AAU needs to be re-enabled by writing a 1 to the AA Enable bit before initiating a new operation.



# 7.12 Power-up/Default Status

Upon power-up, an external hardware reset, the Application Accelerator Registers are initialized to their default values.

# 7.13 Register Definitions

The Application Accelerator Unit contains forty two memory-mapped registers for controlling its operation. There is read/write access only to the Accelerator Control Register, Accelerator Status Register, the Accelerator Next Descriptor Address Register, and the three Extended Descriptor Control Registers. All other registers are read-only and are loaded with new values from the chain descriptor whenever the AA reads a chain descriptor from memory.

#### Table 249. Application Accelerator Unit Registers

Section, Register Name - Acronym (page)
Section 7.13.1, "Accelerator Control Register - ACR" on page 466
Section 7.13.2, "Accelerator Status Register - ASR" on page 467
Section 7.13.3, "Accelerator Descriptor Address Register - ADAR" on page 468
Section 7.13.4, "Accelerator Next Descriptor Address Register - ANDAR" on page 469
Section 7.13.5, "Data / Source Address Register1 - D/SAR1/PQSAR1" on page 470
Section 7.13.6, "Source Address Registers 232 - SAR232" on page 471
Section 7.13.7, "P+Q RAID-6 Source Address Registers 216 - PQSAR216" on page 473
Section 7.13.8, "P+Q RAID-6 Galois Field Multiplier Registers 15 - GFMR15" on page 474
Section 7.13.9, "Destination Address Register - DAR" on page 476
Section 7.13.10, "Accelerator Byte Count Register - ABCR" on page 477
Section 7.13.11, "Accelerator Descriptor Control Register - ADCR" on page 478
Section 7.13.12, "Extended Descriptor Control Register 0 - EDCR0" on page 482
Section 7.13.13, "Extended Descriptor Control Register 1 - EDCR1" on page 484
Section 7.13.14, "Extended Descriptor Control Register 2 - EDCR2" on page 486



## 7.13.1 Accelerator Control Register - ACR

The Accelerator Control Register (ACR) specifies parameters that dictate the overall operating environment. The ACR should be initialized prior to all other AA registers following a system reset. Table 250 shows the register format. This register can be read or written while the AA is active.



#### Table 250. Accelerator Control Register - ACR



## 7.13.2 Accelerator Status Register - ASR

The Accelerator Status Register (ASR) contains status flags that indicate status. This register is typically read by software to examine the source of an interrupt. See Section 7.11 for a description of the error conditions that are reported in the ASR. See Section 7.10 for a description of interrupts caused by the Application Accelerator.

When an AA error occurs, application software should check the status of Accelerator Active flag before processing the interrupt.



#### Table 251. Accelerator Status Register - ASR

## 7.13.3 Accelerator Descriptor Address Register - ADAR

The Accelerator Descriptor Address Register (ADAR) contains the address of the current chain descriptor in local memory. This read-only register is loaded when a new chain descriptor is read. Table 252 depicts the ADAR. Depending on the number of sources, the chain descriptors are required to be aligned on different address boundaries. These include four sources on an eight word address boundary, eight sources on a 16 word address boundary, 16 sources on a 32 word address boundary, and 32 sources on a 64 word address boundary.

*Note:* In the above paragraph, the term "word" refers to a DWORD.



#### Table 252. Accelerator Descriptor Address Register - ADAR
## 7.13.4 Accelerator Next Descriptor Address Register - ANDAR

The Accelerator Next Descriptor Address Register (ANDAR) contains the address of the next chain descriptor in local memory. When starting a transfer, this register contains the address of the first chain descriptor. Table 253 depicts the Accelerator Next Descriptor Address Register.

All chain descriptors are aligned on an eight DWORD boundary. The AA may set bits 04:00 to zero when loading this register.

*Note:* The *Accelerator Enable* bit in the ACR and the *Accelerator Active* bit in the ASR must both be clear prior to writing the ANDAR. Writing a value to this register while the AA is active may result in undefined behavior.



#### Table 253. Accelerator Next Descriptor Address Register - ANDAR



## 7.13.5 Data / Source Address Register1 - D/SAR1/PQSAR1

The Data / Source Address Register (D/SAR1/PQSAR1) contains a 32-bit, local memory address or immediate data to be written in case of Memory Block Fill operations. The ADCR register (Table 260) controls the operation performed on data block referenced by this register. The local memory address space is a 32-bit, byte addressable address space.

Reading the D/SAR1/PQSAR1 register once the AA has started a chain descriptor returns the current source address or immediate data to be written in case of Memory Block Fill operations.

Once an operation is initiated, these registers contain the current source addresses. For example; when the Byte Count is initially 4096 bytes and the AA has completed the operation on the first three 1K-byte data blocks, the value in register SAR1/PQSAR1 is the equal to the programmed descriptor value + 3072 (SAR1 + 3072).

During Memory Block Fills the register always contains the data to be written and does not change.

Table 254 shows the Data / Source Address Register1/P+Q RAID-6 SAR1. This read-only register is loaded when a chain descriptor is read from memory.



Table 254. Data / Source Address Register - SAR1/PQSAR1

## 7.13.6 Source Address Registers 2..32 - SAR2..32

int

The Source Address Registers 2..32 (SAR2..32) contain 32-bit, local memory addresses. There are 31 Source Address Registers (SAR2 - SAR32). Each of these registers is loaded with address of blocks of data to be operated upon by the AA. The ADCR, EDCR0, EDCR1, and EDCR2 registers control the operation performed on each data block referenced by the registers (SAR2 - SAR32). The local memory address space is a 32-bit, byte addressable address space.

Reading SARx registers once AA has started a chain descriptor returns the current source addresses. Once an operation is initiated, these registers contain current source addresses. For example; when Byte Count is initially 4096 bytes and AA has completed operation on the first three 1K-byte data blocks, the value in register SARx is the equal to the programmed descriptor value + 3072 (SARx + 3072).

For *Dual-XOR-transfers* SAR3 is the Horizontal Source Address for the XOR result for Horizontal XOR result, and SAR4 is the Diagonal Source Address for the Diagonal XOR result. Also, SAR5 is the Diagonal Destination Address for the Diagonal XOR result.

Note: For P+Q RAID-6 Mode, refer to section Section 7.13.7, "P+Q RAID-6 Source Address Registers 2..16 - PQSAR2..16" on page 473 for source addresses 2 through 16 and Section 7.13.8, "P+Q RAID-6 Galois Field Multiplier Registers 1..5 - GFMR1..5" on page 474 for P+Q RAID-6 Multiplier Word definitions.



Table 254 shows the Source Address Register2..32. These read-only registers are loaded when a chain descriptor is read from memory.



## 7.13.7 P+Q RAID-6 Source Address Registers 2..16 -PQSAR2..16

*Note:* The following definition applies only when P+Q RAID-6 mode is enabled. When P+Q RAID-6 is NOT enabled, refer to Section 7.13.6, "Source Address Registers 2..32 - SAR2..32" on page 471 for definition and internal bus addresses of Source Address Registers

The P+Q RAID-6 Source Address Register2..16 (PQSAR2..16) contain 32-bit, local memory addresses. There are 16 P+Q RAID-6 Source Address Registers (PQSAR1..PQSAR16). Each of these registers is loaded with address of blocks of data to be operated upon by the AA when P+Q RAID-6 Mode is enabled. The ADCR, EDCR0, and EDCR1 registers control the operation performed on each data block referenced by the registers (PQSAR1..PQSAR16). The local memory address space is a 32-bit, byte addressable address space.

Reading PQSARx registers once AA has started a chain descriptor returns the current source addresses. Once an operation is initiated, these registers contain current source addresses. For example; when Byte Count is initially 4096 bytes and AA has completed operation on the first three 1K-byte data blocks, the value in register PQSARx is the equal to the programmed descriptor value + 3072 (PQSARx + 3072).

Table 256 shows the P+Q RAID-6 Source Address Registers 2..16. These read-only registers are loaded when a chain descriptor is read from memory.

*Note:* See Section 7.13.8 for definition of the Data Multipliers in P+Q RAID-6 mode.



#### Table 256. P+Q RAID-6 Source Address Registers 2..16 - PQSAR2..16



## 7.13.8 P+Q RAID-6 Galois Field Multiplier Registers 1..5 - GFMR1..5

*Note:* The following definition applies only when P+Q RAID-6 mode is enabled. When P+Q RAID-6 is NOT enabled, refer to Section 7.13.6, "Source Address Registers 2..32 - SAR2..32" on page 471 for definition and internal bus addresses of Source Address Registers

The P+Q RAID-6 Galois Field Multiplier Registers 1..5 (GFMR1..5) contain the 8-bit multiplier values. There are 16 Data Multipliers distributed through the five Data Multiplier Words (GFMR1..GFMR5). Each of these registers is loaded with data multiplier values to be used by the AA GF Multiply Function when P+Q RAID-6 Mode is enabled. The ADCR, EDCR0, and EDCR1 registers control the operation performed on each source data block.

Table 257 shows the Galois Field Multiplier Registers GFMR[1:5]. These read-only registers are loaded when a chain descriptor is read from memory.

*Note:* See Section 7.13.7 for definition of the Data Integrity Source Addresses in P+Q RAID-6 Source AddressS.



 Table 257.
 Galois Field Multiplier Registers 1..5- GFMR1..5 (Sheet 1 of 2)



Table 2	Galois Field Multiplier Registers 15- GFMR15 (Sheet 2 of 2)					
Att GFMR1 GFMR2 GFMR3 GFMR4 GFMR5	IOP ributes na na PCI na na Internal FFFF E& FFFF E& FFFF E& FFFF E& FFFF E&	28       24       20       16         1 ro	12       8       4       0         5       10			
Bit	Default		Description			
15:8	00H	<ul> <li>Data Multiplier - Data Multiplier Byte used by t from corresponding PQSARx, when P+Q RAI</li> <li>GFMR1 - Data Multiplier 2 (DMLT2)</li> <li>GFMR2 - Data Multiplier 5 (DMLT5)</li> <li>GFMR3 - Data Multiplier 8 (DMLT8)</li> <li>GFMR4 - Data Multiplier 11 (DMLT11)</li> <li>GFMR5 - Data Multiplier 14 (DMLT14)</li> </ul>	he P+Q RAID-6 function (GF Multiply) with source data D-6 mode is enabled.			
7:0	00H	Data Multiplier - Data Multiplier Byte used by t from corresponding PQSARx, when P+Q RAI • GFMR1 - Data Multiplier 1 (DMLT1) • GFMR2 - Data Multiplier 4 (DMLT4) • GFMR3 - Data Multiplier 7 (DMLT7) • GFMR4 - Data Multiplier 10 (DMLT10) • GFMR5 - Data Multiplier 13 (DMLT13)	he P+Q RAID-6 function (GF Multiply) with source data D-6 mode is enabled.			



## 7.13.9 Destination Address Register - DAR

The Destination Address Register (DAR) contains a 32-bit, local memory address. The DAR may also contain an address targeting the ATU outbound windows for writing the AAU result to the PCI bus. During operations, this address is the destination address in local memory where data will be stored. The 80333 local memory address space is a 32-bit, byte addressable address space. When programming the result to be on the PCI bus, this address is one of the ATU outbound windows, which results in a 32-bit or 64-bit PCI address depending on the window addressed.

During Dual-XOR operations, this address points to the memory block to be written with the Horizontal XOR result.

During Memory Block Fill operations, this address points to the memory block to be written with the constant value contained in the D/SAR1 register.

Reading the DAR once the AA has started a chain descriptor returns the current destination address. For example; during an XOR operation when the Byte Count is initially 4096 bytes and the AA has completed the *XOR-transfer* operation on the first three 1K-byte data blocks, the value in the Destination Address Register (DAR) will be equal to the programmed descriptor value + 3072 (DAR + 3072).

Table 258 shows the Destination Address Register. This read-only register is loaded when a chain descriptor is read from memory



#### Table 258. Destination Address Register - DAR



## 7.13.10 Accelerator Byte Count Register - ABCR

The Accelerator Byte Count Register (ABCR) contains the number of bytes to transfer for an operation. This is a read-only register that is loaded from the Byte Count word in a chain descriptor. It allows for a maximum transfer of 16 Mbytes. A value of zero is a valid byte count and results in no read or write cycles being generated to the Memory Controller Unit. No cycles are generated on the internal bus.



#### Table 259. Accelerator Byte Count Register - ABCR

*Note:* Anytime this register is read, it contains the number of bytes left to transfer on the internal bus. Note that during an operation valid data may be present in the Application Accelerator store queue. This register is decremented by 1 through 8 for every successful transfer from the store queue to the destination location. During *Memory Block Fills* this register is decremented by 1 through 8 for every successful write operation. Table 259 shows the Accelerator Byte Count Register. The byte count value is not required to be aligned to a DWORD boundary (i.e., the byte count value can be a DWORD aligned, short aligned, or byte aligned).



## 7.13.11 Accelerator Descriptor Control Register - ADCR

The Accelerator Descriptor Control Register contains control values for data transfer on a per-chain descriptor basis. This read-only register is loaded when a chain descriptor is read from memory. These values may vary from chain descriptor to chain descriptor. The AAU determines whether a mini-descriptor is appended to the end of the current chain descriptor by examining bits 26:25. Table 260 shows the definition of the Accelerator Descriptor Control Register.







Table 260.         Accelerator Descriptor Control Register - ADCR (Sheet 2 of 4)						
Att	31       28       24       20       16       12       8       4       0         Attributes					
Bit	Default		Description			
26:25	00	Supplemental Block Control Interpreter - Th segments beyond the principle descriptor or 00 Principle Descriptor only - This specifies the principle descriptor to initialize the first ei- or up to 3 sources for P+Q RAID-6. 01 Mini-Descriptor - This specifies that their mini-descriptor to initialize four additional re operation, or up to 6 sources for P+Q RAID- 10 Extended Descriptor 0 - This specifies the therefore reads the mini-descriptor and one registers. Set for up to 16 sources for XOR, 11 Extended Descriptors 1 and 2 - This spe- words. The AA therefore reads the mini-desc total of thirty-nine registers. Set for up to 32 <b>Block 8 Command Control</b> - This bit field spointed at by SAR8 register. 000Null command - This implies that Block at The Application Accelerator does not transfer	s bit field specifies the number of additional descriptor which the operation is executed. that no additional descriptor words exist. The AA only reads ght AA descriptor registers. Set for up to 4 sources for XOR, e are up to 4 additional words. The AA therefore reads the gisters. Set for up to 8 sources for XOR, or Dual-XOR 6. nat there are up to nine additional descriptor words. The AA extended-descriptor to initialize a total of twenty-one or up to 12 sources for P+Q RAID-6 . ecifies that there are up to eighteen additional descriptor riptor and three extended-descriptors to initialize registers a sources for XOR, or up to 16 sources for P+Q RAID-6 . epecifies the type of operation to be carried out on the data 8 Data can be disregarded for the current chain descriptor. er data from this block while processing the current chain			
24.22	0	descriptor. 001XOR command - This implies that Block 8 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved				
21:19 0 Block 7 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR7 register. 000Null command - This implies that Block 7 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 7 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved						
18:16	0	Block 6 Command Control - This bit field s pointed at by SAR6 register. 000Null command - This implies that Block The Application Accelerator does not transfe descriptor. 001XOR command - This implies that Block execute the XOR function. All other values are reserved	pecifies the type of operation to be carried out on the data 5 Data can be disregarded for the current chain descriptor. For data from this block while processing the current chain 6 Data is transferred to the Application Accelerator to			

Att	31       28       b8cc       b7cc       b6cc       b5cc       b4cc       b3cc       b2cc       b1cc       0         Attributes					
	[	1	RS = Read/Set NA = Not Access	sible		
Bit	Default	C	escription			
15:13	0         Block 5 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR5 register.           0         000Null command - This implies that Block 5 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor.           001XOR command - This implies that Block 5 Data is transferred to the Application Accelerator to execute the XOR function.           All other values are reserved			ne data criptor. chain to		
12:10	12:10       0       Block 4 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR4 register.         12:10       000Null command - This implies that Block 4 Data can be disregarded for the current chain descriptor.         000Null command - This implies that Block 4 Data can be disregarded for the current chain descriptor.         001XOR command - This implies that Block 4 Data is transferred to the Application Accelerator to execute the XOR function. (required for Dual-XOR-transfers)         All other values are reserved			ne data criptor. chain to		
09:07	09:07 09:07 09:07 0 09:07 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0			ne data criptor. chain to		
06:04	0	Block 2 Command Control - This bit field spe pointed at by SAR2 register. 000Null command - This implies that Block 2 D The Application Accelerator does not transfer of descriptor. 001XOR command - This implies that Block 2 execute the XOR function. (required for <i>Dual-X</i> All other values are reserved	cifies the type of operation to be carried out on the tata can be disregarded for the current chain des ata from this block while processing the current of Data is transferred to the Application Accelerator <i>OR-transfers</i> )	ne data criptor. chain to		

### Table 260. Accelerator Descriptor Control Register - ADCR (Sheet 3 of 4)







## 7.13.12 Extended Descriptor Control Register 0 - EDCR0

The Extended Descriptor Control Register 0 contains control values for data transfer on a per-chain descriptor basis. This read-only register is loaded when a chain descriptor that requires a minimum of 16 Source Addresses is read from memory. The values in EDCR0 define the command/control value for SAR16 - SAR9. The AAU determines whether an extended descriptor requiring the use of EDCR0 is appended to the end of the current chain descriptor by examining bits 26:25 of the Accelerator Descriptor Control Register. Table 261 shows the definition of the Extended Descriptor Control Register 0.



#### Table 261. Extended Descriptor Control Register 0 - EDCR0 (Sheet 1 of 2)



Table 20	Table 261.Extended Descriptor Control Register 0 - EDCR0 (Sheet 2 of 2)					
Att	Bifec       bifec <th< th=""></th<>					
Bit	Default	Description				
15:13	0	Block 13 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR13 register.         000Null command - This implies that Block 13 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor.         001XOR command - This implies that Block 13 Data is transferred to the Application Accelerator to execute the XOR function.				
12:10	0	<ul> <li>Block 12 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR12 register.</li> <li>000Null command - This implies that Block 12 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor.</li> <li>001XOR command - This implies that Block 12 Data is transferred to the Application Accelerator to execute the XOR function.</li> </ul>				
09:07	9:07 0 Block 11 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR11 register. 000Null command - This implies that Block 11 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 11 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved					
06:04	06:04 0 Block 10 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR10 register. 000Null command - This implies that Block 10 Data can be disregarded for the current chain escriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 10 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved					
03:01	03:01       0       Block 9 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR9 register.         03:01       0       000Null command - This implies that Block 9 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 9 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved         00       0       Reserved					



## 7.13.13 Extended Descriptor Control Register 1 - EDCR1

The Extended Descriptor Control Register 1 contains control values for data transfer on a per-chain descriptor basis. This read-only register is loaded when a chain descriptor that requires 32 Source Addresses is read from memory. The values in EDCR1 define the command/control value for SAR24 - SAR17. The AAU determines whether an extended descriptor requiring the use of EDCR1 is appended to the end of the current chain descriptor by examining bits 26:25 of the Accelerator Descriptor Control Register. Table 262 shows the definition of the Extended Descriptor Control Register 1.



#### Table 262. Extended Descriptor Control Register 1 - EDCR1 (Sheet 1 of 2)



Table 2	62. Ext	Extended Descriptor Control Register 1 - EDCR1 (Sheet 2 of 2)			
At	IOP tributes	b24cc       b23cc       b21cc       b20cc       b19cc       b18cc       b17cc         c       24       20       16       12       8       4       0         rv       rv       rv       rv       rv       rv       ro       ro <t< th=""></t<>			
		RS = Read/Set NA = Not Accessible			
Bit	Default	Description			
15:13	0	<ul> <li>Block 21 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR21 register.</li> <li>000Null command - This implies that Block 21 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor.</li> <li>001XOR command - This implies that Block 21 Data is transferred to the Application Accelerator to execute the XOR function.</li> </ul>			
Block 20 Command Control - This bit field specifies the type of operation to be carried out on pointed at by SAR20 register.           000Null command - This implies that Block 20 Data can be disregarded for the current chain d. The Application Accelerator does not transfer data from this block while processing the current					
		descriptor. 001XOR command - This implies that Block 20 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved			
		Block 19 Command Control - This bit field specifies the type of operation to be carried out on the data			
09:07	09:07 0 pointed at by SAR19 register. 000Null command - This implies that Block 19 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 19 Data is transferred to the Application Accelerator to execute the XOR function.				
		Block 18 Command Control - This bit field specifies the type of operation to be carried out on the data			
06:04	06:04 0 pointed at by SAR18 register. 000Null command - This implies that Block 18 Data can be disregarded for the current chain descripto The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 18 Data is transferred to the Application Accelerator to execute the XOR function.				
	All other values are reserved				
03:01	0	<ul> <li>Block 1 / Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR17 register.</li> <li>000Null command - This implies that Block 17 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor.</li> <li>001XOR command - This implies that Block 17 Data is transferred to the Application Accelerator to the the the Application Accelerator to the the the the the the the the the the</li></ul>			
		execute the XOR function.			
00	0	All other values are reserved			
00	0	Neocived.			



## 7.13.14 Extended Descriptor Control Register 2 - EDCR2

The Extended Descriptor Control Register 2 contains control values for data transfer on a per-chain descriptor basis. This read-only register is loaded when a chain descriptor that requires 32 Source Addresses is read from memory. Values in EDCR2 define the command/control value for SAR32 - SAR25. The AAU determines whether an extended descriptor requiring the use of EDCR2 is appended to the end of the current chain descriptor by examining bits 26:25 of the Accelerator Descriptor Control Register. Table 263 shows the definition of the Extended Descriptor Control Register 2.







Table 263.         Extended Descriptor Control Register 2 - EDCR2 (Sheet 2 of 2)						
Att	31       28       24       20       16       12       8       4       0         Attributes					
		1	RS = Read/Set	NA = Not Accessible		
Bit	Default	De	escription			
12:10	0	Block 28 Command Control - This bit field spec pointed at by SAR28 register. 000Null command - This implies that Block 28 D The Application Accelerator does not transfer da descriptor. 001XOR command - This implies that Block 28 D execute the XOR function. All other values are reserved Block 27 Command Control - This bit field spec	cifies the type of operation to ata can be disregarded for th ta from this block while proce Data is transferred to the App cifies the type of operation to	be carried out on the data ne current chain descriptor. essing the current chain plication Accelerator to		
09:07	09:07 0 Block 27 Command Control - This bit field specifies the type of operation to be carried out on the dat pointed at by SAR27 register. 000Null command - This implies that Block 27 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 27 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved			ne current chain descriptor. essing the current chain		
06:04	06:04 0 Block 26 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR26 register. 000Null command - This implies that Block 26 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 26 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved					
03:01	03:01 0 Block 25 Command Control - This bit field specifies the type of operation to be carried out on the data pointed at by SAR25 register. 000Null command - This implies that Block 25 Data can be disregarded for the current chain descriptor. The Application Accelerator does not transfer data from this block while processing the current chain descriptor. 001XOR command - This implies that Block 25 Data is transferred to the Application Accelerator to execute the XOR function. All other values are reserved					
00	0	Reserved.				



**This Page Left Intentionally Blank** 



# Memory Controller

This chapter describes the integrated Memory Controller Unit (MCU). The operating modes, initialization, external interfaces, and implementation are detailed in this chapter.

## 8.1 Overview

The Intel<sup>®</sup> 80333 I/O processor (80333) integrates a high performance, multi-ported Memory Controller to provide a direct interface between the 80333 and its local memory subsystem. The Memory Controller supports:

- DDR 333 MHz SDRAM with 128/256/512 Mbit and 1Gbit density supported.
- DDR-II 400 MHz SDRAM with 256/512 Mbit density supported.
- Between 64 MBytes and 2 GByte of 64-bit DDR SDRAM (1 GByte for DDR-II SDRAM).
- Between 32 MBytes and 1 GBytes of 32-bit DDR SDRAM for low cost solutions (512 MByte for DDR-II SDRAM).
- Optimized core processor data processing 32-bit region.
- Single-bit error correction, multi-bit detection support (ECC).
- 32-, 40- and 64-, 72-bit wide Memory Interfaces (non-ECC and ECC support).

The DDR SDRAM (DDR and DDR-II SDRAM) interface provides a direct connection to a high bandwidth and reliable memory subsystem. The DDR SDRAM interface consists of a 64-bit wide data path to support up to 3.2 GBytes/sec throughput. An 8-bit Error Correction Code (ECC) across each 64-bit word improves system reliability. The ECC is stored into the DDR SDRAM array along with the data and is checked when the data is read. When the code is incorrect, the MCU corrects the data (when possible) before reaching the initiator of the read. User-defined fault correction software is responsible for scrubbing the memory array.

The MCU supports two banks of DDR SDRAM in the form of one two-bank dual inline memory module (DIMM).

- The MCU has support for Registered and Unbuffered DDR333 and Registered DDR-II 400 DIMMs.
- The MCU supports a 32-bit SDRAM data interface. This mode enables lower-cost solutions at the cost of system performance.
- The MCU responds to internal bus and core processor memory accesses within its programmed address range and issues the memory request to the DDR SDRAM interface.
- The MCU contains transaction queues for each port enabling pipelining of transactions to the DDR SDRAM for maximum performance.
- For 64-bit ECC memory, a 32-bit memory region can be programmed to operate as 32-bit ECC memory for higher performance core write performance by avoiding Read-Modify-Write (RMW) operation of DDR SDRAM.
- The MCU provides two chip enables to the memory subsystem. These two chip enables service the DDR SDRAM subsystem (one per bank).

## 8.2 Glossary

This section lists commonly used terms throughout this chapter:

#### Table 264.Commonly Used Terms

Term	Definition
Bank	A bank is defined as a memory region defined with a base register and a bank size register. Physically, a bank of memory is controlled by a single chip select. A DIMM could be comprised of a single or dual banks.
Column	A column refers to a portion of memory within an DDR SDRAM device. An DDR SDRAM device can be thought of as a grid with rows and columns. Once a row is activated, any column within that row can be accessed multiple times without reactivating the row. Columns are activated with CAS#.
DIMM	A DIMM is an acronym for Dual Inline Memory Module. A DIMM is a physical card comprising multiple DDR SDRAM devices. The card could be populated on one or both sides. A DIMM can be single or dual-bank.
Leaf	DDR SDRAM devices use multiple banks within the device operating in an interleaved mode. The MCU supports 128/256/512 Mbit, 1 Gbit DDR SDRAM and 256/512 Mbit DDR-II SDRAM devices containing four internal banks. An internal bank is defined as a leaf (to avoid confusion with a memory bank).
Page	A page is a row of memory. Once a row is activated, any column within that row can be accessed multiple times without reactivating the row. This is referred to as "keeping the page open." Page size depends on the DDR SDRAM device configuration, and the MCU supports the maximum possible page size (16 Kbytes for 64-bit wide memory and 8 Kbytes for 32-bit wide memory). The MCU breaks up pages across physical memory due to address mapping.
Row	A row refers to a portion of memory within an DDR SDRAM device. An DDR SDRAM device can be thought of as a grid with rows and columns. Once a row is activated, any column within that row can be accessed multiple times without reactivating the row. Rows are activated with RAS#.
Scrubbing	Once an error is detected within the memory array, the MCU must correct the error (when possible) while delivering the data to the initiator. Correcting the memory location is referred to as "scrubbing the array." The MCU relies on software to scrub any errors.
Syndrome	A syndrome is a value which indicates an error in the data read from the memory array. The MCU computes the syndrome with every memory read. Decoding the syndrome indicates: the bit in error for a single-bit error, or a multi-bit error. Table 280 defines the syndrome decoding.



## 8.3 Theory of Operation

The 80333 memory controller translates the internal bus and core processor transactions into the protocol supported by the DDR SDRAM memory subsystem.

## 8.3.1 Functional Blocks

*Note:* The memory controller and Intel XScale<sup>®</sup> core processor bus interface unit (BIU) logically comprises the blocks illustrated in Figure 76. The memory controller is a multiported unit, supporting an inbound path for both the BIU and Internal Bus (IB) to the DDR SDRAM.

Figure 76. Memory Controller Block Diagram





#### 8.3.1.1 Transaction Ports

Note: The MCU provides two transaction ports for DDR SDRAM access.

#### 8.3.1.1.1 Core Processor Port

The Core Processor Port provides a direct connection between the 80333 core processor bus interface and the Memory Controller. This Core Processor Port allows core transactions targeting the DDR SDRAM to pass directly to the DDR SDRAM.

#### 8.3.1.1.2 Internal Bus Port

The Internal Bus Port provides the connection to the DDR SDRAM from the Internal Bus. All peripheral unit transactions targeting the DDR SDRAM are claimed by this port.



#### 8.3.1.2 Address Decode Blocks

Address Decode is performed for transactions from input ports to determine when the MCU should claim the transaction. There are two address ranges the MCU ports can claim transactions: DDR SDRAM memory space and Peripheral Memory-Mapped Register (PMMR) space.

#### 8.3.1.2.1 DDR SDRAM Memory Space

The DDR SDRAM memory space is defined with the DDR SDRAM Base Address Register (SDBR) and the DDR SDRAM Boundary Registers (SBR0, SBR1, S32SR). The transaction is intended for a DDR SDRAM bank when the address is between the base register (SDBR) and between the boundaries programmed with SBR0, SBR1 and S32SR as defined in Section 8.3.3.2, "DDR SDRAM Addressing" on page 504.

*Note:* For DDR SDRAM Bank 0 or Bank 1 overlapping PMMR space (FFFF E000H to FFFF FFFFH), DDR SDRAM is not accessible, and PMMRs are addressed. The read, write, and reserved characteristics of PMMR space is applicable as defined in Chapter 19, "Peripheral Registers".

#### 8.3.1.2.2 Memory-Mapped Register Space

The MCU PMMR memory space is FFFF E500H to FFFF E5FFH and FFFF F500H to FFFF F5FFH. The registers are detailed in Section 8.7, "Register Definitions" on page 545.Each port decodes inbound transactions for these address ranges. The Address Decode blocks determine how the Memory Controller responds to inbound transactions. The details of the address decode for each port is described below.

#### 8.3.1.2.3 Core Processor Port Address Decode

*Note:* The address decode block for the Core Processor transactions resides in the BIU. The Core MCU port therefore does not require any decode, and will process any transaction received from the BIU via the Core MCU Port.

#### 8.3.1.2.4 Internal Bus Port Address Decode

Internal Bus transactions are decoded to determine when they address the DDR SDRAM Memory Space or MCU MMR Space. When the transaction addresses either of these two spaces, the transaction is claimed by the Internal Bus Port. When the transaction addresses the DDR SDRAM Memory Space, the transaction is queued in the Internal Bus Port Transaction Queue. When the transaction addresses the MCU MMR Space, the transaction is serviced by accessing the Configuration Register block.



#### 8.3.1.3 Memory Transaction Queues

There are two transaction queues for transactions which address the MCU DDR Memory Space. One transaction queue for Core Processor transactions and one transaction queue for Internal Bus transactions.

#### 8.3.1.3.1 Core Processor Memory Transaction Queue (CMTQ)

The CMTQ stores memory transactions from the core which have not been processed by the Memory Controller. The CMTQ supports 8 Core Processor read transactions up to 32-Bytes each, which equals the maximum number of outstanding transaction the Core Processor Bus Controller can support. The CMTQ also supports 8 Core Processor posted write transactions up to 16-Bytes each.

The number of CMTQ write transactions is selectable between 8 and 2, where the lower number may find benefit in some applications by throttling back on the core processor, and thereby 'encouraging' coalescing in the Core Processor Bus Controller.

#### 8.3.1.3.2 Internal Bus Memory Transaction Queue (IBMTQ)

*Note:* The IBMTQ stores memory transactions from the Internal Bus that address the DDR SDRAM Memory Space. The IBMTQ can hold 8 outstanding Internal Bus read transaction requests. Of the 8 read transactions in the IBMTQ, 4 can be processed at a time, each with up to 1KBytes of read data. The other 4 read transactions are simply stored in the IBMTQ. The IBMTQ also supports 4 posted write transactions up to 1KBytes each.



#### 8.3.1.4 Configuration Registers

*Note:* The Configuration Registers block contains all of the memory-mapped registers listed in Section 8.7, "Register Definitions" on page 545. These registers define the memory subsystem connected to the 80333. The status registers indicate the current MCU status.

#### 8.3.1.5 Refresh Counter

The Refresh Counter block keeps track of when the DDR SDRAM devices need to be refreshed. The refresh interval is programmed in the "Frequency Register - RFR" on page 562. Once the 12-bit refresh counter reaches the programmed interval, the DDR SDRAM state machine issues a refresh command to the DDR SDRAM devices. When a transaction is currently in progress, the DDR SDRAM State Machine waits for the completion of the transaction to issue the refresh cycle. See Section 8.3.3.12, "DDR SDRAM Refresh Cycle" on page 523 for more details.

- *Note:* When the memory controller is completing transactions from the core processor memory transaction queue (CMTQ) when the refresh counter expires, the MARB will complete the CMTQ tenure based on the "MCU Port Transaction Count Register MPTCR" on page 560, before any refresh cycles are issued.
- *Note:* When the memory interface is busy when the refresh counter expires, it is possible for the MCU to generate more than one refresh cycle when the memory interface becomes available.

#### 8.3.1.6 Memory Controller Arbiter (MARB)

The MARB determines what transaction to issue next to the DDR SDRAM Control Unit. The MARB is configurable to adjust the selection of transactions from the MCU ports and the refresh counter. The MARB selects transactions from all memory transaction queues based on a programmable arbitration scheme as described in Section 8.3.2, "MCU Arbitration and Configuration" on page 497.



#### 8.3.1.7 DDR SDRAM Control Block

The DDR SDRAM Control Block contains all functionality to process the DDR SDRAM data accesses per the transactions issued by the MARB. To process a transaction the DDR SDRAM Control Block employs several sub-blocks. The sub-blocks include the Page Control block, DDR SDRAM State Machine and Pipeline Queues, and Error Correction Logic.

#### 8.3.1.7.1 Page Control Block

The Page Control Block records and maintains the open DDR SDRAM pages. The MCU can keep a maximum of eight pages open simultaneously (4 per bank). This block keeps track of open pages and determines when the transactions hit an open page. For more details about the page hit/miss determination, see Section 8.3.3.5, "Page Hit/Miss Determination" on page 510.

#### 8.3.1.7.2 DDR SDRAM State Machine and Pipeline Queues

Since the MCU generates error correction codes based on the data, the MCU is a pipelined architecture. Pipelining also ensures acceptable AC timings to the memory interfaces. The DDR SDRAM state machine pipelines DDR SDRAM memory operations for several clocks.

#### 8.3.1.7.3 Error Correction Logic

The Error Correction Logic generates the ECC code for DDR SDRAM reads and writes. For reads, this logic compares the ECC codes read with the locally generated ECC code. When the codes mismatch then the Error Correction Logic determines the error type. For a single-bit error, this block determines which bit is in error and corrects the error. For a single-bit or multi-bit error, the Error Correction Logic logs the error in ELOG0 and ELOG1. See Section 8.3.4, "Error Correction and Detection" on page 525 for more details.



## 8.3.2 MCU Arbitration and Configuration

The order transactions are processed by the DDR SDRAM Control block is determined by the Memory Controller Arbiter (MARB). The MARB selects the next transaction from the memory transaction queues and refresh control based on the settings programmed in the MARB configuration registers ("MCU Port Transaction Count Register - MPTCR" on page 560, "MCU Preemption Control Register - MPCR" on page 561). These registers enable the MCU to be tuned for optimal performance for the design.

#### 8.3.2.1 MCU Port Priority

*Note:* The Memory controller for the 80333 has only two inbound ports for memory transactions. Given this, there is no priority programming required, and the MARB will effectively toggle tenure between the two ports when both ports' memory transaction queues have pending transactions.



#### 8.3.2.2 MCU Port Transaction Count

The MCU Port Transaction Count Register (MPTCR) defines the number of transactions the MARB will select from a given port during the corresponding port's tenure. The intent is that ports with frequent small transactions (core processor port) be allowed to complete multiple transactions during a tenure, while ports with less frequent large transactions (IB port) be limited to a small number of transactions (possibly just 1).

#### 8.3.2.3 Core Processor Port Preemption

In order to maximize core processor performance, the MARB can be programmed to interrupt active transactions from other memory transaction queues, in order to process transactions from the Core Processor. By preempting an active transaction, the MARB guarantees a maximum number of data cycles getting processed for the current transaction, therefore reducing the possible latency the Core Processor might see for memory accesses. Core processor preemption is enabled through the MARB Preemption Control Register (MPCR).

*Note:* Preemption control applies to all inbound MCU ports based on pending core transactions. Once the core processor transaction queue is 'activated' by the MARB, a standard tenure will take place. At the completion of the Core Processor transaction queue tenure, the MARB resumes the previously active transaction. The previously active transaction may be interrupted again, provided the conditions for preemption occur again prior to the completion of that transaction. No other effect on the MARB processing will occur.



#### 8.3.2.4 Core Port Transaction Ordering

The Core port maintains order of Intel XScale<sup>®</sup> core requests addressing the DDR SDRAM. Coherency between the core port and the IB port is maintained by the MCU as described in Section 8.3.2.6, "MCU Port Coherency" below.

#### 8.3.2.5 IB Port Ordering

The IB port implements ordering of like transactions where reads do not pass reads, and writes do not pass writes. Write transactions are allowed to pass read transactions. Read transactions are processed by the MCU as described in Section 8.3.2.6, "MCU Port Coherency" below.

#### 8.3.2.6 MCU Port Coherency

With the queueing of DDR SDRAM transactions in multiple ports, coherency of memory must be maintained. The MARB maintains memory coherency by ensuring all writes to a given memory address are completed before any read to the same address is processed. This address comparison is done with a 1 KByte granularity.

The highest priority Read transaction is compared to all pending write transactions from both memory transaction queues. When a write transaction is pending for the same memory location (1 KByte granularity), the write is allowed to complete first, before the read transaction is processed. Also, to maintain ordering rules, all write transactions preceding the 'incoherent write' are also processed.



## 8.3.3 DDR SDRAM Memory Support

The 80333 memory controller supports one or two banks of DDR SDRAM. DDR SDRAM allows zero data-to-data wait-state operation. DDR SDRAM offers an extremely wide range of configuration options emerging from the SDRAMs internal interleaving and bursting capabilities.

The MCU supports both 32-bit and 64-bit data bus width memory implementations (with and without ECC). The data bus width is controlled by the DDR SDRAM Control Register. In addition, a 64-bit data bus width DDR SDRAM implementation can be configured to operate as when a region is 32-bits wide, providing higher performance when the core processor is processing data by eliminating any RMW cycle required for 4-Byte store ECC generation.

The MCU supports DDR SDRAM burst length of four. A burst length of four enables seamless read/write bursting of long data streams as long as the memory transaction does not cross the page boundary. Page boundaries are at naturally aligned boundaries. The MCU ensures that the page boundary is not crossed within a single transaction by initiating a disconnect at next ADB (128 byte address boundary) on the internal bus prior to the page boundary.

#### 8.3.3.1 DDR SDRAM Interface

The DDR SDRAM interface signals generated by the memory controller unit are divided into two sections. The first section lists those signals for DDR SDRAM operation, and the second list those additional signals supporting DDR-II SDRAM.

The MCU SDRAM interface provides a flexible mix of combinations including:

#### Table 265. DDR SDRAM Memory Configuration Options

Data Bus Width	ECC Enabled	Maximum Throughput(1)
64 bit	Yes	3200 Mbyte/s
64 bit	No	3200 Mbyte/s
32 bit	Yes	1600 Mbyte/s
32 bit	No	1600 Mbyte/s

NOTE:

1. Based on DDR-II 400MHz SDRAM



Table 266 shows the DDR SDRAM interface signals.

#### Table 266.DDR SDRAM Interface Signals

Pin Name	Description
M_CK[2:0]	DDR SDRAM Clock Out - Three (positive lines) of the output clocks driven to the Unbuffered DIMMs supported by the 80333.Registered DIMMs will only use <b>M_CK[0]</b> which will drive the input to the on-DIMM PLL. Section 8.3.6, "DDR SDRAM Clocking" on page 536 describes the DDR SDRAM clocking strategy for both Unbuffered and Registered DIMMs.
M_CK[2:0]#	DDR SDRAM Clock Out - Three (negative lines) of the output clocks driven to the Unbuffered DIMMs supported by the 80333.Registered DIMMs will only use <b>M_CK[0]#</b> which will drive the input to the on-DIMM PLL. Section 8.3.6, "DDR SDRAM Clocking" on page 536 describes the DDR SDRAM clocking strategy for both Unbuffered and Registered DIMMs.
M_RST#	DDR Registered DIMM Reset - Used to re-initialize registered DIMM during a <b>PWRGD</b> assertion or whenever internal bus reset bit is asserted in the PCSR. PCI Configuration and Status Register - PCSR describes the internal bus reset sequence.
CKE[1:0]	<i>Clock enables</i> - One clock after <b>CKE[1:0]</b> is de-asserted, data is latched on <b>DQ[63:0]</b> and <b>CB[7:0]</b> . Burst counters within DDR SDRAM device are not incremented. De-asserting this signal places the DDR SDRAM in self-refresh mode. For normal operation, <b>CKE[1:0]</b> must be asserted.
CS[1:0]#	<i>Chip Select</i> - Must be asserted for all transactions to the DDR SDRAM device. One per bank.
BA[1:0]	DDR SDRAM Bank Selects - Controls which of the internal DDR SDRAM banks to read or write. <b>BA[1:0]</b> are used for all technology types supported.
MA[10]	Address bit 10 - When high during a read or write command, auto-precharge occurs after the command. During a <b>row-activate</b> command, this bit is part of the address (Section 8.3.3.2).
MA[13:0]	Address bits 12 through 0 - Indicates the row or column to access depending on the state of RAS# and CAS# (Section 8.3.3.2)
DQ[63:0]	Data Bus - 64-bit wide data bus.
CB[7:0]	ECC Bus - 8-bit error correction code which accompanies the data on DQ[63:0].
DM[8:0]	<i>Data Bus Mask</i> - Controls the DDR SDRAM data input buffers. Asserting <b>WE#</b> causes the data on <b>DQ[63:0]</b> and <b>CB[7:0]</b> to be written into the DDR SDRAM devices. <b>DM[8:0]</b> controls this operation on a per byte basis.
DQS[8:0]	<i>Data Strobes</i> - Strobes that accompany the data to be read or written from the DDR SDRAM devices. Data is sampled on the negative and positive edges of these strobes.
WE#	<i>Write Strobe</i> - Defines whether or not the current operation by the DDR SDRAM is to be a read or a write.
RAS#	Row Address Strobe - Indicates that the current address on MA[13:0] is the row.
CAS#	Column Address Strobe - Indicates that the current address on MA[13:0] is the column.
VREF	SDRAM Voltage Reference - is used to supply the reference voltage to the differential inputs of the memory controller pins.



Table 267 shows the DDR-II SDRAM interface signals.

#### Table 267. DDR-II SDRAM Interface Signals

Pin Name	Description	
DQS[8:0]#	<i>Data Strobes Differential-</i> Strobes that accompany the data to be read or written from the DDR SDRAM devices. Data is sampled on the negative and positive edges of these strobes.	
DDRRES[2:1]	2:1] COMPENSATION FOR DDR OCD (analog) DDR2 mode only	
ODT[1:0]	On Die Termination signals control - turns on SDRAM termination during writes	

Utilizing the DDR SDRAM chip selects **CS[1:0]**# and internal bank selects **BA[1:0]**, the MCU keeps a maximum of eight pages open simultaneously. The number of available pages depends on the memory subsystem population. A single bank allows four pages and two banks allow eight pages.

Open pages allow optimal performance when a read or write occurs to an open page. Multiple open pages allow multiple memory segments to be open simultaneously and is well-suited for the 80333 system environment. The MCUs paging algorithm is detailed in Section 8.3.3.5, "Page Hit/Miss Determination" on page 510. The waveforms illustrating the performance issues are in Section 8.3.3.10, "Intel<sup>®</sup> 80333 I/O ProcessorDDR SDRAM Read Cycle" on page 517 and Section 8.3.3.11, "DDR SDRAM Write Cycle" on page 520.



Figure 77 illustrates how two banks of DDR SDRAM would interface with the 80333 through the MCU.

#### Figure 77. Dual-Bank DDR SDRAM Memory Subsystem



### 8.3.3.2 DDR SDRAM Addressing

MCU supports memory subsystem ranging from 64 Mbytes-2 Gbyte. An ECC or non-ECC system may be implemented using x8, or x16 devices and 32-bit or 64-bit data bus widths. Required information to configure the MCU are: row address size, column address size and size of DIMM module which can be read via the DIMM Serial Presence Detect (SPD).

Table 268 illustrates supported DDR SDRAM configurations. This table reflects the bank size associated with the respective row and column address size. The Bank size reflects 64-bit data bus width. The table also specifies the Page Size corresponding to each column address size.

#### Table 268. Supported DDR SDRAM Bank and Page Sizes<sup>a</sup>

Pow Address Size (bits)	Column Address Size (bits)			
Now Address Size (bits)	9	10	11	
12	64MB	128MB	n/a	
13	128MB	256MB	512MB	
14	n/a	512MB	1GB	
Page Size	4KB	8KB	16KB	

a. Table indicates 64-bit wide memory subsystem sizes. For 32-bit wide memory, memory subsystem and page size are half of the size indicated.

The supported SDRAM devices comprise four internal leaves. The MCU controls the leaf selects within the SDRAM by toggling **BA[0]** and **BA[1]**.

#### Table 269.Bank Address Decode

Bank Size <sup>a</sup>	Leaf Select		
	BA[1]	BA[0]	
64M	I_AD[25]	I_AD[24]	
128M	I_AD[26]	I_AD[25]	
256M	I_AD[27]	I_AD[26]	
512M	I_AD[28]	I_AD[27]	
1G	I_AD[29]	I_AD[28]	

a. Table indicates 64-bit wide memory bank sizes. For 32-bit wide memory, bank size is half of the size indicated.

Two DDR SDRAM chip enables (**CS**[1:0]#) support DDR SDRAM memory subsystem consisting of two banks. Base address for two contiguous banks are programmed in DDR SDRAM Base Register (SDBR) and must be aligned to a 32 Mbyte boundary. Size of each DDR SDRAM bank is programmed with DDR SDRAM boundary registers (SBR0 and SBR1). Bank 0 is configurable into multiple regions. By default bank 0 is a single 64-bit region. For higher core data processing performance, a 32-bit region can be defined within Bank 0 with the DDR SDRAM 32-bit Region Size Register (S32SR).

Table 270 is used to determine which DDR SDRAM address decode is used to program "SDRAM Boundary Register 0 - SBR0" and "SDRAM Boundary Register 1 - SBR1".

#### Table 270. DDR SDRAM Address Decode Summary

Row Address Size (bits)	Column Address Size (bits)		
	9	10	11
12	1	1	n/a
13	2	1	1
14	n/a	3	1


Table 271, Table 272 and Table 273 illustrate the address decodes listed in Table 270. These tables show how the internal address is mapped to the **MA[13:0]** lines for DDR SDRAM devices.

DDR SDRAM Address decode #1 is the normal translation used for most memory configurations. When the column and row address size is less than shown in Table 270, the unused address lines are still driven with the internal bus address shown.

#### Table 271. DDR SDRAM Address Decode #1

MA[13:0]	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Row	I_AD[27]	I_AD[25]	I_AD[23]	I_AD[22]	I_AD[21]	I_AD[20]	I_AD[19]	I_AD[18]	I_AD[17]	I_AD[16]	I_AD[15]	I_AD[14]	I_AD[13]	I_AD[12]
Column	-	-	I_AD[26]	$\mathbf{V}^{1}$	I_AD[24]	I_AD[11]	I_AD[10]	I_AD[9]	I_AD[8]	I_AD[7]	I_AD[6]	I_AD[5]	I_AD[4]	I_AD[3]

NOTES:

1. A10 is used for precharge variations on the read or write command. See Table 278 for more details.

2. For the Leaf Selects, see Section 8.3.3.2.

DDR SDRAM Address translation #2 presents **I\_AD[24]** as bit 12 of the row address instead of **I\_AD[25]** as in Address decode #1.

#### Table 272. DDR SDRAM Address Decode #2

MA[12:0]	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Row	I_AD[27]	I_AD[24]	I_AD[23]	I_AD[22]	I_AD[21]	I_AD[20]	I_AD[19]	I_AD[18]	I_AD[17]	I_AD[16]	I_AD[15]	I_AD[14]	I_AD[13]	I_AD[12]
Column	-	-	I_AD[26]	$V^1$	I_AD[24]	I_AD[11]	I_AD[10]	I_AD[9]	I_AD[8]	I_AD[7]	I_AD[6]	I_AD[5]	I_AD[4]	I_AD[3]

NOTES:

1. A10 is used for precharge variations on the read or write command. See Table 278 for more details.

2. For the Leaf Selects, see Section 8.3.3.2.

DDR SDRAM Address decode #3 presents **I\_AD[26]** as bit 13 of the row address instead of **I\_AD[27]** as in Address decode #1.

#### Table 273. DDR SDRAM Address Decode #3

MA[12:0]	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Row	I_AD[26]	I_AD[25]	I_AD[23]	I_AD[22]	I_AD[21]	I_AD[20]	I_AD[19]	I_AD[18]	I_AD[17]	I_AD[16]	I_AD[15]	I_AD[14]	I_AD[13]	I_AD[12]
Column	-	-	I_AD[26]	$V^1$	I_AD[24]	I_AD[11]	I_AD[10]	I_AD[9]	I_AD[8]	I_AD[7]	I_AD[6]	I_AD[5]	I_AD[4]	I_AD[3]

NOTES:

1. A10 is used for precharge variations on the read or write command. See Table 278 for more details.

2. For the Leaf Selects, see Section 8.3.3.2.

Since the MCU supports DDR SDRAM bursting, the MCU increments the column address based on the burst length of four for each DDR SDRAM read or write burst. The MCU supports a sequential and random burst types. Sequential bursting means that the address issued to the DDR SDRAM is incremented by the DDR SDRAM device in a linear fashion during the burst cycle. Random bursting means that the address issued to the DDR SDRAM is any address in a currently active page.



### 8.3.3.3 DDR SDRAM Configuration

#### Table 274. DDR SDRAM Address Register Summary

DDR SDRAM Address Register	Definition
DDR SDRAM Base Register (SDBR)	Lowest address for DDR SDRAM memory space aligned to a 32-Mbyte boundary.
DDR SDRAM Boundary Register 0 (SBR0)	The upper address boundary for bank 0 of DDR SDRAM memory space. SBR0 must be greater than or equal to the value of SDBR[30:25].
DDR SDRAM Boundary Register 1 (SBR1)	The upper address boundary for bank 1 of DDR SDRAM memory space. SBR1 must be greater than or equal to SBR0.
DDR SDRAM 32-Bit Size Register (S32SR)	The size for the memory space to operate as 32-bit memory (in MBs). S32SR must be less than, or equal to 1/2 of bank 0 size. Ignored with a 32-bit data bus width.

*Note:* DDR SDRAM memory space must be aligned to a 32 Mbyte boundary and must *never* cross a 2 Gbyte boundary.

#### Figure 78. DDR SDRAM 64-bit Memory Address Map



#### Figure 79. FDDR SDRAM 64-bit Physical Map



Figure 78 illustrates the80333 DDR SDRAM 64-bit Memory Address Map, and Figure 79 illustrates the Physical map of the 64-bit DDR SDRAM. The size of each region is configured separately. The 32-bit region is only applicable when the 80333 is connected to 64-bit DDR SDRAM memory. With 32-bit DDR SDRAM attached to 80333, all DDR SDRAM memory space behaves as 32-bit DDR SDRAM and the value in S32SR is ignored.

*Note:* When the use of 32-bit memory or the 32-bit memory region is required, write coalescing must be turned off. For more details see *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #8.



When the 32-bit region defined by S32SR is greater than zero, bank 0's address space is split into three regions as illustrated in Figure 78. The lowest addressable region operates as a 32-bit region of size defined in the S32SR. A second region, contiguous with the 32-bit region, and equal in size, is defined as invalid for the DDR SDRAM memory space. This invalid region reflects the other half of the 64-bit DDR SDRAM which is not used in the 32-bit region. Transactions which address this region will result in an error and interrupt to the Intel XScale<sup>®</sup> core. The third region is contiguous with the second and is the remainder of bank 0 address space, and is a 64-bit region.

The base register defines the upper seven address bits of the DDR SDRAM memory space. The boundary registers define the address limits for each DDR SDRAM bank in 32 Mbyte granularity. Table 275 defines the conditions which must be satisfied to activate a DDR SDRAM memory bank.

#### Table 275. Address Decoding for DDR SDRAM Memory Banks

Condition	DDR SDRAM Bank Selected
AD[31] is not equal to the SDBR[31]	None
AD[31] is equal to the SDBR[31] AD[30:25] is greater than or equal to the SDBR[30:25] AD[30:25] is less than the value in SBR0	Bank 0
AD[31] is equal to the SDBR[31] AD[30:25] is greater than or equal to the value in SBR0 AD[30:25] is less than the value in SBR1	Bank 1
AD[31] is equal to the SDBR[31] AD[30:25] is greater than or equal to the value in SBR1	None

Table 276 shows the correct programming values for the DDR SDRAM Bank Size encoding.

#### Table 276. Programming Codes for the DDR SDRAM Bank Size

Bank Size	Code	Bank Size	Code
Empty	00H	256MB	08H
32MB	01H	512MB	10H
64MB	02H	1GB	20H
128MB	04H		

These Bank Size Codes are used in the calculation of the DDR SDRAM boundary registers programming values. Equation 6 and Equation 7 show the required equations to calculate the programming values for the DDR SDRAM boundary registers.

#### Equation 6. Programming Value for DDR SDRAM Boundary Register 0 (SBR0[6:0])

SBR0[6:0] = Bank 0 Size Code+ SDBR[30:25]

#### Equation 7. Programming Value for DDR SDRAM Boundary Register1 (SBR1[6:0])

SBR1[6:0] = Bank 1 Size Code + SBR0[6:0]



Table 277 shows the correct programming values for the 32-bit DDR SDRAM Size Register.

#### Table 277. Programming Values for the DDR SDRAM 32-bit Size Register (S32SR[29:20])

32-bit Region Size	S32SR[29:20]	32-bit Region Size	S32SR[29:20]		
Empty	000H	32M	020H		
1M	001H	64M	040H		
2M	002H	128M	080H		
4M	004H	256M	100H		
8M	008H	512M	200H		
16M 010H		reserved	all other values		

#### Example 2. Address Register Programming Example 1

The user wants to program the DDR SDRAM memory space to begin at B000 0000H. Bank 0 is 64 Mbytes and Bank 1 is 128 Mbytes yielding in a total memory of 192 Mbytes. There is no 32-bit memory region. The memory space summary is:

Memory Space Limit	Address			
DDR SDRAM Base	B000 0000H			
32-Bit Region Top (end 32-bit region)	n/a			
Invalid Region Top (start 64-bit region)	n/a			
Bank 0 Top	B3FF FFFFH			
Bank 1 Top	BBFF FFFFH			

The registers would be programmed as follows:

Bank 0 Size = 64 MB, code =  $000010_2$ Bank 1 Size = 128 MB, code =  $000100_2$ SDBR = B000 0000H, SDBR[30:25] =  $011000_2$ SBR0[6:0]=  $011010_2$  = 1AH SBR1[6:0]=  $011110_2$  = 1EH S32SR[29:20] = 0H, S32SR = 0H

#### Example 3. Address Register Programming Example 2

The user wants to program the DDR SDRAM memory space to begin at A000 0000H. Bank 0 is 64 Mbytes and Bank 1 is un-populated. There is a 4 MB 32-bit region. The memory space summary is:

Memory Space Limit	Address			
DDR SDRAM Base	A000 0000H			
32-Bit Region Top (end 32-bit region)	A03F FFFFH			
Invalid Region Top (start 64-bit region)	A07F FFFFH			
Bank 0 Top	A3FF FFFFH			
Bank 1 Top	n/a			

The registers would be programmed as follows:

Bank 0 Size = 64 MB, code =  $000010_2$ Bank 1 Size = empty, code =  $000000_2$ SDBR = A000 0000H, SDBR[30:25] =  $000000_2$ SBR0[6:0] =  $000010_2 = 02H$ SBR1[6:0] =  $000010_2 = 02H$ S32SR[29:20] = 004H, S32SR =  $0040\ 0000H$ 



#### Example 4. Address Register Programming Example 3

The user wants to program the DDR SDRAM memory space to begin at C000 0000H. Bank 0 is 512 Mbytes and Bank 1 is 512 Mbytes yielding a total memory of 1 Gbyte. The 32-bit memory region is 256 Mbytes (Bank 0 is entirely 32-bit region). The memory space summary is:

Memory Space Limit	Address			
DDR SDRAM Base	C000 0000H			
32-Bit Region Top (end 32-bit region)	CFFF FFFFH			
Invalid Region Top (start 64-bit region)	DFFF FFFFH			
Bank 0 Top	DFFF FFFFH			
Bank 1 Top	FFFF DFFFH*			

Note: Bank 1 Top is limited by PMMR space overlap at FFFF E000H in this example.

The registers would be programmed as follows:

Bank 0 Size = 512 MB, code =  $010000_2$ Bank 1 Size = 512 MB, code =  $010000_2$ SDBR = C000 0000H, SDBR[30:25] =  $010000_2$ SBR0[6:0] =  $100000_2 = 40$ H SBR1[6:0] =  $110000_2 = 50$ H S32SR[29:20] = 100H, S32SR = 1000 0000H

#### Example 5. Address Register Programming Example 4

The user wants to program the DDR SDRAM memory space to begin at 0000 0000H. Bank 0 is 1 Gbyte and Bank 1 is 1 Gbyte yielding a total memory of 2 Gbytes. There is a 128MB 32-bit region. THis configuration also requires that the ATU Outbound Direct Addressing window be disabled before configuring the MCU for 2 Gbytes. The memory space summary is:

Memory Space Limit	Address			
DDR SDRAM Base	0000 0000H			
32-Bit Region Top (end 32-bit region)	01FF FFFFH			
Invalid Region Top (start 64-bit region)	03FF FFFFH			
Bank 0 Top	3FFF FFFFH			
Bank 1 Top	7FFF FFFFH			

The registers would be programmed as follows:

Bank 0 Size = 1 GB, code = 100000<sub>2</sub> Bank 1 Size = 1 GB, code = 100000<sub>2</sub> SDBR =0000 0000H, SDBR[30:25] = 000000<sub>2</sub> SBR0[6:0] = 0100000<sub>2</sub> = 20H SBR1[6:0] = 1000000<sub>2</sub> = 40H S32SR[29:20] = 080H, S32SR = 0800 0000H



#### 8.3.3.4 32-bit Data Bus Width

Using 64 Mbit SDRAMs, a 64-bit data bus yields a minimum memory size of 64 Mbytes. To address cost-sensitive applications requiring less than 64 Mbytes of local memory, the MCU supports a 32-bit data bus. While 32-bit mode decreases the minimum memory size to 32 Mbytes, the bus throughput also reduces by half.

The MCU does not support switching between 32-bit data bus width and 64-bit data bus width. The data bus width is selected by bit 2 of the SDCR (see Section 8.7.2, "SDRAM Control Register 0 - SDCR0" on page 547). The default is 64-bit bus width.

Reducing the data bus width by half also reduces the page size by half. The page sizes listed in Section 8.3.3.2 are listed for a 64-bit data bus, and must therefore are one half for a 32-bit data bus. The MCU disconnects from the internal bus when the page is crossed during a burst read or write.

*Note:* When the use of 32-bit memory or the 32-bit memory region is required, write coalescing must be turned off. For more details see *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Erratum #8.

#### 8.3.3.5 Page Hit/Miss Determination

The MCU keeps up to eight pages open simultaneously; one page each of Bank0/Leaf0, Bank0/Leaf1, Bank0/Leaf2, Bank0/Leaf3, Bank1/Leaf0, Bank1/Leaf1, Bank1/Leaf2, and Bank1/Leaf3. The page size is based on the DDR technology as specified in Section 8.3.3.2.

The MCU logic determines the hit/miss status for reads and writes. For a new DDR SDRAM transaction, the MCU compares the address of the current transaction with the address stored in the appropriate page address register. Given the supported DDR SDRAM devices and two banks, there are eight pages kept open simultaneously. The DDR SDRAM chip enables (**CS**[1:0]#) and leaf selects (**BA**[1:0]) determine which page address to compare.

When the current transaction misses the open page selected then the MCU closes the open page pointed to by **CS[1:0]**# and **BA[1:0]** by issuing a **precharge** command. The MCU opens the current page with a **row-activate** command and the transaction completes with a **read** or **write** command. When the MCU opens the current page, the row address from the corresponding memory transaction queue is stored in the page address register pointed to by **CS[1:0]**# and **BA[1:0]** so it may be compared for future transactions. See Table 271, Table 272and Table 273 for address mapping to row address.

Once the MARB issues a command to the MCU DDR Control Block, the paging logic makes a hit/miss comparison. The performance is best for page hits and therefore the MCUs behavior is different for the hit and miss scenario.

For a page hit, the MCU does not need to open the page (assert RAS#) and avoids the RAS-to-CAS delay achieving greater performance. The waveform for a write including the row activation in the case of a page miss is illustrated in Figure 87. For a page hit, the two cycles required for row activation are saved resulting in lower first word write latency.

When the current transaction hits the open page, then the page is already active and the **read** or **write** command may be issued without a **row-activate** command. When the next transaction is the same command type as the current transaction, and also a page hit, the MCU does not need to issue the command again, but simply drive column address for an open page.

When the refresh timer expires and the MCU issues an **auto-refresh** command, all pages are closedFigure 86 illustrates the performance benefit of a read hit versus a read miss in Figure 89. Figure 87 illustrates the performance benefit of a write hit versus a write miss in Figure 92.

# intel

#### **Page Address Registers** Page 0 (closed) Page 1 (closed) Page 2 Bank0 Leaf0 Page 2 (OPEN) Leaf 0 Page 3 (closed) Page 0 Bank0 Leaf1 Page 1 Bank0 Leaf2 Page 0 (OPEN) Page 1 Page 1 (closed) Bank0 Leaf3 Page 2 (closed) Leaf 1 Page 3 (closed) Page 2 Bank1 Leaf0 Page 0 Bank1 Leaf1 Bank 0 Page 0 (closed) Page 1 Bank1 Leaf2 Page 1 (OPEN) Page 2 (closed) Page 1 Leaf 2 Bank1 Leaf3 Page 3 (closed) Page 0 (closed) Page 1 (OPEN) Page 2 (closed) Leaf 3 Page 3 (closed) Page 0 (closed) Page 1 (closed) Page 2 (OPEN) Leaf 0 Page 3 (closed) Page 0 (OPEN) Page 1 (closed) Page 2 (closed) Leaf 1 Page 3 (closed) Bank 1 Page 0 (closed) Page 1 (OPEN) Page 2 (closed) Leaf 2 Page 3 (closed) Page 0 (closed) Page 1 (OPEN) Page 2 (closed) Leaf 3 Page 3 (closed)

#### Figure 80. Logical Memory Image of a DDR SDRAM Memory Subsystem

Figure 80 illustrates how the logical memory image is partitioned with respect to open and closed pages. When the above image represents a 128 Mbyte DDR SDRAM memory size, each bank is 64 Mbytes and each leaf is 16 Mbytes.

Only one page may be open within each of the leaf blocks. The page sizes depend on the memory size implemented in the DDR SDRAM memory subsystem as listed in Section 8.3.3.2. The programmer can optimize DDR SDRAM transactions by partitioning code and data across the leaf boundaries to maximize the number of page hits.

#### 8.3.3.6 **On DIMM Termination**

The memory controller supports On DIMM Termination (ODT) when controlling DDR-II SDRAM technology. ODT eliminates the need for on board termination resisters, thereby reducing the implementation cost and area. Control for enabling ODT is performed in Section 8.7.2, "SDRAM Control Register 0 - SDCR0" on page 547.

#### 8.3.3.7 **DDR SDRAM Commands**

The MCU issues specific commands to the DDR SDRAM devices by encoding them on the CS[1:0]#, RAS#, CAS#, and WE# inputs. Table 278 lists all of the DDR SDRAM commands understood by DDR SDRAM devices. The MCU supports a subset of these commands.

**Table 278. DDR SDRAM Commands** 

Commondab			Condit	Commonto		
Command	SCE#	RAS#	CAS#	SWE#	Other	Comments
NOP	0	1	1	1		No Operation
Mode Register Set	0	0	0	0	<b>BA[0]</b> = Sel <sup>c</sup> <b>BA[1]</b> = 0	Load the (Extended) Mode Register from <b>MA[13:0]</b>
Row Activate	0	0	1	1	BA[1:0] = Leaf	Activate a row specified on <b>MA[13:0]</b>
Read	0	1	0	1	<b>BA[1:0]</b> = Leaf <b>MA[10]</b> = 0	Column burst read Column address on <b>MA[13:0]</b>
Read w/ Auto-Precharge	0	1	0	1	BA[1:0] = Leaf MA[10] = 1	Column burst read with row precharge at the end of the transfer
Write	0	1	0	0	<b>BA[1:0]</b> = Leaf <b>MA[10]</b> = 0	Column burst write Column address on <b>MA[13:0]</b>
Write w/ Auto-Precharge	0	1	0	0	<b>BA[1:0]</b> = Leaf <b>MA[10]</b> = 1	Column burst write with row precharge at the end of the transfer
Precharge	0	0	1	0	<b>BA[1:0]</b> = Leaf <b>MA[10]</b> = 0	Precharge a single leaf
Precharge All	0	0	1	0	<b>MA[10]</b> = 1	Precharge all leaves
Auto-Refresh	0	0	0	1		Refresh both banks from on-chip refresh counter
Self-Refresh	0	0	0	1	<b>CKE</b> = 0	Refresh autonomously while <b>CKE</b> = 0
Power Down	х	х	х	х	<b>CKE</b> = 0	Power down when both banks precharged when <b>CKE</b> = 0
Stop	0	1	1	0		Interrupt a read or write burst.

This table copied from New DRAM Technologies by Steven Przybylski. a. b.

Shaded boxes indicate commands not supported by 80333. They are included for completeness.

During a Mode Register Set command, BA[1:0] = 002 selects the Mode Register, BA[1:0] = 012 selects the Extended Mode c. Register, all others are served.

DDR SDRAM commands are synchronous to the clock so the MCU sets up the above conditions prior to the M\_CK[2:0] rising edge.

# intel

## 8.3.3.8 DDR SDRAM Initialization

Since DDR SDRAM devices contain a controller within the device, the MCU must initialize them specifically. Upon the deassertion of **I\_RST#**, software initializes the DDR SDRAM devices with the sequence illustrated with Figure 84:

- 1. The MCU applies the clock (**M\_CK[2:0**]) at power up along with system power (clock frequency unknown).
- 2. The MCU must stabilize M\_CK[2:0] within 100 µs after power stabilizes.
- The MCU holds all the control inputs inactive (RAS#, CAS#, WE#, SCE[1:0]# = 1), places all data outputs and strobes in the High-Z state (DQS[8:0], DQ[71:0]), and deasserts CKE[1:0] for a minimum of 200 us after supply voltage reaches the desired level. Asserting PWRGD achieves this state.
- 4. Software disables the refresh counter by setting the RFR to zero.
- 5. Software issues one **NOP** cycle after the 200 us device deselect. A **NOP** is accomplished by setting the SDIR to 0011<sub>2</sub>. The MCU asserts **CKE[1:0]** with the NOP.
- 6. Software issues a **precharge-all** command to the DDR SDRAM interface by setting the SDIR to 0010<sub>2</sub>.
- 7. Software issues an **extended-mode-register-set** command to enable the DLL by writing  $0100_2$  to the SDIR. The MCU supports the following DDR and DDR-II SDRAM mode parameters:
  - a. DLL = Enable/Disable
  - b. Off-Chip Driver (OCD) Impedance Adjustment (set DCAL registers) applies to DDR-II SDRAM only.
  - c. Additive Latency (AL) is always zero for 80333.

#### Figure 81. Supported DDR SDRAM Extended Mode Register Settings





#### Figure 82. Supported DDR-II SDRAM Extended Mode Register Settings



- 8. After waiting T<sub>mrd</sub> cycles, software issues a **mode-register-set** command by writing 0001<sub>2</sub> to the SDIR to program the DDR SDRAM parameters and to **Reset** the DLL. The MCU supports the following DDR SDRAM mode parameters:
  - a. CAS Latency (tCAS) = two and one-half for DDR, or three<sup>1</sup> or four for DDR-II based on the programmed setting in Section 8.7.2, "SDRAM Control Register 0 SDCR0"
  - b. Burst Type = Sequential
  - c. Burst Length (BL) = four

#### Figure 83. Supported DDR SDRAM Mode Register Settings



- 9. After waiting  $T_{mrd}$  cycles, software issues a **precharge-all** command to the DDR SDRAM interface by setting the SDIR to 0010<sub>2</sub>.
- 10. After waiting  $T_{rp}$  cycles, software provides two **auto-refresh** cycles. An **auto-refresh** cycle is accomplished by setting the SDIR to  $0110_2$ . Software must ensure at least  $T_{rfc}$  cycles between each **auto-refresh** command.
- 11. Following the second **auto-refresh** cycle, software must wait  $T_{rfc}$  cycles. Then, software issues a **mode-register-set** command by writing to the SDIR to program the DDR SDRAM parameters **without** resetting the DLL by writing 0000<sub>2</sub> to the SDIR.
- 12. The MCU may issue a row-activate command  $T_{mrd}$  cycles after the mode-register-set command.
- 13. Software re-enables the refresh counter by setting the RFR to the required value.

<sup>1.</sup> CAS latency of 3 is not supported.



The waveform in Figure 84 illustrates the DDR SDRAM initialization sequence.

#### Figure 84. DDR SDRAM Initialization Sequence (controlled with software)



When the DDR SDRAM subsystem implements ECC (see Section 8.3.4, "Error Correction and Detection" on page 525), then initialization software should initialize the entire memory array with the 80333. It is important that every memory location has a valid ECC byte. The AAU includes a memory block fill mode which can be used to fill the memory array with a constant (Section 7.5.7, "Memory Block Fill Operation" on page 454), thereby initializing the associated ECC bytes in the process. A small portion of memory must be initialized by the Intel XScale<sup>®</sup> core processor software to store the AAU descriptor. When the memory array is not initialized, the BIU may attempt to read memory locations beyond the specified word(s). In this case, the MCU will report an ECC error even though software did not specifically request the un-initialized data.



#### 8.3.3.9 DDR SDRAM Mode Programming

The MCU programs the DDR SDRAM devices through a **mode-register-set** command. During the initialization sequence this command sets the DDR SDRAM mode register (see Section 8.3.3.8, "DDR SDRAM Initialization" on page 513) by programming the SDIR and SDCR[1:0].

The DDR SDRAM state machine ensures that a **row-activate** command is issued no sooner than  $T_{mrd}$  cycles after the **mode-register-set** command.

The values to be programmed in the SDCR[1:0] registers are based on the SDRAM devices being interfaced to 80333. Because the parameters that define the time between allowed commands are programmable, this allows flexibility in the type of DDR device that is selected, in addition to de-coupling the hardware to any frequency dependencies. See Section 8.7.2, "SDRAM Control Register 0 - SDCR0" on page 547 and Section 8.7.3, "SDRAM Control Register 1 - SDCR1" on page 549 for equations and values for programming of the MCU timing parameters.

*Note:* The MCU\_DDRSM will NOT interact properly with the DDR SDRAM until the SDCR[1:0] registers have been programmed.

The timing parameters and protocol for all non-read and non-write commands are derived directly from the *JEDEC Standard Double Data Rate (DDR) SDRAM Specification JESD79*, June 200 and *JEDEC DDR - II SDRAM Specification*, September 2002. Please see the JEDEC specification for the timing parameters specific to the DDR device that is to be implemented in the system.



# 8.3.3.10 Intel<sup>®</sup> 80333 I/O ProcessorDDR SDRAM Read Cycle

The MCU performance is optimized for page hits and the MCUs behavior is different for the hit and miss scenario.

The waveform for a read including the row activation in the case of a page miss is illustrated in Figure 86. For a page hit, the two cycles required for row activation are saved resulting in lower first word read latency.



#### Figure 85. DDR SDRAM Pipelined Reads





#### Figure 86. DDR SDRAM Read, 36 Bytes, ECC Enabled, BL=4

1. Each of the MCU inbound memory transaction ports decodes the address to determine when the transaction should be claimed.

- When the address falls in the DDR SDRAM address range indicated by the SDBR, SBR0, SBR1, and S32SR the MCU claims the transaction and latches the transaction in the respective memory transaction queue.
- 2. Once the MARB selects the highest priority transaction from the memory transaction queues, it forwards the transaction to the DDR SDRAM control block. The DDR SDRAM Control Block decodes the address to determine whether or not any of the open pages are hit.

# intel®

A read that misses the open pages encounters a miss penalty because the currently open page needs to be closed before the read can be issued to the new page. Refer to Section 8.3.3.5, "Page Hit/Miss Determination" on page 510 for the paging algorithm details. When a page hit occurs, steps 3-4 are skipped by the MCU.

- 3. The DDR SDRAM Control Block closes the currently open page by issuing a **precharge** command to the currently open row. (Not depicted in Figure 84).
  - The DDR SDRAM Control Block waits T<sub>rp</sub> cycles after the precharge before issuing the row-activate command for the new read transaction.
- 4. The **row-activate** command enables the appropriate row.
  - The DDR SDRAM Control Block asserts RAS#, de-asserts WE#, and drives the row address on MA[13:0].
- In the following cycle in the case of a page hit or after T<sub>rcd</sub> cycles in the case of a page miss, the DDR SDRAM Control Block asserts CAS#, de-asserts WE#, and places the column address on MA[13:0]. This initiates the burst read cycle.
- 6. After the CAS latency expires, the DDR SDRAM device drives data to the MCU. A CAS latency of 2 is depicted in Figure 86.
- Upon receipt of the data, the DDR SDRAM Control Block calculates the ECC code from the data and compares it with the ECC returned by the DDR SDRAM array. Section 8.3.4, "Error Correction and Detection" on page 525 explains the ECC algorithm in more detail.
- 8. Assuming the calculated ECC matches the read ECC, the DDR SDRAM Control Block drives the data back to the corresponding memory transaction queue.
- For each burst read issued, the memory controller increments the column address by four.

The MCU continues to return data to the corresponding memory transaction queue based on the byte count of the transaction.

### 8.3.3.11 DDR SDRAM Write Cycle

All write transactions to the DDR SDRAM are posted to the MCU in the memory transaction queues. This implies that the transaction completes between a given port and corresponding unit prior to data being written to the SDRAM array. Once the MARB issues a **write** command from a memory transaction queue to the MCU DDR Control Block, the paging logic makes a hit/miss comparison. The performance is best for page hits and therefore the MCUs behavior is different for the hit and miss scenario.

Write transactions require ECC codes to be generated and stored in the SDRAM array with the data being written. The behavior is different depending on the size of the data being written. Section 8.3.4, "Error Correction and Detection" on page 525 explains the ECC algorithm in more detail.

For a page hit, the MCU does not need to open the page (assert RAS#) and avoids the RAS-to-CAS delay achieving greater performance. The waveform for a write including the row activation in the case of a page miss is illustrated in Figure 87. For a page hit, the two cycles required for row activation are saved resulting in lower first word write latency.

#### Figure 87. DDR SDRAM Write, 36 Bytes, ECC Enabled, BL=4



#### Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Memory Controller

# intel

- 1. Each of the MCU inbound memory transaction ports decodes the address to determine when the transaction should be claimed.
  - When the address falls in the DDR SDRAM address range indicated by the SDBR, SBR0, SBR1, and S32SR the MCU claims the transaction and latches the transaction in the respective memory transaction queue.
- 2. Once the MARB selects the highest priority transaction from the memory transaction queues, it forwards the transaction to the DDR SDRAM control block. The DDR SDRAM Control Block decodes the address to determine whether or not any of the open pages are hit.
  - The ECC logic generates the ECC code for the data to be written.

A write that misses the open page encounters a miss penalty because the currently open page needs to be closed before the write can be issued to the new page. Refer to Section 8.3.3.5, "Page Hit/Miss Determination" on page 510 for the paging algorithm details. When a page hit occurs, steps 2-3 are skipped by the MCU.

- 3. The DDR SDRAM Control Block closes the currently open page by issuing a **precharge** command to the currently open row.
  - The DDR SDRAM Control Block waits T<sub>rp</sub>cycles after the precharge before issuing the row-activate command for the new write transaction.
- 4. The **row-activate** command enables the appropriate row.
  - The DDR SDRAM Control Block asserts RAS#, de-asserts WE#, and drives the row address on MA[13:0].
- 5. After T<sub>rcd</sub> cycles in the case of a page miss, the DDR SDRAM Control Block asserts CAS#, asserts WE#, and places the column address on MA[13:0]. This initiates the burst write cycle. The DDR SDRAM Control Block drives the data to be written and its ECC code to the DDR SDRAM devices.
- The DDR SDRAM Control Block drives the new data to the corresponding memory transaction queue each cycle until the transaction is completed with the byte count expiring, or the transaction is interrupted when preemption conditions are met.
- For each burst issued, the DDR SDRAM Control Block increments the address by four.
- When ECC is enabled, when the data to write is not aligned on an 8 byte boundary (4 byte for 32-bit data bus width or 32-bit region), the DDR SDRAM Control Block will perform a read-modify-write of the entire 8 byte aligned quad-word (4 byte aligned double-word for 32-bit data bus width or 32-bit region) and incorporate the new data while regenerating ECC.



#### Figure 88. DDR SDRAM Pipelined Writes

inte



### 8.3.3.12 DDR SDRAM Refresh Cycle

Since the DDR SDRAM is a dynamic memory, the MCU issues a refresh cycle periodically. The interval of these refresh cycles is programmable in the RFR register. The DDR SDRAM device generates the refresh address internally. The MCU initiates two sequential refresh cycles (one per bank) after the MCUs refresh timer expires and any current transaction is complete. The waveform in Figure 89 illustrates the case where the refresh timer expires while the memory bus is not busy.

Figure 89. Refresh While the Memory Bus is Not Busy





- Once the refresh timer expires, the MCU knows that a refresh cycle is necessary.
  - The refresh timer continues to count for the next refresh cycle.
- The MARB allows the current transaction to complete.
  - When the DDR SDRAM Control Block and the DDR SDRAM array are transferring data, or a CMTQ tenure is ongoing, the refresh cycle is queued until the transaction is complete or the CMTQ tenure expires.
- The DDR SDRAM Control Block closes all open pages with a **precharge-all** command to all the populated DDR SDRAM banks.
  - The DDR SDRAM Control Block resets the page register valid bits.
- The DDR SDRAM Control Block issues an **auto-refresh** command to DDR SDRAM bank 0.
  - This command affects all internal leaves.
- In the next cycle, the DDR SDRAM Control Block issues an **auto-refresh** command to DDR SDRAM bank 1.
- After T<sub>rfc</sub> cycles, the DDR SDRAM Control Block can service a new transaction or another refresh cycle.

It is recommended that the RFR ("Frequency Register - RFR" on page 562) is programmed with the value to achieve 7.8 us, though some DDR SDRAM devices may provide for the ability to refresh at a period of 15.6 us. The value is based on the frequency of the DDR SDRAM and Table 279 can provides for these two typical values.

#### Table 279. Typical Refresh Frequency Register Values

DDR Speed	<b>7.8</b> μ <b>s Value</b>	15.6 μs Value
333 MHz	500H	АООН
400 MHz	600H	Соон

The longest possible internal bus transaction is writing a 1 Kbyte burst where each data cycle results in a read-modify-write due to partial writes (see Section 8.3.4.2, "ECC Generation for Partial Writes" on page 527). The longest possible CMTQ tenure is 16 transactions where each of the transaction are page misses and partial writes. Such periods potentially require queueing two refresh cycles.

# intel®

# 8.3.4 Error Correction and Detection

The MCU is capable of correcting any single bit errors and detecting any double bit errors in the 80333 DDR SDRAM memory subsystem. ECC enhances the reliability of a memory subsystem by correcting single bit errors caused by electrical noise or occasional alpha particle hits on the DDR SDRAM devices.

Similar to parity, which simply detects single bit errors, error correction requires an additional 8-bit code word for the 64-bit (32-bit) datum. This means that a memory must have the additional 8-bit error correction code (**CB**[7:0]) per 64-bit (32-bit) datum (**DQ**[63:0]) resulting in a 72-bit (40-bit) wide memory subsystem. During DDR SDRAM read cycles, the DDR SDRAM Control Block detects single bit errors and corrects the data prior to returning the data to the respective memory transaction queue. DDR SDRAM write cycles generate the ECC and sends it with the data to the memories.

In the 32-bit region with 64-bit memory, or with 32-bit wide memory, the 80333 will zero extend the 32-bit datum to a 64-bit datum in order to generate, check and correct ECC. This means that a 32-bit datum memory with ECC will result in a 40-bit wide memory since an 8-bit error correction code is still required.

Scrubbing is the process of correcting an error in the memory array. The chance of an unrecoverable multi-bit error increases when the software does not correct a single-bit error in the array. For the 80333, scrubbing is handled by software. When error reporting is enabled, the MCU logs the error type in ELOG0 or ELOG1 and the address in ECAR0 or ECAR1 when an error occurs.



#### 8.3.4.1 ECC Generation

For write operations, the MCU generates the error correction code which is written along with the data. This section describes the operation of the DDR SDRAM Control Block for ECC generation in a 64-bit wide memory and 64-bit region. The same principles apply for 32-bit wide memory and in the 32-bit region in 64-bit wide memory, however the MCU will generate 8-bit wide ECC by zero extending the data to 64-bits The algorithm for a write transaction is:

```
if data to write is 64 bits wide
  Generate the ECC_with the G-matrix
  Write the new data and ECC
  else {Partial Write}
  Read entire 64-bit data word from memory
  Merge the new data portion with the data from memory
  Generate the new ECC with the G-matrix
  Write new data and ECC
```

Figure 90 shows how the data logically flows through the ECC hardware for a write transaction.





The G-Matrix in Figure 91 generates the ECC. The data to be written is input to the matrix and the output is the ECC code. Each row of the G-Matrix indicates which data bits of **AD[63:0]** needs to be XORed together to form the ECC bit. The resulting ECC bits are driven on **CB[7:0]**.

# intel

### 8.3.4.2 ECC Generation for Partial Writes

# Figure 91. Intel<sup>®</sup> 80333 I/O Processor G-Matrix (generates the ECC)

	32			×	×	1 [	×				1	0	×						×	×
	33			×	×			×				٢		×					×	×
	34			×	×				×	×		2			×				×	×
	35			×	×				×	×	3				×			×	×	
	36		×		×		×					4	×					×		×
	37		×		×			×				5		×				×		×
	38		×		×				×			9			×			×		×
	39		×		×					×		7				×		×		×
	40	×			×		×					8	×				×			×
	41	×			×			×				6		×			×			×
	42	×			×				×			10			×		×			×
	43	×			×					×		11				×	×			×
	44		×	×			×					12	×					×	×	
	45		×	×				х				13		×				×	×	
JS	46		×	×					×		JS	14			×			×	×	
itio	47		×	×						×	itio	15				×		×	×	
Pos	48	×		×			×				Pos	16	×				×		×	
Bit	49	×		×				х	$\uparrow$		Bit	17		×			×		×	
Data	50	×		×					×		Data	18			×		×		×	
-	51	×		×						×		19				×	Х		×	
	52	×	×				$\times$					20	×				×	×		
	53	×	×					×				21		×			×	×		
	54	×	×						×			22			×		×	×		
	55	×	×							×		23				×	×	×		
	56	×	×	×	×					×		24				×	$\times$	$\times$	$\times$	×
	57			×			Х	×	×	×		25	×	×	×	×			×	
	58	×	×	×								26					×	×	×	
	59	×		×	×							27					×		×	×
	60							×	×	×		28		×	×	×				
	61					$\lfloor  brace$	×	×		×		29	×	×		×				
	62	×	×	×	×		×					30	×				×	×	×	×
	63		×				×	×	×	×		31	×	×	×	×		×		
		CB0	CB1	CB2	CB3		CB4	CB5	CB6	CB7			CB0	CB1	CB2	CB3	CB4	CB5	CB6	CB7



When the memory transaction writes less than the data bus width programmed in the SDCR, then the DDR SDRAM Control Block translates the write transaction into a read-modify-write transaction. For a partial write, the DDR SDRAM Control Block calculates the ECC for the modified datum and writes it back. So, when an external unit issues a write cycle with partial data to an MCU port, the MCU:

- 1. Issues a 64-bit (32-bit) read.
- 2. Modifies the value with the new portion to be written.
- 3. Calculates the ECC on the modified value.
- 4. Writes the 64-bit (32-bit) value and ECC.
- *Note:* When the MCU detects a single-bit error during the read, it is corrected BEFORE being merged with the write data so the corrected data is written back to the array. When a multi-bit error is detected, the MCU causes an interrupt to the core by writing to the MCISR. The memory location is overwritten by the MCU with the error data but valid ECC, making the contents of memory invalid. For more details on how the MCU handles error conditions, see Section 8.5, "Interrupts/Error Conditions" on page 542.

Figure 92 shows an example where the data of a write is less than 64-bits wide. The waveform illustrates how the DDR SDRAM Control Block issues a read-modify-write cycle for the data  $(D_1)$ .

*Note:* In 32-bit wide memory and in the 32-bit region in 64-bit wide memory, the DDR SDRAM Control Block will still generate 8-bit wide ECC by zero extending the data to 64-bits. A partial write is a write of less then 4-Bytes.

528



L Write

New ECC for D<sub>0</sub>

Merge New Data and Generate

A7856-02

#### <-tcκ tCL CK CK# tн tıs CKE t<sub>IH</sub> tis Command NOF ACT NOF NOF NOP NOP Writ NO t<sub>IH</sub> t<sub>IS</sub> → A0-A9, RA COL r COLI A11, A12 t<sub>IS</sub>-> t<sub>IS</sub> DIS AF DIS AF A10 RA tiH → t<sub>IS</sub> → -BA0, BA1 BA x BA x CL=2 DQS $\chi$ DQ $\sum$ ∕\_D<sub>in</sub> **↓** D<sub>out</sub> 00 FF DM I L

Read  CAS Latency

ECC Calculation Comparision, and Correction for D<sub>0</sub>

#### Figure 92. Sub 64-bit DDR SDRAM Write (D<sub>0</sub>)

Document Number: 305432001US

### 8.3.4.3 ECC Checking

When enabled, the ECC logic uses the following ECC read algorithm. This algorithm corrects the data before it's driven onto the internal bus. The ECC algorithm for a read transaction is:

```
Read 64-bit data and 8-bit ECC
Compute the syndrome by passing the 64-bit data through the G-Matrix and XORing the
8-bit result with the 8-bit ECC
if the syndrome <> 0 {ECC Error}
   Look up in H-matrix to determine error type
   Register the address where the error occurred
   if error is correctable {single bit}
       if single-bit error correction is enabled
          Correct data
           Send corrected data to internal bus
       if single bit error reporting is enabled
           Interrupt core for software scrubbing
   else {uncorrectable}
       if the read cycle is not part of a RMW cycle {read}
           Target-Abort the Internal Bus read transaction.
       else {write requiring RMW}
           Merge the new data portion with the read data from memory
           Generate the new ECC with the G-matrix
           Write new data and ECC
       if multi-bit error reporting is enabled
               Interrupt the core for uncorrectable error
```

When the MCU reads the ECC code from the memory subsystem, it is compared (XORed) with an ECC that the MCU generates from the data read from the memory. The result is called the syndrome. Table 280 shows how the MCU decodes the syndrome for DDR SDRAM read cycles.

#### Table 280.Syndrome Decoding

Error Type	Symptom
None	The syndrome is 0000 0000.
Single-Bit	Use the H-Matrix in Figure 94 to determine which bit the MCU will invert to fix the error.
Multi-Bit	When the Syndrome does not match an 8-bit value in the H-matrix, the error is uncorrectable



Figure 93 shows how the data flows through the ECC hardware for a read transaction.





Figure 94 illustrates the H-Matrix used for decoding the syndrome. For single-bit errors, the H-Matrix indicates the bit that contains the error and consequently, which bit to fix.



	9		_		_	11	_				1		Γ_							_		
	7 34		-	_	-		-					0	-							-	-	
	3.37		-		-			-				-		-						-	-	
	35		-	-	-				-		-	2			-					-	-	
	35		-		-					-		3				-				-	-	
	40	-			-		1					4	-						-		-	
	41	-			-			-				5		-					٢		-	
	42	-			~				-			9			-				-		-	
	43	-			~					-		7				-			-		-	
	44		-	~			-					8	-					-			-	
	45		~	~				٢				6		~				١			-	
	46		-	-					-			10			-			-			-	
	47		-	-						-		11				٦		-			-	
	48	-		-			-					12	-						-	-		
	49	-		-				۲				13		-					-	-		
	50	٦		-					-		]	14			-				-	-		
	51	-		-		1				۲	1	15				-			٢	-		
"	52	-	-				٦				6	16	-					۲		-		
ions	53	-	-		1	1		-			ion	17		-				٦		-		
ositi	54	-	-		1	1			-		osit	18			-			-		-		
it Pç	55	-	-			1				-	ît Pç	19				-		-		-		
B	56	-	-	-	-	1				-	B	20	-					-	-			
	57			-			-	-	-	-		21		-				٢	-			
	58	-	-	-		1						22			-			٦	-			
	59	-		-	-							23				-		-	-			
	09					Π		-	-	-		24				-		-	-	-	-	
	61					$\square$	-	-		-		25	-	-	-	-				-		
	62	-	-	-	-	1	-					26						-	-	-		
	63		-				-	-	-	-		27						-		-	-	
	∠B									-		28		-	-	-						
	e CB			$\vdash$					-			29	-	-		-						
	5B							-				30	-					٢	-	-	-	
	4 CB			$\vdash$			-					31	-	-	-	-			+			
	338				-							32 3			-	-		-				
	5 8 8 8			-								33			-	-			-			
	8-		-	-								2			-	-				-		
	80	-										35 3			-	-					-	
	0-	0	Σ	2	ñ		4	55	9	1		en en	0	2	2	33		4	5	9	1	
		S S	S	S	S	1	0	S)	ŝ	S	J		S	0)	ŝ	S	I	0	S)	S	3	

### Figure 94. Intel<sup>®</sup> 80333 I/O Processor H-Matrix (indicates the single-bit error location)

# intel®

Referring to Figure 93, the syndrome bits are created by XORing the data bits as indicated by the appropriate row of the G-Matrix in Figure 91 with the corresponding ECC bit. For example, the MCU derives syndrome bit 0 by XORing data bits 0, 4, 8, 12, 16, 20, 25, 29..31, 40..43, 48..56, 58, 59, 62, and ECC bit 0 (physically read on **CB[0]**). The MCU performs eight such XOR operations (one per syndrome bit).

When decoding the syndrome indicates multi-bit error (see Table 280), the transaction results in a target-abort for Internal Bus transactions, or a multi-bit error in the BIU for Core transactions. When an internal bus master detects a target-abort, the master asserts an interrupt to the core. Write cycles are posted to the memory transaction queues, and already completed to the initiating master. For write cycles with a multi-bit error and ECC Error reporting is enabled, the MCU reports the interrupt in the MCISR and interrupts the core.

When the syndrome indicates a single-bit error and single-bit error correction is enabled, the H-Matrix is used to determine the bit in error (see Figure 94). For example, when the syndrome was 1100 0001, the error is with bit 0 of **DQ[63:0]**. The MCU inverts bit 0 before driving the data on **AD[63:0]**.

When error reporting is enabled in the ECCR and the MCU detects a single-bit or multi-bit error, the MCU stores the address in ECARx and the syndrome in ELOGx. Then, the MCU signals an interrupt to the core. Software decides how to proceed through an interrupt handler. By registering the address in ECARx, software can identify the faulty DIMM.

For details about the MCU error conditions and how the MMR registers are affected, refer to Section 8.5, "Interrupts/Error Conditions" on page 542.

*Note:* In 32-bit wide memory and in the 32-bit region in 64-bit wide memory, the DDR SDRAM Control Block will still generate 8-bit wide ECC by zero extending the data to 64-bits. A partial write is a write of less then 4-Bytes.



#### 8.3.4.4 Scrubbing

Fixing the data error in memory is called scrubbing. The 80333 relies on Intel XScale<sup>®</sup> core software to perform the scrubbing. When the MCU detects an error during a read, the MCU logs the address where the error occurred and interrupts the core. The core decides how to fix the error through an interrupt handler. Software could decide to perform the scrubbing on:

- the data location that failed
- the entire row of the data that failed
- the entire memory

For single-bit errors reported on a write transaction scrubbing is not required, as the MCU will have scrubbed the data during the RMW operation. For single-bit errors, the error is fixed by reading the location that failed and writing back the data after the ECC hardware fixed it. The scrubbing routine should read either DWORD of the 64-bit memory space (QWORD aligned location) using a ld instruction and write the data back with a st instruction. Software should isolate activity on the memory location to guarantee atomicity.

*Note:* When the scrubbing routine reads the failed location in order to fix the single-bit error, a second error will be reported. Therefore, software should disable single-bit ECC reporting (ECCR[0]) during the scrubbing routine. Also, the scrubbing routine should be aware that partial writes will automatically scrub the QWORD aligned location when it contains a single-bit ECC error

Multi-bit errors cannot be fixed by the H-Matrix.

#### 8.3.4.4.1 ECC Example Using the H-Matrix

Assume the core writes 1234 5678 9ABC DEF0H to the SDRAM memory space. The Core Address Decoder decodes the address and determines the write should be sent to the Core Memory Transaction Queue. The CMTQ latches the transaction with data 1234 5678 9ABC DEF0H on **AD[63:0]**.

During the next CMTQ tenure, this transaction is processes and the DDR SDRAM Control Block receives the data and must calculate the ECC code.

Using the G-Matrix in Figure 91, the DDR SDRAM Control Block creates each check bit by XORing the appropriate bits in the row. Using 1234 5678 9ABC DEF0H, the ECC code generated is D2H. This code is written with the data to the SDRAM devices on **CB**[7:0].

Assume that bit 17 was corrupted in the array. Therefore, the bit has been inverted from 0 to 1.

At some later point in time, the core wishes to read from the same address. The core issues a read transaction which is latched by the CMTQ after the Core Address Decoder decodes the address and determines the read targets the DDR SDRAM address space. Upon the receipt of 1234 5678 9ABE DEF0H on **DO[63:0]**, the DDR SDRAM Control Block calculates the syndrome

with the G-Matrix in Figure 91. The DDR SDRAM Control Block calculates the syndrome of 52H.

*Note:* During a memory write, ECC code is created by XORing the appropriate data bits indicated by the G-Matrix. The syndrome is created during a memory read by XORing the 8-bit value generated by XORing appropriate data bits (**DQ**[63:0]) indicated by the G-Matrix with the check bits (**CB**[7:0]).

Referring to Table 280, when the syndrome is non-zero and matches a value in the H-Matrix, there is a single-bit error that can be fixed. A syndrome of 52H matches a value in the H-Matrix (see Figure 94) which indicates that bit 17 has an error. The DDR SDRAM Control Block inverts bit 17 prior to returning the corrected data on **AD[63:0]**. The MCU returns 1234 5678 9ABC DEF0H on **AD[63:0]**.

Assuming this was the first error, MCU records the address where the error occurred in ECAR0 and error type in ELOG0. When error reporting is enabled in the ECCR, MCU writes a 1 to MCISR[0] which generates an interrupt to the core. A software interrupt handler scrubs the array and fixes the error in bit 17. Unless more errors occur, future reads from this location do not result in an error.



### 8.3.4.5 ECC Disabled

When software disables ECC, the MCU does generate the ECC byte for writes, but does not check the ECC byte for reads. For writes, the MCU will not perform the Read and Modify steps normally performed for a sub 64-bit write as described in Section 8.3.4.2, "ECC Generation for Partial Writes" on page 527. When the Intel XScale<sup>®</sup> core writes 0's to memory, with ECC disabled, the MCU will pad all sub 64-bit writes with zeros, calculate ECC, and store the ECC value along with the write data.

This mode can be used to initialize ECC and memory from the Intel XScale<sup>®</sup> core. By writing the entire 64-bits of data (two stores by the core), all of the data will be zero, and the ECC value will be updated (twice) with a valid code for zero data.

*Note:* For faster initialization of large memory space, the Block Fill function of the Application Accelerator should be used.

#### 8.3.4.6 ECC Testing

Section 8.3.4.4, "Scrubbing" on page 534 explains how the software is responsible for correcting an error in the memory array once it has been detected by the ECC logic. The MCU implements the ECTST register providing the programmer the ability to test error handling software. For write transactions, the ECTST register value is XORed with the generated ECC. This inverts the bits where the mask is set prior to writing the ECC to memory. When the MCU reads the address later, the ECC mismatches and the error condition occurs (see Section 8.5, "Interrupts/Error Conditions" on page 542).

Int

# 8.3.5 Overlapping Memory Regions

The MCU supports two independent memory regions:

- Memory Mapped Register (MMR) Space
- DDR SDRAM Memory Space

The MMR memory space is fixed at FFFF E000H to FFFF FFFFH. Software programs the DDR SDRAM memory space by providing a base address in SDBR, each of the two bank boundaries in SBR0 and SBR1, and the size of the 32-bit region in S32SR when desired.

While it is not recommended, the two ranges could overlap. In the case of a memory space overlap, refer to Table 281 for the priority rules.

#### Table 281. Overlapping Address Priorities

Priority	Address Region						
Highest	Memory Mapped Register Address Space						
Lowest	DDR SDRAM Address Space						

## 8.3.6 DDR SDRAM Clocking

The MCU provides 6 clocks, three positive (**M\_CK[2:0]**) and three negative (**M\_CK[2:0]#**), to the DDR SDRAM memory subsystem at the selected DDR SDRAM command rate. The 72-bit 2-bank unbuffered *JEDEC Standard Double Data Rate (DDR) SDRAM Specification JESD79*, June 2000 requires 6 clocks to distribute the loading across eighteen x8 DDR SDRAM components.

These metrics do not measure performance impact of preemption capability.

# 8.4 **Power Failure Mode**

The 80333 is an I/O processor used in server applications including networking and storage. Specifically, the storage applications supported utilize the 80333 as the IOP for a SCSI RAID disk subsystem and the local memory is used for disk caching. The local memory is used for the temporary storage of disk writes which greatly improves disk performance.

While the host assumes all written data is stored on the non-volatile disk subsystem, the IOP must ensure that eventually all the data in the disk cache is actually stored onto disk.

The power supply could fail to provide power to the I/O subsystem in the case of a power outage or a failed power supply. It is imperative that the cached data within the IOP's local memory is not lost. When power fails, the local memory subsystem must remain powered with a battery backup and some agent must continue to refresh at the appropriate interval specified by the memory component datasheet.

This section defines a mechanism with which the 80333 memory controller ensures that the data within local memory is not lost during a power failure.

## 8.4.1 Theory of Operation

DDR SDRAM technology provides a simple way of enabling data preservation through the **self-refresh** command. This command is issued by the memory controller and the DDR SDRAM will refresh itself autonomously with internal logic and timers. The **self-refresh** command is defined in Table 278.

The DDR SDRAM device will remain in self-refresh mode as long as:

- The device continues to be powered.
- CKE is held low until the memory controller is ready to control the DDR SDRAM once again.

Power to the DDR SDRAM subsystem is ensured with an adequate battery backup and a reliable method for switching between system power and battery power. The memory controller is responsible for deasserting **CKE[1:0]** when issuing the **self-refresh** command but while power gradually drops, **CKE[1:0] must** remain deasserted regardless of the state of  $V_{cc}$  powering the 80333.



## 8.4.2 **Power Failure Sequence**

Figure 95 illustrates the sequence of events during a power failure as defined by *PCI Local Bus Specification*, Revision 2.2.



#### Figure 95. Power Failure Sequence

#### 8.4.2.1 Power Failure Impact on the System

Upon initial power-up a power supply provides the appropriate voltage to the system. The voltage level will increase at a rate that is dependent on the type of power supply used and the components in the system. These variables are not certain, so the power supply often provides a signal called **PWRGD** which indicates the time when the voltage has reached a reliable level. The power supply deasserts **PWRGD** when the voltage level drops below a certain minimum threshold.

*PCI Local Bus Specification*, Revision 2.2 indicates that once **PWRGD** is deasserted, the PCI reset pin (**P\_RST**#) is asserted in order to float the output buffers. In the specification  $T_{fail}$  is defined as the time when **P\_RST**# is asserted in response to the power rail going out of specification.  $T_{fail}$  is the minimum of:

- 500 ns from either power rail going out of specification (exceeding specified tolerances by more than 500mV)
- 100 ns from the 5V rail falling below the 3.3V rail by more than 300mV

#### 8.4.2.2 System Assumptions

This proposal makes specific assumptions about the system's behavior during a power failure. When the below assumptions are not guaranteed, it is the vendor's responsibility to ensure them.

1. **P\_RST#** is asserted to the 80333 when there is at least 2 us of reliable power remaining. This is required so that the memory controller can execute it's power-failure state machine in response to the assertion of **P\_RST#**.

# 8.4.3 Memory Controller Response to PWRGD

When **PWRDELAY** is asserted indicating that an initial power up sequence has completed, the memory controller assumes a power failure condition whenever**PWRGD** is asserted.**PWRGD** assertion following an initial power up sequence results in the following sequence of events:

- 1. The Clock/Reset Unit (CRU) will request that the MCU run the power fail sequence depicted in Figure 97 on page 540. The MCU will take the following steps to execute the power fail sequence:
  - a. Gracefully terminate the current transaction.
  - b. Deactivate all DDR SDRAM leaves with the precharge-all command.
  - c. After Trp, the MCU will issue a **self-refresh** command to the DDR SDRAM devices one bank at a time and continue to deassert **CKE**[1:0].
- 2. The MCU will notify the CRU that the power fail sequence has completed on the memory bus.
- 3. The CRU will then assert **I\_RST#** to reinitialize all internal bus agents including the MCU.

Note that **I\_RST#** is asserted in response to the assertion of **PWRGD**. If **PWRGD** indicates a true power failure (i.e., **PWRDELAY** is asserted), then battery-backup power is supplied to the DDR SDRAM array.

Refer to Figure 96 on page 539 for a high-level state machine representation illustrating the memory controller's behavior during a power failure condition.





*Note:* Following the request from the CRU to run the power fail sequence, any data that is in the MCUs 1024 byte posted write buffer will be discarded.



Figure 97 on page 540 illustrates the DDR SDRAM waveforms after the assertion of **PWRGD** during a true power failure.



Figure 97. Power Failure Sequence

**CKE[1:0]** must be held low throughout the power-down period. The memory controller drives it low initially with the **self-refresh** command, but an external pull-down is required to continually drive it low when the 80333 loses power. External logic ensures that **CKE[1:0]** is held low after the memory controller initially deasserts it. Likewise, the external logic must stop driving **CKE[1:0]** low once **PWRGD** is deasserted by the system. Figure 97 shows one example of the external logic required for power failure mode.

Due to the high loading on **CKE** and the requirement of PC200 DDR operation, the memory controller must drive two copies to the DDR SDRAM DIMM. The board layout will distribute the two **CKE[1:0]** signals between the two DDR SDRAM banks equally.

As long as the DDR SDRAM memory subsystem is powered with a battery source and **CKE[1:0]** is held low, the DDR SDRAM preserves its memory image.

When power is restored, the system asserts **PWRGD** to the 80333. While the 80333 is reset, **CKE[1:0]** is held low by memory controller. After **PWRGD** is deasserted (and subsequently, **I\_RST#** is deasserted), the 80333 must be re-initialized to reset the DDR SDRAM memory subsystem operating parameters. The first step of DDR SDRAM initialization sequence re-asserts **CKE[1:0]** to ones and the memory controller resumes refreshing. DDR SDRAM initialization sequence does not affect memory contents. For more details about the DDR SDRAM initialization sequence, refer to Section 8.3.3.8, "DDR SDRAM Initialization" on page 513.

*Note:* The power failure mechanism in the memory controller is not responsible for maintaining the 80333 state. The purpose of this mechanism is to maintain the memory so that any data cached in the local memory can be flushed once power is restored. Any data queued within the 80333 components (MCU, ATU, DMAs, etc.) will be lost.


### 8.4.3.1 External Logic Required for Power Failure

#### 8.4.3.1.1 Assertion of PWRGD During Power Failure

The 80333 PWRGD input signal must be deasserted when system power goes below 3 V.

# 8.4.3.1.2 Distinguishing Between a Power Up and a Power Failure PWRGD Assertion

The 80333 provides a dedicated input pin, **PWRDELAY** that will be used to distinguish between and initial power up and a power failure assertion of **PWRGD**. This signal should be driven by external circuitry that will assert **PWRDELAY** following the initial assertion of **PWRGD**. **PWRDELAY** is not deasserted until power has truly failed.

This circuitry could consist of a capacitor that is charged up through a FET enabled via assertion of **PWRGD**. The only discharge path available to the capacitor would be through a Zener diode connected to VCC, thus **PWRDELAY** would be deasserted only when power has truly failed.

The Clock Reset Unit will **not** request the MCU to run the power failure sequence when **PWRGD** is deasserted while **PWRDELAY** is deasserted.

*Note:* PWRDELAY only needs a pull-up resistor. For more details, see *Intel*<sup>®</sup> 80333 I/O Processor Specification Update, Specification Clarification #35.

#### 8.4.3.2 P\_RST# Usage Versus I\_RST#

The memory controller logic is initialized with the assertion of the Internal Bus Reset signal **I\_RST#**. It is important to note that the external logic required for power failure mode uses **PWRGD** since **I\_RST#** is not available. It is possible to assert **I\_RST#** with either **PWRGD** or software using the PCI Configuration and Status Register (see Table 168, "PCI Configuration and Status Register - PCSR" on page 283 located in the ATUs MMR space. When software resets the internal bus, the external logic will not function (but it is not required to anyway).



# 8.5 Interrupts/Error Conditions

The MCU has two conditions which require intervention from the Intel XScale<sup>®</sup> core. When a single-bit error is detected during a read cycle, the MCU can correct the data returned but software needs to fix the error in the memory array. When a multi-bit error is detected, the core decides how to handle the condition. For all ECC errors, the MCU records the requester of the transaction resulting in the error in ELOGx[23:16] and interrupts the core.

When the MCU detects an ECC error during a read or write cycle<sup>1</sup>, MCISR[0] or MCISR[1] is set to 1. Whenever the MCU toggles one of the MCISR bits from 0 to 1, an interrupt is generated to the core.

Table 282 shows how the MCU responds to error conditions.

#### Table 282.MCU Error Response

Error Type	MCU Action <sup>a</sup>
Single-Bit during a read or write	Fix Error (when ECC error correction enabled in the SDCR)
Multi-bit during a read	Target Abort the Internal Bus transaction or Terminate the Core transaction, notify the BIU of multi-bit error
Multi-bit during a write	New ECC is generated with bad data and written to DDR SDRAM array. Data location is no longer valid.

a. The ECC Enable bit in the SDCR needs to be set in order for these actions to occur.

*Note:* When ECC reporting is enabled with ECCR[1] or ECCR[0] and an ECC error occurs, MCISR[1] or MCISR[0] is set and ELOGx/ECARx logs the error in addition to Table 282 actions.

<sup>1.</sup> Any error condition during a write cycle actually occurs while performing the read portion of a read-modify-write on a partial write. See Section 8.3.4.1, "ECC Generation" on page 526 for details.

# intel

## 8.5.1 Single-Bit Error Detection

When enabled, the MCU interrupts the core when the ECC logic detects a single-bit error by setting the appropriate bit in the MCISR register. The core knows the interrupt was caused by a single-bit error by polling the ELOG0 or ELOG1 register. The DDR SDRAM Control Block ensures that correct data is returned but the interrupt handler is responsible for scrubbing the error in the array (refer to Section 8.3.4.4, "Scrubbing" on page 534).

An example flow for a single-bit error with error detection and reporting enabled is:

- A single-bit ECC error is detected on the data bus by the MCU.
- The MCU fixes the error prior to returning the data.
- The MCU clears ELOG0[8] indicating a single-bit error.
- The MCU records the requester of the transaction that resulted in an error in ELOG0[23:16]
- The MCU loads ELOG0[7:0] with the syndrome that indicated the error.
- The MCU loads ECAR0[31:2] with address where the error occurred.
- Since the core needs to scrub the error in the array, the MCU sets MCISR[0] to 1 (assuming it is not already set).
  - Setting any bit in the MCISR causes an interrupt to the core.
- Software polls the interrupt status register. Bit 0 set to 1 indicates that the first error has occurred.
- Software polls ELOG0 and ECAR0 and scrubs the error at the location specified by ECAR0.
- Software writes a 1 to MCISR[0] thereby clearing it.

When software does not perform error scrubbing, the probability of an unrecoverable multi-bit error increases for the memory location containing the single-bit error.

ECARx and ELOGx remain registered until software explicitly clears them.

When a second error occurs before software clears the first by resetting MCISR[0] or MCISR[1], the error is recorded in the remaining ELOGx/ECARx register. When none are available, the error is not logged but the MCU carries out the action described in Table 282.



## 8.5.2 Multi-bit Error Detection

When a multi-bit error occurs during a read or write transaction and error reporting is enabled, the MCU sets MCISR[0] or MCISR[1] which asserts an interrupt to the core. Upon receiving an interrupt, the core knows the interrupt was caused by a multi-bit error by polling the ELOGx registers.

When MCU detects a multi-bit error during a read cycle and ECC calculation is enabled in the ECCR, the MCU target aborts the transaction in the IBMTQ, indicating to the internal bus masters that an unrecoverable error has been detected. For core transactions issued by the CMTQ, when a multi-bit error is detected during a read cycle, the MCU signals a multi-bit error to the BIU. The MCU records the error type in ELOGx and the address in ECARx.

When MCU detects a multi-bit error during a write<sup>1</sup> cycle and error reporting is enabled in the ECCR, the MCU records the first multi-bit error by programming ELOGx and ECARx. The MCU generates new ECC with the data before sending it on **DQ[63:0]** so the contents of memory after the read-modify-write cycle will be corrupted with correct ECC.

When a second error occurs before software clears the first by resetting MCISR[0] or MCISR[1], the error is recorded in the remaining ELOGx/ECARx register. When none are available, the error is not logged but the MCU carries out the action described in Table 282.

It is interrupt handler responsibility to decide how to handle this error condition and clear the MCISR.

# 8.6 Reset Conditions

Once **I\_RST#** is deasserted and 200 us have passed, software must issue the initialization sequence defined in Section 8.3.3.8, "DDR SDRAM Initialization" on page 513. After initialization, the DDR SDRAM devices are ready to be written to or read from. Reads issued prior to a write to the same address results in an ECC error and are not recommended when ECC is enabled.

While **I\_RST#** is asserted, the MCU initializes its MMR registers to the states defined in Section 8.7, "Register Definitions" on page 545.

Note: The operation of any memory transactions are not guaranteed when **PWRGD** is asserted.

<sup>1.</sup> Any error condition during a write cycle actually occurs while performing the read portion of a read-modify-write on a partial write. See Section 8.3.4.1, "ECC Generation" on page 526 for details.

# intel®

# 8.7 Register Definitions

A series of configuration registers control the MCU. Software can determine the status of the MCU by reading the status registers. Table 283 lists all of the MCU registers which are detailed further in proceeding sections.

*Note:* Constant polling of MCU MMRs can result in inducing long latencies in peripheral unit DDR SDRAM transactions, and therefore may negatively impact performance. Polling of MCU MMRs should be avoided.

#### Table 283. Memory Controller Register

Section, Register Name - Acronym (Page)
Section 8.7.1, "SDRAM Initialization Register - SDIR" on page 546
Section 8.7.2, "SDRAM Control Register 0 - SDCR0" on page 547
Section 8.7.3, "SDRAM Control Register 1 - SDCR1" on page 549
Section 8.7.4, "SDRAM Base Register - SDBR" on page 551
Section 8.7.5, "SDRAM Boundary Register 0 - SBR0" on page 552
Section 8.7.6, "SDRAM Boundary Register 1 - SBR1" on page 553
Section 8.7.7, "SDRAM 32-bit Region Size Register - S32SR" on page 554
Section 8.7.8, "ECC Control Register - ECCR" on page 555
Section 8.7.9, "ECC Log Registers - ELOG0, ELOG1" on page 556
Section 8.7.10, "ECC Address Registers - ECAR0, ECAR1" on page 557
Section 8.7.11, "ECC Test Register - ECTST" on page 558
Section 8.7.12, "Memory Controller Interrupt Status Register - MCISR" on page 559
Section 8.7.13, "MCU Port Transaction Count Register - MPTCR" on page 560
Section 8.7.14, "MCU Preemption Control Register - MPCR" on page 561
Section 8.7.15, "Frequency Register - RFR" on page 562
Section 8.7.16, "DCAL Control and Status Register - DCALCSR" on page 563
Section 8.7.17, "DCAL Address Register - DCALADDR" on page 565
Section 8.7.18, "DCAL Data Registers 17:0 - DCALDATA[17:0]" on page 566
Section 8.7.19, "Receive Enable Delay Register - RCVDLY" on page 571
Section 8.7.20, "Slave Low Mix 0 - SLVLMIX0" on page 572
Section 8.7.21, "Slave Low Mix 1 - SLVLMIX1" on page 573
Section 8.7.22, "Slave High Mix 0 - SLVHMIX0" on page 574
Section 8.7.23, "Slave High Mix 1 - SLVHMIX1" on page 575
Section 8.7.24, "Slave Length - SLVLEN" on page 576
Section 8.7.25, "Master Mix - MASTMIX" on page 577
Section 8.7.26, "Master Length - MASTLEN" on page 578
Section 8.7.27, "DDR Drive Strength Status Register - DDRDSSR" on page 579
Section 8.7.28, "DDR Drive Strength Control Register - DDRDSCR" on page 580
Section 8.7.29, "DDR Miscellaneous Pad Control Register - DDRMPCR" on page 581



## 8.7.1 SDRAM Initialization Register - SDIR

The DDR SDRAM Initialization Register (SDIR) is responsible for programming the operation of the DDR SDRAM device state machines. The SDIR provides a method for software to execute the DDR SDRAM initialization sequence (see Section 8.3.3.8, "DDR SDRAM Initialization" on page 513).





# intel

## 8.7.2 SDRAM Control Register 0 - SDCR0

The SDRAM Control Registers (SDCR[1:0]) are responsible for programming the operation of the DDR SDRAM state machines. The SDCR0 specifies the DIMM type, data bus width, and some SDRAM timing parameters required by the DDR SDRAM state machine as defined in Section 8.3.3.8, "DDR SDRAM Initialization" on page 513 and Section 8.3.3.9, "DDR SDRAM Mode Programming" on page 516. The remaining SDRAM timing parameters required by the DDR SDRAM state machine are set in SDCR1.

IOP       31       28       24       20       16       12       8       4       0         Attributes						
Intel XScale <sup>®</sup> Core Local Bus Address       Attribute Legend:       RW = Read/Wri         FFFF E504H       RV = Reserved       RC = Read Cleat         PR = Preserved       RO = Read Onl       RS = Read/Set         NA = Not Access       NA = Not Access						
Bit	Default	De	escription			
		RAS: Active to Precharge duration in MCLK per	iods			
31:28	ОH	Equation 8. RAS = tRAS - 1				
	I	where tRAS is from SPD				
27	0 <sub>2</sub>	Reserved				
		RP: Precharge Command Period in MCLK perio	ods			
26:24	26:24 000 <sub>2</sub> Equation 9. RP = tRP -1					
	I	where tRP is from SPD				
23	0 <sub>2</sub>	Reserved	Reserved			
		RCD: Active to Read, Active to Write Period in N	VICLK periods			
22:20	22:20 000 <sub>2</sub> Equation 10. RCD = tRCD - 1					
	I	where tRCD is from SPD				
19:18	002	Reserved				
17.16	010	EDP: Data Path Latency in MCLK periods				
17.10	012	This field should be programmed to 10 <sub>2.</sub>				
15:14	002	Reserved				
13:12	<ul> <li>WDL: Write Data Latency in MCLK periods used by the Memory Controller state machine. See Equation 11.</li> <li>00 = 0 MCLK periods (tCAS = 2.5) - (DDR-I Type only)</li> <li>01 = 1 MCLK period (tCAS = 3) - (DDR-II Type only) - not supported</li> <li>10 = 2 MCLK periods (tCAS = 4) - (DDR-II Type only)</li> <li>11 = reserved</li> <li>where tCAS is from SPD</li> </ul>					
11:10	002	Reserved				

#### Table 285. DDR SDRAM Control Register 0 - SDCR0 (Sheet 1 of 2)

able 28	35. DDR 9	SDRAM Control Register 0 - SDCR0 (Sheet 2 of 2)				
Att Att	IOP       31       28       24       20       16       12       8       4       0         Attributes					
Bit	Default	Description				
09:08	00 <sub>2</sub>	<ul> <li>CAS: CAS Latency: Indicates the CAS Latency used by the Memory Controller state machine. Encoded value based on tCAS from SPD</li> <li>00 = reserved</li> <li>01 = 2.5 MCLK periods (DDR-I Type only)</li> <li>10 = 3 MCLK periods (DDR-II Type only) - not supported</li> <li>11 = 4 MCLK periods (DDR-II Type only)</li> </ul>				
07:06	002	Reserved.				
05:04	00 <sub>2</sub>	00       ODT Termination Value: Determines the termination value of the On Die Termination for both Banks (controlled by <b>ODT[1:0]</b> ). Applies to DDR-II SDRAM memory type only.         002       00 Disabled         01 75 ohm         10 150 ohm         11 reserved				
03	0 <sub>2</sub>	Reserved				
02	Varies with external state of <b>MEM_TYPE</b> at PCI bus reset	<b>DDR Type:</b> Identifies the selected DDR generation of SDRAM based on the <b>MEM_TYPE</b> reset strap. 0 = DDR-II (supported speed of 400 MHz) - <b>MEM_TYPE</b> Deasserted. 1 = DDR (supported speed of 333 MHz) - <b>MEM_TYPE</b> Asserted.				
01	02 Data Bus Width: Indicates the width of the data bus. See Section 8.3.3.4, "32-bit Data Bus Width" on page 510. 0 = 64 bits 1 = 32 bits					
00	DIMM Type: Selects unbuffered or registered DIMM operating modes for the MCU. 0 = Unbuffered* 1 = Registered NOTE: Unbuffered DDB SDBAM memory subsystems will use the Unbuffered mode					

### Ta

intel



## 8.7.3 SDRAM Control Register 1 - SDCR1

The SDRAM Control Registers (SDCR[1:0]) are responsible for programming the operation of the DDR SDRAM state machines as defined in Section 8.3.3.8, "DDR SDRAM Initialization" on page 513 and Section 8.3.3.9, "DDR SDRAM Mode Programming" on page 516. The SDCR1 specifies the remaining SDRAM timing parameters required by the DDR SDRAM state machine not specified in SDCR0.





intel
-------

IOP Attributes       31       28       24       20       16       12       8       4       0         Matributes						
Bit	Default	Description	Description			
19:17	0002	Reserved				
16:12	0 0000 <sub>2</sub>	tRFC: Refresh-to-Active and Refresh-to-Refresh period in MCLK periods.         Equation 13.       RFC = tRFC -1         where tRFC is from SPD				
11:09	000 <sub>2</sub>	<ul> <li>tWR: Write Recovery time in MCLK periods.</li> <li>000 = 0 - for DDR333</li> <li>010 = 3 - for DDR-II 400 (encoding per JEDEC spec) all other values reserved</li> </ul>				
08:04	4     0 00002     RC: Active-to-Active and Active-to-Refresh period in MCLK periods.       Equation 14.     RC = tRC -1 where tRC is from SPD					
03:00	03:00       0H       WTRD: Write-to-Read turnaround period in MCLK periods.         Equation 15.       WTRD = tCAS + tWTR + tREG         where tREG = 1 for registered DIMM and 0 for unbuffered DIMM, tCAS and tWTR are from SPD.					

### Table 286. DDR SDRAM Control Register 1 - SDCR1 (Sheet 2 of 2)



## 8.7.4 SDRAM Base Register - SDBR

This register indicates the beginning of SDRAM space. See Section 8.3.3.2, "DDR SDRAM Addressing" on page 504 for usage details. There can be two contiguous physical banks defined by SBR0 and SBR1 in the DDR SDRAM subsystem starting at this address.

- Note: DDR SDRAM memory space must never cross a 2 Gbyte boundary.
- *Note:* This register should be read back after being written, before the Intel XScale<sup>®</sup> core performs transactions which address the DDR SDRAM.



#### Table 287. SDRAM Base Register - SDBR



## 8.7.5 SDRAM Boundary Register 0 - SBR0

This register indicates the upper boundary of SDRAM bank 0 and its memory technology. When bank 0 is unpopulated, SBR0[6:0] is programmed either with all zeros or a number equal to the value in SDBR[30:25]. See Section 8.3.3.2, "DDR SDRAM Addressing" on page 504 for more details and programming examples.

Bank 0 may have multiple regions based on the value in "SDRAM 32-bit Region Size Register - S32SR" on page 554. See Section 8.3.3.2, "DDR SDRAM Addressing" on page 504 for more details

- Note: DDR SDRAM memory space must never cross a 2 Gbyte boundary.
- *Note:* This register should be read back after being written, before the Intel XScale<sup>®</sup> core performs transactions which address the DDR SDRAM.



Table 288. SDRAM Boundary Register 0 - SBR0



## 8.7.6 SDRAM Boundary Register 1 - SBR1

This register indicates the upper boundary of SDRAM bank 1 and its memory technology. When bank 1 is unpopulated, SBR1[6:0] is programmed either with all zeroes or a value equal to SBR0[6:0]. When bank 1 is populated, SBR1[6:0] must be programmed greater than or equal to SBR0[6:0]. See Section 8.3.3.2, "DDR SDRAM Addressing" on page 504 for more details and programming examples.

- Note: DDR SDRAM Memory Space must never cross a 2 Gbyte boundary.
- *Note:* This register should be read back after being written, before the Intel XScale<sup>®</sup> core performs transactions which address the DDR SDRAM.



Table 289. SDRAM Boundary Register - SBR1



## 8.7.7 SDRAM 32-bit Region Size Register - S32SR

.Defines the size of the 32-bit region located at the base of SDRAM Bank 0.This register must be programmed with a size that is less than or equal to one half of the size of DDR SDRAM Bank 0. Sizes are limited to binary sizes with a minimum size of 1MB. Also, the DDR SDRAM type must be 64-bit, as defined by "SDRAM Control Register 0 - SDCR0" on page 547. See also "SDRAM Boundary Register 0 - SBR0" on page 552 and Section 8.3.3.2, "DDR SDRAM Addressing" on page 504.

This register should be read back after being written, before the Intel XScale<sup>®</sup> core performs transactions which address the DDR SDRAM.



#### Table 290. DDR SDRAM 32-bit Region Size Register - S32SR

# intel

## 8.7.8 ECC Control Register - ECCR

This register programs the MCU error correction and detection capabilities. The configuration depends on the application's needs but a typical configuration is:

- ECC Mode Enabled
- Enable multi-bit error reporting
- Disable single-bit error reporting
- Enable single-bit error correcting

For more details, see Section 8.3.4, "Error Correction and Detection" on page 525 and Section 8.5, "Interrupts/Error Conditions" on page 542.

#### Table 291. ECC Control Register - ECCR





## 8.7.9 ECC Log Registers - ELOG0, ELOG1

The ECC Log Registers are responsible for logging the error types detected on the local memory bus. Two errors can be detected and logged. The error type is logged (single-bit or multi-bit) along with the syndrome that indicated the error. For a single-bit error, software can read this syndrome and determine which bit had the error in order to perform scrubbing. For a multi-bit error, the syndrome will not match an entry in the H-Matrix and thus, is uncorrectable (see Table 280, "Syndrome Decoding" on page 530).

The error recorded in ELOG0 corresponds to the address in ECAR0. ELOG1 corresponds to ECAR1.

The ELOGx registers comprise read-only bits and only have meaning when MCISR[0] or MCISR[1] is non-zero. For more details on error handling, see Section 8.3.4, "Error Correction and Detection" on page 525.



#### Table 292. ECC Log Registers - ELOG0, ELOG1



## 8.7.10 ECC Address Registers - ECAR0, ECAR1

These registers are responsible for logging the addresses where the errors were detected on the local memory bus. Two errors can be detected and logged. The software knows which DDR SDRAM address had the error by reading these registers and decoding the syndrome in the log registers. For error details, see Section 8.3.4, "Error Correction and Detection" on page 525).



#### Table 293. ECC Address Registers - ECAR0, ECAR1



## 8.7.11 ECC Test Register - ECTST

This register allows testing between the ECC logic and the memory subsystem (Section 8.3.4.6, "ECC Testing" on page 535). To test error handling software, the programmer writes this register with a non-zero masking function. Any subsequent writes to memory stores a masked version of the computed ECC. Therefore, any subsequent reads to these locations result in an ECC error.



#### Table 294. ECC Test Register - ECTST

## 8.7.12 Memory Controller Interrupt Status Register - MCISR

Setting the MCISR asserts an interrupt to the core. Upon an interrupt, the Intel XScale<sup>®</sup> core polls the interrupt status register for each unit. The interrupt status register tells the core the reason for the interrupt. The MCU has three interrupt conditions: first ECC error (MCISR[0]), second ECC error (MCISR[1]), and more than two ECC errors (MCISR[2]).

When the MCU detects an ECC error and both MCISR[0] and MCISR[1] are cleared, the error is logged in ELOG0 and MCISR[0] is set to 1. When one of the MCISR bits are not clear and the MCU detects an error, the error is logged in the unused ELOGx register and the appropriate MCISR bit is set to 1. When both MCISR[0] and MCISR[1] are not clear, any additional ECC errors are not logged and MCISR[2] is set.

Bits 2:0 are read/clear bits which means that to clear them, software must write a one to these bits.



#### Table 295. Memory Controller Interrupt Status Register - MCISR



## 8.7.13 MCU Port Transaction Count Register - MPTCR

Sets the number of transactions a given port can have processed during a single tenure. The 4-bit fields for each port allow up to 16 transactions to be processed by a port before the MCU arbiter selects a different port for DDR SDRAM transactions. This register along with the "MCU Preemption Control Register - MPCR" on page 561 used to optimize the memory controller operation.



#### Table 296. MCU Port Transaction Count Register - MPTCR



## 8.7.14 MCU Preemption Control Register - MPCR

Enables Preemption of IB port transaction when a core processor transaction is pending.



#### Table 297. MCU Preemption Control Register - MPCR



## 8.7.15 Frequency Register - RFR

The Refresh Frequency Register is programmed for refreshing the DDR SDRAM subsystem at the specified interval. Writing to the RFR programs the refresh counter with the Refresh Interval. Reading from the RFR results in the value currently within the refresh counter. Refer to Section 279, "Typical Refresh Frequency Register Values" on page 524 for recommended programmed values.



#### Table 298. Refresh Frequency Register - RFR



## 8.7.16 DCAL Control and Status Register - DCALCSR

This 32 bit register controls and shows status for the DCAL operation.Table 299.DCAL Control and Status Register - DCALCSR (Sheet 1 of 2)

Ą	Attributes Attributes Intel XScale <sup>(1)</sup>	28       24       20       16       12       8       4       0         w/rw/rv/rv/rv/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/		
Bit	Default	Description		
31	0	When written: 0 = no action 1 = Start operation defined in this register When read: 0 = Operation completed 1 = Operation in progress		
30:28	D:28 000 Pass Fail Indicators 000 - Pass x01 - SDRAM Access Denied x10 - Unpopulated Row Select x11 - Unsupported Opcode 1xx - Operation Completed with a Failure			
27:25	000	Reserved		
24	1	SDRAM I/F Select 0 = Not Selected - will result in Fail Indicator 001 in bits 30:28 1 = Selected		
23	0	Operation Mode: applicable to the Receive Enable Calibration and DQS Calibration operations. 0 = One Pass - uses value in operation modifier field 1 = All Passes - cycles through all possible encodings		
22:21	00	Reserved		
20	20 0 Row Select (Chip Select) 0 = CS0# selected 1 = CS1# selected			
19	0	Reserved 1 =		



Ą	Attributes Attributes Intel XScale	28       24       20       16       12       8       4       0         rw/rw/rv/rv/rw/rv/rv/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/
Bit	Default	Description
18:16	000	Fixed Data Pattern Selection: Only applicable to DQS Calibration operation. • 000 - F -> 0 -> F -> 0 • 001 - 0 -> F -> 0 -> F • 010 - A -> 5 -> A -> 5 • 011 - 5 -> A -> 5 -> A • 100 - C -> 3 -> C -> 3 • 101 - 3 -> C -> 3 -> C • 110 - 9 -> 6 -> 9 -> 6 • 111 - 6 -> 9 -> 6 -> 9
15	0	Reserved
14:4	000H	<ul> <li>Opcode Modifiers</li> <li>EMRS OCD Calibration: <ul> <li>14 - Reserved</li> <li>13:4 - Number of clock cycles to wait before collecting OCD Drive(0)/Drive(1) Samples.</li> </ul> </li> <li>Receive Enable <ul> <li>14:10 - Reserved</li> <li>9:4 - Receive Enable Delay - contain a 6-bit receive enable delay. Applicable for Single Pass operation only</li> </ul> </li> <li>DQS Calibration <ul> <li>14:12 - DLL Slave Length: used to set the coarse DLL delay adjustment, used in both single pass and all-pass modes</li> <li>11:8 - DLL Slave Mix: used to set the fine DLL delay adjustment. Applicable for Single Pass operation only.</li> <li>7:4 - Reserved</li> </ul> </li> <li>Physical Address used for operation <ul> <li>Logical Address used for operation</li> </ul> </li> </ul>
3	0	Reserved
2:0	он	<ul> <li>Opcode</li> <li>011 - EMRS OCD Calibration: requires the BA field (001b) and OCD Drive[0] or Drive[1] command is selected (MA[9:7] = 010b or 001b) in the DCALADDR register.</li> <li>100 - Receive Enable Calibration</li> <li>101 - DQS Calibration</li> <li>All other values are reserved</li> </ul>

### Table 299. DCAL Control and Status Register - DCALCSR (Sheet 2 of 2)



## 8.7.17 DCAL Address Register - DCALADDR

This 32 bit register supplies address for the DCAL operation selected in Section 8.7.16, "DCAL Control and Status Register - DCALCSR" on page 563. The opcodes use this register to specify the row, column and bank address that should be driven to the DIMMs with the command.



A	IOP ttrributes	28       24       20       16       12       8       4         rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/r	0 rw hna /rite lear nly essible
Bit	Default	Description	
31:30	00	Reserved	
29:16	0000H	Row Address: 14-bit address driven by the MCU on MA[13:0] pins for the activate (row) of for the Receive Enable Calibration and DQS Calibration operations and for the EMRS OC Command.	command CD
15:14	00	Reserved	
13:4	13:4Column Address: Used to construct the 14-bit address to be driven on the MA[13:0] pins for the read/write (column) command for the Receive Enable Calibration and DQS Calibration operations and for the EMRS OCD Calibration Command. Note, the MA[13, 1:0] pins are always driven to 0.		
3:2	00	Reserved	
1:0	00	Bank Address: 2-bit Bank address to be driven on the BA[1:0] pins for all operations.	



## 8.7.18 DCAL Data Registers 17:0 - DCALDATA[17:0]

These ten 32-bit registers are used to support the different opcodes. The definition of these bits as with the address register changes depending on the operation.



#### Table 301.DCAL Data Registers 17-0 - DCALDATA[17:0]



## 8.7.18.1 Opcode: EMRS OCD Adjust/Drive Commands

The DCAL Data Registers consist of Adjust and Drive fields for each DQS line. The fields are identical for each DQS line and are defined in Table 302. The location of the fields within the DCAL Data Registers is provided in Table 303 following the field definitions.

**Adjust:** A 4-bit value stored as 0000  $D_{T3}D_{T2}D_{T1}D_{T0}$  within the byte (MSB to LSB), which is the 4-bit data burst for the OCD adjust command as documented in the latest version of the JEDEC DDR-II Component specification.

*Note:* The bits are reversed, bit 0 in left hand column.

4-bit burst codes inputs to all DQs		Operation			
DT0	DT1	DT2	DT3	Pull-up driver strength	Pull-down driver strength
0	0	0	0	NOP	NOP
0	0	0	1	Increase by 1 step	NOP
0	0	1	0	Decrease by 1 step	NOP
0	1	0	0	NOP	Increase by 1 step
1	0	0	0	NOP	Decrease by 1 step
0	1	0	1	Increase by 1 step	Increase by 1 step
0	1	1	0	Decrease by 1 step	Increase by 1 step
1	0	0	1	Increase by 1 step	Decrease by 1 step
1	0	1	0	Decrease by 1 step	Decrease by 1 step
Other combinations		Res	served		

#### Table 302.OCD Adjust Field Encoding

**Drive:** An 8-bit value stored as  $D_{1S1}D_{1S2}D_{1S3}D_{0S4}D_{0S1}D_{0S2}D_{0S3}D_{0S4}$  within the byte. The  $D_{1SX}$  values represent the 4 samples collected during the EMRS OCD Drive[1] operation and the  $D_{0SX}$  values represent the 4 samples collected during the EMRS OCD Drive[0] operation. After running the Drive[1] and Drive[0] commands, this 8-bit value should be 0xFF to guarantee that the drive strengths are sufficient. When Drive[1] samples do not equal 0xF, then the pull-up driver strength needs to be increased by (at least) one step with the OCD adjust command. When the Drive[0] samples do not equal 0xF, then the pull-down driver strength needs to be increased by (at lease) one step with the OCD adjust command.



Table 303 defines the usage of the DCAL Data Registers for the OCD Adjust/Drive fields.

*Note:* The Adjust fields must be duplicated per DQS as indicated. The Drive fields are also duplicated, and either field (odd or even byte) can be read for the corresponding DQS value.

#### Table 303. OCD Definition of DCALDATA[17:0] Registers

DCALDATA Register	Bits[31:24] - Byte3	Bits[23:16] - Byte2	Bits[15:8] - Byte1	Bits[7:0] - Byte0	
17	Deserved		DQS8 - Adjust	DQS8 - Adjust	
16	Reserved		DQS8 - Drive	DQS8 - Drive	
15:12	Reserved				
11	DQS7 - Adjust	DQS7 - Adjust	DQS5 - Adjust	DQS5 - Adjust	
10	DQS3 - Adjust	DQS3 - Adjust	DQS1 - Adjust	DQS1 - Adjust	
9	DQS6 - Adjust	DQS6 - Adjust	DQS4 - Adjust	DQS4 - Adjust	
8	DQS2 - Adjust	DQS2 - Adjust	DQS0 - Adjust	DQS0 - Adjust	
7:4	Reserved				
3	DQS7 - Drive	DQS7 - Drive	DQS5 - Drive	DQS5 - Drive	
2	DQS3 - Drive	DQS3 - Drive	DQS1 - Drive	DQS1 - Drive	
1	DQS6 - Drive	DQS6 - Drive	DQS4 - Drive	DQS4 - Drive	
0	DQS2 - Drive	DQS2 - Drive	DQS0 - Drive	DQS0 - Drive	

#### 8.7.18.2 Opcode: Receive Enable Calibration

The DCAL Data Registers consist of Write pointer and Vector fields for each DQS line. The fields are identical for each DQS line and are defined below. The location of the fields within the DCAL Data Registers is provided following the field definitions.

Write Pointers (wrptr): A 5-bit value stored as 000x xxxx within the byte. The write pointers should be one-hot encoded value. The legal values are 0x01, 0x02, 0x04, 0x08, and 0x10.

**Sampled High/Low Vectors:** A 64-bit value stored in 2 consecutive registers that contain 1-bit per pass of the receive enable operation. Bit 0 represents the results from pass 0, bit 1 represents the results from pass 1, and so forth. By comparing the sampled high vector to the sampled low vector, the rising/falling edges of DQS can be located. In general, the pattern in each vector should repeat every 8 bits, except where the preamble is located. Legal values: bit x of one (or both) of the sampled high and low vectors should be equal to zero (both should never be equal to one). In other words, a bit-wise AND of the 2 vectors should result in all zeros.

#### Table 304. Receive Enable Calibration of DCALDATA[9] Register

Bits	Description
31:24	Expected write Pointer when DQS is sampled high. This is when the 1 cycle wide rcvcal pulse ANDed with the DQS signal produces a two small pulses that causes the write pointers to advance twice. The initial value of the write pointer is 0x01, so this value should be programmed to 0x04.
23:16	Expected write Pointer when DQS is sampled low. This is when the 1 cycle wide rcvcal pulse ANDed with the DQS signal produces a single pulse that causes the write pointers to advance once. The initial value of the write pointer is 0x01, so this value should be programmed to 0x02.
15:8	Reserved
7:0	DRAM I/F ORed wrptr: All of the channel A write pointers are logically ORed to produce this value, which is then compared to the expected sampled high/low write pointers to produce the sampled high/low vectors stored in DCALDATA[8:0]

# intel

Table 305 defines the usage of the DCAL Data Registers for the Receive Enable fields.

#### Table 305. Receive Enable Calibration Definition of DCALDATA[17:0] Registers

DCALDATA Register	Bits[31:24] - Byte3	Bits[23:16] - Byte2	Bits[15:8] - Byte1	Bits[7:0] - Byte0		
17	Expected Write Pointer High Sample	Expected Write Pointer Low Sample	Reserved	DRAM I/F ORed wrptr		
16		DQS8 - wrptr				
15:12	Reserved					
11	reserved	DQS7 - wrptr	reserved	DQS6 - wrptr		
10		DQS5 - wrptr		DQS4 - wrptr		
9		DQS3 - wrptr		DQS2 - wrptr		
8		DQS1 - wrptr		DQS0 - wrptr		
7:4	Reserved					
3	Cumulative "Sampled High" Vector for passes (63:32					
2	Cumulative "Sampled High" Vector for passes (31:0)					
1	Cumulative "Sampled Low" Vector for passes (63:32)					
0	Cumulative "Sampled Low" Vector for passes (31:0)					

## 8.7.18.3 Opcode: DQS Calibration

The DCAL Data Registers consist of 16-bit field for each DQS line in DQS Calibration mode. The fields are identical for each DQS line and are defined in Table 306. The location of the fields within the DCAL Data Registers is provided following the field definitions in Table 307.

*Note:* Each nibble may have a different left and right edge value, but the calibration will be done across the byte.

#### Table 306. DQS Calibration Field Encodings for DCALDATA[17:0] Registers

Bits	Description				
15:14	Right Edge Detection Indicator 11 - right edge found				
13:8	Right edge slavelen[2:1] and slavemix[3:0] values				
7:6	Left Edge Detection Indicator 11 - left edge found				
5:0	Left edge slavelen[2:1] and slavemix[3:0] values				

Table 307 defines the usage of the DCAL Data Registers for the Receive Enable fields.

#### Table 307. DQS Definition of DCALDATA[17:0] Registers

DCALDATA Register	Bits[31:24] - Byte3	Bits[23:16] - Byte2	Bits[15:8] - Byte1	Bits[7:0] - Byte0	
17	reserved		DQS8 - High Nibble		
16			DQS8 - Low Nibble		
15:12	Reserved				
11	DQS7 - High Nibble		DQS5 - High Nibble		
10	DQS3 - High Nibble		DQS1 - High Nibble		
9	DQS6 - High Nibble		DQS4 - High Nibble		
8	DQS2 - High Nibble		DQS0 - High Nibble		
7-4	Reserved				
3	DQS7 - L	ow Nibble	DQS5 - Low Nibble		
2	DQS3 - L	ow Nibble	DQS1 - Low Nibble		
1	DQS6 - L	ow Nibble	DQS4 - Low Nibble		
0	DQS2 - Low Nibble		DQS0 - Low Nibble		



## 8.7.19 Receive ENABLE Delay Register - RCVDLY

This 32 bit register consists of a packed 3-bit field for DQS Receive Enable Calibration.



#### Table 308. Receive Enabled Delay Register - RCVDLY

*Note:* The Receive Enable Delay Register is used for DQS receive enable calibration. In other words, RCVDLY adjusts the memory controllers relationship of DQS to an internal M\_CLK.

The RCVDLY value is highly dependant on the board layout and DIMM characteristics. Also, the memory controller only supports a non integer CAS latency (tCAS = 2.5, SDCR0.9:8) for DDR-I, which means that RCVDLY may need to be adjusted because DQS is no longer synchronized with M\_CLK.

Therefore, when using DDR-I memory the RCVDLY default setting of 5, may need to be changed to 6 or 7 to operate correctly with a specific DIMM based on the board layout. For example, the Intel reference code uses a value to 7 in order to allow for a wider compatibility with various DIMMs.



## 8.7.20 Slave Low Mix 0 - SLVLMIX0

This 32 bit register consists of the first 8 of 9 packed 4-bit fields for dynamic DQS DLL delay for the low nibble.



#### Table 309. Slave Low Mix 0 - SLVLMIX0



## 8.7.21 Slave Low Mix 1 - SLVLMIX1

This 32 bit register consists of the last 4-bit field for dynamic DQS DLL delay for the low nibble.



Table 310. Slave Low Mix 1 - SLVLMIX1



## 8.7.22 Slave High Mix 0 - SLVHMIX0

This 32 bit register consists of the first 8 of 9 packed 4-bit fields for dynamic DQS DLL delay for the high nibble.



#### Table 311. Slave High Mix 0 - SLVHMIX0



## 8.7.23 Slave High Mix 1 - SLVHMIX1

This 32 bit register consists of the last 4-bit field for dynamic DQS DLL delay for the high nibble.







## 8.7.24 Slave Length - SLVLEN

This 32 bit register consists of a packed 3-bit field for Static DQS Slave DLL length.



#### Table 313. Slave Length - SLVLEN


## 8.7.25 Master Mix - MASTMIX

This 32-bit register consists of a 4-bit field for master DQS DLL delay loop mixer control.



#### Table 314. Master Mix - MASTMIX



## 8.7.26 Master Length - MASTLEN

This 32 bit register consists of the nine 2-bit field for master DQS DLL length control.



#### Table 315. Master Length- MASTLEN



## 8.7.27 DDR Drive Strength Status Register - DDRDSSR

This 32-bit register consists of a 4-bit field for drive strength value.







## 8.7.28 DDR Drive Strength Control Register - DDRDSCR

This 32-bit register consists of a 5-bit field to override the DDR pad automatic drive strength control mechanism.



#### Table 317. DDR Drive Strength Control Register - DDRDSCR



## 8.7.29 DDR Miscellaneous Pad Control Register - DDRMPCR

This 32-bit register contains miscellaneous control signals for the DDR pads



	IOP Attributes PCI Attributes Intel XScale <sup>®</sup> FFFF F578H	28 24 20 16 12 8 4 0 v/rv/rv/rv/rv/rv/rv/rv/rv/rv/rv/rv/rv/rv					
Bit	Default	Description					
31:18	0000H	Reserved					
17:16	10	read pointer delay: Determines the read pointer delay for DCAL, which is based on DIMM topology and technology. Command Clocks per Frequency DDR-I 333 DDR-II 400 • 00 0 (0 ns) 0 (0 ns) • 01 1 (6 ns) 1 (5 ns) • 10 2 (12 ns) 2 (10 ns) • 11 3 (18ns) 3 (15 ns)					
15	0	Fast slew rate control: allows extended range on slew rate: Set high for DDR-II, low for DDR-I					
14	0	Half gain control: Select for DQS differential amplifier gain. Set 1 to cut gain in half for differential strobe mode associated with DDR-II.					
13	02	Leg test mode: When set, Places output buffer strength control in a test mode which will allow testing of each leg of a buffer driver independently. This can be done by placing the DDR pads in slew rate override mode (DDRDSCR register). Each value of drive strength override (bits 2:0 in DDRDSCR) will enable an individual driver leg in each DDR pad.					
12	02	VOX Start: Back-up override for the voltage output crossing control loop.					
11	02	Slew Rate Override: Slew rate override adjustment for analog validation					
10	02	DLL Bypass: Bypasses DLL for calibration on DDR bus					
9	02	OCD Load Enable: Calibration buffer load placed on incoming signals to perform calibration. Required for OCD calibration when DDR-II ODT mode is in use					



	Attributes Attributes Attributes	28 24 20 16 12 8 4 0 / rv/rv/rv/rv/rv/rv/rv/rv/rv/rv/rv/rv/rv/r
Bit	Default	Description
8	02	Analog Validation Mode (behavior TBD)
7:4	00002	Reserved.
3:0	0000 <sub>2</sub>	VREF Select: Controls the VREF selection/generation mechanism for DDR pads. VREF calibration is a required BIOS function 0 External VREF from bump (default) 1 Internal VREF - 250mV 2 Internal VREF - 200mV 3 Internal VREF - 100mV 4 Internal VREF - 100mV 5 Internal VREF - 100mV 6 Internal VREF - 50mV 6 Internal VREF - 10mV 7 Internal precision VREF 8 Internal VREF + 10mV 9 Internal VREF + 10mV 10 Internal VREF + 50mV 10 Internal VREF + 50mV 11 Internal VREF + 150mV 12 Internal VREF + 200mV 13 Internal VREF + 250mV 14 Pull-down Calibration Mode 15 Pull-up Calibration Mode

#### Table 318. DDR Miscellaneous Pad Control Register - DDRMPCR

## **Peripheral Bus Interface Unit**

This chapter describes the Peripheral Bus Interface Unit (PBI) of the Intel<sup>®</sup> 80333 I/O processor (80333). It explains the following:

- Peripheral Bus signals, which consist of address/data, control/status.
- Peripheral Bus Read, and write transactions.
- Peripheral Bus configuration and Flash Memory Support.
- Support for Intel XScale<sup>®</sup> core (ARM\* architecture compliant) boot from the PCI Bus.
- Registers.

intel

This chapter also serves as a starting point for the hardware designer when interfacing typical flash components to the 80333 Peripheral Bus.

Figure 98. The Peripheral Bus Interface Unit



## 9.1 **Overview**

The Peripheral Bus Interface Unit (PBI) is a data communication path to the flash memory components and peripherals of a 80333 hardware system. The PBI allows the processor to read and write data to these supported flash components and other peripherals. To perform these tasks at high bandwidth, the bus features a burst transfer capability which allows successive 8- or 16-bit data transfers.

The peripheral bus is controlled by the on-chip bus masters: the Intel XScale<sup>®</sup> core, the ATU, AAU and DMA units.

The address/data path is multiplexed for economy, and the bus width is programmable to 8-, and 16-bit widths. The PBI performs the necessary packing and unpacking of bytes to communicate properly across the 80333 Internal Bus.

The PBI unit includes two chip enables. The PBI chip enables activate the appropriate peripheral device when the address falls within one of the PBI's two programmable address ranges. Both address ranges incorporate functionality that optimizes an interface for Flash Memory devices.



## 9.2 Peripheral Bus Signals

Bus signals consist of two groups: address/data, and control/status.

## 9.2.1 Address/Data Signal Definitions

The address/data signal group consists of 26 lines. 16 of these signals multiplex within the processor to serve a dual purpose. During and address cycle ( $T_A$ ), the processor drives A[22:16] and AD[15:0] with the address of the bus access. At all other times, the AD[15:0] lines are defined to contain data. A[2:0] are demultiplexed address pins providing incrementing byte addresses during burst cycles.

## 9.2.2 Control/Status Signal Definitions

The control/status signals control peripheral device enables and direction. All output control/status signals are three-state.

The PBI pulses **ALE** (address latch enable) active high for one clock during  $T_A$  to latch the multiplexed address on **AD**[15:2] in external address latches.

A peripheral read may be either non-burst or burst. A non-burst read ends after one data transfer to a single location.

When the data bus is configured for 16 bits, demultiplexed address bits A[2:1] are used to burst across up to four short-words. For an 8-bit data bus, demultiplexed address bits A[1:0] are used to burst across up to four bytes.

The Output Enable, **POE#**, is used for burst or non-burst read accesses to a peripheral device and is asserted during the  $T_{A/T_W}/T_D$  states.

The Write Enable, **PWE#**, is used for non-burst write accesses to a peripheral device and is asserted during the  $T_W/T_D$  states.

Note: Burst write accesses to Flash Devices are not supported.



## 9.2.3 Bus Width

Each address range's attributes are programmed in the PBIs boundary registers. The PBI allows an 8-, or 16-bit data bus width for each range. The PBI places 8- and 16-bit data on low-order data signals, simplifying the interface to narrow bus external devices. As shown in Figure 99, 8-bit data is placed on lines **AD**[7:0]; 16-bit data is placed on lines **AD**[15:0].

#### Figure 99. Data Width and Low Order Address Lines



The user needs to wire up the flash memories in a manner consistent with the programmed bus width:

- 8-bit region: A[1:0] provide the demultiplexed byte address for a read burst.
- 16-bit region: A[2:1] provide the demultiplexed short-word address for a read burst.

During initialization, bus width is selected for each of the two address ranges in the Peripheral Base Address Registers (PBBAR0 - PBBAR1). In addition, the PBBAR0-PBBAR5 can be used to configure these ranges as Peripheral Windows and to set a Wait state profile.

The PBI drives determinate values on all address/data signals during  $T_W/T_D$  write operation states. For an 8-bit bus, the PBI continues to drive address on unused data signals **AD**[15:8].



## 9.2.4 Detailed Signal Descriptions

Bus signal descriptions are detailed in Table 320.

#### Table 319.Bus Signal Descriptions

NAME	DESCRIPTION				
A[22:16]	<b>ADDRESS BUS 22:16</b> carries a demultiplexed version of address bits <b>A</b> [ <b>22:16</b> ]. During address ( $T_a$ ), wait state ( $T_w$ ) and data cycles ( $T_d$ ) cycles, <b>A</b> [ <b>22:16</b> ] represents the upper 7 address bits for the current access. <b>A</b> [ <b>22:16</b> ] allows the PBI interface to address up to 8 MBytes per peripheral device.				
AD[15:0]	ADDRESS / DATA BUS carries 16-bit physical addresses and 8-, or 16-bit data to and from memory. During an address ( $T_a$ ) cycle, bits 2-15 contain a physical word address (bits 0-1 indicate SIZE; see below). During a data ( $T_d$ ) cycle, bits 0-7, or 0-15 contain read or write data, depending on the corresponding bus width. During write operations to 8-bit wide memory regions, the PBI drives unused bus pins high or low. SIZE, which comprises bits 0-1 of the AD lines during a $T_a$ cycle specifies the number of data transfers during the bus transaction.				
	AD1 AD0         0       0       1 Transfer         0       1       2 Transfers         1       0       3 Transfers         1       1       4 Transfers				
A[2:0]	<b>ADDRESS BUS 2:0</b> carries a demultiplexed version of bits 2:0 of the <b>AD[15:0]</b> bus. During a bursted read data ( $T_d$ ) cycle, <b>A[2:0]</b> represents the current byte address in the bursted transaction.				
	<b>A[2:1]</b> are used for an 16-bit wide peripheral while A[1:0] are used for an 8-bit wide peripheral.				
ALE	<b>ADDRESS LATCH ENABLE</b> indicates the transfer of a physical address. ALE is asserted during a $T_a$ cycle and deasserted before the beginning of the $T_d$ state.				
POE#	<b>PERIPHERAL OUTPUT ENABLE</b> specifies, during a $T_a$ cycle, whether the operation is a write (1) or read (0). It is latched on-chip and remains valid during $T_d$ cycles.				
PCE[1:0]#	Address Ranges are associated with the current bus access. It remains valid during T <sub>d</sub> cycles				
PWE#	<b>PERIPHERAL WRITE ENABLE</b> indicates to a peripheral device whether or not to use the data on the AD15:0 bus to write the addressed space. It is low during $T_w$ cycles and deasserts during the $T_d$ cycle for a write; it is high during $T_a$ and $T_w/T_d$ cycles for a read.				



## 9.2.5 Flash Memory Support

PBI peripheral bus interface supports 8-, or 16- bit Flash devices.

The PBI provides programmable wait state functionality for peripheral memory windows.

*Note:* Potentially, programmable wait state functionality could be connected to any peripheral device that has a deterministic wait state profile. However, data valid and turn-around times would need to fit within parameters provided by programmable wait state profiles to support Flash devices.

Any write transactions issued to a Flash address space window must always represent a single flash bus data cycle (**strb**, **strh**).

The peripheral chip enables, **PCE[1:0]**#, activate the appropriate Peripheral window when the address falls within one of the Peripheral address ranges.

*Note:* By default, bank 0 is enabled with the maximum number of Address-to-Data and Recovery Wait states. The width of the interface can be strapped for either 8-bit wide Flash or 16-bit wide flash. Thus, **PCE0#** is the Peripheral Bus chip enable to be used for booting purposes.

Figure 100 on page 588 illustrates how two 8-bit Flash devices would interface with the 80333 through the PBI Interface.

#### Figure 100. Four Mbyte Flash Memory System<sup>a</sup>



a. 16-bit wide flash devices would require two latches.

### 9.2.5.1 Flash Read Cycle

Reading a Flash device involves driving the address, output enable, and chip enable. Depending on the speed of the Flash device, the data returns several cycles later.

The definition of address-to-data wait states are the number of cycles between the assertion of **PCE[1:0]**#, and the arrival of data from the Flash device on **AD[7:0]** (8-bit Flash). The definition of recovery wait states are the number of cycles between the data arrival on **AD[7:0]** and the address for the next Peripheral transaction.

Address-to-data and recovery wait states are programmed in PBBAR0 and PBBAR1 and are identical for reads and writes. Since the read wait state requirement is typically greater, the write wait state requirement is guaranteed to be met.

Refer to Figure 101 illustrates a read cycle for a 120 ns Flash device.

#### Figure 101. 120 ns Flash Read Cycle



Refer to Table 320 for the programmable address-to data and recovery wait states. These numbers are based on a 66 MHz internal clock for the PBI interface.

 Table 320.
 Flash Wait State Profile Programming<sup>a</sup>

Flash Speed	Address-to-Data Wait States	Recovery Wait States		
<= 55 ns	5	0		
<= 115 ns	9	3		
<= 150 ns	13	3		

a. Each Wait State Represents 15 ns.



#### 9.2.5.2 Flash Write Cycle

Address-to-data and recovery wait states for reads and writes are identical and programmed in FBBAR0 and FBBAR1. Refer to Table 320 for the programmable address-to data wait states. However, Any write transactions issued to a Peripheral address space window must always represent a single peripheral bus data cycle (**strb**, **strh**) depending on the bus width selected in PBBAR0 - PBBAR1. Bursting is not supported to a Flash device since the Flash device has no write buffers to support bursting data.

Figure 102 illustrates a write cycle to a 120 ns Flash device.

#### Figure 102. 120 ns Flash Write Cycle



## 9.3 Intel XScale<sup>®</sup> Core PCI Memory Boot Support

When PBI Window 0 is disabled (Section 9.4.4, "PBI Limit Register 0 - PBLR0" on page 596) and the Intel XScale<sup>®</sup> core PCI Bus Boot Enable in the PBCR (Section 9.4.1, "PBI Control Register - PBCR" on page 593) is set, the PBI provides the ability for the Intel XScale<sup>®</sup> core to boot (Section 17.3.3, "Exception Priorities and Vectors" on page 753) from PMMR register space. When the mode is enabled, three registers is available at the Internal Bus address 0000 0000H, 0000 0020H, and 0000 0024H, otherwise these registers are accessible only from their normal PMMR register addresses FFFF E6C0H, FFFF E6E0H, and FFFF E6E4H, respectively.

To boot from PCI memory, the user needs to configure the 80333 power-on-reset straps to hold the Intel XScale<sup>®</sup> core in reset (PCSR bit 1 is set) and to allow a host processor to configure the 80333 PCI interface (PCSR bit 2 is clear) (see Table 168, "PCI Configuration and Status Register - PCSR" on page 283 for details).

Following the host configuration of the 80333 PCI interface, the user can boot using code that resides on the PCI bus (typically host memory) through an Outbound Memory window. When the Intel XScale<sup>®</sup> core is released from reset (PCSR bit 1 is cleared), the PBI provides facilities for the core to execute a short jump or a long jump from the reset vector to any of the Outbound Memory windows including the direct addressing window.

#### Example 6. Memory-less Boot through by Primary ATU Outbound Memory Window 0

- 1. Disable or Remap PBI Memory Window 0 which is the default window used to boot from Flash. Write to PBLR0 to disable the PBI Memory Window 0. Remap PBI Memory Window 0 to an address above 0x3FH by writing to PBBAR0.
- 2. Set the Intel XScale<sup>®</sup> core PCI Bus Boot Enable bit (PBCR bit 3).
- 3. Write a short branch into PMBR0 (Intel XScale<sup>®</sup> core Reset exception vector). Typically, PMBR0 is written with a short branch to 0000 0020H (PMBR1) (see Section 9.4.7, "PBI Memory-less Boot Registers 2:0 PMBR[2:0]" on page 599).
- 4. Write a long branch into the PMBR1/2 registers. PMBR1 holds the branch instruction while PMBR2 represents the 32-bit branch vector.
- 5. Write the 64 Mbyte aligned Host Memory Address where the Intel XScale<sup>®</sup> core boot code resides into the OMWTVR0 register (Section 3.10.34, "Outbound Memory Window Translate Value Register 0 OMWTVR0" on page 273).
- When the Host Memory Address resides above the 4 Gbyte address boundary, write the upper 32 bits of the Host Memory Address into the OUMWTVR0 register (Section 3.10.35, "Outbound Upper 32-bit Memory Window Translate Value Register 0 - OUMWTVR0" on page 274).
- 7. Release the Intel XScale<sup>®</sup> core to boot from host memory by clearing the core processor reset bit in the PCSR (PCSR bit 1).
- *Note:* The Long Branch at 0000 0020H (PMBR1/2) should branch to the 64 Mbytes of address space that is claimed by the ATUs Primary Outbound Memory Window 0 (8000 0000H to 83FF FFFFH). With the MMU enabled following boot-up, the exception vector's physical addresses can be moved from the address range 0000 0000H to 0000 001CH to an address range contained in the 64 Mbytes of Primary Outbound Memory Window 0. Typically, the exception vectors is locked into the cache following boot-up.

## 9.4 Register Definitions

A series of configuration registers control the PBI. Software can determine the status of the PBI by reading the status register. Table 321 lists all of the PBI registers which are detailed further in proceeding sections.

#### Table 321. Peripheral Bus Interface Register

Section, Register Name - Acronym (Page)
Section 9.4.1, "PBI Control Register - PBCR" on page 593
Section 9.4.3, "PBI Base Address Register 0 - PBBAR0" on page 595
Section 9.4.4, "PBI Limit Register 0 - PBLR0" on page 596
Section 9.4.5, "PBI Base Address Register 1 - PBBAR1" on page 597
Section 9.4.6, "PBI Limit Register 1 - PBLR1" on page 598
Section 9.4.7, "PBI Memory-less Boot Registers 2:0 - PMBR[2:0]" on page 599
Section 9.4.8, "PBI Drive Strength Control Register - PBDSCR" on page 600



## 9.4.1 PBI Control Register - PBCR

The PBI Control Register (PBCR) is responsible for enabling the operation of the PBI state machines.



#### Table 322. PBI Control Register - PBCR



## 9.4.2 Determining Block Sizes for Memory Windows

The memory window size can be determined by writing ones to the appropriate upper bits of the limit register. The binary-weighted value of the first non-zero bit set in the limit register indicates the size of the memory window. Table 323 describes the relationship between limit register values and the byte sizes of the memory window.

#### Table 323. Memory Block Size Limit Register Value

Limit Register Value <sup>a</sup>	Size (in Bytes)	Limit Register Value	Size (in Bytes)
FFFFFF0H	16	FFF00000H	1 M
FFFFFE0H	32	FFE00000H	2 M
FFFFFC0H	64	FFC00000H	4 M
FFFFF80H	128	FF800000H	8 M
FFFFF00H	256	FF000000H	
FFFFE00H	512	FE000000H	
FFFFC00H	1K	FC000000H	
FFFF800H	2K	F8000000H	
FFFF000H	4K	F000000H	
FFFE000H	8K	E000000H	Address
FFFFC000H	16K	С000000Н	Closed.
FFFF8000H	32K	8000000H	
FFFF0000H	64K		
FFFE0000H	128K	00000000	
FFFC0000H	256K	000000011	
FFF80000H	512K		

a. Shaded Limit Register Values and Memory Block Sizes are not supported by the PBI.

As an example, assume that FFF0 0004H is written to the PBI Limit Register 0 (FBLR0). Scanning upwards starting at bit 12, bit 20 is the first one bit found. The binary-weighted value of this bit is 1,048,576, indicating a 1 Mbyte of memory window.

When programming the Base and Limit Registers for a memory window, the Base Address always needs to be aligned the size of the memory window set in a limit register. For a 1 Mbyte memory window, only bit 20 through bit 31 of the base address from the PBI Base Address Register 0 (FBBAR0) is relevant to the PBI when decoding Memory Window 0.

Warning: A given PBI Base (PBBAR0-PBBAR1) and Limit (PBLR0-PBLR1) register pair should not be modified during the time there is activity on the peripheral bus associated with that particular peripheral memory window. For instance, following boot-up, code executing from Peripheral Memory Window 0 may be used to modify the PBI Base and Limit registers for Peripheral Memory Window 1, but not for Peripheral Memory Window 0.



## 9.4.3 PBI Base Address Register 0 - PBBAR0

The PBI Base Address Register 0 (PBBAR0) defines the block of memory addresses where PBI Memory Window 0 begins. The PBBAR0 defines the base address and describes the required memory block size; see Section 9.4.2, "Determining Block Sizes for Memory Windows" on page 594. The selected base address needs to be naturally aligned to the granularity of the memory block size. For instance, when a 64 Kbyte memory window size is selected, the base address needs to be 64 Kbyte address aligned (i.e., bits 15:12 of the base address are required to be 000b).

Bits 1:0 define the PBI bus width for PBI Memory Window 0.



#### Table 324.PBI Base Address Register 0 - PBBAR0



## 9.4.4 PBI Limit Register 0 - PBLR0

The 80333 limit register (PBLR0) programmed value must be naturally aligned with the base address register (PBBAR0) programmed value. The limit register is used as a mask when the address decode for memory window 0 is performed.



#### Table 325. PBI Limit Register 0 - PBLR0



## 9.4.5 PBI Base Address Register 1 - PBBAR1

The PBI Base Address Register 1 (PBBAR1) defines the block of memory addresses where PBI Memory Window 1 begins. The PBBAR1 defines the base address and describes the required memory block size; see Section 9.4.2, "Determining Block Sizes for Memory Windows" on page 594. The selected base address needs to be naturally aligned to the granularity of the memory block size. For instance, when a 64 Kbyte memory window size is selected, the base address needs to be 64 Kbyte address aligned (i.e., bits 15:12 of the base address are required to be 000b).

Bits 1:0 define the PBI bus width for PBI Memory Window 1.



#### Table 326. PBI Base Address Register 1- PBBAR1



## 9.4.6 PBI Limit Register 1 - PBLR1

The 80333 limit register's (PBLR1) programmed value must be naturally aligned with the base address register's (PBBAR1) programmed value. The limit register is used as a mask when the address decode for memory window 1 is performed.







## 9.4.7 PBI Memory-less Boot Registers 2:0 - PMBR[2:0]

The three PBI Memory-less Boot Registers are used to hold the Intel XScale<sup>®</sup> core Reset exception vector and a long branch during a PCI Memory Boot sequence.

These three registers is made accessible at the addresses 0000 0000H, 0000 0020H, and 0000 0024H, when the 80333 performs a PCI Memory Boot sequence (see Section 9.3, "Intel XScale<sup>®</sup> Core PCI Memory Boot Support" on page 591 for more details).

Table 328 shows the PBI Memory-less Boot Registers 2:0.



#### Table 328. PBI Memory-less Boot Registers 2:0 - PMBR[2:0]



## 9.4.8 **PBI Drive Strength Control Register - PBDSCR**

The PBI drive strength control register is used to manually control the slew rate and drive strength of the peripheral bus interface, the **TDO** pin, and the **GPIO**[7:0] pins.



#### Table 329. PBI Drive Strength Control Register - PBDSCR

# int<sub>el.</sub> *P*C Bus Interface Units

## 10

This chapter describes the two  $I^2C$  (Inter-Integrated Circuit) bus interface units, including the operation modes and setup. Throughout this manual, these peripherals are referred to as the  $I^2C$  units.

## 10.1 Overview

Both I<sup>2</sup>C Bus Interface Units allows the Intel<sup>®</sup> 80333 I/O processor (80333) to serve as a master and slave device residing on the I<sup>2</sup>C bus. The I<sup>2</sup>C bus is a serial bus developed by Philips Corporation consisting of a two-pin interface. **SDA** is the data pin for input and output functions and **SCL** is the clock pin for reference and control of the I<sup>2</sup>C bus.

The I<sup>2</sup>C bus allows the 80333 to interface to other I<sup>2</sup>C peripherals and microcontrollers for system management functions. The serial bus requires a minimum of hardware for an economical system to relay status and reliability information on the 80333 subsystem to an external device.

The I<sup>2</sup>C Bus Interface Unit is a peripheral device that resides on a 80333 internal bus. Data is transmitted to and received from the I<sup>2</sup>C bus via a buffered interface. Control and status information is relayed through a set of memory-mapped registers. Refer to the I<sup>2</sup>C - *Bus Specification*, version 2.1 for complete details on I<sup>2</sup>C bus operation.

## 10.2 I<sup>2</sup>C Interface

 $I^2C$  channel 1 bus interface is multiplexed with the SMBus interface, and is by default disabled. To enable  $I^2C$  channel 1 requires both the SMBus interface to be disabled in the Section 11.4.10, "SMBus Enable Register - SMBER" on page 650 and the  $I^2C$  channel 1 interface unit to be enabled in the " $I^2C$  Control Register x - ICRx" on page 626.

Table 330.I<sup>2</sup>C Interface Pins

Signal	Description			
SCL0	I <sup>2</sup> C Clock: Provides synchronous operation of I <sup>2</sup> C bus zero.			
SDA0	C Data: Is used for data transfer and arbitration of $I^2C$ bus zero.			
SCL1/SCLK	1/SCLK I <sup>2</sup> C Clock: Provides synchronous operation of I <sup>2</sup> C bus one.			
SDA1/SDTA	$I^2C$ Data: Is used for data transfer and arbitration of $I^2C$ bus one.			
Total	4			

## 10.3 Theory of Operation

The I<sup>2</sup>C bus defines a serial protocol for passing information between agents on the I<sup>2</sup>C bus using only a two pin interface. The interface consists of a Serial Data/Address (**SDA**) line and a Serial Clock Line (**SCL**). Each device on the I<sup>2</sup>C bus is recognized by a unique 7-bit address and can operate as a transmitter or as a receiver. In addition to transmitter and receiver, the I<sup>2</sup>C bus uses the concept of master and slave. Table 331 lists the I<sup>2</sup>C device types.

#### Table 331.I<sup>2</sup>C Bus Definitions

I <sup>2</sup> C Device	Definition				
Transmitter	Sends data to the I <sup>2</sup> C bus.				
Receiver	Receives data from the I <sup>2</sup> C bus.				
Master	Initiates a transfer, generates the clock signal, and terminates the transactions.				
Slave	The device addressed by a master.				
Multi-master	More than one master can attempt to control the bus at the same time without corrupting the message.				
Arbitration	Procedure to ensure that, when more than one master simultaneously tries to control the bus, only one is allowed. This procedure ensures that messages are not corrupted.				

As an example of  $I^2C$  bus operation, consider the case of the 80333 acting as a master on the bus (see Figure 103). The 80333, as a master, addresses an EEPROM as a slave to receive data. The 80333 is a master-transmitter and the EEPROM is a slave-receiver. When the 80333 reads data, the 80333 is a master-receiver and the EEPROM is a slave-transmitter. In both cases, the master generates the clock, initiates the transaction and terminates it.

#### Figure 103. I<sup>2</sup>C Bus Configuration Example



The I<sup>2</sup>C bus allows for a multi-master system, which means more than one device can initiate data transfers at the same time. To support this feature, the I<sup>2</sup>C bus arbitration relies on the wired-AND connection of all I<sup>2</sup>C interfaces to the I<sup>2</sup>C bus. Two masters can drive the bus simultaneously provided they are driving identical data. The first master to drive **SDA** high while another master drives **SDA** low loses the arbitration. The **SCL** line consists of a synchronized combination of clocks generated by the masters using the wired-AND connection to the **SCL** line.

The I<sup>2</sup>C bus serial operation uses an open-drain wired-AND bus structure, which allows multiple devices to drive the bus lines and to communicate status about events such as arbitration, wait states, error conditions and so on. For example, when a master drives the clock (**SCL**) line during a data transfer, it transfers a bit on every instance that the clock is high. When the slave is unable to accept or drive data at the rate that the master is requesting, the slave can hold the clock line low between the high states to insert a wait interval. The master's clock can only be altered by a slow slave peripheral keeping the clock line low or by another master during arbitration.

 $I^2C$  transactions are either initiated by the 80333 as a master or are received by the processor as a slave. Both conditions may result in the processor doing reads, writes, or both to the  $I^2C$  bus.



## 10.3.1 Operational Blocks

The I<sup>2</sup>C Bus Interface Unit is a slave peripheral device that is connected to the internal bus. The 80333 interrupt mechanism can be used for notifying the 80333 that there is activity on the I<sup>2</sup>C bus. Polling can be also be used instead of interrupts, although it would be very cumbersome. Figure 104 shows a block diagram of the I<sup>2</sup>C Bus Interface Unit and its interface to the internal bus.

The I<sup>2</sup>C Bus Interface Unit consists of the two wire interface to the I<sup>2</sup>C bus, an 8-bit buffer for passing data to and from the 80333, a set of control and status registers, and a shift register for parallel/serial conversions.







The I<sup>2</sup>C interrupts are signalled through a single pin which provides a level sensitive interrupt to the 80333 interrupt control unit. The I<sup>2</sup>C Bus Interface Unit can cause and interrupt when a buffer is full, buffer empty, slave address detected, arbitration lost, or bus error condition occurs. All interrupt conditions must be cleared explicitly by software. See Section 10.9.2, "I<sup>2</sup>C Status Register x - ISRx" on page 628 for details.

The I<sup>2</sup>C Data Buffer Register (IDBR) is an 8-bit data buffer that receives a byte of data from the shift register interface of the I<sup>2</sup>C bus on one side and parallel data from the 80333 internal bus on the other side. The serial shift register is not user accessible.

The control and status registers are located in the I<sup>2</sup>C memory-mapped address space (FFFF F680H to FFFF F694H). The registers and their function are defined in Section 10.9.

The I<sup>2</sup>C Bus Interface Unit supports fast mode operation of 400 Kbits/sec. Fast mode logic levels, formats, and capacitive loading, and protocols are exactly the same as the 100 Kbits/sec standard mode. Because the data setup and hold times differ between the fast and standard mode, the I<sup>2</sup>C is designed to meet the slower, standard mode requirements for these two specifications. Refer to the I<sup>2</sup>C Bus Specification for details.



## **10.3.2** I<sup>2</sup>C Bus Interface Modes

The I<sup>2</sup>C Bus Interface Unit can be in different modes of operation to accomplish a transfer. Table 332 summarizes the different modes.

#### Table 332.Modes of Operation

Mode	Definition				
Master - Transmit	<ul> <li>I<sup>2</sup>C Bus Interface Unit acts as a master.</li> <li>Used for a write operation.</li> <li>I<sup>2</sup>C Bus Interface Unit sends the data.</li> <li>I<sup>2</sup>C Bus Interface Unit is responsible for clocking.</li> <li>Slave device is in slave-receive mode</li> </ul>				
Master - Receive	<ul> <li>I<sup>2</sup>C Bus Interface Unit acts as a master.</li> <li>Used for a read operation.</li> <li>I<sup>2</sup>C Bus Interface Unit receives the data.</li> <li>I<sup>2</sup>C Bus Interface Unit is responsible for clocking.</li> <li>Slave device is in slave-transmit mode</li> </ul>				
Slave - Transmit	<ul> <li>I<sup>2</sup>C Bus Interface Unit acts as a slave.</li> <li>Used for a read (master) operation.</li> <li>I<sup>2</sup>C Bus Interface Unit sends the data.</li> <li>Master device is in master-receive mode.</li> </ul>				
Slave - Receive (default)	<ul> <li>I<sup>2</sup>C Bus Interface Unit acts as a slave.</li> <li>Used for a write (master) operation.</li> <li>I<sup>2</sup>C Bus Interface Unit receives the data.</li> <li>Master device is in master-transmit mode.</li> </ul>				

While the I<sup>2</sup>C Bus Interface Unit is in idle mode (neither receiving or transmitting serial data), the unit defaults to Slave-Receive mode. This allows the interface to monitor the bus and receive any slave addresses that might be intended for the 80333.

When the I<sup>2</sup>C Bus Interface Unit receives an address that matches the 7-bit address found in the I<sup>2</sup>C Slave Address Register (ISAR) or the General Call Address (00H), the interface either remains in Slave-Receive mode or transitions to Slave-Transmit mode. This is determined by the Read/Write (R/W#) bit (the least significant bit of the byte containing the slave address). When the R/W# bit is low, the master initiating the transaction intends to do a write and the I<sup>2</sup>C Bus Interface Unit remains in Slave-Receive mode. When the R/W# is high, the initiating master wants to read data and the slave transitions to Slave-Transmit mode. Slave operation is further defined in Section 10.4.6, "Slave Operations" on page 618.

When the 80333 wants to initiate a read or write on the  $I^2C$  bus, the  $I^2C$  Bus Interface Unit transitions from the default Slave-Receive mode to Master-Transmit mode. When the 80333 wants to write data, the interface remains in Master-Transmit mode after the address transfer has completed. (see Section 10.3.3.1, "START Condition" on page 607) for START information). When the 80333 wants to read data, the  $I^2C$  Bus Interface Unit transmits the start address, then transition to Master-Receive mode. Master operation is further defined in Section 10.4.5, "Master Operations" on page 614.



## **10.3.3 Start and Stop Bus States**

The I<sup>2</sup>C bus defines a transaction START and a transaction STOP bus state that are used at the beginning and end of the transfer of one to an unlimited number of bytes on the bus.

The 80333 uses the START and STOP bits in the I<sup>2</sup>C Control Register (ICR) to:

- initiate an additional byte transfer
- initiate a START condition on the I<sup>2</sup>C bus
- enable Data Chaining (repeated START)
- initiate a STOP condition on the  $I^2C$  bus

Table 333 summarizes the definition of the START and STOP bits in the ICR.

#### Table 333. START and STOP Bit Definitions

STOP bit	STAR T bit	Condition	Notes			
0	0	No START or STOP	<ul> <li>No START or STOP condition is sent by the I<sup>2</sup>C Bus Interface Unit. This is used when multiple data bytes need to be transferred.</li> </ul>			
0	1	START Condition and Repeated START	<ul> <li>The I<sup>2</sup>C Bus Interface Unit sends a START condition and transmit the contents of the 8 bit IDBR after the START. The IDBR must contain the 7-bit address and the R/W# bit before a START is initiated.</li> </ul>			
			<ul> <li>For a repeated start, the IDBR contents contains the target slave address and the R/W# bit. This enables multiple transfers to different slaves without giving up the bus.</li> </ul>			
			• The interface stays in Master-Transmit mode when a write is used or transition to master-receive mode when a read is requested.			
1	x	STOP Condition	<ul> <li>In Master-Transmit mode, the I<sup>2</sup>C Bus Interface Unit transmits the 8-bit IDBR and then send a STOP on the I<sup>2</sup>C bus.</li> </ul>			
			<ul> <li>In Master-Receive mode, the Ack/Nack Control bit in the ICR must be changed to a negative Ack (see Section 10.4.3). The I<sup>2</sup>C Bus Interface Unit writes the Nack bit (Ack/Nack Control bit must be 1), receive the data byte in the IDBR, then send a STOP on the I<sup>2</sup>C bus.</li> </ul>			

Figure 105 shows the relationship between the SDA and SCL lines for a START and STOP condition.

#### Figure 105. Start and Stop Conditions



# intel

#### 10.3.3.1 START Condition

The START condition (bits 1:0 of the ICR set to  $01_2$ ) initiates a master transaction or repeated START. Software must load the target slave address and the R/W# bit in the IDBR (see Section 10.9.4, "I<sup>2</sup>C Data Buffer Register x - IDBRx" on page 631) before setting the START ICR bit. The START and the IDBR contents are transmitted on the I<sup>2</sup>C bus when the ICR Transfer Byte bit is set. The I<sup>2</sup>C bus stays in master-transmit mode when a write is requested or enters master-receive mode when a read is requested. For a repeated start (a change in read or write or a change in the target slave address), the IDBR contains the updated target slave address and the R/W# bit. A repeated start enables multiple transfers to different slaves without giving up the bus.

The START condition is not cleared by the  $I^2C$  unit. When arbitration is lost while initiating a START, the  $I^2C$  unit may re-attempt the START when the bus becomes free. See Section 10.4.4, "Arbitration" on page 612 for details on how the  $I^2C$  unit functions under those circumstances.

#### 10.3.3.2 No START or STOP Condition

No START or STOP condition (bits 1:0 of the ICR set to  $00_2$ ) is used in master-transmit mode while the 80333 is transmitting multiple data bytes (see Figure 105). Software writes the data byte, sets the IDBR Transmit Empty bit in the ISR (and interrupt when enabled), and clears the Transfer Byte bit in the ICR. The software then writes a new byte to the IDBR and sets the Transfer Byte ICR bit, which initiates the new byte transmission. This continues until the software sets the START or STOP bit. The START and STOP bits in the ICR are not automatically cleared by the I<sup>2</sup>C unit after the transmission of a START, STOP or repeated START.

After each byte transfer (including the Ack/Nack bit) the  $I^2C$  unit holds the **SCL** line low (inserting wait states) until the Transfer Byte bit in the ICR is set. This action notifies the  $I^2C$  unit to release the **SCL** line and allow the next information transfer to proceed.

#### 10.3.3.3 STOP Condition

The STOP condition (bits 1:0 of the ICR set to  $10_2$ ) terminates a data transfer. In master-transmit mode, the STOP bit and the Transfer Byte bit in the ICR must be set to initiate the last byte transfer (see Figure 105). In master-receive mode, to initiate the last transfer the 80333 must set the Ack/Nack bit, the STOP bit, and the Transfer Byte bit in the ICR. Software must clear the STOP condition after it is transmitted.

#### Figure 106. START and STOP Conditions

No STAR D	T or STOP Con	Ack/ Nack				
START C START	ondition Target Slave	Address	R/W#	Ack/ Nack		
STOP Co Data	ndition Byte R/	W# Ack/ Nack	STOP	,		



## **10.4** I<sup>2</sup>C Bus Operation

The I<sup>2</sup>C Bus Interface Unit transfers in 1 byte increments. A data transfer on the I<sup>2</sup>C bus always follows the sequence:

- 1) START
- 2) 7-bit Slave Address
- 3) R/W# Bit
- 4) Acknowledge Pulse
- 5) 8 Bits of Data
- 6) Ack/Nack Pulse
- 7) Repeat of Step 5 and 6 for Required Number of Bytes
- 8) Repeated START (Repeat Step 1) or STOP

### **10.4.1** Serial Clock Line (SCL) Generation

The 80333  $I^2C$  unit is required to generate the  $I^2C$  clock output when in master mode (either receive or transmit). **SCL** clock generation is accomplished through the use of the Fast Mode Enable bit, which is programmed at initialization. The following equation is used to determine the **SCL** transition period:

*Note:* An I<sup>2</sup>C bus lock condition can be cleared by software doing a toggle of the GPOD [11:10] to toggle SCL[1:0]. Bits 11 and 10 of the GPOD register are writable. When GPOD[11] is high, SCL[1] is driven low. When GPOD[10] is driven high, SCL[0] is driven low.

#### **Equation 16. SCL Transition Period**

SCL Transition Period = (30 ns) \* (167 - (Fast Mode Enable \* 125))

March 2005

## 10.4.2 Data and Addressing Management

Data and slave addressing is managed via the I<sup>2</sup>C Data Buffer Register (IDBR) and the I<sup>2</sup>C Slave Address Register (ISAR). The IDBR (see Section 10.9.4, "I<sup>2</sup>C Data Buffer Register x - IDBRx" on page 631) contains data or a slave address and R/W# bit. The ISAR contains the 80333 programmable slave address. Data coming into the I<sup>2</sup>C unit is received into the IDBR after a full byte is received and acknowledged. To transmit data, the processor writes to the IDBR, and the I<sup>2</sup>C unit passes this onto the serial bus when the Transfer Byte bit in the ICR is set. See Section 10.9.1, "I<sup>2</sup>C Control Register x - ICRx" on page 626.

When the  $I^2C$  unit is in transmit mode (master or slave):

- 1. Software writes data to the IDBR over the internal bus. This initiates a master transaction or sends the next data byte, after the IDBR Transmit Empty bit is sent.
- 2. The I<sup>2</sup>C unit transmits the data from the IDBR when the Transmit Empty bit in the ICR is set.
- 3. When enabled, an IDBR Transmit Empty interrupt is signalled when a byte is transferred on the I<sup>2</sup>C bus and the acknowledge cycle is complete.
- 4. When the I<sup>2</sup>C bus is ready to transfer the next byte before the processor has written the IDBR (and a STOP condition is not in place), the I<sup>2</sup>C unit inserts wait states until the processor writes a new value into the IDBR and sets the ICR Transfer Byte bit.

When the  $I^2C$  unit is in receive mode (master or slave):

- 1. The processor reads the IDBR data over the internal bus after the IDBR Receive Full interrupt is signalled.
- 2. The  $I^2C$  unit transfers data from the shift register to the IDBR after the Ack cycle completes.
- 3. The I<sup>2</sup>C unit inserts wait states until the IDBR is read. Refer to Section 10.4.3, "I<sup>2</sup>C Acknowledge" on page 611 for acknowledge pulse information in receiver mode.
- 4. After the Intel XScale<sup>®</sup> core (ARM\* architecture compliant) reads the IDBR, the I<sup>2</sup>C unit writes the ICR's Ack/Nack Control bit and the Transfer Byte bit, allowing the next byte transfer to proceed.



#### 10.4.2.1 Addressing a Slave Device

As a master device, the I<sup>2</sup>C unit must compose and send the first byte of a transaction. This byte consists of the slave address for the intended device and a R/W# bit for transaction definition. The slave address and the R/W# bit are written to the IDBR (see Figure 107).

#### Figure 107. Data Format of First Byte in Master Transaction



The first byte transmission must be followed by an Ack pulse from the addressed slave. When the transaction is a write, the  $I^2C$  unit remains in master-transmit mode and the addressed slave device stays in slave-receive mode. When the transaction is a read, the  $I^2C$  unit transitions to master-receive mode immediately following the Ack and the addressed slave device transitions to slave-transmit mode. When a Nack is returned, the  $I^2C$  unit aborts the transaction by automatically sending a STOP and setting the ISR bus error bit.

When the  $I^2C$  unit is enabled and idle (no bus activity), it stays in slave-receive mode and monitors the  $I^2C$  bus for a START signal. Upon detecting a START pulse, the  $I^2C$  unit reads the first seven bits and compares them to those in the  $I^2C$  Slave Address Register (ISAR) and the general call address (00H). When the bits match those of the ISAR register, the  $I^2C$  unit reads the eighth bit (R/W# bit) and transmits an Ack pulse. The  $I^2C$  unit either remains in slave-receive mode (R/W# = 0) or transitions to slave-transmit mode (R/W# = 1). See Section 10.4.7, "General Call Address" on page 620 for actions when a general call address is detected.

## 10.4.3 I<sup>2</sup>C Acknowledge

Every I<sup>2</sup>C byte transfer must be accompanied by an acknowledge pulse, which is always generated by the receiver (master or slave). The transmitter must release the **SDA** line for the receiver to transmit the acknowledge pulse (see Figure 108).

In master-transmit mode, when the target slave receiver device cannot generate the acknowledge pulse, the **SDA** line remains high. This lack of acknowledge (Nack) causes the I<sup>2</sup>C unit to set the bus error detected bit in the ISR and generate the associated interrupt (when enabled). The I<sup>2</sup>C unit aborts the transaction by generating a STOP automatically.

In master-receive mode, the I<sup>2</sup>C unit signals the slave-transmitter to stop sending data by using the negative acknowledge (Nack). The Ack/Nack bit value driven by the I<sup>2</sup>C bus is controlled by the Ack/Nack bit in the ICR. The bus error detected bit in the ISR is not set for a master-receive mode Nack (as required by the I<sup>2</sup>C bus protocol). The I<sup>2</sup>C unit automatically transmits the Ack pulse, based on the Ack/Nack ICR bit, after receiving each byte from the serial bus. Before receiving the last byte, software must set the Ack/Nack Control bit to Nack. Nack is then sent after the next byte is received to indicate the last byte.

In slave mode, the  $I^2C$  unit automatically acknowledges its own slave address, independent of the Ack/Nack bit setting in the ICR. As a slave-receiver, an Ack response is automatically given to a data byte, independent of the Ack/Nack bit setting in the ICR. The  $I^2C$  unit sends the Ack value after receiving the eighth data bit of the byte.

In slave-transmit mode, receiving a Nack from the master indicates the last byte is transferred. The master then sends either a STOP or repeated START. The ISR's unit busy bit (2) remains set until a STOP or repeated START is received.



#### Figure 108. Acknowledge on the I<sup>2</sup>C Bus



#### **10.4.4** Arbitration

Arbitration on the  $I^2C$  bus is required due to the multi-master capabilities of the  $I^2C$  bus. Arbitration is used when two or more masters simultaneously generate a START condition within the minimum  $I^2C$  hold time of the START condition.

Arbitration can continue for a long period. When the address bit and the R/W# are the same, the arbitration moves to the data. Due to the wired-AND nature of the  $I^2C$  bus, no data is lost when both (or all) masters are outputting the same bus states. When the address, the R/W# bit, or the data are different, the master which outputted the high state (master's data is different from **SDA**) loses arbitration and shut its data drivers off. When losing arbitration, the  $I^2C$  Bus Interface Unit shuts off the **SDA** or **SCL** drivers for the remainder of the byte transfer, set the Arbitration Loss Detected bit, then return to idle (Slave-Receive) mode.

#### 10.4.4.1 SCL Arbitration

Each master on the  $I^2C$  bus generates its own clock on the **SCL** line for data transfers. With masters generating their own clocks, clocks with different frequencies may be connected to the **SCL** line. Since data is valid when the clock is in the high period, a defined clock synchronization procedure is needed during bit-by-bit arbitration.

Clock synchronization is accomplished by using the wired-AND connection of the I<sup>2</sup>C interfaces to the **SCL** line. When a master's clock transitions from high to low, this causes the master to hold down the **SCL** line for its associated period (see Figure 109). The low to high transition of the clock may not change when another master has not completed its period. Therefore, the master with the longest low period holds down the **SCL** line. Masters with shorter periods are held in a high wait-state during this time. Once the master with the longest period completes, the **SCL** line transitions to the high state, masters with the shorter periods can continue the data cycle.

#### Figure 109. Clock Synchronization During the Arbitration Procedure


# intel

# 10.4.4.2 SDA Arbitration

Arbitration on the **SDA** line can continue for a long period, starting with address and R/W# bits and continuing with data bits. Figure 110 shows the arbitration procedure for two masters (more than two may be involved depending on how many masters are connected to the bus). When the address and R/W# are the same, arbitration moves to the data. Due to the wired-AND nature of the I<sup>2</sup>C bus, no data is lost when both (or all) masters are outputting the same bus states. When address, R/W#, or data is different, the master that output the first low data bit loses arbitration and shuts its data drivers off. When the I<sup>2</sup>C unit loses arbitration, it shuts off the **SDA** or **SCL** drivers for remainder of byte transfer, sets arbitration loss detected ISR bit, then returns to idle (Slave-Receive) mode.





When the I<sup>2</sup>C unit loses arbitration during transmission of the seven address bits and the 80333 is not being addressed as a slave device, the I<sup>2</sup>C unit re-sends the address when the I<sup>2</sup>C bus becomes free. This is possible because the IDBR and ICR registers are not overwritten when arbitration is lost.

When the arbitration loss is to due to another bus master addressing the 80333 as a slave device, the  $I^2C$  unit switches to slave-receive mode and the original data in the  $I^2C$  data buffer register is overwritten. Software is responsible for clearing the start and re-initiating the master transaction at a later time.

*Note:* Software must not allow the  $I^2C$  unit to write to its own slave address. This can cause the  $I^2C$  bus to enter an indeterminate state.

Boundary conditions exist for arbitration when an arbitration process is in progress and a repeated START or STOP condition is transmitted on the I<sup>2</sup>C bus. To prevent errors, the I<sup>2</sup>C unit, acting as a master, provides for the following sequences:

- No arbitration takes place between a repeated START condition and a data bit
- No arbitration takes place between a data bit and a STOP condition
- No arbitration takes place between a repeated START condition and a STOP condition

These situations arise only when different masters write the same data to the same target slave simultaneously and arbitration is not resolved after the first data byte transfer.

*Note:* Typically, software is responsible for ensuring arbitration is lost soon after the transaction begins. For example, the protocol might insist that all masters transmit their  $I^2C$  address as the first data byte of any transaction ensuring arbitration is ended. A restart is then sent to begin a valid data transfer (the slave can then discard the master's address).



# **10.4.5 Master Operations**

When software initiates a read or write on the  $I^2C$  bus, the  $I^2C$  unit transitions from the default slave-receive mode to master-transmit mode. The start pulse is sent followed by the 7-bit slave address and the R/W# bit. After the master receives an acknowledge, the  $I^2C$  unit has the option of two master modes:

- Master-Transmit The 80333 writes data
- Master-Receive The 80333 reads data

The 80333 initiates a master transaction by writing to the ICR register. Data is read and written from the  $I^2C$  unit through the memory-mapped registers.

Table 334 describes the I<sup>2</sup>C Bus Interface Unit responsibilities as a master device.

## Table 334. Master Transactions (Sheet 1 of 3)

I <sup>2</sup> C Master Action	Mode of Operation	Definition	
Generate clock output	Master-transmit Master-receive	<ul> <li>The master always drives the SCL line.</li> <li>The SCL Enable bit must be set.</li> <li>The Unit Enable bit must be set.</li> </ul>	
Write target slave address to IDBR	Master-transmit Master-receive	<ul> <li>The Intel XScale<sup>®</sup> core writes to IDBR bits 7-1 before a START condition is enabled.</li> <li>First 7 bits sent on bus after START.</li> <li>See Section 10.3.3.</li> </ul>	
Write R/W# Bit to IDBR	Master-transmit Master-receive	<ul> <li>The Intel XScale<sup>®</sup> core writes to the least significant IDBR bit with the target slave address.</li> <li>When low, the master remains a master-transmitter. When high, the master transitions to a master-receiver.</li> <li>See Section 10.4.2.</li> </ul>	
Signal START Condition	Master-transmit Master-receive	<ul> <li>See "Generate clock output" above.</li> <li>Performed after the target slave address and the R/W# bit are in the IDBR.</li> <li>Intel XScale<sup>®</sup> core sets the START bit.</li> <li>Intel XScale<sup>®</sup> core sets the Transfer Byte bit which initiates the start condition.</li> <li>See Section 10.3.3.</li> </ul>	
Initiate first data byte transfer	Master-transmit Master-receive	<ul> <li>Intel XScale<sup>®</sup> core writes byte to IDBR</li> <li>I<sup>2</sup>C Bus Interface Unit transmits the byte when the Transfer Byte bit is set.</li> <li>I<sup>2</sup>C Bus Interface Unit clears the Transfer Byte bit and sets the IDBR Transmit Empty bit when the transfer is complete.</li> </ul>	

# intel

## Table 334.Master Transactions (Sheet 2 of 3)

I <sup>2</sup> C Master Action	Mode of Operation	Definition
Arbitrate for I <sup>2</sup> C Bus	Master-transmit Master-receive	<ul> <li>When two or more masters signal a start within the same clock period, arbitration must occur.</li> </ul>
		<ul> <li>The I<sup>2</sup>C Bus Interface Unit arbitrates for as long as necessary. Arbitration takes place during slave address, R/W# bit, and data transmission and continues until all but one master loses the bus. No data is lost during arbitration.</li> </ul>
		<ul> <li>When the I<sup>2</sup>C Bus Interface Unit loses arbitration, it sets the Arbitration Loss Detect ISR bit after byte transfer is complete and transition to slave-receive (default) mode.</li> </ul>
		<ul> <li>When I<sup>2</sup>C Bus Interface Unit loses arbitration while attempting to send the target address byte, the I<sup>2</sup>C Bus Interface Unit attempts to resend it when the bus becomes free.</li> </ul>
		• The system designer must ensure the boundary conditions described in Section 10.4 do not occur.
		Data transmit mode of I <sup>2</sup> C master operation.
Write one data	Master-transmit only Master-transmit only	<ul> <li>Occurs when the IDBR Transmit Empty ISR bit is set and the Transfer Byte bit is clear. When enabled, the IDBR Transmit Empty Interrupt is signalled to the Intel XScale<sup>®</sup> core.</li> </ul>
byte to the IDBR		<ul> <li>Intel XScale<sup>®</sup> core writes 1 data byte to the IDBR, set the appropriate START/STOP bit combination, and then set the Transfer Byte bit to send the data. Eight bits are written on the serial bus followed by a STOP when requested.</li> </ul>
Wait for Acknowledge from		<ul> <li>As a master-transmitter, the I<sup>2</sup>C Bus Interface Unit generates the clock for the acknowledge pulse. The I<sup>2</sup>C Bus Interface Unit is responsible for releasing the SDA line to allow slave-receiver Ack transmission.</li> </ul>
slave-receiver		• See Section 10.4.3.
	Master-receive only	Data receive mode of I <sup>2</sup> C master operation.
		• Eight bits are read from the serial bus, collected in the shift register then transferred to the IDBR after the Ack/Nack bit is read.
		<ul> <li>The Intel XScale<sup>®</sup> core reads the IDBR when the IDBR Receive Full bit is set and the Transfer Byte bit is clear. When enabled, a IDBR Receive Full Interrupt is signalled to the Intel XScale<sup>®</sup> core.</li> </ul>
Read one byte		<ul> <li>When the IDBR is read, when the Ack/Nack Status is clear (indicating Ack), the Intel XScale<sup>®</sup> core writes the Ack/Nack Control bit and set the Transfer Byte bit to initiate the next byte read.</li> </ul>
of I <sup>2</sup> C Data from the IDBR		<ul> <li>When the Ack/Nack Status bit is set (indicating Nack), Transfer Byte bit is clear, STOP bit in the ICR is set, and Unit Busy bit in the ISR is set, then the last data byte has been read into the IDBR and the I<sup>2</sup>C Bus Interface Unit is sending the STOP.</li> </ul>
		<ul> <li>When the Ack/Nack Status bit is set (indicating Nack), Transfer Byte bit is clear, but the STOP bit is clear, then the Intel XScale<sup>®</sup> core has two options: 1. set the START bit, write a new target address to the IDBR, and set the Transfer Byte bit which sends a repeated start condition, 2. set the Master Abort bit and leave the Transfer Byte clear which sends a STOP only.</li> </ul>



# Table 334.Master Transactions (Sheet 3 of 3)

I <sup>2</sup> C Master Action	Mode of Operation	Definition
Transmit Acknowledge to slave-transmitter	Master-receive only	<ul> <li>As a master-receiver, the I<sup>2</sup>C Bus Interface Unit generates the clock for the acknowledge pulse. The I<sup>2</sup>C Bus Interface Unit is also responsible for driving the SDA line during the Ack cycle.</li> <li>When the next data byte is to be the last transaction, the Intel XScale<sup>®</sup> core sets the Ack/Nack Control bit for Nack generation.</li> <li>See Section 10.4.3.</li> </ul>
Generate a Repeated START to chain I <sup>2</sup> C transactions	Master-transmit Master-receive	<ul> <li>When data chaining is desired, a repeated START condition is used instead of a STOP condition.</li> <li>This occurs after the last data byte of a transaction has been written to the bus.</li> <li>The Intel XScale<sup>®</sup> core writes the next target slave address and the R/W# bit to the IDBR, set the START bit, and set the Transfer Byte bit.</li> <li>See Section 10.3.3.</li> </ul>
Generate a STOP	Master-transmit Master-receive	<ul> <li>Generated after the Intel XScale<sup>®</sup> core writes the last data byte on the bus.</li> <li>Intel XScale<sup>®</sup> core generates a STOP condition by setting the STOP bit in the ICR.</li> <li>See Section 10.3.3.</li> </ul>

# intel

When the 80333 needs to read data, the  $I^2C$  unit transitions from slave-receive mode to master-transmit mode to transmit the start address and immediately following the ACK pulse transitions to master-receive mode to wait for the reception of the read data from the slave device (see Figure 111). It is also possible to have multiple transactions during an  $I^2C$  operation such as transitioning from master-receive to master-transmit through a repeated start or Data Chaining (see Figure 112). Figure 113 shows the wave forms of **SDA** and **SCL** for a complete data transfer.

### Figure 111. Master-Receiver Read from Slave-Transmitter



#### Figure 112. Master-Receiver Read from Slave-Transmitter / Repeated Start / Master-Transmitter Write to Slave-Receiver



Figure 113. A Complete Data Transfer



# **10.4.6** Slave Operations

Table 335 describes the I<sup>2</sup>C Bus Interface Unit's responsibilities as a slave device.

#### Table 335.Slave Transactions

I <sup>2</sup> C Slave Action	Mode of Operation	Definition	
Slave-receive (default mode)	Slave-receive only	<ul> <li>I<sup>2</sup>C Bus Interface Unit monitors all slave address transactions.</li> <li>The I<sup>2</sup>C Bus Interface Unit Enable bit must be set.</li> <li>I<sup>2</sup>C Bus Interface Unit monitors bus for START conditions. When a START is detected, the interface reads the first 8 bits and compares the most significant 7 bits with the 7 bit I<sup>2</sup>C Slave Address Register and the General Call address (00H). When there is a match, the I<sup>2</sup>C Bus Interface Unit sends an Ack.</li> <li>When the first 8 bits are all zero's, this is a general call address. When the General Call Disable bit is clear, both the General Call Address Detected bit and the Slave Mode Operation bit in the ISR is set. See Section 10.4.7.</li> <li>When the 8th bit of the first byte (R/W# bit) is low, the I<sup>2</sup>C Bus Interface Unit stays in slave-receive mode and the Slave Mode Operation bit is cleared. When the R/W# bit is high, the I<sup>2</sup>C Bus Interface Unit transitions to slave-transmit mode and the Slave Mode Operation bit is set.</li> </ul>	
Setting the Slave Address Detected bit	Slave-receive Slave-transmit	<ul> <li>Indicates the interface has detected an I<sup>2</sup>C operation that addresses the 80333 (this includes general call address). The Intel XScale<sup>®</sup> core can distinguish an ISAR match from a General Call by reading the General Call Address Detected bit.</li> <li>An interrupt is signalled (when enabled) after the matching slave address is received and acknowledged.</li> </ul>	
Read one byte of I <sup>2</sup> C Data from the IDBR	Slave-receive only	<ul> <li>Data receive mode of I<sup>2</sup>C slave operation.</li> <li>Eight bits are read from the serial bus into the shift register. When a full byte has been received and the Ack/Nack bit has completed, the byte is transferred from the shift register to the IDBR.</li> <li>Occurs when the IDBR Receive Full bit in the ISR is set and the Transfer Byte bit is clear. When enabled, the IDBR Receive Full Interrupt is signalled to the Intel XScale<sup>®</sup> core.</li> <li>Intel XScale<sup>®</sup> core reads 1 data byte from the IDBR. When the IDBR is read, the Intel XScale<sup>®</sup> core writes the desired Ack/Nack Control bit and set the Transfer Byte bit. This causes the I<sup>2</sup>C Bus Interface Unit to stop inserting wait states and let the master transmitter write the next piece of information.</li> </ul>	
Transmit Acknowledge to master-transmitter	Slave-receive only	<ul> <li>As a slave-receiver, the I<sup>2</sup>C Bus Interface Unit is responsible for pulling the SDA line low to generate the Ack pulse during the high SCL period.</li> <li>The Ack/Nack Control bit controls the Ack data the I<sup>2</sup>C Bus Interface Unit drives. See Section 10.4.3.</li> </ul>	
Write one byte of I <sup>2</sup> C data to the IDBR	Slave-transmit only	<ul> <li>Data transmit mode of I<sup>2</sup>C slave operation.</li> <li>Occurs when the IDBR Transmit Empty bit is set and the Transfer Byte bit is clear. When enabled, the IDBR Transmit Empty Interrupt is signalled to the Intel XScale<sup>®</sup> core.</li> <li>Intel XScale<sup>®</sup> core writes a data byte to the IDBR and set the Transfer Byte bit to initiate the transfer.</li> </ul>	
Wait for Acknowledge from master-receiver	Slave-transmit only	<ul> <li>As a slave-transmitter, the I<sup>2</sup>C Bus Interface Unit is responsible for releasing the <b>SDA</b> line to allow the master-receiver to pull the line low for the Ack.</li> <li>See Section 10.4.3.</li> </ul>	



Figure 114 through Figure 116 are examples of  $I^2C$  transactions. These show the relationships between master and slave devices.

Figure 114. Master-Transmitter Write to Slave-Receiver



### Figure 115. Master-Receiver Read to Slave-Transmitter



#### Figure 116. Master-Receiver Read to Slave-Transmitter / Repeated START / Master-Transmitter Write to Slave-Receiver





# 10.4.7 General Call Address

The  $I^2C$  unit supports both sending and receiving general call address transfers on the  $I^2C$  bus. When sending a general call message from the  $I^2C$  unit, software must set the General Call Disable bit in the ICR to keep the  $I^2C$  unit from responding as a slave. Failure to set this bit causes the  $I^2C$  Bus to enter an indeterminate state.

A general call address is defined as a transaction with a slave address of 00H. When a device requires the data from a general call address, it acknowledges the transaction and stays in slave-receiver mode. Otherwise, the device can ignore the general call address. The second and following bytes of a general call transaction are acknowledged by every device using it on the bus. Any device not using these bytes must not Ack. The meaning of a general call address is defined in the second byte sent by the master-transmitter. Figure 117 shows a general call address transaction. The least significant bit (B) of the second byte defines the transaction. Table 336, "General Call Address Second Byte Definitions" on page 620 shows the valid values and definitions when B=0.

When the 80333 is acting as a slave, and the  $I^2C$  unit receives a general call address and the ICR General Call Disable bit is clear the  $I^2C$  unit:

- Sets the ISR general call address detected bit
- Sets the ISR slave address detected bit
- Interrupts (when enabled) the 80333

When the  $I^2C$  unit receives a general call address and the ICR General Call Disable bit is set, the  $I^2C$  unit ignores the general call address.

#### Figure 117. General Call Address



#### Table 336. General Call Address Second Byte Definitions

Least Significant Bit of Second Byte (B)	Second Byte Value	Definition
0	06H	2-byte transaction where the second byte tells the slave to reset and then store this value in the programmable part of their slave address.
0	04H	2-byte transaction where the second byte tells the slave to store this value in the programmable part of their slave address. No reset.
0	00H	Not allowed as a second byte

When directed to reset, the  $I^2C$  Bus Interface Unit returns to its default reset condition with the exception of the ISAR. The 80333 is responsible for ensuring this occurs, not the  $I^2C$  Bus Interface Unit hardware.

When B=1, the sequence is used as a hardware general call by hardware masters only they cannot transmit a slave address, only their own address. The I<sup>2</sup>C Bus Interface Unit does not support this mode of operation.

I<sup>2</sup>C 10-bit addressing and CBUS compatibility are not supported.

# intel

# **10.5 Slave Mode Programming Examples**

# **10.5.1** Initialize Unit

- 1. Write ISAR: Set slave address
- 2. Write ICR: Enable all interrupts, set Unit Enable

# 10.5.2 Write 1 Byte as a Slave

- Wait for Slave Address Detected interrupt. Read ISR: Slave Address Detected (set), Unit Busy (set), R/W# bit (1), Ack/Nack (Clear - Ack)
- 2. Write IDBR: Load data byte to transfer
- 3. Write ICR: Set Transfer Byte bit
- 4. Wait for IDBR Transmit Empty interrupt. Read ISR: IDBR Transmit Empty (set), Ack/Nack (set - indicates last byte write), R/W# bit (0)
- 5. Clear interrupt by clearing the IDBR Transmit Empty Interrupt bit.
- Wait for interrupt. Read ISR: Unit Busy (clear), Slave STOP Detected (set)
- 7. Clear interrupt by clearing Slave STOP Detected Interrupt bit.

# 10.5.3 Read 2 Bytes as a Slave

- Wait for Slave Address Detected interrupt. Read ISR: Slave Address Detected (set), Unit busy (set), R/W# bit (0)
- Read byte 1 on I<sup>2</sup>C bus Write ICR: Set Transfer Byte bit to initiate the transfer
- Wait for interrupt. Read ISR: IDBR Receive Full (set), Ack/Nack (clear), R/W# bit (0) Clear interrupt by clearing IDBR Receive Full bit. Read IDBR: To get the data.
- Read byte 2 on I<sup>2</sup>C bus Write ICR: Set Transfer Byte bit to initiate the transfer
- Wait for interrupt.
   Read ISR: IDBR Receive Full (set), Ack/Nack (clear), R/W# bit (0) Clear interrupt by clearing IDBR Receive Full bit.
   Read IDBR: To get the data.
   Write ICR: Set Transfer Byte bit (to release I<sup>2</sup>C bus allowing next transfer)
- Wait for interrupt.
   Read ISR: Unit busy (clear), Slave STOP Detected (set) Clear interrupt by clearing Slave STOP Detected bit.

# **10.6 Master Programming Examples**

# 10.6.1 Initialize Unit

- 1. Write ISAR: Set slave address
- 2. Write ICR: Enable all interrupts (except Arb Loss), set SCL Enable, set Unit Enable

# 10.6.2 Write 1 Byte as a Master

- 1. Write IDBR: Target slave address and R/W# bit (0 for write)
- 2. Write ICR: Set START bit, Clear STOP bit, Set Transfer Byte bit to initiate the access
- 3. Wait for IDBR Transmit Empty interrupt. When interrupt arrives: Read status register: IDBR Transmit Empty (set), Unit Busy (set), R/W# bit (clear) Clear IDBR Transmit Empty Interrupt bit to clear the interrupt.
- *Note:* Arbitration Loss Detected bit may be set. When arbitration was lost, because Arb Loss interrupt was disabled, an address retry occurs when bus becomes free. Clear Arbitration Loss Detected bit when set.
  - 4. Send byte with STOP Write IDBR: With data byte to send Write ICR: Clear START bit, Set STOP bit, Enable Arb Loss interrupt, Set Transfer Byte bit to initiate the access
  - 5. Wait for Buffer empty interrupt. When interrupt arrives (Note: Unit is sending STOP): Read status register: IDBR Transmit Empty (set), Unit busy (set - maybe), R/W# bit (clear) Clear IDBR Transmit Empty Interrupt bit to clear the interrupt. Clear ICR STOP bit (optional) Wait until Unit busy is clear before clearing the ICR SCL Enable bit.

# 10.6.3 Read 1 Byte as a Master

- 1. Write IDBR: Target slave address and R/W# bit (1 for read)
- 2. Write ICR: Set START bit, Clear STOP bit, Disable Arb loss interrupt, Set Transfer Byte bit to initiate the access
- 3. Wait for IDBR Transmit Empty interrupt. When interrupt arrives: Read status register: IDBR Transmit Empty (set), Unit busy (set), R/W# bit (set) Clear IDBR Transmit Empty bit to clear the interrupt.
- Read byte with STOP Write ICR: Clear START bit, Set STOP bit, Enable arb loss interrupt, Set Ack/Nack bit (Nack), Set Transfer Byte bit to initiate the access
- 5. Wait for Buffer full interrupt. When interrupt arrives (Note: Unit is sending STOP): Read status register: IDBR Receive Full (set), Unit Busy (set - maybe), R/W# bit (Set), Ack/Nack bit (Set) Clear IDBR Receive Full bit to clear the interrupt. Read IDBR data. Clear ICR STOP bit (optional), Clear ICR Ack/Nack Control bit (optional) Wait until Unit busy is clear before clearing the ICR SCL Enable bit. (optional)

# 10.6.4 Write 2 Bytes and Repeated Start Read 1 Byte as a Master

- 1. Write IDBR: Target slave address and R/W# bit (0 for write)
- 2. Write ICR: Set START bit, Clear STOP bit, Set Transfer Byte bit to initiate the access
- 3. Wait for IDBR Transmit Empty interrupt. When interrupt arrives: Read status register: IDBR Transmit Empty (set), Unit busy (set), R/W# bit (clear) Clear IDBR Transmit Empty bit to clear the interrupt.
- Send byte 1 Write IDBR: With data byte to send Write ICR: Clear START bit, Clear STOP bit, Enable Arb Loss interrupt, Set Transfer Byte bit to initiate the access
- 5. Wait for Buffer empty interrupt. Read status register: IDBR Transmit Empty (set), Unit busy (set), R/W# bit (clear) Clear IDBR Transmit Empty bit to clear the interrupt.
- Send byte 2 Write IDBR: With data byte to send Write ICR: Clear START bit, Clear STOP bit, Set Transfer Byte bit to initiate the access
- Wait for Buffer empty interrupt. Read status register: IDBR Transmit Empty (set), Unit busy (set), R/W# bit (clear) Clear IDBR Transmit Empty bit to clear the interrupt.
- Send repeated start as a master Write IDBR: Target slave address and R/W# bit (1 for read) Write ICR: Set START bit, Clear STOP bit, Disable Arb Loss interrupt, Set Transfer Byte bit the initiate the access
- 9. Wait for IDBR Transmit Empty interrupt. When interrupt comes. Read status register: IDBR Transmit Empty (set), Unit busy (set), R/W# bit (set) Clear IDBR Transmit Empty bit to clear the interrupt.
- Read byte with STOP Write ICR: Clear START bit, Set STOP bit, Enable arb loss interrupt, Set Ack/Nack bit (Nack), Set Transfer Byte bit to initiate the access
- 11. Wait for Buffer full interrupt. When interrupt comes (Note: Unit is sending STOP). Read status register: IDBR Receive Full (set), Unit busy (set - maybe), R/W# bit (Set), Ack/Nack bit (Set)
  Clear IDBR Receive Full bit to clear the interrupt. Read IDBR data.
  Clear ICR STOP bit (optional), Clear ICR Ack/Nack Control bit (optional)
  Wait until Unit busy is clear before clearing the ICR SCL Enable bit. (optional)



# **10.6.5** Read 2 Bytes as a Master - Send STOP Using the Abort

- 1. Write IDBR: Target slave address and R/W# bit (1 for read)
- 2. Write ICR: Set START bit, Clear STOP bit, Disable Arb loss interrupt, Set Transfer Byte bit to initiate the access
- 3. Wait for IDBR Transmit Empty interrupt. When interrupt comes. Read status register: IDBR Transmit Empty (set), Unit Busy (set), R/W# bit (set) Clear IDBR Transmit Empty bit to clear the interrupt.
- Read byte 1 Write ICR: Clear START bit, Clear STOP bit, Enable Arb Loss interrupt, Clear Ack/Nack bit (Ack), Set Transfer Byte bit to initiate the access
- 5. Wait for Buffer full interrupt. Read status register: IDBR Receive Full (set), Unit busy (set), R/W# bit (Set), Ack/Nack bit (Clear) Clear IDBR Receive Full bit to clear the interrupt. Read IDBR data.
- Read byte 2 with Nack (STOP is not set because STOP or Repeated START is decided on the byte read)
   Write ICR: Clear START bit, Clear STOP bit, Enable Arb Loss interrupt, Set Ack/Nack bit (Nack), Set Transfer Byte bit to initiate the access
- 7. Wait for Buffer full interrupt. Read status register: IDBR Receive Full (set), Unit Busy (set), R/W# bit (Set), Ack/Nack bit (Set)
  Clear IDBR Receive Full bit to clear the interrupt. Read IDBR data.

There are now two options based on the byte read:

- Send a repeated START
- Send a STOP only

Here, a STOP abort is sent.

- Note: Had a NACK not been sent, the next transaction *must* involve another data byte read.
  - 8. Send STOP abort condition. (STOP with no data transfer.) Write ICR: Set Master abort.

# **10.7 Glitch Suppression Logic**

The I<sup>2</sup>C Bus Interface Unit has built-in glitch suppression logic. Glitches is suppressed according to:  $2 * I^2C$  clock period. For example, with the 33 MHz (30. ns period) I<sup>2</sup>C clock glitches of 60ns or less is suppressed. This is within the 50 ns glitch suppression specified.

# **10.8 Reset Conditions**

The I<sup>2</sup>C unit is reset with **I\_RST#**. Software is responsible for ensuring the I<sup>2</sup>C unit is not busy (ISR[3]) before asserting reset. Software is also responsible for ensuring the I<sup>2</sup>C bus is idle when the unit is enabled after reset. When directed to reset, the I<sup>2</sup>C unit returns to its default reset condition with the exception of the ISAR. ISAR is not affected by a reset.

When the Unit Reset bit in the ICR is set, only the 80333  $I^2C$  unit resets, the associated  $I^2C$  MMRs remain intact. When resetting the  $I^2C$  unit with the ICR units reset, use the following guidelines:

- 1. In the ICR register, set the reset bit and clear the remainder of the register.
- 2. Clear the ISR register.
- 3. Clear reset in the ICR.

# 10.9 Register Definitions

The following registers are associated with the  $I^2C$  Bus Interface Units. Each  $I^2C$  Bus Interface Unit has five memory-mapped control registers for independent operation. In register titles, x is 0 or 1 for unit 0 or 1, respectively.

They are all located within the peripheral memory- mapped address space of the 80333. See Section 5.8, "Register Definitions" on page 370 for the register addresses

## Table 337.I<sup>2</sup>C Register Summary

Section, Register Name, Acronym, Page
Section 10.9.1, "I <sup>2</sup> C Control Register x - ICRx" on page 626
Section 10.9.2, "I <sup>2</sup> C Status Register x - ISRx" on page 628
Section 10.9.3, "I <sup>2</sup> C Slave Address Register x - ISARx" on page 630
Section 10.9.4, "I <sup>2</sup> C Data Buffer Register x - IDBRx" on page 631
Section 10.9.5, "I <sup>2</sup> C Bus Monitor Register x - IBMRx" on page 632



# 10.9.1 I<sup>2</sup>C Control Register x - ICRx

The 80333 uses the bits in the I<sup>2</sup>C Control Register (ICRx) to control the I<sup>2</sup>C unit.



IOP       31       28       24       20       16       12       8       4       0         Attributes       / rv				
Unit #Intel XScale® Core internal bus addressAttribute Legend:RW = Read/Write0FFFF F680HRV = ReservedRC = Read Clear1FFFF F6A0HPR = PreservedRO = Read Only				
Dit	Defeult			
		Description		
31:16	0000H	Reserved		
15	0	0 = 100 KBit/sec operation 1 = 400 KBit/sec operation		
		Unit Reset:		
14	02	0 = No reset. 1 = Reset the I <sup>2</sup> C unit only.		
		Slave Address Detected Interrupt Enable:		
13	02	<ul> <li>0 = Disable interrupt.</li> <li>1 = Enables the l<sup>2</sup>C unit to interrupt the 80333 upon detecting a slave address match or a general call address.</li> </ul>		
		Arbitration Loss Detected Interrupt Enable:		
12	02	<ul> <li>0 = Disable interrupt.</li> <li>1 = Enables the I<sup>2</sup>C unit to interrupt the 80333 upon losing arbitration while in master mode.</li> </ul>		
11	0	Slave STOP Detected Interrupt Enable:		
11	02	<ul> <li>0 = Disable interrupt.</li> <li>1 = Enables the I<sup>2</sup>C unit to interrupt the 80333 when it detects a STOP condition while in slave mode.</li> </ul>		
		Bus Error Interrupt Enable:		
		0 = Disable interrupt. 1 = Enables the $l^2C$ unit to interrupt the 80333 for the following $l^2C$ bus errors:		
10	02	<ul> <li>As a master transmitter, no Ack was detected after a byte was sent.</li> </ul>		
		<ul> <li>As a slave receiver, the I<sup>2</sup>C unit generated a Nack pulse.</li> <li>NOTE: Software is responsible for guaranteeing that misplaced START and STOP conditions do not occur. See Section 10.7, "Glitch Suppression Logic" on page 625.</li> </ul>		
		IDBR Receive Full Interrupt Enable:		
09	02	0 = Disable interrupt. 1 = Enables $I^2C$ unit to interrupt the 80333 when the IDBR has received a data byte from the $I^2C$ bus.		
		IDBR Transmit Empty Interrupt Enable:		
08	02	0 = Disable interrupt. 1 = Enables the $l^2C$ unit to interrupt the 80333 after transmitting a byte onto the $l^2C$ bus.		
		General Call Disable:		
07	02	<ul> <li>0 = Enables the I<sup>2</sup>C unit to respond to general call messages.</li> <li>1 = Disables I<sup>2</sup>C unit response to general call messages as a slave.</li> </ul>		
		This bit must be set when sending a master mode general call message from the I <sup>2</sup> C unit.		



Table 3	38. I <sup>2</sup> C C	ontrol Register x - ICRx (Sheet 2 of 2)	
31       28       24       20       16       12       8       4       0         IOP Attributes			
Unit # Intel XScale <sup>®</sup> Cor 0 FFFF F680H 1 FFFF F6A0H		cale® Core internal bus addressAttribute Legend: RV = ReservedRW = Read/Write RC = Read Clear580H 5A0HPR = Preserved RS = Read/SetRO = Read Only NA = Not Accessible	
Bit	Default	Description	
06	02	$I^2C$ <b>Unit Enable</b> : 0 = Disables the unit and does not master any transactions or respond to any slave transactions. 1 = Enables the $I^2C$ unit (defaults to slave-receive mode). Software must guarantee the $I^2C$ bus is idle and that before setting this bit.	
05	02	SCL Enable:0 = Disables the I <sup>2</sup> C unit from driving the SCL line.1 = Enables the I <sup>2</sup> C clock output for master mode operation.	
04	02	Master Abort: used by the I <sup>2</sup> C unit when in master mode to generate a STOP without transmitting another data byte.         0 = The I <sup>2</sup> C unit transmits STOP using the STOP ICR bit only.         1 = The I <sup>2</sup> C unit sends STOP without data transmission.         When in Master transmit mode, after transmitting a data byte, the ICR's Transfer Byte bit is clear and IDBR Transmit Empty bit is set. When no more data bytes need to be sent, setting master abort bit sends the STOP. The Transfer Byte bit (03) must remain clear.         In master-receive mode, when a Nack is sent without a STOP (STOP ICR bit was not set) and the 80333 does not send a repeated START, setting this bit sends the STOP. Once again, the Transfer Byte bit (03) must remain clear.	
03	02	0 <sub>2</sub> Transfer Byte: used to send/receive a byte on the I <sup>2</sup> C bus. 0 = Cleared by I <sup>2</sup> C unit when the byte is sent/received. 1 = Send/receive a byte. The 80333 can monitor this bit to determine when the byte transfer has completed. In master or slave mode, after each byte transfer including Ack/Nack bit, the I <sup>2</sup> C unit holds the SCL line low (inserting wait states) until the Transfer Byte bit is set.	
02	02	Ack/Nack Control: defines the type of Ack pulse sent by the $I^2C$ unit when in master receive mode. $0 = \text{The } I^2C$ unit sends an Ack pulse after receiving a data byte. $1 = \text{The } I^2C$ unit sends a negative Ack (Nack) after receiving a data byte. The $I^2C$ unit automatically sends an Ack pulse when responding to its slave address or when responding in slave-receive mode, independent of the Ack/Nack control bit setting.	
01	02	<ul> <li>STOP: used to initiate a STOP condition after transferring the next data byte on the I<sup>2</sup>C bus when in master mode. In master-receive mode, the Ack/Nack control bit must be set in conjunction with this bit. See Section 10.3.3.3, "STOP Condition" on page 607 for more details on the STOP state.</li> <li>0 = Do not send a STOP.</li> <li>1 = Send a STOP.</li> </ul>	
00	02	<ul> <li>START: used to initiate a START condition to the I<sup>2</sup>C unit when in master mode. See Section 10.3.3.1, "START Condition" on page 607 for more details on the START state.</li> <li>0 = Do not send a START.</li> <li>1 = Send a START.</li> </ul>	



# 10.9.2 I<sup>2</sup>C Status Register x - ISRx

 $I^2C$  interrupts are signalled to the 80333 interrupt controller by the  $I^2C$  Interrupt Status Register (ISRx). Software uses the ISR bits to check the status of the  $I^2C$  unit and bus. ISRx bits (bits 9-5) are updated after the Ack/Nack bit has completed on the  $I^2C$  bus.

The ISRx is also used to clear interrupts signalled from the I<sup>2</sup>C Bus Interface Unit. These are:

- IDBRx Receive Full
- IDBRx Transmit Empty
- Slave Address Detected
- Bus Error Detected
- STOP Condition Detect
- Arbitration Lost

## Table 339. I<sup>2</sup>C Status Register x - ISRx (Sheet 1 of 2)

31       28       24       20       16       12       8       4       0         IOP Attributes			
Unit # 0	Intel XS	Cale® Core internal bus addressAttribute Legend:RW = Read/Write84HRV = ReservedRC = Read Clear84HPR = PreservedRO = Read Only	
1		RS = Read/Set NA = Not Accessible	
Bit	Default	Description	
31:11	000000H	Reserved	
10	02	<ul> <li>Bus Error Detected:</li> <li>0 = No error detected.</li> <li>1 = The I<sup>2</sup>C unit sets this bit when it detects one of the following error conditions:</li> <li>As a master transmitter, no Ack was detected on the interface after a byte was sent.</li> <li>As a slave receiver, the I<sup>2</sup>C unit generates a Nack pulse.</li> <li>NOTE: When an error occurs, I<sup>2</sup>C bus transactions continue. Software must guarantee that misplaced START and STOP conditions do not occur. See Section 10.4.4, "Arbitration" on page 612.</li> </ul>	
09	02	<ul> <li>Slave Address Detected:</li> <li>0 = No slave address detected.</li> <li>1 = I<sup>2</sup>C unit detected a 7-bit address that matches the general call address or ISAR. An interrupt is signalled when enabled in the ICR.</li> </ul>	
08	02	General Call Address Detected:         0 = No general call address received.         1 = 1 <sup>2</sup> C unit received a general call address.	
07	02	<ul> <li>IDBR Receive Full:</li> <li>0 = The IDBR has not received a new data byte or the I<sup>2</sup>C unit is idle.</li> <li>1 = The IDBR register received a new data byte from the I<sup>2</sup>C bus. An interrupt is signalled when enabled in the ICR.</li> </ul>	









# **10.9.3** I<sup>2</sup>C Slave Address Register x - ISARx

The I<sup>2</sup>C Slave Address Register (ISARx) (see Table 340) defines the I<sup>2</sup>C unit 7-bit slave address to which the 80333 responds when in slave-receive mode. This register is written by the 80333 before enabling I<sup>2</sup>C operations. The register is fully programmable (no address is assigned to the I<sup>2</sup>C unit) so it can be set to a value other than those of hard-wired I<sup>2</sup>C slave peripherals that might exist in the system. The ISAR is not affected by the 80333 being reset. The ISAR register default value is  $0000000_2$ .



## Table 340. I<sup>2</sup>C Slave Address Register x - ISARx

# **10.9.4** I<sup>2</sup>C Data Buffer Register x - IDBRx

The I<sup>2</sup>C Data Buffer Register (IDBRx) is used by the 80333 to transmit and receive data from the I<sup>2</sup>C bus. The accesses the IDBR by the 80333 on one side and by the I<sup>2</sup>C shift register on the other. Data coming into the I<sup>2</sup>C Bus Interface Unit is received into the IDBR after a full byte has been received and acknowledged. Data going out of the I<sup>2</sup>C Bus Interface Unit is written to the IDBRx by the Intel XScale<sup>®</sup> core and sent to the serial bus.

When the I<sup>2</sup>C Bus Interface Unit is in transmit mode (master or slave), the 80333 writes data to the IDBRx over the internal bus. This occurs when a master transaction is initiated or when the IDBRx Transmit Empty Interrupt is signalled. Data is moved from the IDBRx to the shift register when the Transfer Byte bit is set. The IDBR Transmit Empty Interrupt is signalled (when enabled) when a byte has been transferred on the I<sup>2</sup>C bus and the acknowledge cycle is complete. When the IDBRx is not written by the 80333 (and a STOP condition was not in place) before the I<sup>2</sup>C bus is ready to transfer the next byte packet, the I<sup>2</sup>C Bus Interface Unit inserts wait states until the Intel XScale<sup>®</sup> core writes the IDBRx and sets the Transfer Byte bit.

When the I<sup>2</sup>C Bus Interface Unit is in receive mode (master or slave), the processor reads IDBRx data over the internal bus. This occurs when the IDBRx Receive Full Interrupt is signalled. The data is moved from the shift register to the IDBRx when the Ack cycle is complete. The I<sup>2</sup>C Bus Interface Unit inserts wait states until the IDBR has been read. Refer to Section 10.4.3, "I<sup>2</sup>C Acknowledge" on page 611 for acknowledge pulse information in receiver mode. After the 80333 reads the IDBRx, the Ack/Nack Control bit is written and the Transfer Byte bit is written, allowing the next byte transfer to proceed on the I<sup>2</sup>C Bus. The IDBRx register is 00H after reset.

### Table 341. I<sup>2</sup>C Data Buffer Register x - IDBRx





# 10.9.5 I<sup>2</sup>C Bus Monitor Register x - IBMRx

The I<sup>2</sup>C Bus Monitor Register (IBMRx) tracks the status of the SCL and SDA pins. The values of these pins are recorded in this read-only register so that software may determine when the I<sup>2</sup>C bus is hung and the I<sup>2</sup>C unit must be reset.



## Table 342.I<sup>2</sup>C Bus Monitor Register x - IBMRx

# intel

# SMBus Interface Unit

# 11

This chapter describes the System Management Bus (SMBus) interface unit, including the operation modes and setup. Throughout this manual, this peripheral is referred to as the SMBus unit.

# 11.1 Overview

The SMBus Interface Units allows the Intel<sup>®</sup> 80333 I/O processor (80333) to serve as a slave device residing on the SMBus. The SMBus is a two-pin interface. **SDTA** is the data pin for input and output functions and **SCLK** is the clock pin for reference and control of the SMBus.

The SMBus allows the system to interface to 80333 for system management functions. The serial bus requires a minimum of hardware for an economical system to relay status and reliability information of the 80333 to the system.

The SMBus Interface Unit is a peripheral device that resides on a 80333 internal bus. Data is transmitted to and received from the SMBus via a buffered interface. Control and status information is relayed through a set of memory-mapped registers. Refer to the SMBus Specification for complete details on SMBus operation.

# 11.2 SMBus Interface

SMBus provides for full access to registers in 80333 bridges, IOAPICs and SHPC, including configuration and memory registers. 80333 has no internal I/O registers. Systems so configured can use the SMBus to access the registers. 80333 supports a slave-only SMBus mode.

- System Management Bus Specification, Revision 2.0 (SMBus) Compliant.
- Slave mode operation only.
- Full read/write access to configuration and memory spaces in 80333 bridges, IOAPICs and SHPC.

The SMBus interface is multiplexed with one  $I^2C$  interface ( $I^2C$  ch 1), and is by default enabled for SMBus operation. To change the interface to  $I^2C$  requires both the SMBus interface to be disabled in the "SMBus Enable Register - SMBER" on page 650 and the  $I_2C$  to be enabled in the Section 10.9.1, " $I^2C$  Control Register x - ICRx" on page 626.

#### Table 343.SMBus Interface Pins

Signal	Pad Type	
SCL1/SCLK	SCL1/SCLK SMBus Clock: Provides synchronous operation of the SMBus.	
SDA1/SDTA	<b>SDA1/SDTA</b> SMBus Data: Is used for data transfer and arbitration of the SMBus.	
Total	2	



# **11.3 System Management Bus Interface**

This interface has no configuration registers associated with it. The SMBus address is set upon PWRGD by sampling the Peripheral Bus Interface Reset Strap inputs A[19:16]. When the pins are sampled, the resulting 80333 address will be stored in the Reset Strap Status Register and assigned as follows:

Bit	Value
7	'1'
6	'1'
5	A[19]
4	ʻ0'
3	A[18]
2	A[17]
1	A[16]

The values for these bits are stored in. The SMBus controller has access to all internal registers of the PCI Express-to-PCI Bridges and associated IOAPICs and Hot-Plug Controller. It also physically allows accesses to PCI through Type 1 cycles, but this is not a supported usage model and may not function properly. It can perform reads and writes from all registers through the particular interface configuration space. IOAPIC memory space is accessible through its configuration space. SHPC memory space is directly accessible from the SMBus controller via the SMBus Memory Command.



# 11.3.1 SMBus Controller

The 80333 SMBus slave port interfaces to the configuration and memory spaces of each bridge, the IOAPICs, and the Hot-Plug unit. This will give SM (server management) visibility into configuration and memory registers in the 80333 bridges, IOAPICs and SHPC.

## 11.3.1.1 SMBus Commands

80333 supports six SMBus commands:

- Block Write
   Word Write
   Byte Write
- Block Read
   Word Read
   Bytes Read

Sequencing these commands will initiate internal accesses to 80333 configuration and memory registers. For high reliability, 80333 also supports the optional Packet Error Checking feature (CRC-8) and is enabled or disabled with each transaction.

Every configuration and memory read or write first consists of an SMBus write sequence which initializes the Bus Number, Device, function number, memory address offset etc. The term sequence is used since these variables can be initialized by the SMBus master with a single block write or multiple word or byte writes. The last write in the sequence that completes the initialization performs the internal configuration/memory read or write. The SMBus master can then initiate a read sequence which returns the status of the internal read or write command and also the data in case of a read.

Each SMBus transaction has an 8-bit command driven by the master. The command encodes the following information:

#### Table 344.SMBus Command Encoding

Bit	Description
7	Begin: The Begin bit when set indicates the first transaction of the read or write sequence.
6	End: The End bit when set indicates the last transaction of the read or write sequence.
5	Memory/Configuration: Indicate whether memory or configuration space is being accesses in this SMBus sequence. Value of '1' indicates memory and a value of '0' indicate configuration.
4	PEC Enable: Indicates that PEC is enabled when set. When set, each transaction in the sequence ends with an extra CRC byte. 80333 would check for CRC on writes and generate CRC on reads. PEC does not include the command byte itself.
3:2	Internal Command: 00 - Read DWord 01 - Write Byte 10 - Write Word 11 - Write Dword All access are naturally aligned to the access width. This field specifies the internal command to be issued by the SMBus slave logic to the 80333 core.
1:0	SMBus command: 00 - Byte 01 - Word 10 - Block 11 - Reserved This field specifies the SMBus command to be issued on the SMBus. This field is used as an indication of the length of transfer so that the slave knows when to expect the PEC packet (when enabled).



## 11.3.1.2 Initialization Sequence

All Configuration and memory read and writes are accomplished through an SMBus write(s) and later followed by an SMBus read (for a read command). The SMBus write sequence is used to initialize the for the configuration access:

- Bus Number,
- Device/Function and
- 12-bit Register Number (in 2 separate bytes on SMBus)

Each of the parameters above is sent on SMBus in separate bytes. The register number parameter is initialized with two bytes and 80333 ignores the most significant 4 bits of the second byte that initializes the register number. For memory reads and writes, the write sequence initializes:

- Destination memory
- 24-bit memory address offset (in 3 separate bytes on SMBus)

The initialization of the information can be accomplished through any combination of the supported SMBus write commands (Block, Word or Byte). The Internal Command field for each write should specify the same internal command every time (read or write). After all the information is set up, the last write (End bit is set) initiates an internal read or write command. On an internal read when the data is not available before the slave interface acknowledges this last write command (ACK), the slave will "clock stretch" until the data returns to the SMBus interface unit. On a internal write, when the write is not complete before the slave interface acknowledges this last write command (ACK), the salve will "clock stretch" until the write completes internally. When an error occurs (internal time-out, target or master abort on the internal switch) during the internal access, the last write command will receive a NACK.

# 11.3.2 SMBus Signaling

## 11.3.2.1 **Overview**

The SMBus interface includes a pair of signals: **SCLK** (clock) and **SDTA** (serial data). **SCLK** provides the timing mechanism for data transfers. The SMBus master always drives **SCLK**. **SDTA** carries the data, as driven by the sending device (sender), which can be the initiator or the target.

An initiator starts a transfer over the SMBus when it is free. Details of how initiators arbitrate are not described here. The current initiator communicates to the desired target through a unique 7-bit address to the target, sent MSb to LSb. All devices monitor the generated address after detecting the start condition. Once seven address bits are received, all targets compare the received address with their own and the target slave finds a match.

The next data bit from the initiator indicates the transfer direction. A value of '1' indicates that the target needs to transfer data to the initiator (read). Data transfers over SMBus are performed in 8-bit chunks. Data is transferred from MSb to LSb.

## 11.3.2.2 Waveforms

The timing relationship between **SDTA** and **SCLK** is defined such as the **SDTA** value must be valid through the duration of **SCLK** being in High state. The following diagram illustrates data transfer:

### Figure 118. Basic SMBus Transfer Waveform



#### 11.3.2.2.1 Start Phase

A start condition is generated when SMBus is idle to indicate that its state is changing to busy. The start condition occurs when **SDTA** transitions from High to Low while **SCLK** remains High. The SMBus protocol also allows a master to "Repeat Start", meaning that a new transfer is started by the same master, without a stop condition.

#### Figure 119. Start (S) / Repeat Start (Sr) Signaling





#### 11.3.2.2.2 Stop Phase

A stop condition is generated when SMBus is busy to indicate that its state is changing to idle. The Stop condition occurs when **SDTA** transitions from Low to High while **SCLK** remains High.

#### Figure 120. Stop (P) Signaling

SCLK	
SDTA	

A stop bit can occur at any point in a data stream. It is not guaranteed to occur after an ACK from a target (as later waveforms will show). The 80333 must be able to accept a stop condition at any time and clean up.

## 11.3.2.2.3 ACK/NACK

For every 8 bits of data transfer (including address and direction), the receiving agent must respond with ACK or NACK. An ACK is requires **SDTA** = 0 during **SCLK** = 1. A NACK requires **SDTA** = 1 during **SCLK** = 1 as shown below.

#### Figure 121. ACK (A) Signaling



### Figure 122. NACK (N) Signaling

SCLK	
SDTA	

During a write cycle, the 80333 must drive an ACK after the address/direction phase, and after the data phase. During a read cycle, the 80333 must drive an ACK/NACK after the address/direction phase, and (when ACKed) the initiator must drive an ACK/NACK after the 80333 returns its 8 bits of data.

#### 11.3.2.2.4 Wait States

The receiver (initiator or target) can add wait states, after driving ACK for receiving the last byte, by driving the **SCLK** line low. Further data transfers are delayed until the receiver stops driving **SCLK** low. It is expected the 80333 will drive the **SCLK** line low after receiving data on writes until the write is complete, and after receiving the direction bit on reads until the read data is ready.

# 11.3.3 Data Transfer Examples

For Figure 124 through Figure 134, the following terminology is used:

- S Start Bit
- Sr Start Repeat Bit
- W Write Command
- R Read Command
- A Acknowledge
- N Retry / not Acknowledge
- P Stop Bit

Clear boxes indicate phases of the cycle driven by the initiator, and shaded boxes indicate phases of the cycle driven by the target.

In Figure 123, an initiator targets the 80333 with a write, which retries the cycle. Since all read cycles will first be set up with a write to the register index stack pointer, the 80333 will only be busy on a write cycle. For example, when a command had been written indicating a DWORD read from a bus/dev-func/reg number, the 80333 will be busy doing the read when the subsequent write comes into the register index.

## Figure 123. Intel<sup>®</sup> 80333 I/O Processor Busy





## 11.3.3.1 Configuration and Memory Reads

80333 supports only read dword to internal register space. All Configuration and memory reads are accomplished through an SMBus write(s) and later followed by an SMBus read to read the status and the read data. For SMBus read transactions, the last byte of data (or the PEC byte when enabled), is NACKed by the master to indicate the end of the transaction. The SMBus memory read command returns the status of the previous internal command and the data associated previous internal read command. The status field encoding is:

#### Table 345. SMBus Status Byte Encoding

Bit	Description
7	Internal Time-out. This bit is set when an SMBus Request is not completed in 2ms internally.
6	Reserved
5	Internal Master Abort
4	Internal Target Abort
3:1	Reserved
0	Successful

Examples of configuration and memory reads are shown in Figure 124 through Figure 129. For the definition of the diagram conventions below, refer to the *System Management Bus Specification*, Revision 2.0.

#### Figure 124. DWORD Configuration Read Protocol (SMBus Block Write/Block Read, PEC Enabled)



## Figure 125. DWORD Memory Read Protocol (SMBus Block Write/Block Read, PEC Enabled)





### Figure 126. DWORD Configuration Read Protocol (SMBus Word Write/Word Read, PEC Enabled)

S	11X0_XXX	W A	Cmd = 10010001	А	Bus Number	А	Device/Function	A	PEC	AP
	<b>`</b>	111	<b>`</b>	11	<b>`</b>	11		11		
S	11X0_XXX	W A	Cmd = 01010001	А	Register Num[15:8]	Α	Register Num[7:0]	A	PEC	CLOCK STRETCH A P
$\Box$	<b>`</b>	111		1	· · · ·	$\sim$		<u>H</u>		
S	11X0_XXX	WA	Cmd = 10010001	Α	)					
Sr	11X0_XXX	RA	Status	Α	Data[31:24]	Α	PEC	NP		
$\sim$				~ `		11		) $)$ $)$		
S	11X0_XXX	WA	Cmd = 00010001	Α	)					
Sr	11X0_XXX	RA	Data[23:16]	Α	Data[16:8]	Α	PEC	NP		
$\sim$	<b>`</b>	111	<b>`</b>	1	<b>`</b>	11		111		
_										
S	11X0_XXX	W A	Cmd = 01010000	А			_			
Sr	11X0_XXX	RA	Data[7:0]	Α	PEC	NF	רב			
		11		1		11				

### Figure 127. DWORD Configuration Read Protocol (SMBus Block Write/Block Read, PEC Disabled)

Reg Number [7:0]         CLOCKSTREETCH A P           S         11X0_XXX         WA         Circl = 11000010         A	S 11X0_XXX WA Cmd=11000010	A Byte Count = 4	A Bus Number	A Device/Runction A F	ag Nurba(15.8)
S 11X0_XXX WA Cirid=11000010 A	Reg Number [7:0] CLOCK STREET ON A P				
<u>3° 11X0_XXX R A Byte-Count = 5 A Status A Data(31:24) A Data(23:16) A Data(158) A Data(7:0) NP</u>	S 11X0_XXX WA Crnd = 11000010 ≌ 11X0_XXX RA ByteCount = 5	A Status	A Data(31:24)	A Data(2316) A	Data(158) A Data(7:0) NP

### Figure 128. DWORD Memory Read Protocol (SMBus Block Write/Block Read, PEC Disabled)



#### Figure 129. DWORD Configuration Read Protocol (SMBus Word Write/Word Read, PEC Disabled)





## 11.3.3.2 Configuration and Memory Writes

Configuration and memory writes are accomplished through a series of SMBus writes. As with reads, a write sequence is first used to initialize the Bus Number, Device, Function, and Register Number for the configuration access and the destination memory, address offset for the memory write. The writing of this information can be accomplished through any combination of the supported SMBus write commands (Block, Word or Byte).

Note: On SMBus, there is no concept of byte enables. Therefore, the Register Number written to the slave is assumed to be aligned to the length of the Internal Command. In other words, for a Write Byte internal command, the Register Number specifies the byte address. For a Write DWord internal command, the two least-significant bits of the Register Number are ignored. This is different from PCI where the byte enables are used to indicate the byte of interest.

After all the information is set up, the SMBus master initiates one or more writes which sets up the data to be written. The final write (End bit is set) initiates an internal configuration or memory write. The slave interface could potentially clock stretch the last data write until the write completes without error. When an error occurred, the SMBus interface NACKs the last write operation just before the stop bit. Examples of configuration writes are illustrated below. All the figures are with PEC Enabled. When PEC is disabled, there is no PEC byte in any of the sequences and the PEC enable bit in the command field is 0.

For the definition of the diagram conventions below, refer to the *System Management Bus Specification*, Revision 2.0.

#### Figure 130. DWORD Configuration Write Protocol (SMBus Block Write, PEC Enabled)

s	11XQ_XXX	MA	2md - 11011110	Ą	Byte Count - 8	Ą	BusNumber	A	DesiceFundion	A	Reg Ninter(154)	A	Reg Number (74)	A	Data(31:24)	Ŋ
A	Dalaį2316j	Ą	Cata(168)	A	Data(7:0)	A	REC .	0								_

#### Figure 131. DWORD Memory Write Protocol (SMBus Word Write, PEC Enabled)

S 11X0_XXX	W A Cmd = 1011110	01 A Dest Mem	A Add Offset[23:16]	A PEC	AP
S 11X0_XXX	W A Cmd = 001111	01 A Add Onsel[15	8] A Add Offset[7:0]	A PEC	AP
S 11X0_XXX	W A Cmd = 00111	01 A Data(31:24	4 Data[23:16]	A PEC	AP
S 11X0_XXX	WA Crind=01111	101 A Data[15:8	n A Data(7:0)	A PEC	CLOCK STRETCH A P



Figure 132.	<b>DWORD Confid</b>	guration Write	<b>Protocol (S</b>	SMBus By	te Write,	PEC Enabled)



Figure 133. DWORD Memory Write Protocol (SMBus Word Write/(Word, Byte) Read, PEC Enabled)

S 11X0_XXX	WA Cand = 101100	1 A DestMem	A Add Offset[23:16] A	PEC A P
s 11X0_XXX	W A Circl =011100	01 A Add Offsel [15:8	Add Offset[7:0] A	PEC CLOCK STRETCH A P
S 11X0_XXX	W A Cand = 1011000	л А		_
a 11X0_XXX	R A Status	A Data[31:24]		<u>-</u>
S 11X0_XXX	W A Crnd = 0011000	л А		~
a 11x0_x0x	R A Data[23:16]	A Data[15:8]		<u> </u>
S 11X0_XXX	WA Crnd =011100	00 A.		
3 11X0_XXX	RA Data[7:0]	A PEC	NP	

Figure 134. DWORD Memory Read Protocol (SMBus Word Write/Byte Read, PEC Enabled)



# 11.3.4 Error Handling

The SMBus slave interface handles two types of errors: internal and PEC. Internal errors can occur for example when the SMBus tries to access the APIC or SHPC configure/memory space and these units in 80333 are stuck servicing the PCI Express interface which is broke. 80333 internally times out in such a case and this error manifests itself as a Not-Acknowledge (NACK) for the read or write command (End bit is set). Other internal errors include the read or write command receiving a master or target abort on the internal interface and these conditions also cause a NACK on the SMBus. When the master receives a NACK, the entire transaction should be reattempted.

When the master supports packet error checking (PEC) and the PEC enable bit in the command is set, then the PEC byte is checked in the slave interface. When the check indicates a failure, then the slave will NACK the PEC packet and not issue the command on the internal interface.

An SMBus master must either do PEC on all transactions in a sequence or not do it at all i.e. it cannot turn on PEC in the middle of a sequence.

A PEC error in the middle of a sequence must be re-started from the beginning of the sequence i.e. the begin bit set

# 11.3.5 SMBus Interface Reset

The master in two ways can reset the slave interface state machine in 80333:

The master holds **SCLK** low for 25ms cumulative. Cumulative in this case means that all the "low time" for **SCLK** is counted between the Start and Stop bit. When this totals 25ms before reaching the Stop bit, the interface is reset.

The master holds SCLK continuously high for 50ms.

Besides these, the SMBus interface in 80333 is also reset on a **PWRGD**, **RSTIN**# or an in-band warm reset from PCI Express.

# **11.3.6 Configuration Access Arbitration**

When the CPU is currently accessing a unit, SMBus cannot access it. Whoever gets in first wins arbitration. The other agent is stalled until the first agent finishes. The micro-architecture of this area is critical. The reason for the SMBus interface is to access registers when the system may be unstable or locked, which can result with broken queues. Any register access through SMBus must be able to proceed while the system is stuck.



# 11.4 Register Definitions

The following registers are associated with the SMBus Interface Unit. The SMBus Interface Unit has only one memory-mapped control register ("SMBus Enable Register - SMBER"), and all other registers are accessed in stack order through the SMBus interface.

#### Table 346. SMBus Register Summary Table

Section, Register Name, Acronym, Page
Section 11.4.2, "SMBus Controller Command - SM_CMD" on page 646
Section 11.4.3, "SMBus Controller Byte Count Register- SM_BC" on page 646
Section 11.4.4, "SMBus Controller Address Register 3 - SM_ADDR3" on page 647
Section 11.4.5, "SMBus Controller Address Register 2 - SM_ADDR2" on page 647
Section 11.4.6, "SMBus Controller Address Register 1 - SM_ADDR1" on page 647
Section 11.4.7, "SMBus Controller Address Register 0 - SM_ADDR0" on page 648
Section 11.4.8, "SMBus Controller Data Registers - SM_DATA" on page 648
Section 11.4.9, "SMBus Controller Status - SM_STS" on page 649
Section 11.4.10, "SMBus Enable Register - SMBER" on page 650

# 11.4.1 Register Summary

#### Table 347. Register Summary

Symbol	Configuration Description	Memory Access Description
SM_CMD	Command	Command
SM_BC	Byte Count	Byte Count
SM_ADDR3	Bus Number	Device Select
SM_ADDR2	Device / Function Number	Reserved - zero
SM_ADDR1	Extended RNUM register	Memory Address [11:8]
SM_ADDR0	Register Number	Memory Address [7:0]
SM_DATA3	Data Byte 3 [31:24]	Data Byte 3 [31:24]
SM_DATA2	Data Byte 2 [23:16]	Data Byte 2 [23:16]
SM_DATA1	Data Byte 1 [15:8]	Data Byte 1 [15:8]
SM_DATA0	Data Byte 0 [7:0]	Data Byte 0 [7:0]
SM_STS	Status (Reads only)	Status (Reads only)



# 11.4.2 SMBus Controller Command - SM\_CMD

When written, this is the Command Register. All configuration accesses from the SMBus port are initiated by writing to this register. While a command is in progress, all future writes or reads will be NACKd by the 80333 to avoid having registers overwritten.

### Table 348. SMBus Controller Command Register - SM\_CMD

Bit	Description
07	Begin: The Begin bit when set indicates the first transaction of the read or write sequence.
06	End: The End bit when set indicates the last transaction of the read or write sequence.
05	Memory/Configure: Indicate whether memory or configuration space is being accesses in this SMBus sequence. 0 = Configuration space 1 = Memory space
04	PEC Enable: Indicates that PEC is enabled when set. When set, each transaction in the sequence ends with an extra CRC byte. 80333 would check for CRC on writes and generate CRC on reads. PEC does not include the command byte itself.
03:02	Internal Command: 00 - Read DWord 01 - Write Byte 10 - Write Word 11 - Write Dword All access are naturally aligned to the access width. This field specifies the internal command to be issued by the SMBus slave logic to the 80333 core.
01:00	SMBus command:         00 - Byte         01 - Word         10 - Block         11 - Reserved         This field specifies the SMBus command to be issued on the SMBus. This field is used as an indication of the length of transfer so that the slave knows when to expect the PEC packet (when enabled).

# 11.4.3 SMBus Controller Byte Count Register- SM\_BC

This register should be programmed with the Byte Count of the transfer request.

 Table 349.
 SMBus Controller Byte Count Register - SM\_BC

Bit	Туре	Description
07:00	RW	Byte Count (BC): Indicates the byte count of the transfer.



# 11.4.4 SMBus Controller Address Register 3 - SM\_ADDR3

This register should be programmed with the Bus Number of the desired configuration register.

### Table 350. SMBus Controller Address Register 3- SM\_ADDR3

Bit	Description	
07:02	Reserved.	
01	<ul> <li>Bus Number (BNUM): Indicates the bus number to access for register access.</li> <li>BAR selection. bit encoding for memory BAR selection for memory accesses.</li> <li>00010 - B-Segment SHPC Memory Space</li> <li>All other values are reserved.</li> <li>The destination memory is a byte of information that indicates the internal memory space to access in the 80333. A '1' written in this bit location selects the SHPC B-segment memory space. The 24-bit address offset is used to address any internal memory with up to an offset of 24 bits. The 80333 only uses 12 bits of address, and ignores the most significant 12 bits of the 24-bit address. The 80333 slave interface always expects 24 bits of address from the SMBus master though it uses only 12 bits</li> </ul>	
00	Reserved.	

# 11.4.5 SMBus Controller Address Register 2 - SM\_ADDR2

This register should be programmed with the Device Number and Function Number of the desired configuration register. This is set to zero for memory accesses.

#### Table 351. SMBus Controller Address Register 2\_SM\_ADDR2

Bit	Description
07:03	Device Number (DEV): Device number of device to access.
02:00	Function Number (FNUM): Function number of device to access.

# 11.4.6 SMBus Controller Address Register 1 - SM\_ADDR1

This register should be programmed with the Extended Register Number of the desired configuration register, or the upper four bits of the memory address.

#### Table 352. SMBus Controller Address Register 1 - SM\_ADDR1

Bit	Description	
07:04	Reserved	
03:00	Number (XRNUM): Indicates the upper 4 address bits of the PCI Express enhanced configuration space. Register address [11:8] bits. Memory Address [11:8] (ADDR1).	



# 11.4.7 SMBus Controller Address Register 0 - SM\_ADDR0

This register should be programmed with the Register Number of the desired configuration register or the lower 8 bits of memory address.

#### Table 353. SMBus Controller Address Register 0 - SM\_ADDR0

Bit	Description
07:00	Number (RNUM): Indicates the register number to access. Register address [7:0] bits.
	Memory Address [7:0] (ADDR0)

# 11.4.8 SMBus Controller Data Registers - SM\_DATA

This register is used to read or write data to the desired Configuration Register.

At the completion of a Read command, this register will contain the data from the selected configuration register. For reads the data register will always return 32 bits and will always be aligned on a DWORD boundary.

Before issuing a write command this register should be written with the desired write data. For a byte, write only the Byte0 [7:0] data will be written to the desired configuration register. For a word write, only the Byte1 & Byte 0 [15:0] data will be written to the desired configuration register. The register number must be word aligned for word writes. For a DWORD write, all 32 bits of data will be used. The register number must be DWORD aligned.

The Status Register should be checked to make sure that there is not a command currently in progress, before writing to this register. Writing to this register when the Busy bit in the Status Register is asserted will have indeterminate effects.

#### Table 354. SMBus Controller Data Registers - SM\_DATA

Bit	Description
31:24	Byte 3 (B3): Data bits [31:24].
23:16	Byte 2 (B2): Data bits [23:16].
15:08	Byte 1 (B1): Data bits [15:8].
07:00	Byte 0 (B0): Data bits [7:0].


## 11.4.9 SMBus Controller Status - SM\_STS

On reads, this Status Register precedes the data.

#### Table 355. SMBus Controller Status Register - SM\_STS

Bit	Description
07	Internal Time-out - error condition detected during transaction access.
06	Reserved
05	Internal Master Abort - register access transaction resulted in a master abort.
04	Internal Target Abort - register access transaction resulted in a target abort.
03:01	Reserved
00	Successful - transaction completed successfully



## 11.4.10 SMBus Enable Register - SMBER

The SMBus Enable Register controls when the SMBus Unit is enabled on the multiplexed  $I^2C$  pins. The  $I^2C$  units have their own enable in the Section 10.9.1, " $I^2C$  Control Register x - ICRx" on page 626.

*Note:* This register is memory mapped to an address with the GPIO control registers and is read/write accessible by the Intel XScale<sup>®</sup> Core processor.



 Table 356.
 SMBus Enable Register - SMBER



## 12

This chapter describes the Universal Asynchronous Receiver/Transmitter (UART) serial ports. The Intel<sup>®</sup> 80333 I/O processor (80333) UARTs are controlled via programmed I/O through memory-mapped registers.

## 12.1 Overview

Each asynchronous serial port supports all the functions of a 16550 UART. Each UART performs serial-to-parallel conversion on data characters received from a peripheral device or a modem and parallel-to-serial conversion on data characters received from the processor. The processor can read the complete status of a UART at any time during the functional operation. Available status information includes the type and condition of the transfer operations being performed by a UART and any error conditions (parity, overrun, framing, or break interrupt).

Each serial port 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.

Each UART includes a programmable baud rate generator which is capable of dividing the input clock by divisors of 1 to  $(2^{16}-1)$  and producing a 16X clock to drive the internal transmitter and receiver logic. Interrupts can be programmed to the user requirements, minimizing the computing required to handle the communications link. Each UART operates in a polled or an interrupt driven environment which is selected by software.



The UART hardware is responsible for executing serial protocol communication and for providing the programming interface. The UART features include:

- Registers are compatible with the 16550 and 16750
- Adds or deletes standard asynchronous communications bits (start, stop, and parity) to or from the serial data
- Independently controlled transmit, receive, line status and data set interrupts
- Baud-rate generator allows division of clock by 1 to (2 16–1) and generates an internal 16X clock; baud-rate can be manually or automatically programmed via auto-baud-rate detection circuitry
- Modem control functions (CTS#, RTS#)
- Autoflow capability controls data I/O without generating Interrupts:
  - RTS# (output) controlled by UART Receiver FIFO
  - CTS# (input) from modem controls UART transmitter
- Fully programmable serial-interface characteristics:
  - 5, 6, 7 or 8-bit characters
  - Even, odd, or no parity detection
  - 1, 1-1/2, or 2 stop bit generation
  - Baud rate generation (up to 115kbps)
- False start bit detection
- 64-byte Transmit FIFO
- 64-byte Receive FIFO with programmable threshold
- Complete status reporting capability
- Break generation and detection
- Internal diagnostic capabilities include:
  - Loopback controls for communications link fault isolation
  - Break, parity, overrun, and framing error simulation
- Fully prioritized interrupt system controls

### 12.1.1 Compatibility with 16550 and 16750

The UARTs can be programmed to be functionally compatible with industry standard 16550 and 16750. Each UART supports most of the 16550 and 16750 functions and has additional features, as listed below.

- DMA requests for transmit and receive data services
- NRZ encoding/decoding function
- 64-byte Transmit/Receive FIFO buffers
- Programmable Receive FIFO threshold
- Auto baud-rate detection
- Auto flow

# intel

## 12.2 Signal Descriptions

The name and description of external signals connected to a UART module are shown in Table 357. These signals are multiplexed with the GPIO signals as specified in Chapter 18, "General Purpose I/O Unit". The selection between GPIO and UART function for these multiplexed pins is controlled by the UART Unit Enable bit (UUE bit 6) of the "UART x Interrupt Enable Register" on page 663.

Name*	Туре	Description
Ux_RXD	Input	SERIAL INPUT: Serial data input from device pin to the receive shift register.
Ux_TXD	Output	SERIAL OUTPUT: Composite serial data output to the communications link-peripheral, modem, or data set. The TXD signal is set to the MARKING (logic 1) state upon a Reset operation.
		CLEAR TO SEND: When low, this pin indicates that the receiving UART is ready to receive data. When the receiving UART deasserts <b>CTS#</b> high, the transmitting UART should stop transmission to prevent overflow of the receiving UARTs buffer. The <b>CTS#</b> signal is a modem-status input whose condition can be tested by the host processor or by the UART when in Autoflow mode as described below:
		Non-Autoflow Mode:
Ux_CTS#	Input	When not in Autoflow mode, bit 4 (CTS) of the Modem Status register (MSR) indicates the state of <b>CTS#</b> . Bit 4 is the complement of the <b>CTS#</b> signal. Bit 0 (DCTS) of the Modem Status register indicates whether the <b>CTS#</b> input has changed state since the previous reading of the Modem Status register. <b>CTS#</b> has no effect on the transmitter. The user can program the UART to interrupt the processor when DCTS changes state. The programmer can then stall the outgoing data stream by starving the transmit FIFO or disabling the UART with the IER register.
		Note: When UART transmission is stalled by disabling the UART, the user does not receive an MSR interrupt when <b>CTS#</b> reasserts. This is because disabling the UART also disables interrupts. To get around this, the user can use Auto CTS in Autoflow Mode, or program the <b>CTS#</b> pin to interrupt.
		Autoflow Mode:
		In Autoflow mode, the UART Transmit circuity checks the state of <b>CTS#</b> before transmitting each byte. IF <b>CTS#</b> is high, no data is transmitted. See Section 12.4.7, UART x Modem Control Register for more information on Auto CTS mode.
		REQUEST TO SEND: When low, this informs the remote device that the UART is ready to receive data. A reset operation sets this signal to its Inactive (high) state. LOOP mode operation holds this signal in its Inactive state.
		Non-Autoflow Mode:
Ux_RTS#	Output	The <b>RTS#</b> output signal can be asserted by setting bit 1 (RTS) of the Modem Control register to a 1. The RTS bit is the complement of the <b>RTS#</b> signal.
		Autoflow Mode:
		<b>RTS#</b> is automatically asserted by the autoflow circuitry when the Receive buffer exceeds its programmed threshold. It is deasserted when enough bytes are removed from the buffer to lower the data level back to the threshold. See Section 12.4.7, UART x Modem Control Register for more information on Auto RTS mode.

#### Table 357. UART Signal Descriptions

*Note:* \* "x" in signal name replaced with either "0" or "1" for UART-0 or UART-1 respectively.



## 12.3 Theory of Operation

The format of a UART data frame is shown in Figure 135.

#### Figure 135. Example UART Data Frame

	Start Bit	Data< 0>	Data< 1>	Data< 2>	Data< 3>	Data< 4>	Data< 5>	Data< 6>	Data< 7>	Parity Bit	Stop Bit 1	Stop Bit 2	
тхі	D or RXD	pin											
		LSB							MSB				

Shaded bits are optional and can be programmed by user. -->

Each data frame is between 7 bits and 12 bits long, depending on the size of data programmed and when parity and stop bits are enabled. The frame begins with a start bit that is represented by a high-to-low transition. Next, five to eight bits of data are transmitted, beginning with the least significant bit. An optional parity bit follows, which is set when even parity is enabled and an odd number of ones exist within the data byte; or when odd parity is enabled and the data byte contains an even number of ones. The data frame ends with one, one-and-one-half or two stop bits (as programmed by the user), which is represented by one or two successive bit periods of a logic one.

NRZ coding can be used by the UART to represent individual bit values. NRZ coding is enabled when Interrupt Enable Register (IER) bit-5 is set to high. A one is represented by a line transition and a zero is represented by no line transition. Figure 136 shows the NRZ coding of the data byte 8b 0100 1011. Note that the byte's LSB is transmitted first.

#### Figure 136. NRZ Bit Encoding Example – (0100 1011)



The unit is disabled upon reset, and users need to program GPIO registers first and then enable the unit by setting the UART Unit Enable bit (UUE, bit-6) of Interrupt Enable Register. When the unit is enabled, the receiver starts looking for the start bit of a frame; the transmitter sends data to the transmit data pin when there is data available in the transmit FIFO. Transmit data can be written to the FIFO before the unit is enabled. When the UART is disabled, the transmitter/receiver finishes the current byte being transmitted/received (when it is in the middle of transmitting/receiving a byte), and stops transmitting/receiving more data. Disabling the UART with the UUE bit does not clear transmission/reception with the original data.

Each UART has a Transmit FIFO and a Receive FIFO each holding 64 characters of data. There

are two methods for moving data into/out of the FIFOs: Interrupts and Polling.



## **12.3.1 FIFO Interrupt Mode Operation**

### 12.3.1.1 Receiver Interrupt

When the Receive FIFO and receiver interrupts are enabled (FCR[0]=1 and IER[0]=1), receiver interrupts occur as follows:

- The Receive Data Available Interrupt is asserted when the FIFO has reached its programmed trigger level. The interrupt is cleared when the FIFO drops below the programmed trigger level.
- The 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 Data Ready bit (DR in 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.

#### 12.3.1.2 Transmit Interrupt

When the transmitter FIFO and transmitter interrupt are enabled (FCR[0]=1, IER[1]=1), transmit interrupts occur as follows:

- When the Flow Control Register Transmitter Interrupt Level (TIL) bit (FCR[3]) is clear (0), The Transmit Data Request interrupt occurs when the transmit FIFO is half empty or more than half empty. The interrupt is cleared when the data level exceeds the half-empty mark. The interrupt is cleared as soon as the Transmit Holding Register is written or the IIR is read. 1 to 32 characters may be written to the transmit FIFO while servicing the interrupt when TIL=0.
- When the Flow Control Register Transmitter Interrupt Level (TIL) bit is set (1), The Transmit Data Request Interrupt occurs when the Transmit FIFO is empty. The interrupt is cleared as soon as the Transmit Holding Register is written or the IIR is read. 1 to 64 characters may be written to the Transmit FIFO while servicing the interrupt when TIL = 1.

Users could cause the UART Transmit FIFO to overflow when too many characters are written. FIFO underflow does not cause an error as the UART waits for the Transmit FIFO to be serviced.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual UARTs



## 12.3.2 Removing Trailing Bytes In Interrupt Mode

When the number of entries in the Receive FIFO is less than its trigger level, and no additional data is received, the remaining bytes are called trailing bytes. When the receive FIFO is being serviced by processor interrupts, trailing bytes need to be removed via the processor using the 16550 compliant character timeout interrupt: Time Out Detected (TOD) bit of Interrupt Identification Register. To enter this mode, users need to insure that the character timeout interrupt is enabled via IER[4].

To remove trailing bytes in Interrupt mode, the user must wait for the character timeout interrupt and then read all remaining bytes as indicated in the FIFO Occupancy Register (FOR), or read one byte at a time until the FIFO is empty. This can be determined by polling the Line Status Register bit 0 through programmed I/O.

#### 12.3.2.1 Character Timeout Interrupt

When the Receiver FIFO and Receiver Time-out Interrupt are enabled, a character time-out interrupt (TOD) occurs to signal the presence of trailing bytes. The Interrupt is cleared and the timer is reset when a character is read from the Receiver FIFO. When 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.

When enabled via IER[4], a character time-out occurs under the following conditions:

- At least one character is in the FIFO.
- A character has not been received for the amount of time it takes to receive four or more characters at the current baud rate.
- The FIFO has not been read for the amount of time it takes to receive four or more characters

### 12.3.3 FIFO Polled Mode Operation

With the FIFOs enabled (TRFIFOE bit of FCR set to 1), clearing IER[4:0] 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 Operation mode. In this mode, software checks Receiver and Transmitter status via the LSR. The processor polls the following bits for Receive and Transmit Data Service.

#### 12.3.3.1 Receive Data Service

• Processor should check *Data Ready bit of LSR* which is set when 1 or more bytes remains in the Receive FIFO or Receive Buffer register (RBR).

#### 12.3.3.2 Transmit Data Service

- Processor should check *Transmit Data Request bit of LSR* which is set when transmitter needs data.
- • Processor can also check *Transmitter Empty bit of LSR*, which is set when the Transmit FIFO or Holding register is empty.

## intel®

## 12.3.4 Autoflow Control

Autoflow Control uses the Clear-to-Send (**CTS#**) and Request-to-Send (**RTS#**) signals to automatically control the flow of data between the UART and external modem. When autoflow is enabled, the remote device is not allowed to send data unless the UART asserts nRTS low. When the UART deasserts **RTS#** while the remote device is sending data, the remote device is allowed to send one additional byte after **RTS#** is deasserted. An overflow could occur when the remote device violates this rule. Likewise, the UART is not allowed to transmit data unless the remote device asserts **CTS#** low. This feature increases system efficiency and eliminates the possibility of a Receive FIFO Overflow error due to long Interrupt latency.

Autoflow mode can be used in two ways: **Full autoflow**, automating both **CTS#** and **RTS#**, and **half autoflow**, automating only **CTS#**. Full Autoflow is enabled by writing a 1 to bits 1 and 5 of the Modem Control register (MCR). Auto-CTS-Only mode is enabled by writing a 1 to bit 5 and a 0 to bit 1 of the MCR register.

### 12.3.4.1 RTS Autoflow

When in full autoflow mode, **RTS**# is asserted when the UART FIFO is ready to receive data from the remote transmitter. This occurs when the amount of data in the Receive FIFO is below the programmable threshold value. When the amount of data in the Receive FIFO reaches the programmable threshold, **RTS**# is deasserted. It is asserted once again when enough bytes are removed from the FIFO to lower the data level below the threshold.

### 12.3.4.2 CTS Autoflow

When in Full or Half-Autoflow mode, **CTS#** is asserted by the remote receiver when the receiver is ready to receive data from the UART. The UART checks **CTS#** before sending the next byte of data and does not transmit the byte until **CTS#** is low. When **CTS#** goes high while the transfer of a byte is in progress, the transmitter completes this byte.

Note: Autoflow mode can be used only in conjunction with FIFO mode.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual UARTs



## 12.3.5 Auto-Baud-Rate Detection

Each UART supports auto-baud-rate detection. When enabled, the UART counts the number of 33.334 MHZ clock cycles within the *start*-bit pulse. This number is then written into the Auto-Baud-Count register (ACR) and is used to calculate the baud rate. When the ACR is written, a Auto-Baud-Lock Interrupt is generated (when enabled), and the UART automatically programs the Divisor Latch registers with the appropriate baud rate. When preferred, the processor can read the Auto-Baud-Count register and use this information to program the Divisor-Latch registers with a baud rate calculated by the processor. After the baud rate has been programmed, it is the responsibility of the processor to verify that the predetermined characters (usually **AT** or **at**) are being received correctly.

When the UART programs Divisor Latch registers, users can choose between two auto-baud calculation methods: **table-** and **formula-based**. The method is selected via bit *ABT* of the Auto-Baud Control register (ABR). When the formula method is used, any baudrate allowed in Equation 17 can be programmed by the UART. This method works well for higher baud rates, but could possibly fail below 28.8 kbps when the remote transmitter's actual baud rate differs by more than one percent of its target.

#### **Equation 17. Baud-Rate Equation**

## $BaudRate = \frac{33.334 \text{Mhz}}{(16XDivisor)}$

The table method is more immune to such errors as the table rejects uncommon baud rates and rounds to the common ones. The table method allows any baud rate in Equation 17 above 28.8 kbps. Below 28.8 kbps the only baud rates which can be programmed by the UART are 19200, 14400, 9600, 4800, 1200, and 300 baud. Some typical values for Divisor and corresponding baud rates are provided in Table 358. Baud rates above 3600 baud require only Divisor Latch Low Register to be programmed, as Divisor Latch High Register would be 0.

#### Table 358.Divisor Values for Typical Baud Rates

Baud Rate	Divisor	UART Rate	Error
115.2K	18	115.74K	0.47%
57.6K	36	57.87K	0.47%
38.4K	54	38.58K	0.47%
33.6K	62	33.60K	0.01%
28.8K	72	28.94K	0.47%
19.2K	109	19.29K	0.47%
14.4K	145	14.47K	0.47%
9600	217	9645	0.47%
4800	434	4800	0.01%
3600	579	3600	0.01%
2400	868	2400	0.01%
1200	1736	1200	0.01%
600	3472	600	0.01%
300	6944	300	0.01%

When the baud rate is detected, auto-baud circuitry disarms itself by clearing bit *ABE* of the Auto-Baud Control register (ABR). When users want to rearm the circuitry, the *ABE* bit must be rewritten.

*Note:* For the auto-baud-rate detection circuit to work correctly, the first data bit transmitted after the start bit must be a logic '1'. When a logic '0' is transmitted instead, the autobaud circuit counts the zero as part of the start bit, resulting in an incorrect baud rate being programmed into the DLL and DLH registers.

## 12.3.6 Manual Baud Rate Selection

Each UART contains a programmable Baud Rate Generator that is capable of taking the fixed input clock of 33.334 MHz and dividing it by any divisor from 1 to  $(2^{16}-1)$ . The baud-rate generator output frequency is 16 times the baud rate. Two 8-bit registers store the divisor in a 16-bit binary format. These Divisor Registers must be loaded during initialization to ensure proper operation. When both Divisor Latches are loaded with 0, the 16X output clock is stopped. Access to the Divisor latch can be done with a word write. Equation 17 or Table 358 are used by the programmer to select the Divisor Latch value for the desired baud rate.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual UARTs

## 12.4 Register Descriptions

There are 15 registers in each UART. The registers are all 32 bit registers, but only lower 8 bits have valid data. The 12 UART registers share eight address locations in the MMR address space. Table 359 shows the registers and their addresses as offsets of a base address. The base address for each UART is 32 bits and is internal bus address FFFF F700H for UART 0, and FFFF F740H for UART 1. Note that the state of the Divisor Latch Bit (DLAB), which is the MOST significant bit of the Serial Line Control Register, affects the selection of certain of the UART registers. The DLAB bit must be set high by the system software to access the Baud Rate Generator Divisor Latches.

UART Register Addresses	DLAB Bit Value	Name	Register Accessed
Base	0	UxRBR	UART x Receive BUFFER (read only)
Base	0	UxTHR	UART x Transmit BUFFER (write only)
Base + 04H	0	UxIER	UART x Interrupt Enable (R/W)
Base + 08H	Х	UxIIR	UART x Interrupt I.D. (read only)
Base + 08H	Х	UxFCR	UART x FIFO Control (write only)
Base + 0CH	Х	UxLCR	UART x Line Control (R/W)
Base + 10H	Х	UxMCR	UART x Modem Control (R/W)
Base + 14H	Х	UxLSR	UART x Line Status (Read only)
Base + 18H	Х	UxMSR	UART x Modem Status (Read only)
Base + 1CH	Х	UxSPR	UART x Scratch Pad (R/W)
Base	1	UxDLL	UART x Divisor Latch (Low Byte, R/W)
Base + 04H	1	UxDLH	UART x Divisor Latch (High Byte, R/W)
Base + 24H	Х	UxFOR	UART x FIFO Occupancy Register (R/W)
Base + 28H	Х	UxABR	UART x Autobaud Control Register (R/W)
Base + 2CH	Х	UxACR	UART x Autobaud Count Register (read only)

#### Table 359. UART Register Addresses as Offsets of a Base

#### Table 360.UART Unit Registers

Section, Register Name, Acronym (page)
Section 12.4.1, "UART x Receive Buffer Register" on page 662
Section 12.4.2, "UART x Transmit Holding Register" on page 662
Section 12.4.3, "UART x Interrupt Enable Register" on page 663
Section 12.4.4, "UART x Interrupt Identification Register" on page 664
Section 12.4.5, "UART x FIFO Control Register" on page 666
Section 12.4.6, "UART x Line Control Register" on page 668
Section 12.4.7, "UART x Modem Control Register" on page 670
Section 12.4.8, "UART x Line Status Register" on page 672
Section 12.4.9, "UART x Modem Status Register" on page 675
Section 12.4.10, "UART x Scratchpad Register" on page 676
Section 12.4.11, "Divisor Latch Registers" on page 677
Section 12.4.13, "UART x Auto-Baud Control Register" on page 679
Section 12.4.14, "UART x Auto-Baud Count Register" on page 680

# intel

UART Register Addresses	DLAB Bit Value	Name	Register Accessed
	0	U0RBR	UART 0 Receive BUFFER (read only)
	0	U0THR	UART 0 Transmit BUFFER (write only)
FFFF F704H	0	U0IER	UART 0 Interrupt Enable (R/W)
	Х	U0IIR	UART 0 Interrupt I.D. (read only)
	Х	U0FCR	UART 0 FIFO Control (write only)
FFFF F70CH	Х	U0LCR	UART 0 Line Control (R/W)
FFFF F710H	х	U0MCR	UART 0 Modem Control (R/W)
FFFF F714H	Х	U0LSR	UART 0 Line Status (Read only)
FFFF F718H	Х	U0MSR	UART 0 Modem Status (Read only)
FFFF F71CH	Х	U0SPR	UART 0 Scratch Pad (R/W)
FFFF F700H	1	U0DLL	UART 0 Divisor Latch (Low Byte, R/W)
FFFF F704H	1	U0DLH	UART 0 Divisor Latch (High Byte, R/W)
FFFF F720H	Х	n/a	Reserved
FFFF F724H	X	U0FOR	UART 0 FIFO Occupancy Register (R/W)
FFFF F728H	Х	U0ABR	UART 0 Autobaud Control Register (R/W)
FFFF F72CH	Х	U0ACR	UART 0 Autobaud Count Register (read only)
FFFF F730H through FFFF F73FH	x	n/a	Reserved
	0	U1RBR	UART 1 Receive BUFFER (read only)
	0	U1THR	UART 1 Transmit BUFFER (write only)
FFFF F744H	0	U1IER	UART 1 Interrupt Enable (R/W)
	Х	U1IIR	UART 1 Interrupt I.D. (read only)
FFFF F740F1	Х	U1FCR	UART 1 FIFO Control (write only)
FFFF F74CH	X	U1LCR	UART 1 Line Control (R/W)
FFFF F750H	Х	U1MCR	UART 1 Modem Control (R/W)
FFFF F754H	Х	U1LSR	UART 1 Line Status (Read only)
FFFF F758H	Х	U1MSR	UART 1 Modem Status (Read only)
FFFF F75CH	Х	U1SPR	UART 1 Scratch Pad (R/W)
FFFF F740H	1	U1DLL	UART 1 Divisor Latch (Low Byte, R/W)
FFFF F744H	1	U1DLH	UART 1 Divisor Latch (High Byte, R/W)
FFFF F760H	Х	n/a	Reserved
FFFF F764H	Х	U1FOR	UART 1 FIFO Occupancy Register (R/W)
FFFF F768H	х	U1ABR	UART 1 Autobaud Control Register (R/W)
FFFF F76CH	Х	U1ACR	UART 1 Autobaud Count Register (read only)
FFFF F770H through FFFF F77FH	х	n/a	Reserved

### Table 361. UART Register MMR Addresses

#### 12.4.1 **UART x Receive Buffer Register**

In non-FIFO mode, this register holds the character(s) received by the UART Receive Shift register. When it receives fewer than eight bits, the bits are right-justified and the leading bits are zeroed. Reading the register empties the register and resets the *data ready* (DR) bit in the Line Status register to 0. Other (error) bits in the Line Status register are not cleared. In FIFO mode, this register latches the value of the data byte(s) at the bottom of the FIFO.

When the UART is in eight-bit Peripheral Bus mode, the 24 most significant bits must be ignored and not used. Reading these bits returns unpredictable results.



#### **Table 362.** UART x Receive Buffer Register - (UxRBR)

#### **UART x Transmit Holding Register** 12.4.2

This register holds the next data byte(s) to be transmitted. When the Transmit Shift register becomes empty, the contents of the Transmit Holding register are loaded into the Shift register and the Transmit Data Request (TDRQ) bit in the Line Status register is set to one (see Table 366).

In FIFO mode, writing to THR puts data to the top of the FIFO. The data at the bottom of the FIFO is loaded to the Shift register when it is empty. In eight-bit Peripheral mode, the 24 most significant bits are ignored and is not transmitted.



#### Table 363. UART x Transmit Holding Register - (UxTHR)

## 12.4.3 UART x Interrupt Enable Register

This register enables six types of interrupts which set a value in the Interrupt Identification register. Each of the six interrupt types can be disabled by clearing the appropriate bit of the IER register. Similarly, by setting the appropriate bits, selected interrupts can be enabled.

This register also has the control bits of the unit enable and NRZ coding enable. The use of bit 7 to bit 4 is different from the register definition of standard 16550.

- *Note:* A global interrupt enable/disable exists in the Modem Control Register bit 3 (IE). After reset, this bit must be set or no interrupts occurs, regardless of the state of the IER bits. See Section 12.4.7, "UART x Modem Control Register" on page 670.
- *Note:* Users need to program the GPIO registers before enabling the UART. See "GPIO Pin Multiplexing" on page 824.

#### 12 24 20 16 8 Λ 0 ttribut Q PCI Intel XScale<sup>®</sup> Core internal bus address RW = Read/Write Attribute Legend: Unit # RV = Reserved RC = Read Clear 0 FFFF F704H (DLAB=x) PR = Preserved RO = Read Only FFFF F744H (DLAB=x) 1 NA = Not Accessible RS = Read/Set Bit Default Description 00 0000h 31:8 Reserved 7 Preserved $0_{2}$ UART Unit Enable (UUE): Controls UART operation and pin multiplexing with the GPIO. Refer to "GPIO Pin Multiplexing" on page 824 for details. 6 $0_{2}$ 0 = the unit is disabled and the shared pins operate as GPIO signals. 1 = the unit is enabled and the shared pins operate as UART signals NRZ coding Enable (NRZE): 5 $0_{2}$ 0 = NRZ coding disabled 1 = NRZ coding enabled Receiver Time Out Interrupt Enable: (RTOIE) 4 02 0 = Receiver data Time out Interrupt disabled 1 = Receiver data Time out Interrupt enabled Modem Interrupt Enable (MIE): 3 $0_{2}$ 0 = Modem Status interrupt disabled 1 = Modem Status interrupt enabled Receiver Line Status Interrupt Enable (RLSE): 2 02 0 = Receiver Line Status interrupt disabled 1 = Receiver Line Status interrupt enabled Transmit Data request Interrupt Enable (TIE): 1 02 0 = Transmit FIFO Data Request interrupt disabled 1 = Transmit FIFO Data Request interrupt enabled Receiver Data Available Interrupt Enable (RAVIE): 0 $0_{2}$ 0 = Receiver Data Available (Trigger level reached) interrupt disabled Receiver Data Available (Trigger level reached) interrupt enabled 1 =

#### Table 364. UART x Interrupt Enable Register - (UxIER)



## **12.4.4 UART x Interrupt Identification Register**

The IIR register is read to determine the type and source of UART interrupts. To be 16550 compatible, the lower 4 bits (0-3) of the IIR register are priority encoded as shown in Table 366. When two or more interrupts represented by bits (0-3) occur, only the interrupt with the highest priority is displayed. The upper 4 bits, (4-7) are not priority encoded. These bits asserts/deasserts independently of the lower 4 bits.

Bit 0 (nIP) is used to indicate the existence of an interrupt in the priority encoded bits (0-3) of the IIR register. A low signal on this bit indicates an encoded interrupt is pending. When this bit is high, no encoded interrupt is pending, regardless of the state of the other 3 bits. IP# has no effect or association with the upper bits four bits (4-7) which assert/deassert independently of IP#.

In order to minimize software overhead during data character transfers, the UART prioritizes interrupts into four levels (listed in Table 366) and records these in the Interrupt Identification register. The Interrupt Identification register (IIR) stores information indicating that a prioritized interrupt is pending and the source of that interrupt.

#### Table 365. UART x Interrupt Identification Register - (UxIIR)

Unit # 0 1	31 28 rv rv rv rv rv rv na na na na na na Intel XSG FFFF F7 FFFF F7	24       20       16       12       8       4       0         rv       r				
Bit	Default	Description				
31:8	00 0000h	Reserved				
7:6	00 <sub>2</sub>	FIFO Mode Enable Status (FIFOES[1:0]): 00 = Non-FIFO mode is selected 01 = Reserved 10 = Reserved 11 = FIFO mode is selected (TRFIFOE = 1)				
5	02	Reserved				
4	02	Autobaud Lock (ABL) 0 = Autobaud circuitry has not programmed Divisor Latch registers (DLL/DLH) 1 = Divisor Latch registers (DLL/DLH) programmed by autobaud circuitry				
3	02	Time Out Detected (TOD): 0 = No time out interrupt is pending 1 = Time out interrupt is pending. (FIFO mode only)				
2:1	0 <sub>2</sub>	Interrupt Source Encoded (IID[1:0]): indicates a Modem Status Interrupt when the IP# bit is low. When IP# bit is high, there is no Interrupt. 00 = Modem Status (CTS, DSR, RI, DCD modem signals changed state) 01 = Transmit FIFO requests data 10 = Received Data Available 11 = Receive error (Overrun, parity, framing, break, FIFO error)				
0	1 <sub>2</sub>	Interrupt Pending (IP#): 0 = Interrupt is pending. (Active low) 1 = No interrupt is pending				

# intel

	Interrupt ID bits					Interrupt SET/RESET Function				
	3	2	1	0	Priority	Туре	Source	RESET Control		
IP#	0	0	0	1	-	None	No Interrupt is pending.	-		
IID[11]	0	1	1	0	Highest	Receiver Line Status	Overrun Error, Parity Error, Framing Error, Break Interrupt.	Reading the Line Status Register.		
							Non-FIFO mode: Receive Buffer is full.	Non-FIFO mode: Reading the Receiver Buffer Register.		
IID[10]	0	1	0	0	Second Highest	Received Data Available.	FIFO mode: Trigger level was reached.	FIFO mode: Reading bytes until Receiver FIFO drops below trigger level or setting RESETRF bit in FCR register.		
TOD	1	1	0	0	Second Highest	Character Timeout indication.	FIFO Mode only: At least 1 character is in receiver FIFO and there was no activity for a time period.	Reading the Receiver FIFO or setting RESETRF bit in FCR register.		
IID[01]	0	0	1	0	Third	Transmit FIFO	Non-FIFO mode: Transmit Holding Register Empty	Reading the IIR Register (when the source of the interrupt) or writing into the Transmit Holding Register.		
					Hignest	Data Nequest	FIFO mode: Transmit FIFO has half or less than half data.	Reading the IIR Register (when the source of the interrupt) or writing to the Transmitter FIFO.		
IID[00]	0	0	0	0	Fourth Highest	Modem Status	Clear to Send, Data Set Ready, Ring Indicator, Received Line Signal Detect	Reading the Modem Status Register.		
	Non Prioritized Interrupts									
ABL 4 None					None	Autobaud Lock Indication	Autobaud circuitry has locked onto the baud rate.	Reading the IIR Register		

### Table 366. Interrupt Identification Register Decode



## 12.4.5 UART x FIFO Control Register

FCR is a write only register that is located at the same address as the IIR (IIR is a read only register). FCR enables/disables the transmitter/receiver FIFOs, clears the transmitter/receiver FIFOs, and sets the receiver FIFO trigger level.

#### Table 367. UART x FIFO Control Register - (UxFCR) (Sheet 1 of 2)



# intel®

<i>.</i>	UARTX	FIFO Contr	of Register - (UXFCR) (Sh	eet 2 0f 2)		
	PCI IOP Attributes Attributes	31   28     rv   rv   rv     na   na     na   na	24 20 10 ry ry ry na na	5 12 8 TV TV TV TV TV TV TV TV na na	4 0 wowo/pr/pr/wo/wo/wo/wo/ na na na na na na na na na	
	Unit #	Intel XSo	ale <sup>®</sup> Core internal bus address	Attribute Legend: RV = Reserved	RW = Read/Write RC = Read Clear	
	1	FFFF F7	08H (DLAB=x) 48H (DLAB=x)	PR = Preserved RS = Read/Set	RO = Read Only NA = Not Accessible	
	Bit	Default		Description		
	2	0 <sub>2</sub>	<ul> <li>Reset Transmitter FIFO (RESETTF): When set, the Transmitter FIFO counter is reset to clear all the bytes in the FIFO. The <i>TDRQ</i> bit of LSR is set generating a Transmitter Requests Data Interrupt IID field of IIR when the <i>TIE</i> bit in the IER register is set. The Transmitter Shift register is not reset; it completes the current transmission. Any transmit FIFO Service-Request Interrupts are cleared.</li> <li>Note: After the FIFO is cleared, RESETTF is automatically reset to 0.</li> <li>0 = no effect</li> <li>1 = The transmitter FIFO is cleared (FIFO counter set to 0). After clearing, bit is automatically reset to 0</li> </ul>			
	1	02	Reset Receiver FIFO (RESETR clear all the bytes in the FIFO. Th FIFO and the FIFOE bit in LSR a had been set in LSR are still set. Receive FIFO Service Request I Note: After the FIFO is cleared, I 0 = no effect 1 = The receiver FIFO is cleare automatically reset to 0	F): When set, the receiver the DR bit in LSR is reset to are cleared. Any error bits ( The receiver shift register interrupts are cleared. RESETRF is automatically d (FIFO counter set to 0).	FIFO counter is reset to 0. All the error bits in the (OE, PE, FE or BI), that is not cleared. Any reset to 0. After clearing, bit is	
	0	02	Transmit and Receive FIFO En and receiver FIFOS. When TRFI When TRFIFOE = 0, the FIFOS a this bit clears all bytes in both FI mode and vice versa, data is aut Service Request Interrupts are c Note: This bit must be 1 when ot are not programmed. 0 = FIFOS are disabled 1 = FIFOS are enabled	able (TRFIFOE): Enables, FOE = 1, both FIFOs are e are both disabled (non-FIF FOs. When changing from omatically cleared from the leared when TRFIFOE is o her bits in this register are	/disables the transmitter nabled (FIFO Mode). O Mode). Writing a 0 to FIFO mode to non-FIFO e FIFOs. Any FIFO leared. written, or the other bits	

### Table 367. UART x FIFO Control Register - (UxFCR) (Sheet 2 of 2)



## 12.4.6 UART x Line Control Register

In the Line Control Register, the system programmer specifies the format of the asynchronous data communications exchange. The serial data format consists of a start bit (logic 0), five to eight data bits, an optional parity bit, and one or two stop bits (logic 1). The LCR has bits for accessing the Divisor Latch registers and causing a Break condition. The programmer can also read the contents of the Line Control Register. The read capability simplifies system programming and eliminates the need for separate storage in system memory.



PCI IOP Attributes	31 28 rv/rv/rv/rv/rv/rv/ na na na na na	24       20       16       12       8       4       0         1v       rv       r
Unit # 0 1	Intel XSc FFFF F7 FFFF F7	cale® Core internal bus addressAttribute Legend:RW = Read/Write0CH (DLAB=x)RV = ReservedRC = Read Clear4CH (DLAB=x)PR = PreservedRO = Read OnlyRS = Read/SetNA = Not Accessible
Bit	Default	Description
31:8	00 0000h	Reserved
7	0 <sub>2</sub>	<ul> <li>Divisor Latch register Access Bit (DLAB): This bit must be set (1) to access the Divisor Latches of the Baud Rate Generator during a READ or WRITE operation. It must be clear (0) to access the Receiver Buffer, the Transmit-Holding Register, or the Interrupt-Enable Register. This bit does not have to be set when using autobaud.</li> <li>0 = access Transmit Holding register (THR), Receive Buffer Register (RBR) and Interrupt Enable Register.</li> <li>1 = access Divisor Latch Registers (DLL and DLH).</li> </ul>
6	0 <sub>2</sub>	Set break (SB): Causes a Break condition to the receiving UART. When SB is set (1), the serial output (TXD) is forced to the spacing (logic 0) state and remains there until SB is clear (0). This bit acts only on the TXD pin and has no effect on the transmitter logic. In FIFO mode, wait for the transmitter to be idle (TEMT=1) to set and clear the break bit. 0 = no effect on TXD output. 1 = forces TXD output to 0 (space).
5	0 <sub>2</sub>	Sticky Parity (STKYP): Can be used in multiprocessor communications. When PEN and STKYP are set (1), the bit that is transmitted in the parity bit location (the bit just before the stop bit) is the complement of the EPS bit. When EPS is 0, then the bit at the parity bit location is transmitted as a 1. In the receiver, when STKYP and PEN are 1, then the receiver compares the bit that is received in the parity bit location with the complement of the EPS bit. When the values being compared are not equal, the receiver sets the Parity Error bit in LSR and causes an error interrupt when line status interrupts were enabled. For example, when EPS is 0, the receiver expects the bit received at the parity bit location to be 1. When it is not, then the parity error bit is set. By forcing the bit value at the parity bit location, rather than calculating a parity value, a system with a master transmitter and multiple receivers can identify some transmitted characters as receiver addresses and the rest of the characters as data. When PEN = 0, STKYP is ignored. 0 = no effect on parity bit. 1 = Forces parity bit to be opposite of EPS bit value.

# intel®

	Line Contr			
PCI IOP Attributes Attributes	31 28 rv/rv/rv/rv/rv/ na na na na na	24     20     16       rv     rv <t< th=""><th>12     8       rv     rv     rv     rv     rv     rv     rv     rv       na     na     na     na     na     na     na     na</th><th>8 4 0 rw/rw/rw/rw/rw/rw/rw/rw/ na na na na na na na na na</th></t<>	12     8       rv     rv     rv     rv     rv     rv     rv     rv       na     na     na     na     na     na     na     na	8 4 0 rw/rw/rw/rw/rw/rw/rw/rw/ na na na na na na na na na
Unit # 0	Intel XSo FFFF F7	cale <sup>®</sup> Core internal bus address /0CH (DLAB=x)	Attribute Legend: RV = Reserved PR = Preserved	RW = Read/Write RC = Read Clear RO = Read Only
1		4CH (DLAB=X)	RS = Read/Set	NA = Not Accessible
Bit	Default		Description	
4	02	Even Parity Select (EPS): Wher number of logic ones is transmitte bit. When PEN is set (1) and EPS transmitted or checked in the dat ignored. 0 = sends or checks for odd pari 1 = sends or checks for even pa	PEN is set (1) and EPS ed or checked in the data 5 is also set (1), an even a word bits and parity bit. ity. urity.	is clear (0), an odd word bits and the parity number of logic ones is When PEN = 0, EPS is
3	02	<b>Parity Enable (PEN):</b> When set checked (receive data) between the top roduce word bits and the parity bit are su $0 = 100$ parity function. 1 = allows parity generation and	(1), a parity bit is generate the last data word bit and an even or odd number o immed.) checking.	ed (transmit data) or Stop bit of the serial data. of ones when the data
2	02	Stop bits (STB): This bit specifie in each serial character. When ST transmitted data. When STB is set WLS[1:0], then 1 and one half sto either a 6, 7, or 8-bit word is selec checks the first stop bit only, rega 0 = 1 stop bit 1 = 2 stop bits, except for 5-bit of	s the number of stop bits TB is clear (0), one stop b at (1) when a 5-bit word le op bits are generated. Wh ted, then two stop bits are ardless of the number of s	e transmitted and received bit is generated in the ength is selected via nen STB is set (1) when e generated. The receiver stop bits selected.
1:0	002	Word Length Select (WLS[1:0]) transmitted or received serial cha 00 = 5-bit character (default) 01 = 6-bit character 10 = 7-bit character 11 = 8-bit character	: Specifies the number of rracter.	f data bits in each

### Table 368. UART x Line Control Register - (UxLCR) (Sheet 2 of 2)

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual UARTs



## 12.4.7 UART x Modem Control Register

This register controls the interface with the modem or data set (or a peripheral device emulating a modem). The contents of the Modem Control register are described below:

#### Table 369.UART x Modem Control Register - (UxMCR) (Sheet 1 of 2)

PCI IOP Attributes Attributes	31 28 rv/rv/rv/rv/rv/ na na na na na	24       20       16       12       8       4       0         rv       r				
Unit # 0 1	Intel XSc FFFF F7 FFFF F7	ale® Core internal bus addressAttribute Legend:RW = Read/Write10H (DLAB=x)RV = ReservedRC = Read Clear50H (DLAB=x)PR = PreservedRO = Read OnlyRS = Read/SetNA = Not Accessible				
Bit	Default	Description				
31:6	000 0000h	Reserved				
5	0 <sub>2</sub>	<ul> <li>Autoflow Control Enable (AFE): When set, autoflow control is enabled. Only auto-CTS is enabled when <i>RTS</i> is cleared in MCR while AFE is set. Both auto-CTS and auto-RTS are enabled when <i>AFE</i> and <i>RTS in MCR</i> are set. Auto-RTS is not enabled when <i>AFE</i> is not set regardless of the state of <i>RTS</i>.</li> <li>Autoflow automates the flow of data between the UART and the remote device. See Section 12.3.4, "Autoflow Control" on page 657 for more details.</li> <li>0 = Auto-RTS and Auto-CTS are disabled</li> <li>1 = Auto-CTS is enabled. IF RTS in MCR is also set, both auto-CTS and auto-RTS is enabled</li> </ul>				
4	02	<ul> <li>is enabled</li> <li>Loop back test mode (LOOP): This bit provides a local Loopback feature for diagnostic testing of the UART. In the Diagnostic mode, data that is transmitted is immediately received. This feature allows the processor to verify the UART Transmit and Receive data paths. The Transmit, Receive and Modem Control Interrupts are operational. The modem-control input CTS# is activated by MCR bit 1 instead of the modem-control input. A Break signal can also be transferred from the transmitter section to the receiver section in Loop-Back mode.</li> <li>When LOOP is set (1), the following occurs: <ul> <li>The TXD (transmitter output) pin is set to a logic-1 state.</li> <li>The RXD (receiver input) pin is disconnected.</li> <li>The output of the Transmitter Shift register is "looped back" into the Receiver-Shift register input.</li> <li>The modem-control output pin RTS# is forced to the inactive state.</li> </ul> </li> <li>The RTS bit of the Modem Control register is connected to bits of the Modem Status register bits. Flow control can be tested; when autoflow is enabled the RTS bit of the Modem Control can be tested; when autoflow is enabled the RTS bit of the Modem Control can be tested; when autoflow is enabled the RTS bit of the autoflow logic.: <ul> <li>RTS = 1 forces CTS to a 1</li> </ul> </li> <li>Note: Note: Coming out of the Loop-Back Test mode may result in unpredictable activation of the delta bit (bit 0) in the Modem Status register (MSR). It is recommended that MSR is read once to clear the delta bits in the MSR.</li> </ul>				

# intel®

- 303.								
	31       28       24       20       16       12       8       4         01       12							
	Bit	Default	Description					
	3	02	Interrupt Enable (IE): Global control all UART interrupts. 0 = interrupts disabled. 1 = interrupts enabled. NOTE: This bit is not valid when in Loopback mode.					
	2	02	Preserved.					
	1	02	<ul> <li>Request to Send (RTS):</li> <li>Non-Autoflow mode: When not in Autoflow mode (AFE bit of MCR is clear), this is controls the Request-to-Send (RTS#) output pin.</li> <li>0 = RTS# pin is 1</li> <li>1 = RTS# pin is 0</li> <li>Autoflow mode: When in Autoflow mode (AFE bit of MCR is set), auto-RTS is enabled. RTS# behaves as follows:</li> <li>Auto-RTS disabled. Autoflow works only with auto-CTS.</li> <li>Auto-RTS enabled Autoflow works with both auto-CTS and auto-RTS</li> </ul>					
	0	02	Reserved					

## Table 369. UART x Modem Control Register - (UxMCR) (Sheet 2 of 2)



## 12.4.8 UART x Line Status Register

This register provides status information to the processor concerning the data transfers. Bits 5 and 6 show information about the transmitter section. The remainder of the bits contain information about the receiver.

In non-FIFO mode, three of the LSR register bits, parity error, framing error, and break interrupt, show the error status of the character that has just been received. In FIFO mode, these three status bits are stored with each received character in the FIFO. LSR shows the status bits of the character at the bottom of the FIFO. When the character at the bottom of the FIFO has errors, the LSR error bits are set and are not cleared until software reads LSR, even when the character in the FIFO is read and a new character is now at the bottom of the FIFO.

Bits 1 through 4 are the error conditions that produce a Receiver Line Status Interrupt when any of the corresponding conditions are detected and the interrupt is enabled. These bits are not cleared by reading the erroneous byte from the FIFO or receive buffer. They are cleared only by reading LSR. In FIFO mode, the Line Status Interrupt occurs only when the erroneous byte is at the bottom of the FIFO. When the erroneous byte being received is not at the bottom of the FIFO, an interrupt is generated only after the previous bytes are read and the erroneous byte is moved to the bottom of the FIFO.



#### Table 370. UART x Line Status Register - (UxLSR) (Sheet 1 of 3)

# intel®

P ∃ Butes	31 28	24 20 16 12 8 4 0				
PCI Attributes	na na na na	na n				
Unit # 0 1	Intel XS FFFF F FFFF F	Scale® Core internal bus addressAttribute Legend: RV = ReservedRW = Read/Write RC = Read Clear RO = Read Only RS = Read/Set714H (DLAB=x)PR = Preserved RS = Read/SetRC = Read Only RA = Not Accessible				
Bit	Default	Description				
5	12	Transmit Data Request (TDRQ): Indicates that the UART is ready to accept data for transmission. The assertion of this bit causes the UART to issue an interrupt when the Transmit Data Request Interrupt Enable is set. In non-FIFO mode, the TDRQ bit is set (1) when a character is transferred from the Transmit-Holding register. The bit is cleared (0) concurrently with the loading of the Transmit Holding register by the processor.				
		In FIFO mode, TDRQ is set (1) when the FIFO is less than half full. It is cleared when the FIFO is more than half full. When more than 64 characters are loaded into the FIFO, the excess characters are lost. 0 = The UART is NOT ready to receive data for transmission. 1 = The UART is ready to receive data for transmission.				
	02	<b>Break Indicator (BI):</b> Set (1) when the received data input is held in the Spacing (logic 0) state for longer than a full character transmission time (that is, the total time of <i>start</i> bit + <i>data</i> bits + <i>parity</i> bit + <i>stop</i> bits). The Break Indicator is reset (cleared to 0) when the processor reads the Line-Status register.				
4		In FIFO mode, only one Break character (equal to 0x00), is loaded into the FIFO regardless of the length of the Break condition. <i>BI</i> shows the Break condition for the character at the bottom of the FIFO, not the most recent character received.				
		<ul><li>0 = No break signal has been received.</li><li>1 = Break signal has been received.</li></ul>				
	02	<b>Framing Error (FE):</b> Indicates that the received character did not have a valid stop bit. FE is set (1) when the bit following the last data bit or parity bit is detected as a logic 0 bit (spacing level). When the Line Control register had been set for two stop bits, the receiver does not check for a valid second stop bit. The FE indicator is reset when the processor reads the Line Status Register.				
3		The UART resynchronizes after a framing error by assuming that the framing error was due to the next start bit. Therefore it samples this start bit twice and then takes in the data. In FIFO mode, FE shows a framing error for the character at the bottom of the FIFO, not for the most recently received character.				
		0 = No Framing error. 1 = Invalid stop bit has been detected.				

#### Table 370. UART x Line Status Register - (UxLSR) (Sheet 2 of 3)



PCI IOP Attributes Attributes	31 28 rv/rv/rv/rv/rv na na na na na	24     20     16       rv     rv <t< th=""><th>12 8 rv rv rv rv rv rv rv na na na na na na na na</th><th>4 0 ro ro ro ro ro ro ro na na na na na na na na</th></t<>	12 8 rv rv rv rv rv rv rv na na na na na na na na	4 0 ro ro ro ro ro ro ro na na na na na na na na		
Unit # 0 1	Intel XS FFFF F <sup>-</sup> FFFF F <sup>-</sup>	cale <sup>®</sup> Core internal bus address 714H (DLAB=x) 754H (DLAB=x)	Attribute Legend: RV = Reserved PR = Preserved RS = Read/Set	RW = Read/Write RC = Read Clear RO = Read Only NA = Not Accessible		
Bit	Default	Description				
2	02	Parity Error (PE): Indicates that the received data character does not have the correct even or odd parity, as selected by the even parity select bit. The PE is set (1) upon detection of a parity error and is cleared (0) when the processor reads the Line Status register. In FIFO mode, PE shows a parity error for the character at the bottom of the FIFO, not the most recently received character. 0 = No Parity error. 1 = Parity error has occurred.				
1	02	Overflow Error (OE): In non-FIFO mode, OE indicates that data in the Receiver Buffer register was not read before the next character was received, the new character is lost.         In FIFO mode, OE indicates that all 64 bytes of the FIFO are full and the most recently received byte has been discarded. The OE indicator is set (1) upon detection of an overflow condition and reset when the processor reads the Line Status register 0 = No overflow error. Data has not been lost.         1       = Overflow error. Receive data has been lost.				
0	02	Data Ready (DR): Set to a logic 1 wh received and transferred into the received, DR is reset to 0 when the received a logic 0 when the FIFO is empty (las RESETRF bit is set in FCR.         0 = No data has been received.         1 = Data is available in RBR or the F	en a complete incomi siver buffer register or ive buffer is read. In F t character has been i	ng character has been the FIFO. In non-FIFO IFO mode, DR is reset to read from RBR) or the		

#### Table 370.UART x Line Status Register - (UxLSR) (Sheet 3 of 3)

## 12.4.9 UART x Modem Status Register

This register provides the current state of the control lines from the modem or data set (or a peripheral device emulating a modem). In addition to this current state information, the Modem Status register also provides change information. The change bit is set to a logic 1 when the control input from the Modem changes state. The change bit is reset to a logic 0 when the processor reads the Modem Status register.

*Note:* When the change bit (bit 0) is set to logic 1, a Modern Status interrupt is generated when bit 3 of the Interrupt Enable Register is set.



#### Table 371. UART x Modem Status Register - (UxMSR)

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual UARTs



## 12.4.10 UART x Scratchpad Register

This read/write register has no effect on the UART. It is intended as a scratchpad register for use by programmers.

#### Table 372. UART x Scratchpad Register - (UxSCR)





## 12.4.11 Divisor Latch Registers

The description of use for the Divisor Latch Registers are provided in Section 12.3.5, Auto-Baud-Rate Detection and Section 12.3.6, Manual Baud Rate Selection. Refer to those sections for details on how to program these registers.

Bit DLAB in the LCR register must be set high before the Divisor Latch registers can be accessed.

A Divisor value of 0 in the Divisor Latch Register is not allowed. A value of 0 has the affect of disabling the UART. The reset value of the divisor is 02.

#### Table 373. UART x Divisor Latch Low Register - (UxDLL)



#### Table 374.UART x Divisor Latch High Register - (UxDLH)





## 12.4.12 UART x FIFO Occupancy Register

This register shows the number of bytes currently remaining the Receive FIFO. It can be used by the processor to determine the number of trailing bytes to remove from the receive FIFO when the Character Time-out Interrupt is detected. Refer to Section 12.3.2, "Removing Trailing Bytes In Interrupt Mode" on page 656. The FOR register is incremented once for each byte of data written to the Receive FIFO and decremented once for each byte read.

#### 20 12 28 24 16 8 4 31 Attribu ğ RW = Read/Write Intel XScale<sup>®</sup> Core internal bus address Attribute Legend: Unit # RV = Reserved RC = Read Clear FFFF F724H (DLAB=x) 0 PR = Preserved RO = Read Only FFFF F764H (DLAB=x) 1 RS = Read/Set NA = Not Accessible Bit Default Description 31:7 000 0000h Reserved 000 00002 FOR[6:0]: Number of bytes (0-63) remaining in the Receiver FIFO 6:0

#### Table 375. UART x FIFO Occupancy Register - (UxFOR)

## 12.4.13 UART x Auto-Baud Control Register

This read/write register has no effect on the UART. It is intended as a scratchpad register for use by programmers.

#### Table 376. UART x Auto-Baud Control Register - (UxABR)



Intel<sup>®</sup> 80333 I/O Processor Developer's Manual **UARTs** 



#### 12.4.14 **UART x Auto-Baud Count Register**

The Auto-Baud Count register stores the number of 33.334 MHZ clock cycles within a start bit pulse. This value is then used by the processor or auto-baud circuitry within the UART to calculate the baud rate. When Auto-Baud mode and Auto-Baud Interrupts are enabled, the UART interrupt the processor with the Auto-Baud Lock Interrupt after it has written the count value into the ACR. The value is written regardless of the state of the auto-baud UART program bit.



#### Table 377.



## 13

This chapter describes the components comprising the Intel<sup>®</sup> 80333 I/O processor (80333) arbitration, including one Internal Bus Arbiter, one PCI Selector, and two Multi-Transaction Timers. The operation modes, setup, and implementation of these components are described in this chapter.

## 13.1 Arbitration Overview

The 80333 interfaces to one PCI bus and contains an internal PCI-like bus. Therefore, there are two buses which need an arbitration mechanism. Figure 137 illustrates all the potential bus initiators and which arbitration components are responsible for them.

#### Figure 137. Intel<sup>®</sup> 80333 I/O Processor Arbitration Block Diagram



The four components which comprise 80333 arbitration are:

- Internal Bus Arbitrator (Section 13.2, "Internal Bus Arbiter Overview" on page 682) -Arbitrates between multiple initiators. The arbitration scheme is a round-robin with priority/promotion capabilities. The 80333 requires one PCI arbiter: the Internal Bus Arbiter.
  - Internal Bus Arbiter (IARB) arbitrates between six potential internal bus initiators (ATU, two DMA Channels, Application Accelerator, the PBI, and the Bus Interface Unit for the core).
- Multi-Transaction Timer (Section 13.4.2, "Multi-Transaction Timer Register 1 MTTR1" on page 691Section 13.4.3, "Multi-Transaction Timer Register 2 MTTR2" on page 692) Two muti-transaction timers (MTTR[2:1]) control the arbitration control for the internal bus initiators. The MTTRs counts internal bus cycles, an initiator uses across back-to-back transactions. The arbiter continues to assert GNT# to current initiator until the respective MTTR expires, or the initiator deasserts its **REQ**#. The 80333 implements one MTTR for the internal Intel XScale<sup>®</sup> core Bus Interface Unit and one for all peripheral initiators on the internal bus.
- Arbitration Configuration Registers (Section 13.4, "Register Definitions" on page 689) -Priority and latency timer values for arbitration mechanism are programmable, as defined in Arbitration Configuration Registers.



## **13.2** Internal Bus Arbiter Overview

The *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0 requires a central arbitration resource for each PCI bus within a system environment. This section details the operation of the Internal Bus Arbiter block.

The Internal Bus Arbiter supports:

- Three priority levels for each bus initiator
- A "fairness" algorithm which ensures that each potential bus initiator is granted access to the PCI bus independent of other requests
- Hidden, access-based arbitration

PCI uses the concept of access-based arbitration rather than the traditional time slot approach. When a bus initiator requires the PCI bus for a transaction, the device requests the arbitration logic for the PCI bus. PCI arbitration consists of a simple **REQ**# and GNT# handshake protocol. When a device requires the Internal Bus, it asserts its **REQ**# output. The arbitration unit allows the requesting agent access to the bus by asserting that agent's GNT# input.

PCI arbitration is a hidden arbitration scheme where the arbitration sequence occurs in the background while another bus initiator may currently control the bus. Hidden arbitration has the advantage of not consuming any bus bandwidth for arbitration overhead.

The arbiter is required by the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0 to implement a "fair" arbitration algorithm. The Internal Bus Arbiter's algorithm guarantees there is only one GNT# active on the Internal bus at any one time.



## 13.2.1 Theory of Operation

The arbiter's behavior is fully compliant with the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0.

### 13.2.1.1 **Priority Mechanism**

The Internal Bus Arbiter supports six bus initiators on the Internal Bus. Each request can be programmed to one of three priority levels or be disabled. Application software programs the Internal Arbiter Control Register (IACR) to set the initial priority for each bus initiator. The arbiter promotes the bus initiator priority levels using a round-robin scheme.

Figure 138 is an example showing the three priority levels and the reserved slots for the promoted requester.



#### Figure 138. Arbitration Example

In Figure 138, the bus initiators are initially programmed to the priorities shown in Table 378. The IACR defines the initial priority levels for the IARB.

March 2005



Table 385 shows the 2-bit values that correspond to each priority level. A priority level of  $11_2$  effectively disables the associated device by removing it from the arbitration sequence. A device programmed with a  $11_2$  priority does not receive a grant to gain access to the bus.

#### Table 378.Priority Programming Example

Bus Initiator	Programmed Priority
Intel XScale <sup>®</sup> core (BIU)	High - 00 <sub>2</sub>
Memory Controller Unit	High - 00 <sub>2</sub>
DMA Channel 0	Medium - 01 <sub>2</sub>
ATU	Medium - 01 <sub>2</sub>
Application Accelerator	Low - 10 <sub>2</sub>
DMA Channel 1	Disabled - 11 <sub>2</sub>

The priority of the individual bus initiator determines the level to which the device is placed in the round-robin scheme. The programmed priority determines the starting priority or the lowest priority the device is. When the application programs the device for low priority, the device may be promoted up to medium and then high priority until it is granted the Internal Bus. Once the IARB grants the bus and the device asserts FRAME#, the device is reset to its initially programmed priority.

*Note:* When a low priority initiator requests the bus, with no other higher priority agent requesting the bus, that initiator is granted the bus the following clock. The promotion mechanism does not consume bus cycles.

The round-robin arbitration scheme supports three levels of round-robin arbitration: low, medium, and high priority. Using a round-robin mechanism ensures there is a winner for each priority level. To enforce the concept of fairness, a slot is reserved for the winner of each priority level (except the highest) in the next higher priority level. When the winner of a priority level is not granted the bus during that particular arbitration sequence, it is promoted to the next higher level of priority.

#### 13.2.1.2 Priority Example with Three Bus Initiators

Table 379 illustrates an example of bus arbitration with three bus initiators:

#### Table 379. Bus Arbitration Example – Three Bus Initiators

Priority Level	Initial State	Winning Bus Initiator							
		Α	В	Α	С	Α	В	Α	С
High	А	В	A	С	A	В	A	С	Α
Medium	В	С	С	В	В	-	С	В	В
Low	С	_	-	_	-	С	-	-	-

NOTE: In this example, all bus initiators are continually requesting the bus.

Each of the bus initiators (A, B, and C) are constantly requesting the bus and each is at a different priority level. The top row of Table 379 lists the current bus initiator/winner of the highest priority group. The three rows labelled as high, medium and low represent the actual priority levels that devices are currently at based on either their initial programmed priority or promotion through the levels. For example, device C starts out at low priority. Because it is the only device at this priority, it is the winner at low priority and is promoted to medium priority. Later, it wins at the medium priority level (against device B) and is promoted to high priority where it wins the level (against device A) and the bus. Device C is then put back at its programmed priority of low and starts the cycle over.

Continuing with Table 379, the winning bus initiator pattern would follow as: ABACABACABACABAC


# 13.2.1.3 **Priority Example with Six Bus Initiators**

Table 380 illustrates an example of bus arbitration with six bus initiators:

#### Table 380. Bus Arbitration Example – Six Bus Initiators

Priority Level	Initial State	Winning Bus Initiator								
		А	В	С	Α	В	D	А	В	E
High	AB	BC	AC	AB	BD	AD	AB	BE	AE	AB
Medium	CD	DE	DE	DE	CE	CE	CE	CDF	CDF	CDF
Low	EF	F	F	F	F	F	F	_	-	_

**NOTE:** In this example, all bus initiators are continually requesting the bus.

Each of the six bus initiators (A through F) are constantly requesting the bus. There are two initiators programmed at each priority level. The top row of Table 380 lists the current bus initiator/winner of the highest priority group. The three rows labelled as high, medium and low represent the actual priority levels that devices are currently at based on either their initial programmed priority or promotion through the levels.

Continuing with Table 380, the winning bus initiator pattern would follow as: ABCABDABFABCABDABEABCABDABF

# 13.2.1.4 Arbitration Signalling Protocol

The Internal Bus Arbiter interfaces to all requesting agents on the bus through the **REQ**#/GNT# handshaking protocol. A bus initiator asserts its **REQ**# to request ownership of the Internal Bus. When the arbiter determines an agent may use the bus, it asserts the agent GNT# input. Agents must only assert **REQ**# to signal a true need for the bus and not to *reserve* the bus. Figure 139 illustrates arbitration between initiators of equal priority. This Internal Bus Arbiter only operates in the PCI-X mode as defined in *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0.



#### Figure 139. Arbitration Between Three Initiators

An agent can be granted the bus while a previous bus owner still has control of the Internal Bus (hidden arbitration). The arbiter is responsible for deciding which device is granted the bus next while each initiator is responsible for determining when the bus actually becomes free and is allowed to initiate its transaction by asserting FRAME#.

*PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0 states an initiator may deassert its **REQ**# pin before the arbiter grants the Internal bus to that initiator. When an initiator deasserts its **REQ**# pin, the Internal Bus Arbiter re-arbitrates, giving bus ownership to the next initiator based on priority algorithm defined in Section 13.2.1.1, "Priority Mechanism" on page 683.

# *Note:* The Internal Bus Arbiter arbitrates the Internal Bus by checking REQ[8:0]# on every cycle independent of any transactions on the bus.

The Internal Bus Arbiter may deassert an agent's GNT# on any clock. An agent must ensure its GNT# is asserted two clocks prior to the clock edge where it initiates a transaction by asserting FRAME#. When GNT# is deasserted, the transaction may not proceed.



By monitoring **REQ[8:0]**#, the arbiter can control the arbitration algorithm described in Section 13.2.1.1, "Priority Mechanism" on page 683. The arbiter asserts GNT# two clocks after **REQ**# is asserted when the agent has won the bus. An example arbitration flow is:

#### Table 381.Arbitration Flow

Cycle	Event
0	The arbiter is currently driving Initiator_As GNT#. The arbitration flow is independent of whether or not Initiator_A is involved with a transaction. For example, the Internal Bus could be parked with Initiator_A.
1	Initiator_B asserts its <b>REQ#</b> for PCI bus ownership. The arbitration logic calculates that Initiator_B has a higher priority than Initiator_A.
3	The arbiter deasserts GNT# for Initiator_A since Initiator_B is higher priority.
4	The arbiter asserts GNT# for Initiator_B.
7	When Initiator_B drives FRAME#, any of the priority winners that were not granted the bus are promoted to a higher priority level when the reserved promotion slot is unoccupied (see Section 13.2.1.1, "Priority Mechanism" on page 683).

### 13.2.1.5 Internal Bus Arbitration Parking

Arbitration parking occurs when the arbiter asserts GNT# to a selected Internal Bus agent and no agent is currently using or requesting the bus.

Upon reset, the IARB parks the internal bus with the Intel XScale<sup>®</sup> core Bus Interface Unit (BIU). After an initiator requests, and is granted the bus, the arbiter parks the bus with that initiator. In other words, the last initiator that was granted the bus is responsible for parking.

When the Internal Bus is parked during an idle state, the parked agent loses the bus when the arbiter asserts another agent's GNT#. The parked agent relinquishes the bus and stops driving the address and command signals in one clock. When the arbiter deasserts an agent's GNT# on clock N-2, the agent can still initiate a normal bus transaction by driving FRAME# during clock N.



# 13.2.2 Intel XScale<sup>®</sup> Core Arbitration

The Intel XScale<sup>®</sup> core has an inherently small burst size. For this reason, a busy internal bus could inhibit data traffic for the core. To address this issue, the IARB implements a Multi-Transaction Timer (MTT1) which allocates a minimum timeslice where the Intel XScale<sup>®</sup> core BIU can request the IARB to keep the core processor's GNT# asserted potentially across multiple Internal Bus transactions. Refer to Section 13.2.2.1, "Multi-Transaction Timers" on page 688 for details.

#### 13.2.2.1 Multi-Transaction Timers

The Internal Arbiter incorporates two Multi-Transaction Timers (MTT1 and MTT2) allowing the BIU or the other internal bus agents a guaranteed time-slice before losing arbitration to another internal bus agent. PCI is a transaction based protocol. In a system with long bursting agents, an agent such as the BIU with a small burst size could get starved.

The MTT overcomes these potential bottlenecks by guaranteeing a programmed timeslice with which the BIU (MTTR1) or other internal bus agents (MTTR2) are granted the internal bus. For example, once the IARB grants the internal bus to the BIU and the BIU initially asserts **I\_FRAME#**, MTT1 is loaded with the value programmed in the Multi-Transaction Timer Register 1 (MTTR1) and begins to decrement. The arbiter does not remove the BIU grant unless:

- BIU no longer requests bus by deasserting core processor **REQ**#.
- BIU continues to drive core processor **REQ**# and MTT expires.
- *Note:* Even when a higher-priority initiator requests the internal bus, the arbiter does not deassert the BIU grant unless either of the above conditions occur.

Figure 140 illustrates an example of how the BIU uses MTT1 for efficient back-to-back transactions. For this example, the MTTR1 is programmed for 13 cycles.

#### Figure 140. Intel XScale<sup>®</sup> Core Back-to-Back Transactions with MTT1 Enabled



- *Note:* Multi-Transaction Timers keep track of the maximum time that an internal bus agent may continue to initiate new transactions on the internal bus per arbitration cycle. The MTT governs the internal arbiter.
- *Warning:* When the MTTRx are programmed with zero, the MTTx are effectively disabled. However, this is not recommended to avoid potential resource lock out conditions on the internal bus; this is due to the one-deep transaction queue of the MCU. The minimum (default) value for MTT1 is recommended to be 98H (152 clocks) and the minimum (default) value for MTT2 is recommended to be 38H (56 clocks).



# 13.3 Reset Conditions

Table 382 shows all the arbitration blocks and the signal responsible for resetting its logic:

#### Table 382. Arbitration Reset

Arbitration Block	Reset With
Internal Arbiter (IARB)	I_RST#
Multi-Transaction Timer 1	I_RST#
Multi-Transaction Timer 2	I_RST#

**I\_RST#** moves all the internal agents to their programmed priority levels and starts the round robin arbitration sequence on the lowest number device at each priority level.

# 13.4 Register Definitions

Table 383 lists the Arbitration configuration registers which are detailed further in proceeding sections.

#### Table 383. Arbiter Register

Section, Register Name - Acronym (Page)
Section 13.4.1, "Internal Arbitration Control Register - IACR" on page 690
Section 13.4.2, "Multi-Transaction Timer Register 1 - MTTR1" on page 691
Section 13.4.3, "Multi-Transaction Timer Register 2 - MTTR2" on page 692



# 13.4.1 Internal Arbitration Control Register - IACR

The Internal Arbitration Control Register (IACR) sets the arbitration priority of each device that uses the internal bus. This register is part of the local arbitration configuration register space and is accessible from the 80333 core.



#### Table 384. Internal Arbitration Control Register - IACR

Each device is given a 2-bit priority shown in Table 385. The default values for the IACR give all the internal bus initiators the highest priority.

#### Table 385. Programmed Priority Control

2-Bit Programmed Value	Priority Level
002	High Priority
012	Medium Priority
10 <sub>2</sub>	Low Priority
11 <sub>2</sub>	Disabled



# 13.4.2 Multi-Transaction Timer Register 1 - MTTR1

The Multi-Transaction Timer Register 1 defines the duration with which the Intel XScale<sup>®</sup> core Bus Interface Unit retains the core processor GNT# across back-to-back transactions. This is an 8-bit value allowing up to 255 dedicated internal bus cycles for as long as the core processor **REQ**# is asserted. A value of zero effectively disables the MTTR1. This register is part of the local arbitration configuration register space and is accessible from the 80333 core.



#### Table 386. Multi-Transaction Timer Register - MTTR1



# 13.4.3 Multi-Transaction Timer Register 2 - MTTR2

The Multi-Transaction Timer Register 2 defines the duration with which agents other than the Intel XScale<sup>®</sup> core retain its respective GNT# across back-to-back transactions. This is an 8-bit value allowing up to 255 dedicated internal bus cycles for as long as the internal agent's **REQ**# is asserted. A value of zero effectively disables the MTTR2. This register is part of the local arbitration configuration register space and is accessible from the 80333 core.



Table 387. Multi-Transaction Timer Register - MTTR2

# **Standard Hot-Plug Controller**

# 14.1 Introduction

Intel<sup>®</sup> 80333 I/O processor(80333) integrates a standard Hot-Plug controller on the B PCI segment only, that is compliant with the *PCI Standard Hot-Plug Controller and Subsystem Specification*, Revision 1.0. This standard applies to both PCI and PCI-X modes of operation. The new standard is motivated by the concept that standardization of more features of the *PCI Hot-Plug Specification*, Revision 1.0 will reduce both the cost of hardware and software development time. Differences between the PCI Hot-Plug model and the standard Hot-Plug model include standardization of the following features:

- Power and attention indicators
- Manually-operated Retention Latch (MRL)
- MRL Sensor
- S/W Programming model
- Slot Numbering
- Attention button

Two new register sets are defined for the SHPC Capabilities List and the SHPC Working Register Set. The new registers are defined as a PCI-to-PCI Bridge capability of the B Segment PCI-to-PCI Bridge of 80333 and not as a separate PCI controller device. The new specification also fixes the architectural proximity of the SHPC function to the PCI slots it controls. Per the new specification, the 80333 may only control slots on its secondary bus. The standard Hot-Plug controller is a capability of the PCI-to-PCI Bridge. For details on the operation of the standard Hot-Plug controller, refer to the *PCI Standard Hot-Plug Controller and Subsystem Specification*, Revision 1.0.

# 14.1.1 PCI Standard Hot-Plug Features

- PCI Standard Hot-Plug Controller and Subsystem Specification, Revision 1.0 compliant
- One controller for the B PCI bus segment.
- Parallel mode operation for one slot systems. Slot interface logic not needed.
- In-box and out-of-box PCI Express system solution support with firmware initialization of the Hot-Plug controller.
- ACPI support for 80333-SHPC with in-band PCI Express messaging.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Standard Hot-Plug Controller



# 14.2 System Architecture

*Note:* The Standard Hot-Plug Controller in 80333 provides what is called the 1-slot-no-glue mode where the on-board FET switches, that are otherwise needed to isolate the 80333s PCI buffers from the slot, are no longer needed.

# 14.2.1 Slot Control Signals

This section discusses the signals needed to control the slot side interface. This is a brief discussion of their operation. For a more detailed discussion, refer to the standard *PCI Hot-Plug Specification*, Revision 1.0. These signals are direct inputs and outputs from 80333.

#### 14.2.1.1 Output Control

These four signals are output from 80333 and control the slot side interface to do various functions from connecting the slot to the bus to turning on/off the LEDs.

- **PWREN**: Power enable signal connected to on-board slot-specific power controller to regulate current and voltage of the slot.
- **RST**#: Connected to the PCIRST# pin of the slot.
- ATNLED#: Connected to attention LED which is yellow or amber in color.
- **PWRLED#:** Connected to the power LED which is green in color.

#### 14.2.1.2 Input Control

These six signals are used to poll the status of the Hot-Plug slots for such things are MRL switch closure and power fault in the slot power controller.

- **PWRFLT#:** Power controller fault indication for over-current / under-volt indication. When asserted, the 80333, when enabled, immediately asserts slot reset and disconnects slot from the bus.
- **PRSNT1#, PRSNT2#:** These signals are used to indicate to 80333 whether a card is installed in the slot or not and its power requirements. These signals are directly connected to the present bits on the PCI card.

PRSNT1#	PRSNT2#	Meaning
1	1	No expansion board present
0	1	Expansion board present, 25 W maximum
1	0	Expansion board present, 15 W maximum
0	0	Expansion board present, 7.5 W maximum

- **M66EN:** Determines when a card is capable of running at 66 MHz in the conventional PCI mode.
- MRL#: Manually operated retention latch sensor input. A logic low input that is connected directly to the MRL sensor. When asserted it indicates that the MRL latch is closed. When a platform does not support MRL sensors, this must to be wired to a low logic level (MRL closed).
- **ATTN#:** Attention button input signal connected to the attention button of the slot. When low, indicates that the operator has requested attention. When attention button is not implemented, then this input must be wired to a high logic level.



# 14.2.2 Parallel Mode 1-slot No-Glue Operation

In the parallel 1-slot no-glue mode, 80333 provides 4-slot control outputs and 6-slot control inputs for the slot it can control. 80333 operates in this mode, when the number of slots implemented is 1, as programmed into the slot configuration register. Those signals required for 1-slot no-glue parallel mode are listed in Table 388.

#### Table 388. Parallel 1-slot No-Glue Mode Hot-Plug Pins

Signal	Description					
	Inputs					
B_HMRL#	Manually operated retention latch sensor input. A logic low input that is connected directly to the MRL sensor. When asserted it indicates that the MRL latch is closed. When a platform does not support MRL sensors, this must to be wired to a low logic level (MRL closed).					
B_HPWRFLT#	Power controller fault indication for over-current / under-volt indication. When asserted, the 80333, when enabled, immediately asserts reset to the slot and disconnects the slot from the bus.					
B_HPRSNT1# B_HPRSNT2#	These signals are used to indicate to 80333 whether a card is installed in the slot or not and its power requirements. These signals are directly connected to the present bits on the PCI card.					
B_HBUTTON#	Attention button input signal connected to the slot's attention button. When low, indicates that the operator has requested attention. When attention button is not implemented, then this input must be wired to a high logic level.					
B_HM66EN	Determines when a card is capable of running at 66MHz in the conventional PCI mode.					
	Outputs					
B_HPWRLED#	Connected to the power LED which is green in color.					
B_HATNLED#	Connected to attention LED which is yellow or amber in color.					
B_HPWREN	Power enable signal connected to on-board slot-specific power controller to regulate current and voltage of the slot.					
B_HRESET#	Connected to the PCIRST# pin of the slot.					



# 14.2.3 One-Slot-No-Glue Parallel Mode Operation

When only 1 slot is implemented, the bus isolation switches are not required as there are no other devices on the PCI bus other than that one slot. There is no reason to isolate the card from 80333s I/O buffers. 80333 enters this mode when **B\_HSLOT[3:0]** reset straps are a "1111". There are special requirements for how the PCI bus signals are handled in this mode.

### 14.2.3.1 Driving Bus To Ground When PCI Card is Disconnected

When in 1-slot no-glue mode, all PCI signals are to be driven to ground whenever PCI card is disconnected. The signals that must be driven to ground are the following:

- B\_AD[63:0], B\_C/BE[7:0]#, B\_PAR, B\_PAR64, B\_REQ64#, B\_ACK64#.
- B\_FRAME#, B\_IRDY#, B\_TRDY#, B\_STOP#, B\_DEVSEL#, B\_LOCK#.
- B\_RST#
- B\_GNT[3:0]#<sup>1</sup>, B\_REQ[4,3,1,0]#<sup>2</sup>
- B\_PERR#, B\_SERR#.
- B\_CLKO[4:0] (B\_CLKOUT is not driven to ground because it is connected back to B\_CLKIN).
- B\_PME# (for those inputs routed to the B-segment IOAPIC)

These signals will be driven back to their normal PCI levels at various times in the clock connection process. When a card is reconnected to the bus, it follows the following algorithm...

- Power is applied to the card. This does not affect any of the PCI signals which are now being driven to ground.
- After a fixed (refer to PCI Standard Hot-Plug Controller and Subsystem Specification, Revision 1.0) period of time, the clock is connected to the card. When this occurs, B\_CLKO[4:0] will no longer be driven to ground, but will toggle normally (assuming that BIOS has not disabled that particular B\_CLKO pin). In Hot-Plug terms, this is the equivalent of the "CLKEN#" signal.
- After another fixed (refer to SHPC specification) period of time, the bus is connected to the card. When this occurs, the remaining signals listed above, which were driven to ground, are driven to their default values, except for reset, which continues to be driven to ground. The new signal values are listed below:
- B\_AD[63:0], B\_C/BE[7:0]#, B\_PAR, B\_PAR64, B\_REQ64#, B\_ACK64# driven.
- B\_FRAME#, B\_IRDY#, B\_TRDY#, B\_STOP#, B\_DEVSEL#, B\_LOCK# driven to V<sub>CC</sub> for one clock, then tri-stated.
- B\_GNT[3:0] $\#^1$ , B\_REQ[4,3,1,0] $\#^2$  driven to V<sub>CC</sub> for one clock, then tri-stated.
- B\_PERR#, B\_SERR# driven to V<sub>CC</sub> for one clock, then tri-stated.
- B\_PME# and XINT[7:0]# tri-states. (for those inputs routed to the B-segment IOAPIC)
- In Hot-Plug terms, this is the equivalent of the "BUSEN#" signal.
- After a final fixed period of time, the card is taken out of reset. When this occurs, PCIRST# pin is continuously driven to V<sub>CC</sub>. In Hot-Plug terms, this is the equivalent of "PCIRST#" signal.

<sup>1.</sup> B\_GNT[4]# is multiplexed with B\_HRST# which is a valid output for Parallel Hot Plug modes.

<sup>2.</sup> B\_REQ[2]# is multiplexed with B\_HM66EN# which is a valid input for Parallel Hot Plug modes.

intel

# 14.2.3.2 Aborting Outbound PCI Cycles When Card is Disconnected

When a PCI card is not present in a multi-slot system, it has been isolated. This means that all cycles destined for that particular card (peer traffic or other CPU based traffic) will master abort on the PCI bus because no DEVSEL# will be driven. To be consistent in a single-slot system, the 80333 must master abort cycles that are destined for that PCI bus when the card is disconnected. Therefore, the buffer interface will have to internally master abort all outbound transactions destined for that PCI bus until a card is connected again.



# 14.2.4 Switch Debounce

For the switch inputs, 80333 provides a 8 ms debounce timer internally (board does not have to provide one). A change on any of the switch inputs initiates an 8 ms glitch filter timer. When a change is detected on any slot switch input, the switch input is at an 8 ms time interval. When the switches remain stable, then a change is state is confirmed. When any switch is still changing state, the glitch filter counter is cleared and the sequence restarts.

# 14.2.5 Enable/Disable Sequencing

The SHPC in the 80333 follows the enable and disable sequence as specified by the PCI SHPC specification.

# 14.2.6 Secondary Bus Initialization and SHPC Role

#### 14.2.6.1 In-box Architecture

With the in-box solution, it is assumed 80333 would be an embedded part of the system board and system BIOS has complete knowledge of PCI buses below the 80333, like bus loading characteristics, the slot numbering scheme on the bus etc. In such an architecture, whenever the SHPC in 80333 is initialized (power-on or a PCI Express bus reset) the system BIOS/firmware would be invoked to initialize the SHPC working register set with board-specific information (Hardware-Initiation Registers) and also initialize the PCI bus beneath 80333 to the proper mode and frequency. Whenever the standard Hot-Plug controller is reset, the slot interface outputs are reset to the following:

- RST# is asserted.
- PWREN is de-asserted (slot power is removed).
- All PWRLED# and ATNLED# outputs are set to OFF.

When the **B\_HSLOT[3]** reset strap is "1" (Hot-Plug is enabled), then whenever the SHPC is initialized, the PCI bus will power up operating at 33 MHz PCI and the Hot-Plug slot isolated from the bus. The platform BIOS/firmware could later determine the capabilities of the non-Hot-Plug PCI cards (reading the PCI-X and 66 MHz capability bits in the PCI register space of the cards) and also the capabilities of the inserted Hot-Plug card, for PCI-X capability, or PCI capability at 66 MHz, and then could reset the PCI bus to operate in the new mode. The software could execute a set bus frequency/mode command to achieve the mode.

#### 14.2.6.2 Remote-I/O-Box Architecture

In a cabled environment where 80333 can be hot-added onto PCI Express, and in systems that support ACPI, the initialization of 80333 SHPC can be handled with the OSHP method described in the SHPC specifications.

# 14.2.7 M66EN Pin Handling

It should be noted that when the 80333 is in a 1-Slot mode and the slot is disconnected, it will never drive the **B\_M66EN** pin. This is to allow the Hot-Plug controller the ability to correctly sample the M66EN pin on the slot when the PCI bus is grounded (not connected) but the PCI card is powered on. In this mode, it is recommended that the M66EN pin be pulled up to the 3.3 V PCI slot power rail which is controlled by the Hot-Plug controller.

# intel

# 14.3 Interrupts

# 14.3.1 MSI and Pin Interrupts

SHPC in 80333 can either be enabled to generate an MSI interrupt or can generate an interrupt to the internal IOAPIC to be routed through it to the MCH. The SHPC interrupt is routed to the interrupt input 23 of the IOAPIC. These two interrupts are mutually exclusive. The message for MSI comes out of the message address and data registers in 80333 MSI capability registers. The message enable bit in the message control register of the capability registers either enables MSI or interrupt routing through IOAPIC.

When MSI is enabled, SHPC generates MSI messages only when there are no other interrupt events pending (interrupt pending bit is set). When there is already an interrupt pending when another event happened, that event is inhibited from generating an MSI message. Software needs to comprehend this to do a status read of pending interrupts, after it cleared an interrupt source that it currently serviced.

# 14.3.2 ACPI Support

On platforms where the platform ACPI-compliant firmware controls the SHPC, instead of native-OS support for SHPC, the SHPC interrupts need to be converted to an ACPI interrupt (that goes to an SCI interrupt to the processor). In the presence of native OS-support, this interrupt steering is not needed. In previous platforms where Hot-Plug was firmware controlled, this was done by converting the Hot-Plug interrupt into a side-band pin interrupt which was then directly routed to the GPE# pin in the ICH. In the cabled PCI Express world, things are done in-band and hence 80333 generates an in-band message. The interrupt steering is controlled through a BIOS-specific bit SGME (bit 5 in 80333\_CONFIG). This bit programs the Hot-Plug controller to generate an SCI message to MCH instead of an MSI or a pin interrupt to the internal IOAPIC. This SCI message is ultimately routed by MCH to the ICH over the HL inter-connect (DO\_SCI message). On assertion of the GPE# pin to the ICH or the ICH receiving the DO\_SCI special cycle, the ICH raises an SCI interrupt to the processor, which in turn wakes up the ACPI handler.

All logic in the SHPC function the same as for normal MSI or pin interrupts. PCI Express defines the specifics of the in-band signaling mechanism through a pair of vendor-specific messages ASSERT/DEASSERT GPE. 80333 generates the ASSERT/DEASSERT GPE message (when enabled) based on the SHPC interrupt. The asserting edge of the interrupt signal creates the ASSERT GPE message and the deasserting edge causes the DEASSERT GPE message. The Requester ID of the message is constructed using the primary bus number (0).



# 14.4 Error Handling

The standard Hot-Plug controller can detect a variety of error conditions (refer to SHPC specification for details) and it can be programmed to either do a SERR cycle on PCI Express or raise an interrupt.

# 14.5 Assumptions and Intel<sup>®</sup> 80333 I/O Processor Requirements

# 14.5.1 MRL Opening During the Sequence

While executing a enable or disable sequence, when the MRL of the card is opened then the 80333 performs the auto power down for that slot after executing the current enable/disable operation. As the maximum time required to enable is disable is 319 ms, the maximum delay between MRL open and auto-power down would be less than 320 ms.

# 14.5.2 Power Fault

The power controller/slot control logic is responsible for removing power from the slot when a power fault happen. In the 1-slot-no-glue mode, 80333 will assert reset and ground the bus when the power fault happens. The PWREN signal from 80333 is not asynchronously deasserted on a slot power fault and deasserted only under software control.

80333 would keep the slot that has the power fault isolated from the bus, until the bits 17 and 20 in the Slot Event latch field of the Slot Configure register are cleared and the software issues a enable slot command.

# 14.6 SHPC Working Register Set

The software determines that the 80333 supports the Standard Hot-Plug controller by locating the capability list item in the configuration space of the 80333 PCI Express-to-PCI Bridge for PCI bus segment B, at offset 48h. 80333 does not exhibit the SHPC capability when Hot-Plug is disabled (**B\_HSLOT[3]** is "0"). 80333 implements the Capability list item and the Working register set required to support the five slots for PCI bus segment B. The working register set is accessible either from memory space through the 64-bit BAR (SHPC\_BAR) located at offset 0x10h in the 80333 PCI Express-to-PCI Bridge for PCI bus segment B configuration space or from 80333s configuration space through the SHPC index and data registers at offset 0x4Ah (SHPC\_DWORD, SHPC\_DWSEL).

Upon a Power Good reset or Primary reset, the platform firmware is required to initialize the hardware registers in the standard Hot-Plug working register set, to setup platform specific information for OS/driver use later and also to enable the Hot-Plug slot. Writes to initialize these registers must be a full DWORD, writes of any other size (Byte, WORD) are not defined. The SHPC requests 4 K of memory for accesses to the working register set and the working register set in the 80333 is located at offset 0h. This is accomplished by setting the SHPC base offset register value in the working register set to 0h. Therefore software can use the value in the Base address register to access the working register set of the SHPC. Table 389 gives the layout of 80333 standard Hot-Plug working register set. 80333 supports working register set corresponding to PI=1 (compliant with *SHPC Specification*, Revision 1.0).

# 14.6.1 Access Disposition

Only aligned-DWORD reads and writes are allowed to the standard Hot-Plug working register set. All other read accesses receive a completor abort response on PCI Express. Burst writes are allowed to the SHPC working register set.



# 14.6.2 Register Description

#### Table 389. SHPC Working Register Set Layout

31		C	Byte Offset		
	SHPC Base C	Offset Register	00h		
	Slots Availab	ole 1 Register	04h		
	Slots Availab	ole 2 Register	08h		
	Slot Configura	ation Register	0Ch		
SHPC Program I/F	SHPC Program I/F SHPC MSI Control Secondary Bus Configuration Register				
SHPC Command	SHPC Command Status Register SHPC Command Register				
	Interrupt Loc	ator Register	18h		
	SERR Locator Register				
Controller SERR-INT Register					
Logical Slot Register					
Reserved					

#### 14.6.2.1 Standard Hot-Plug Controller Base Offset Register -SHPC\_BASEOFF (Offset 00)

This register specifies the offset from SHPC\_BAR to access working register set.

#### Table 390. Standard Hot-Plug Controller Base Offset Register - SHPC\_BASEOFF

Bits	Туре	Reset	Description
31:0	RO	0000 0000H	SHPC Base Offset – This field contains the byte offset that must be added to the 64-bit Base Address register SHPC_BAR in 80333 configuration space, to access the SHPC Working Register set using memory-mapped accesses. 80333 has the working register set starting at offset 0.

#### 14.6.2.2 Slots Available Register 1 - SLOTS\_AVAIL1 (Offset 04)

This register is initialized by platform firmware with platform specific loading information to be later used by OS/driver.

#### Table 391. Slots Available Register 1 - SLOTS\_AVAIL1

Bits	Туре	Reset	Description
31:29	PR	0002	Preserved.
28:24	RWO	0 0000 <sub>2</sub>	Number of Slots Available (133 MHz PCI-X) – Maximum number of Hot-Plug slots available to be enabled when the bus is running at 133 MHz PCI-X mode.
23:21	PR	0002	Preserved.
20:16	RWO	0 0000 <sub>2</sub>	Number of Slots Available (100 MHz PCI-X) – Maximum number of Hot-Plug slots available to be enabled when the bus is running at 100 MHz PCI-X mode.
15:13	PR	0002	Preserved.
12:8	RWO	0 0000 <sub>2</sub>	Number of Slots Available (66 MHz PCI-X) – Maximum number of Hot-Plug slots available to be enabled when the bus is running at 66 MHz PCI-X mode.
7:5	PR	0002	Preserved.
4:0	RWO	оH	Number of Slots Available (33 MHz conventional) – Maximum number of Hot-Plug slots available to be enabled when the bus is running at 33 MHz conventional mode.



## 14.6.2.3 Slots Available Register 2 - SLOTS\_AVAIL2 (Offset 08)

This register is initialized by platform firmware with platform specific loading information to be later used by OS/driver.

#### Table 392. Slots Available Register 2 - SLOTS\_AVAIL2

Bits	Туре	Reset	Description
31:5	PR	000 0000H	Preserved.
4:0	RWO	0 0000 <sub>2</sub>	Number of Slots Available (66 MHz conventional.) – Maximum number of Hot-Plug slots available to be enabled when the bus is running at 66 MHz conventional mode

### 14.6.2.4 Slot Configuration Register - SLOT\_CONFIG (Offset 0C)

Bits	Туре	Reset	Description
31	RWO	02	Attention Button Implemented – This bit specifies whether the Hot-Plug slots controlled by this SHPC implement the optional Attention Button. When this bit is set, Attention Buttons are implemented on every slot controlled by this SHPC.
30	RWO	02	<ul> <li>MRL Sensor Implemented – This bit specifies whether MRL Sensors are implemented on the Hot-Plug slots controlled by the SHPC.</li> <li>1 = the platform provides an MRL Sensor for each slot controlled by this SHPC</li> <li>0 = MRL Sensors not implemented</li> </ul>
29	RWO	0 <sub>2</sub>	<ul> <li>Physical Slot Number Up/Down – This bit specifies the direction of enumeration of external slot labels, beginning with the value in the Physical Slot Number field of the Slot Configuration Register.</li> <li>1 = each external slot label increments by 1 from the value in the Physical Slot Number field.</li> <li>0 = each external slot label decrements by 1 from the value in the Physical Slot Number</li> </ul>
			field.
28:27	PR	002	Preserved.
26:16	RWO	000H	Physical Slot Number – This field specifies the physical slot number of the device addressed by the First Device Number. This field must be hardware initialized to a value that assigns all slots (controlled by this SHPC) a slot number that is globally unique within the chassis.
15:13	PR	0002	Preserved.
12:8	RWO	0 0000 <sub>2</sub>	First Device Number – This field contains the device number assigned to the first Hot-Plug slot on this bus segment.
7:5	PR	0002	Preserved.
4:0	RO	0 0000 <sub>2</sub>	NOTE: Number of Slots Implemented – This field contains the number of Hot-Plug slots connected to the SHPC (that is, the number of slots controlled by the SHPC). This field must not return a value of 0. (When the controller does not control any slots in the system, the SHPC Capabilities List Item must not appear in the Capabilities List.)

#### Table 393. Slot Configuration Register - SLOT\_CONFIG



## 14.6.2.5 Secondary Bus Configuration Register - SBUS\_CFG (Offset 10)

Reflects the current speed and mode of PCI bus.

#### Table 394. Secondary Bus Configuration Register - SBUS\_CFG

Bits	Туре	Reset	Description	
15:3	PR	0000H	Preserved.	
2:0	RO	000 <sub>2</sub>	Current Bus Segment Speed/Mode – Indicates the current speed and mode at which the PCI bus segment is operating. • 000 : 33 MHz Conventional Mode • 001 : 66 MHz Conventional Mode • 010 : 66 MHz PCI-X Mode1 • 011 : 100 MHz PCI-X Mode1 • 100 : 133 MHz PCI-X Mode1 • Others : Reserved.	

### 14.6.2.6 SHPC MSI Control Register - SHPC\_MSI\_CNTRL (Offset 12)

SHPC Interrupt message number register.

#### Table 395. SHPC MSI Control Register - SHPC\_MSI\_CNTRL

Bits	Туре	Reset	Description
7:5	RV	0002	Reserved.
4:0	RO	0 0000 <sub>2</sub>	SHPC Interrupt Message Number: Reflects the multiple message enable field of 80333 MSI capability control register (bits 6:4)



## 14.6.2.7 SHPC Programming Interface Register - SHPC\_PIF (Offset 13)

Identifies the format of the working register set.

#### Table 396. SHPC Programming Interface Register - SHPC\_PIF

Bits	Туре	Reset	Description
7:0	RO	01H	SHPC Programming Interface – Identifies the format of the SHPC Working Register set. A value of 01h identifies the SHPC Working Register set format defined in the SHPC 1.0 specification.

### 14.6.2.8 SHPC Command Register - SHPC\_CMD (Offset 14)

Hot-Plug Controller Command register.

#### Table 397. SHPC Command Register - SHPC\_CMD

Bits	Туре	Reset	Description	
15:13	PR	0002	Preserved.	
12:8	RW	0 0000 <sub>2</sub>	Target Slot – This field selects the target slot for a Slot Operation command. Software is permitted to write the Command Code and Target Slot fields simultaneously. However, software is not required to write these fields simultaneously. When the fields are not written simultaneously, the Slot Operation command targets the slot associated with the current value in this register. When the command is not a Slot Operation command, this field is ignored. When this field is read, it returns the value that was last written to it, even after the command has completed.	
7:0	RW	00H	Command Code – Command to be executed by the SHPC. Writing to this field triggers the SHPC to begin executing the command. Refer to the specification for command encodings. When read, this field returns the command code that was last written to it, even after the command has completed.	

### 14.6.2.9 SHPC Command Status Register - SHPC\_CMDSTS (Offset 16)

Command status register.

#### Table 398. SHPC Command Status Register - SHPC\_CMDSTS

Bits	Туре	Reset	Description
15:4	PR	000H	Preserved.
3:1	RO	000 <sub>2</sub>	Controller Command Error Code – This field shows the result of the last command completed by the SHPC. This field is updated when the Controller Busy bit transitions from 1 to 0 (indicating a command completion). When the command failed, the appropriate bit is set. When none of the bits in this field are set, the command completed successfully.
0	RO	0 <sub>2</sub>	<ul> <li>Controller Busy – Reflects the state of the SHPC.</li> <li>1 = a command code has been written to the Controller Command register and is being executed. The SHPC ignores writes to the Controller Command register while this bit is set.</li> <li>0 = the SHPC has finished executing a command.</li> </ul>



# 14.6.2.10 Interrupt Locator Register - INT\_LOC (Offset 18)

Interrupt locator register for software to easily identify the source of interrupt.

#### Table 399. Interrupt Locator Register - INT\_LOC

Bits	Туре	Reset	Description
31:2	RV	0000 0000H	Reserved.
1	RO	00 0000 <sub>2</sub>	Slot Interrupt Pending – A set bit in this field indicates an interrupt pending condition on the Hot-Plug slot. An interrupt pending condition occurs when the SHPC detects a Slot Event, and the event interrupt mask bit in the Slot SERR-INT Mask field is cleared. Clearing all bits in the Slot Event Latch field of the slot Logical Slot register clears that slot bit in this field.
0	RO	02	Command Complete Interrupt Pending – The state of this bit is 1 when the Command Completion Detected bit is set indicating a command completion and the Command Complete Interrupt Mask bit located in the Controller SERR-INT register is cleared.

### 14.6.2.11 SERR Locator Register - SERR\_LOC (Offset 1C)

System Interrupt locator register for software to easily identify the source of interrupt.

#### Table 400. SERR Locator Register - SERR\_LOC

Bits	Туре	Reset	Description
31:2	RV	0000 0000H	Reserved.
1	RO	02	Slot SERR Pending – A set bit in this field indicates an SERR pending condition on the Hot-Plug slot. An SERR pending condition occurs when the SHPC detects a slot event capable of generating an SERR and that event SERR Mask bit in the Slot SERR-INT Mask field is cleared. Clearing all bits in the slot Slot Event Latch field that are capable of generating an SERR clears this field.
0	RO	02	Arbiter SERR Pending – The state of this bit is 1 when the Arbiter Time-out Detected bit in the Controller SERR-INT register is set and the Arbiter SERR Mask bit is cleared.



# 14.6.2.12 SERR Interrupt Register - SERR-INT (Offset 20)

The Controller SERR-INT register enables and disables SERR and System generation and reports global controller events.

Bits	Туре	Reset	Description	
31:18	PR	0000H	Preserved.	
17	RWC	02	Arbiter Time-out Detected – This bit is set when the SHPC detects an arbiter time-out.	
16	RWC	02	Command Completion Detected – This bit is set when the Controller Busy bit in the Controller Command Status register transitions from 1 to 0 (indicating a command completion).	
15:4	PR	000H	Preserved.	
3	RW	1 <sub>2</sub>	Arbiter SERR Mask – When this bit is set, arbiter time-out SERRs are masked. This bit is a mask and does not affect whether the Arbiter Time-out Detected bit is set. When this mask is cleared and the Global SERR Mask (bit 1) is clear, arbiter time-out error will cause ERR_NONFATAL message on PCI Express, provided the SERR enable bit is set in the PCICMD register or the NONFATAL message enable bit is set in the PCI Express capability.	
2	RW	1 <sub>2</sub>	Command Complete Interrupt Mask – When this bit is set, command Completion Interrupts are masked. This bit is a mask and does not affect whether the Command Completion Detected bit is set.	
1	RW	1 <sub>2</sub>	Global SERR Mask – When this bit is set, SERR generation from the SHPC is masked.	
0	RW	1 <sub>2</sub>	Global Interrupt Mask – When this bit is set, System Interrupt generation by the SHPC is masked. This bit is a mask and does not affect any bits in the Interrupt Locator register. This bit has no effect on whether the Walk-up Signal is asserted.	

#### Table 401. SERR Interrupt Register - SERR-INT



# 14.6.2.13 Logical Slot Register - (Offset 24)

Gives the status and control for the Hot-Plug slot.

• Slot SERR-INT Mask Field

The Slot SERR-INT Mask field controls masking and unmasking of System Interrupts and system errors generated from events detected by the SHPC.

#### Table 402. Logical Slot Register- Slot SERR-INT Mask Field

Bits	Туре	Reset	Description
31	RV	02	Reserved.
			Connected Power Fault SERR Mask – SERR assertions from Connected Power Fault Detected
30	RW	10	1 = masked.
00		12	The state of this bit has no effect on the state of the Connected Power Fault Detected bit.
			When this bit is clear, then connected power faults can cause ERR_FATAL message on PCI Express provided the SERR enable bit in the PCICMD register is set or the ERR_FATAL enable bit is set in the PCI Express capability.
			MRL Sensor SERR Mask – SERR assertions from MRL Sensor Change Detected
29	RW	1 <sub>2</sub>	0 = not masked. 1 = masked.
			The state of this bit has no effect on the state of the MRL Sensor Change Detected bit.
			Connected Power Fault Interrupt Mask – System Interrupts from Connected Power Fault Detected
29		1 <sub>2</sub>	0 = not masked. 1 = masked.
20			The state of this bit has no effect on the state of the Connected Power Fault Detected bit.
			When this bit is clear, then MRL sensor error condition can cause ERR_FATAL message on PCI Express provided the SERR enable bit in the PCICMD register is set or the ERR_FATAL enable bit is set in the PCI Express capability.
			MRL Sensor Interrupt Mask – System Interrupts from MRL Sensor Change Detected
27	RW	1 <sub>2</sub>	0 = not masked. 1 = masked.
			The state of this bit has no effect on the state of the MRL Sensor Change Detected bit.
			Attention Button Interrupt Mask – System Interrupts from Attention Button Press Detected
26	RW	1 <sub>2</sub>	0 = not masked. 1 = masked.
			The state of this bit has no effect on the state of the Attention Button Press Detected bit.
			Isolated Power Fault Interrupt Mask – System Interrupts from Isolated Power Fault Detected
25	RW	1	0 = not masked. 1 = masked.
			The state of this bit has no effect on the state of the Isolated Power Fault Detected bit.
			Card Presence Interrupt Mask – System Interrupts from Card Presence Change Detected
24	RW	1	0 = not masked. 1 = masked.
			The state of this bit has no effect on the state of the Card Presence Change Detected bit.



• Slot Event Latch Field

The Slot Event Latch field reports all latched events detected by the SHPC.

#### Table 403. Logical Slot Register- Slot Event Latch Field

Bits	Туре	Reset	Description
23:21	RV	000b	Reserved.
20	RWC	0b	Connected Power Fault Detected – 0 = no connected power fault has been detected by the power control circuitry for this slot. 1 = a connected power fault has been detected by the power control circuitry for this slot.
19	RWC	Оb	<ul> <li>MRL Sensor Change Detected –</li> <li>0 = no change in the position of the MRL has been detected.</li> <li>1 = the MRL Sensor bit in the Slot Status field changed state indicating a change in the position of the MRL.</li> </ul>
18	RWC	Оb	<ul> <li>Attention Button Press Detected –</li> <li>0 = the Attention Button has not been pressed.</li> <li>1 = the Attention Button bit in the Slot Status field transitioned from 0 to 1 indicating the Attention Button has been pressed.</li> </ul>
17	RWC	0n	Isolated Power Fault Detected – 0 = no isolated power fault has been detected by the power control circuitry for this slot. 1 = an isolated power fault has been detected by the power control circuitry for this slot.
16	RWC	0n	Card Presence Change Detected – 0 = no change is detected on the PRSNT1#/PRSNT2# bits defined in the Slot Status field. 1 = a change is detected on the PRSNT1#/PRSNT2# bits defined in the Slot Status field.

intel®

• Slot Status Field

Status information about the slot

#### Table 404. Logical Slot Register- Slot Status Field

Bits	Туре	Reset	Description
15:14	RV	00	Reserved.
13:12	RO	-	<ul> <li>PCI-X Capability – Report the current PCI-X capability of the add-in card installed in the slot. These bits are not valid when the slot is empty. When the slot is occupied, these bits are valid regardless of the state of the slot or speed/mode of the bus.</li> <li>00 : PCI Conventional Mode</li> <li>01 : PCI-X Mode1 66 MHz</li> <li>10 : Reserved</li> <li>11 : PCI-X Mode1 133 MHz</li> </ul>
11:10	RO	-	<ul> <li>PRSNT1#/PRSNT2# - Reports current debounced state of PRSNT1# and PRSNT2# pins on the slot. These bits are valid regardless of the state of the slot or speed/mode of the bus.</li> <li>00b - Card Present; 7.5W</li> <li>01b - Card Present; 25W</li> <li>10b - Card Present; 15W</li> <li>11b - Slot Empty</li> </ul>
9	RO	-	<ul> <li>66 MHz Capable – Reports whether add-in card is capable of running at 66 MHz conventional mode and is latched as slot is powered up or enabled, regardless of current speed/mode of the bus.</li> <li>0 = the card is only capable of 33 MHz conventional mode operation.</li> <li>1 = the card is capable of running at 66 MHz conventional mode.</li> <li>This bit is valid only when the slot is occupied and powered or enabled.</li> </ul>
8	RO	-	<ul> <li>MRL Sensor – Reports current MRL state reported by debounced MRL Sensor input signal.</li> <li>0 = the MRL Sensor is reporting that the MRL is closed.</li> <li>1 = the MRL Sensor is reporting that the MRL is open.</li> </ul>
7	RO	-	Attention Button – Reports current state of debounced Attention Button input signal for slot. 0 = the Attention Button is not being pressed. 1 = the Attention Button is being pressed.
6	RO	-	<ul> <li>Power Fault – Reports current state of power fault latch in power controller circuitry for slot.</li> <li>0 = No power fault detected.</li> <li>1 = Power fault (either isolated or connected) detected by the power controller circuitry.</li> </ul>
5:4	RO		Attention Indicator State- Reports current state of Attention Indicator associated with slot.         • 00b - reserved         • 01b - On         • 10b - Blink         • 11b - Off
3:2	RO	-	<ul> <li>Power Indicator State – Reports current state of the Power Indicator associated with the slot.</li> <li>00b – reserved</li> <li>01b – On</li> <li>10b – Blink</li> <li>11b – Off</li> </ul>
1:0	RO	-	<ul> <li>Slot State – This field reports the current state of the slot.</li> <li>00b – reserved</li> <li>01b – Powered Only</li> <li>10b – Enabled</li> <li>11b – Disabled</li> </ul>



# Intel XScale<sup>®</sup> Core and Core Performance Monitoring

This chapter describes the Intel XScale<sup>®</sup> core (ARM\* architecture compliant) and its associated Performance Monitoring facilities within the Intel<sup>®</sup> 80333 I/O processor (80333). The events that are monitored can provide performance information for compiler writers, system application developers and software programmers.

# 15.1 High-Level Overview of Intel XScale<sup>®</sup> Core

The second generation Intel XScale<sup>®</sup> core is designed for high performance and low-power; leading the industry in mW/MIPs. Many of the architectural features added to the Intel XScale<sup>®</sup> core help hide memory latency which 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 which can be configured different for each application.
- · and cache locking.

All these features improve the efficiency of the external bus.

### **15.1.1 ARM Compatibility**

ARM Version 5 (V5) Architecture added floating point instructions to ARM Version 4. The Intel XScale<sup>®</sup> core implements the integer instruction set architecture of ARM V5, but does not provide hardware support of the floating point instructions. The Intel XScale<sup>®</sup> core provides the Thumb instruction set (ARM V5T) and the ARM V5E DSP extensions. Backward compatibility with the first generation of Intel<sup>®</sup> StrongARM\* products is maintained for user-mode applications. Operating systems may require modifications to match the specific hardware features of the 80333 and to take advantage of the performance enhancements added to the second generation Intel XScale<sup>®</sup> core.



# 15.1.2 Features

Figure 141 shows the major functional blocks of the Intel XScale<sup>®</sup> core. The following sections give a brief, high-level overview of these blocks.

#### Figure 141. The Intel XScale<sup>®</sup> Core Architecture Features



### 15.1.2.1 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 data.

#### 15.1.2.2 Memory Management

The Intel XScale<sup>®</sup> core 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



### 15.1.2.3 Instruction Cache

The Intel XScale<sup>®</sup> 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.

### 15.1.2.4 Branch Target Buffer

The Intel XScale<sup>®</sup> core 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.

#### 15.1.2.5 Data Cache

The Intel XScale<sup>®</sup> 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, 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<sup>®</sup> core allows applications to re-configure a portion of the data cache as data RAM. Software may place special tables or frequently used variables in this RAM.



# 15.2 CP14 Registers

Table 405 lists the CP14 registers implemented in the I/O processor.

#### Table 405. CP14 Registers

Register (CRn)	Access	Description
0-3	Read / Write	Performance Monitoring Registers
4-5	Unpredictable	Reserved
6	Read Only	Core Frequency
7	Unpredictable	Reserved
8-15	Read / Write	Software Debug

# 15.2.1 Registers 0-3: Performance Monitoring

The performance monitoring unit contains the following:

- Control Register (PMNC)
- Clock Counter (CCNT)
- Interrupt Enable Register (INTEN)
- Overflow Flag Register (FLAG)
- Event Selection Register (EVTSEL)
- Four Event Counters (PMN0 PMN3)

Opcode\_2 should be zero on all accesses.

These registers can not be accessed by LDC and STC coprocessor instructions.

#### Table 406. Accessing the Performance Monitoring Registers

Description	CRn Register#	CRm Register#	Instruction
(PMNC) Performance Monitor Control Register	0b0000	0b0001	Read: MRC p14, 0, Rd, c0, c1, 0 Write: MCR p14, 0, Rd, c0, c1, 0
(CCNT) Clock Counter Register	0b0001	0b0001	Read: MRC p14, 0, Rd, c1, c1, 0 Write: MCR p14, 0, Rd, c1, c1, 0
(INTEN) Interrupt Enable Register	0b0100	0b0001	Read: MRC p14, 0, Rd, c4, c1, 0 Write: MCR p14, 0, Rd, c4, c1, 0
(FLAG) Overflow Flag Register	0b0101	0b0001	Read: MRC p14, 0, Rd, c5, c1, 0 Write: MCR p14, 0, Rd, c5, c1, 0
(EVTSEL) Event Selection Register	0b1000	0b0001	Read: MRC p14, 0, Rd, c8, c1, 0 Write: MCR p14, 0, Rd, c8, c1, 0
(PMN0) Performance Count Register 0	0b0000	0b0010	Read: MRC p14, 0, Rd, c0, c2, 0 Write: MCR p14, 0, Rd, c0, c2, 0
(PMN1) Performance Count Register 1	0b0001	0b0010	Read: MRC p14, 0, Rd, c1, c2, 0 Write: MCR p14, 0, Rd, c1, c2, 0
(PMN2) Performance Count Register 2	0b0010	0b0010	Read: MRC p14, 0, Rd, c2, c2, 0 Write: MCR p14, 0, Rd, c2, c2, 0
(PMN3) Performance Count Register 3	0b0011	0b0010	Read: MRC p14, 0, Rd, c3, c2, 0 Write: MCR p14, 0, Rd, c3, c2, 0



# 15.2.2 Register 1: Clock Count Register -- CCNT

The Clock counter is reset to '0' by Performance Monitor Control Register (PMNC) or set to a predetermined value by directly writing to it. When CCNT reaches its maximum value 0xFFFF,FFFF, the next clock cycle causes it to roll over to zero and set the overflow flag (bit 6) in PMNC. An IRQ or FIQ is reported when PMNC register bit 6 is enabled.



#### Table 407. Clock Count Register -- CCNT

# 15.2.3 Register 6: Core Clock Configuration Register -- CCLKCFG

Intel XScale<sup>®</sup> core frequency can be determined by reading this register. Opcode\_2 and CRm should be zero.



#### Table 408. Core Clock Count Register -- CCNT



# 15.2.4 Registers 8-15: Software Debug

Software debug is supported by address breakpoint registers (Coprocessor 15, register 14), serial communication over the JTAG interface and a trace buffer. Registers 8 and 9 are used for the serial interface and registers 10 through 13 support a 256 entry trace buffer. Register 14 and 15 are the debug link register and debug SPSR (saved program status register).

Opcode\_2 and CRm are zero.

#### Table 409.Accessing the Debug Registers

Function	CRn (Register #)	Instruction
Access Transmit Debug Register (TX)	0b1000	MRC p14, 0, Rd, c8, c0, 0 MCR p14, 0, Rd, c8, c0, 0
Access Receive Debug Register (RX)	0b1001	MCR p14, 0, Rd, c9, c0, 0 MRC p14, 0, Rd, c9, c0, 0
Access Debug Control and Status Register (DBGCSR)	0b1010	MCR p14, 0, Rd, c10, c0, 0 MRC p14, 0, Rd, c10, c0, 0
Access Trace Buffer Register (TBREG)	0b1011	MCR p14, 0, Rd, c11, c0, 0 MRC p14, 0, Rd, c11, c0, 0
Access Checkpoint 0 Register (CHKPT0)	0b1100	MCR p14, 0, Rd, c12, c0, 0 MRC p14, 0, Rd, c12, c0, 0
Access Checkpoint 1 Register (CHKPT1)	0b1101	MCR p14, 0, Rd, c13, c0, 0 MRC p14, 0, Rd, c13, c0, 0
Access Transmit and Receive Debug Control Register	0b1110	MCR p14, 0, Rd, c14, c0, 0 MRC p14, 0, Rd, c14, c0, 0
Debug Saved Program Status Register	0b1111	MCR p14, 0, Rd, c15, c0, 0 MRC p14, 0, Rd, c15, c0, 0

# 15.3 CP15 Registers

Table 410 lists the CP15 registers implemented in the I/O processor.

### Table 410. CP15 Registers

Register (CRn)	Opcode_2	Access	Description
0	0	Read / Write-Ignored	ID
0	1	Read / Write-Ignored	Cache Type
1	0	Read / Write	Control
1	1	Read / Write	Auxiliary Control
2	0	Read / Write	Translation Table Base
3	0	Read / Write	Domain Access Control
4	-	Unpredictable	Reserved
5	0	Read / Write	Fault Status
6	0	Read / Write	Fault Address
7	0	Read-unpredictable / Write	Cache Operations
8	0	Read-unpredictable / Write	TLB Operations
9	0	Read / Write	Cache Lock Down
10	0	Read / Write	TLB Lock Down
11 - 12	-	Unpredictable	Reserved
13	0	Read / Write	Process ID (PID)
14	0	Read / Write	Breakpoint Registers
15	0	Read / Write	(CRm = 1) CP Access



# 15.3.1 Register 0: ID and Cache Type Registers

Register 0 houses two read-only registers that are used for part identification: an ID register and a cache type register.

The ID Register is selected when *opcode\_2=0*. This register returns the code for the A0 stepping/revision. The low order four bits of the register are the chip revision number and is incremented for future steppings.F

#### Table 411.ID Register

31	30	29	<b>28</b>	27	<b>26</b>	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	1	1	0	1	0	0	1	0	0	0	0	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	1		Revi	sio	n
re	set v	valu	le:	As	Sho	own	1																								
	Bi	its						Α	cce	SS											0	)es	crip	otio	n						
	31	:24		Read / Write Ignored     Implementation trademark (0x69 = 'i'= Intel Corporation)																											
	23	:16		Re	ad /	/Wi	rite I	lgno	orec	1					Arc	chite	ectu	ire v	/ers	ion	= A	RM	* Ve	ersi	on 5	5					
	15	5:4		Read / Write Ignored       Part Number (Implementation Specified)         I/O processor:       Bits[15:12] refer to the processor generation. = 0x4         Bits[11:8] refer to the implementation = 0x0       Bits[7:4] used for implementation derivatives = 0x1																											
	3:	:0		Re	ad /	/ Wi	rite I	lgno	orec	ł					Re Sp A0	visi ecifi ste	on r ied) ppir	num ng =	iber = 0b	for 000	the	pro	oces	sor	. (Iu	nplei	me	nta	tion		

The Cache Type Register is selected when *opcode\_2*=1 and describes the present I/O processor cache.

#### Table 412. Cache Type Register

31	30	29	<b>28</b>	27	<b>26</b>	25	24	23	22	21	<b>20</b>	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	0	1	1	0	0	0	1	1	0	1	0	1	0	1	0	0	0	0	1	1	0	1	0	1	0	1	0
ree	set	valı	16.	Δs	She	wn																									

Bits	Access	Description
31:29	Read-as-Zero / Write Ignored	Reserved
28:25	Read / Write Ignored	Cache class = 0b0101 The caches support locking, write back and round-robin replacement. They do not support address by index.
24	Read / Write Ignored	Harvard Cache
23:21	Read-as-Zero / Write Ignored	Reserved
20:18	Read / Write Ignored	Data Cache Size = 0b110 = 32 kB
17:15	Read / Write Ignored	Data cache associativity = 0b101 = 32
14	Read-as-Zero / Write Ignored	Reserved
13:12	Read / Write Ignored	Data cache line length = 0b10 = 8 words/line
11:9	Read-as-Zero / Write Ignored	Reserved
8:6	Read / Write Ignored	Instruction cache size = 0b110 = 32 kB
5:3	Read / Write Ignored	Instruction cache associativity = 0b101 = 32 kB
2	Read-as-Zero / Write Ignored	Reserved
1:0	Read / Write Ignored	Instruction cache line length = 0b10 = 8 words/line



# 15.3.2 Register 1: Control and Auxiliary Control Registers

Register 1 is made up of two registers, one that is compliant with ARM Version 5 and is referenced by  $opcode_2 = 0x0$ , and the other which is specific to Intel XScale<sup>®</sup> microarchitecture and is referenced by  $opcode_2 = 0x1$ .

The Exception Vector Relocation bit (bit 13 of the ARM control register) allows the vectors to be mapped into high memory rather than their default location at address 0. This bit is readable and writable by software. When the MMU is enabled, the exception vectors are accessed via the usual translation method involving the PID register and the TLBs. To avoid automatic application of the PID to exception vector accesses, software may relocate the exceptions to high memory.

Table 413.         ARM* Control Register
--

31 30 29 28	27 26 25 24 23 22 21 20 19 18 17	16 15 14	13	12 1	1 10	9	•	• •	0	Э	4		2	1 0	
			۷	1	2 0	R	S	S B	1	1	1	1	С	AM	
reset value:	writeable bits set to 0										<u>.</u>				
Bits	Access	Description													
31:14	Read-Unpredictable / Write-as-Zero	Reserved													
13	Read / Write	Exception Vector Relocation (V). 0 = Base address of exception vectors is 0x0000,0000 1 = Base address of exception vectors is 0xFFFF,0000													
12	Read / Write	Instruction Cache Enable/Disable (I) 0 = Disabled 1 = Enabled													
11	Read / Write	Branch T 0 = Disab 1 = Enab	<b>arg</b> led ed	et Bu	ffer	Ena	able	e (Z)							
10	Read-as-Zero / Write-as-Zero	Reserved													
9	Read / Write	ROM Pro This select managen <i>Manual</i> - for more	tect ts t nent ARM nfor	t <b>ion (</b> he ac unit. /I Lim matic	<b>R)</b> See ted. n.	the Orc	eck <i>AF</i> der	ks pei R <i>M A</i> num	rfori <i>rchi</i> ber	mec itect : AF	d b <i>tur</i> RM	oy the re Re 1 DD	e me efere I 010	emory ence 00E.	
8	Read / Write	System I This select managen Manual - for more	Prot cts t nent ARN nfor	ectio he ac unit. /I Lim matic	n <b>(S</b> cess See ted. n.	che the Orc	eck <i>AF</i> der	ks pei R <i>M A</i> num	rfori <i>rchi</i> ber	mec itect : AF	d b <i>tur</i> RM	oy the re Re 1 DD	e me efere I 010	emory ence 00E.	
		<b>Big/Little</b>	En	dian	B)										
7	Read / Write	0 = Little 1 = Big-e	-enc endia	dian c an op	pera erati	ation on	ı								
6:3	Read-as-One / Write-as-One	= 0b1111													
2	Read / Write	Data cac 0 = Disa 1 = Enat	he e bled bled	enabl	e/dis	abl	e (	C)							
		Alignme	nt fa	ult e	nabl	e/di	sa	ble (	A)						
1	Read / Write	0 = Disa 1 = Enat	bled	l											
		Memory	mar	nager	nent	uni	it e	enabl	e/d	isal	ble	e (M)	)		
0	Read / Write	0 = Disabled 1 = Enabled													



The mini-data cache attribute bits, in the I/O processor Control Register, are used to control the allocation policy for the mini-data cache and whether it uses write-back caching or write-through caching.

The configuration of the mini-data cache is setup before any data access is made that may be cached in the mini-data cache. Once data is cached, software must ensure that the mini-data cache has been cleaned and invalidated before the mini-data cache attributes can be changed.

#### Table 414. Auxiliary Control Register

31 30 29 28	27 26 25 24 23 22 21 20 19 18 17	16 15 14 13 12 11 10 9 8 7 6	5 4	3 2	1	0								
			MD		Ρ	K								
reset value:	writeable bits set to 0													
Bits	Access	Description												
31:6	Read-Unpredictable / Write-as-Zero	Reserved												
		Mini Data Cache Attributes (MD)												
5:4	Read / Write	All configurations of the Mini-data cache are cacheable, stores are buffered in the write buffer and stores are coalesced in the write buffer as long as coalescing is globally enabled (bit 0 of this register).												
		0b00 = Write back, Read allocate 0b01 = Write back, Read/Write allocate 0b10 = Write through, Read allocate 0b11 = Unpredictable												
3:2	Read-Unpredictable / Write-as-Zero	Reserved												
		Page Table Memory Attribute (P)												
1	Read / Write	When set, page table accesses are protected by ECC. S Section 11, "Bus Controller" on page 11-1 of the Intel® 80200 processor based on Intel <sup>®</sup> XScale Microarchitect Developer's Manual for more information.												
		Write Buffer Coalescing Disable (K)	<)											
0	Read / Write	This bit globally disables the coalescing of all stores in the write buffer no matter what the value of the Cacheable and Bufferable bits are in the page table descriptors.												
		0 = Enabled 1 = Disabled												


# 15.3.3 Register 2: Translation Table Base Register

#### Table 415. Translation Table Base Register

31 30 29 2	3 27 26 25 24 23 22 21 20 19 18 17	16     15     14     13     12     11     10     9     8     7     6     5     4     3     2     1     0									
	Translation Table Base										
reset value	unpredictable										
Bits	Access	Description									
31:14	Read / Write	Translation Table Base - Physical address of the base of the first-level table									
13:0	Read-unpredictable / Write-as-Zero	o Reserved									

# 15.3.4 Register 3: Domain Access Control Register

#### Table 416. Domain Access Control Register

31 30	29 28	27 26	25 24	23 22	21 20	19 18	19 18 17 16		13 12	11 10	9 8	76	5 4	3 2	1 0
D15	D14	D13	D12	D11	D10	<b>D9</b>	<b>D8</b>	D7	<b>D6</b>	D5	<b>D4</b>	D3	<b>D2</b>	D1	<b>D0</b>
reset	value:	unpredictable													
Bi	its		Access Description												
31	:0	Read /	/ Write				A of <i>R</i> D	<b>ccess p</b> f each fie <i>eference</i> DI 0100	ermiss eld can e <i>Manu</i> E	be four be four al - AR	or all 1 nd in th M Limit	6 doma le ARM led. Orc	ains - T Archite der num	he mea ecture ber: Al	aning RM



# 15.3.5 Register 5: Fault Status Register

The Fault Status Register (FSR) indicates which fault has occurred, which is either a prefetch abort or a data abort. Bit 10 extends the encoding of the status field for prefetch aborts and data aborts. Bit 9 indicates that a debug event occurred and the exact source of the event is found in the debug control and status register (CP14, register 10). When bit 9 is set, the domain and extended status field are undefined.

Upon entry into the prefetch abort or data abort handler, hardware updates this register with the source of the exception. Software is not required to clear these fields.

#### Table 417.Fault Status Register

31 30 29 2	8 27 26 25 24 23 22 21 20 19 18 17	16 15 14 13 12 11	10	9	8	7 6 5 4	3 2 1 0				
			X	D	0	Domain	Status				
reset value	: unpredictable										
Bits	Access			Des	crip	ition					
31:11	Read-unpredictable / Write-as-Zero	Reserved									
		Status Field Extens	sion	(X)	)						
10	Read / Write	This bit is used to ex when there is a prefe abort.	ten etch	d th ab	ie ei ort a	ncoding of the and when the	Status field, re is a data				
		Debug Event (D)									
9	Read / Write	This flag indicates a cause of the debug of the debug control registered	deb evei er ((	ug o nt is CP1	evei s fou 4, r	nt has occurre Ind in the MO egister 10)	d and that the E field of the				
8	Read-as-zero / Write-as-Zero	= 0									
7:4	Read / Write	<b>Domain</b> - Specifies which of the 16 domains was being accessed when a data abort occurred									
3:0	Read / Write	Status - Type of data access being attempted									

# 15.3.6 Register 6: Fault Address Register

#### Table 418.Fault Address Register

31	30	<b>29</b>	<b>28</b>	27	<b>26</b>	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
													Fau	lt V	irtu	al A	١dd	res	S												
res	et v	valu	le:	unp	rec	licta	able	•																							
	Bi	its						Α	cce	SS												Des	crip	otio	n						
	31	:0		Re	ad /	/Wi	rite								Fa aco	ult cess	Virt s th	ual at c	Ad aus	dre ed	<b>ss</b> - the	- Co mei	onta mor	ins y ał	the port	ΜV	Ao	f the	e da	ita	

# 15.3.7 Register 7: Cache Functions

All the functions defined in the first generation of Intel XScale<sup>®</sup> microarchitecture appear here. The I/O processor adds other functions as well. This register is accessed as write-only. Reads from this register, as with an MRC, have an undefined effect.

The Drain Write Buffer function not only drains the write buffer but also drains the fill buffer.

The I/O processor does not check permissions on addresses supplied for cache or TLB functions. Because only privileged software may execute these functions, full accessibility is assumed. Cache functions do not generate any of the following:

- translation faults
- domain faults
- permission faults

The invalidate instruction cache line command does not invalidate the BTB. When software invalidates a line from the instruction cache and modifies the same location in external memory, it needs to invalidate the BTB also. Not invalidating the BTB in this case may cause unpredictable results.

Disabling/enabling a cache has no effect on contents of the cache: valid data stays valid, locked items remain locked. All operations defined in Table 419 work regardless of whether the cache is enabled or disabled.

Since the Clean D Cache Line function reads from the data cache, it is capable of generating a parity fault. The other operations do not generate parity faults.

Function	opcode_2	CRm	Data	Instruction
Invalidate I&D cache & BTB	0b000	0b0111	Ignored	MCR p15, 0, Rd, c7, c7, 0
Invalidate I cache & BTB	0b000	0b0101	Ignored	MCR p15, 0, Rd, c7, c5, 0
Invalidate I cache line	0b001	0b0101	MVA	MCR p15, 0, Rd, c7, c5, 1
Invalidate D cache	0b000	0b0110	Ignored	MCR p15, 0, Rd, c7, c6, 0
Invalidate D cache line	0b001	0b0110	MVA	MCR p15, 0, Rd, c7, c6, 1
Clean D cache line	0b001	0b1010	MVA	MCR p15, 0, Rd, c7, c10, 1
Drain Write (& Fill) Buffer	0b100	0b1010	Ignored	MCR p15, 0, Rd, c7, c10, 4
Invalidate Branch Target Buffer	0b110	0b0101	Ignored	MCR p15, 0, Rd, c7, c5, 6
Allocate Line in the Data Cache	0b101	0b0010	MVA	MCR p15, 0, Rd, c7, c2, 5

#### Table 419.Cache Functions

The line-allocate command allocates a tag into the data cache specified by bits [31:5] of Rd. When a valid dirty line (with a different MVA) already exists at this location, it is evicted. The 32 bytes of data associated with the newly allocated line are not initialized and therefore generates unpredictable results when read.

This command may be used for cleaning the entire data cache on a context switch and also when re-configuring portions of the data cache as data RAM. In both cases, Rd is a virtual address that maps to some non-existent physical memory. When creating data RAM, software must initialize the data RAM before read accesses can occur.



Other items to note about the line-allocate command are:

- It forces all pending memory operations to complete.
- Bits [31:5] of Rd is used to specific the virtual address of the line to allocated into the data cache.
- When the targeted cache line is already resident, this command has no effect.
- This command cannot be used to allocate a line in the mini Data Cache.
- The newly allocated line is not marked as "dirty" so it never gets evicted. However, when a valid store is made to that line it is marked as "dirty" and gets written back to external memory when another line is allocated to the same cache location. This eviction produces unpredictable results when the line-allocate command used a virtual address that mapped to non-existent memory.

To avoid this situation, the line-allocate operation is only used when one of the following can be guaranteed:

- The virtual address associated with this command is not one that is generated during normal program execution. This is the case when line-allocate is used to clean/invalidate the entire cache.
- The line-allocate operation is used only on a cache region destined to be locked. When the region is unlocked, it must be invalidated before making another data access.



# 15.3.8 Register 8: TLB Operations

Disabling/enabling the MMU has no effect on the contents of either TLB: valid entries stay valid, locked items remain locked. All operations defined in Table 420 work regardless of whether the TLB is enabled or disabled.

This register is accessed as write-only. Reads from this register, as with an MRC, have an undefined effect.

#### Table 420.TLB Functions

Function	opcode_2	CRm	Data	Instruction
Invalidate I&D TLB	0b000	0b0111	Ignored	MCR p15, 0, Rd, c8, c7, 0
Invalidate I TLB	0b000	0b0101	Ignored	MCR p15, 0, Rd, c8, c5, 0
Invalidate I TLB entry	0b001	0b0101	MVA	MCR p15, 0, Rd, c8, c5, 1
Invalidate D TLB	0b000	0b0110	Ignored	MCR p15, 0, Rd, c8, c6, 0
Invalidate D TLB entry	0b001	0b0110	MVA	MCR p15, 0, Rd, c8, c6, 1

#### 15.3.9 Register 9: Cache Lock Down

Register 9 is used for locking down entries into the instruction cache and data cache.

Table 421 shows the command for locking down entries in the instruction cache, instruction TLB, and data TLB. The entry to lock is specified by the virtual address in Rd. The data cache locking mechanism follows a different procedure than the others. The data cache is placed in lock down mode such that all subsequent fills to the data cache result in that line being locked in, as controlled by Table 422.

Lock/unlock operations on a disabled cache have an undefined effect. This register is accessed as write-only. Reads from this register, as with an MRC, have an undefined effect.

#### Table 421.Cache Lockdown Functions

Function	opcode_2	CRm	Data	Instruction
Fetch and Lock I cache line	0b000	0b0001	MVA	MCR p15, 0, Rd, c9, c1, 0
Unlock Instruction cache	0b001	0b0001	Ignored	MCR p15, 0, Rd, c9, c1, 1
Read data cache lock register	0b000	0b0010	Read lock mode value	MRC p15, 0, Rd, c9, c2, 0
Write data cache lock register	0b000	0b0010	Set/Clear lock mode	MCR p15, 0, Rd, c9, c2, 0
Unlock Data Cache	0b001	0b0010	Ignored	MCR p15, 0, Rd, c9, c2, 1

#### Table 422. Data Cache Lock Register

1 30	29	<b>28</b>	27	<b>26</b>	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
																														L
eset v	valı	ie:	writ	eat	ole I	bits	set	t to	0																					
Bi	its						A	cce	SS												Des	crip	otio	n						
31	1:1		Re	ad-	unp	redi	ctal	ble	/ W	rite-	as-	Zer	0	Re	ser	ved														
(	D		Re	ad /	′ Wr	ite								Data Cache Lock Mode (L)           0 = No locking occurs           1 = Any fill into the data cache while this bit is set ge locked in								ets								



# 15.3.10 Register 10: TLB Lock Down

Register 10 is used for locking down entries into the instruction TLB, and data TLB. Lock/unlock operations on a TLB when the MMU is disabled have an undefined effect.

This register is accessed as write-only. Reads from this register, as with an MRC, have an undefined effect.

Table 423 shows the command for locking down entries in the instruction TLB, and data TLB. The entry to lock is specified by the virtual address in Rd.

Table 423.TLB Lockdown Functions

Function	opcode_2	CRm	Data	Instruction
Translate and Lock I TLB entry	0b000	0b0100	MVA	MCR p15, 0, Rd, c10, c4, 0
Translate and Lock D TLB entry	0b000	0b1000	MVA	MCR p15, 0, Rd, c10, c8, 0
Unlock I TLB	0b001	0b0100	Ignored	MCR p15, 0, Rd, c10, c4, 1
Unlock D TLB	0b001	0b1000	Ignored	MCR p15, 0, Rd, c10, c8, 1



# 15.3.11 Register 13: Process ID

The I/O processor supports the remapping of virtual addresses through a Process ID (PID) register. This remapping occurs before the instruction cache, instruction TLB, data cache and data TLB are accessed. The PID register controls when virtual addresses are remapped and to what value.

The PID register is a 7-bit value that is ORed with bits 31:25 of the virtual address when they are zero. This effectively remaps the address to one of 128 "slots" in the 4 Gbytes of address space. When bits 31:25 are not zero, no remapping occurs. This feature is useful for operating system management of processes that may map to the same virtual address space. In those cases, the virtually mapped caches on the I/O processor do not require invalidating on a process switch.

#### Table 424.Accessing Process ID

Function	opcode_2	CRm	Instruction
Read Process ID Register	0b000	0b0000	MRC p15, 0, Rd, c13, c0, 0
Write Process ID Register	0b000	0b0000	MCR p15, 0, Rd, c13, c0, 0

#### Table 425.Process ID Register

31 30 29 28	27 26 2	25 24	23	22 21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Proces	ss ID																								
reset value:	0x0000,0	0000																							
Bits			Ac	cess											[	)es	crip	otio	n						
31:25	Read / V	Nrite							Pro ad	oce dres	<b>ss</b> I SS V	I <b>D</b> - /her	Thi 1 bi	s fie ts 3	eld i 1-28	s us 5 of	sed the	for virt	rem ual	app adc	oing Ires	the s ar	virt re zo	ual ero.	
24:0	Read-as	s-Zero	) / W	rite-a	s-Ze	ero			Re coi	<b>ser</b> mpa	<b>vec</b> atibi	l - S lity	shou	ıld I	pe b	rog	ram	nme	d to	ze	ro fo	or fu	iture	Э	

#### **15.3.11.1** The PID Register Affect On Addresses

All addresses generated and used by User Mode code are eligible for being "PIDed" as described in the previous section. Privileged code, however, must be aware of certain special cases in which address generation does not follow the usual flow.

The PID register is not used to remap the virtual address when accessing the Branch Target Buffer (BTB). Any writes to the PID register invalidate the BTB, which prevents any virtual addresses from being double mapped between two processes.

A breakpoint address must be expressed as an MVA when written to the breakpoint register. This means the value of the PID must be combined appropriately with the address before it is written to the breakpoint register. All virtual addresses in translation descriptors are MVAs.



# 15.3.12 Register 14: Breakpoint Registers

The I/O processor contains two instruction breakpoint address registers (IBCR0 and IBCR1), one data breakpoint address register (DBR0), one configurable data mask/address register (DBR1), and one data breakpoint control register (DBCON). The I/O processor also supports a 256 entry, trace buffer that records program execution information. The registers to control the trace buffer are located in CP14.

#### Table 426. Accessing the Debug Registers

Function	opcode_2	CRm	Instruction
Access Instruction Breakpoint Control Register 0 (IBCR0)	06000	0b1000	MRC p15, 0, Rd, c14, c8, 0; read MCR p15, 0, Rd, c14, c8, 0; write
Access Instruction Breakpoint Control Register 1(IBCR1)	06000	0b1001	MRC p15, 0, Rd, c14, c9, 0; read MCR p15, 0, Rd, c14, c9, 0; write
Access Data Breakpoint Address Register (DBR0)	0b000	0b0000	MRC p15, 0, Rd, c14, c0, 0; read MCR p15, 0, Rd, c14, c0, 0; write
Access Data Mask/Address Register (DBR1)	0b000	0b0011	MRC p15, 0, Rd, c14, c3, 0; read MCR p15, 0, Rd, c14, c3, 0; write
Access Data Breakpoint Control Register (DBCON)	0b000	0b0100	MRC p15, 0, Rd, c14, c4, 0; read MCR p15, 0, Rd, c14, c4, 0; write



# 15.3.13 Register 15: Coprocessor Access Register

Sharing resources among different applications requires a state saving mechanism. Two possibilities are:

- The operating system, during a context switch, saves the state of the coprocessor when the last executing process had access rights to the coprocessor.
- The operating system, during a request for access, saves off the old coprocessor state and saves it with last process to have access to it.

Under both scenarios, the OS needs to restore state when a request for access is made. This means the OS has to maintain a list of what processes are modifying CP0 and their associated state.

#### Example 7. Disallowing Access to CP0

;;	The following code clears bit 0 of	the CPAR.
;;	This causes the processor to fault	if software
;;	attempts to access CP0.	
	LDR R0, =0x3FFE	; bit 0 is clear
	MCR P15, 0, R0, C15, C1, 0	; move to CPAR
	CPWAIT	; wait for effect

#### Table 427. Coprocessor Access Register

31	30	<b>29</b>	<b>28</b>	27	<b>26</b>	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
																0	0	C P 13	C P 12	С Р 11	C P 10	С Р 9	C P 8	C P 7	C P 6	C P 5	C P 4	C P 3	C P 2	<b>C</b> <b>P</b> 1	C P 0

|--|--|

Bits	Access	Description				
31:16	Read-unpredictable / Write-as-Zero	Reserved - Program to zero for future compatibility				
15:14	Read-as-Zero/Write-as-Zero	Reserved - Program to zero for future compatibility				
		<b>Coprocessor Access Rights</b> - Each bit in this field corresponds to the access rights for each coprocessor.				
13:0	Read / Write	<ul> <li>0 = Access denied. Attempts to access corresponding coprocessor generates an undefined exception.</li> <li>1 = Access allowed. Includes read and write accesses.</li> <li>Setting any of bits 12:1 has an undefined effect.</li> </ul>				



# **15.4 Core Performance Monitoring Unit (CPMON)**

Refer to the Intel XScale® Core Developer's Manual (Order Number: 273473), for more details.



#### Intel<sup>®</sup> 80333 I/O Processor Developer's Manual **Timers**

This chapter describes the Intel<sup>®</sup> 80333 I/O processor (80333) dual-programmable 32-bit timers and Watch Dog Timer. Topics include timer registers (TMRx, TCRx and TRRx), timer operation, timer interrupts, and timer register values at initialization.

Each timer is programmed by the timer registers. These registers are mapped into Intel XScale® microarchitecture Coprocessor 6, registers 0 to 5. They may be accessed/manipulated with the MCR, MRC, STC, and LDC instructions. The CRn field of the instruction denotes the register number to be accessed. The opcode\_1 and opcode\_2 fields of the instruction should be zero. The *CRm* field of the instruction should be one. Most systems restricts access to CP6 to privileged processes. To control access to CP6, use the Coprocessor Access Register.

Figure 142 shows a diagram of the timer functions. See also Figure 143 for the Programmable Timer state diagram.







When enabled, a timer decrements the user-defined count value with each Timer Clock (TCLOCK) cycle. The countdown rate is also user-configurable to be equal to the internal bus frequency, or the internal bus clock rate divided by 4, 8 or 16. The timers can be programmed to either stop when the count value reaches zero (single-shot mode) or run continuously (auto-reload mode). When a timer's count reaches zero, the timer's interrupt unit signals the processor's interrupt controller.

#### Table 428.Timer Performance Ranges

Internal Bus Frequency (MHz)	Max Resolution (ns)	Max Range (mins)
333	3.75	4.30

# intel

# 16.1 Timer Operation

This section summarizes the programmable timer and Watch Dog Timer operation and describes load/store access latency for the timer registers.

# 16.1.1 Basic Programmable Timer Operation

Each timer has a programmable enable bit in its control register (TMRx.enable) to start and stop counting. This allows the programmer to prevent user mode tasks from enabling or disabling the timer. Once the timer is enabled, the value stored in the Timer Count Register (TCRx) decrements every Timer Clock (TCLOCK) cycle. TCLOCK is determined by the Timer Input Clock Select (TMRx.csel) bit setting. The countdown rate can be set to equal the internal bus clock frequency, or the internal bus clock rate divided by 4, 8 or 16. Setting TCLOCK to a slower rate lets the user specify a longer count period with the same 32-bit TCRx value.

Software can read or write the TCRx value whether the timer is running or stopped. This lets the user monitor the count without using hardware interrupts.

When the TCRx value decrements to zero, the unit's interrupt request signals the processor interrupt controller. See Section 16.2, "Timer Interrupts" on page 736 for more information. The timer checks the value of the timer reload bit (TMRx.reload) setting. When TMRx.reload. = 1, the processor:

- Automatically reloads TCRx with the value in the Timer Reload Register (TRRx).
- Decrements TCRx until it equals 0 again.

This process repeats until software clears TMRx.reload or TMR.enable.

When TMRx.reload = 0, the timer stops running and sets the terminal count bit (TMRx.tc). This bit remains set until user software reads or writes the TMRx register. Either access type clears the bit. The timer ignores any value specified for TMRx.tc in a write request.

 Table 429.
 Timer Mode Register Control Bit Summary

TRRX	TCRX	Bit 2 (TMRx.reload)	Bit 1 (TMRx.enable)	Action
Х	Х	Х	0	Timer disabled.
NOTE	- V	م ماد مام		

**NOTE:** X = don't care

N = a number between 1H and FFFF FFFH



# 16.1.2 Watch Dog Timer Operation

The Watch Dog Timer (WDT) is a 32-bit down counter that can be used to reset the Internal Bus and the Intel XScale<sup>®</sup> core when software gets stuck in an infinite loop or generate an interrupt to the Intel XScale<sup>®</sup> core. A reset of the Internal Bus also results in the **M\_RST#** output to be asserted which can be used to reset the system or as an external indicator of the WDT expiration when Reset generation is enabled. The WDT Mask bit (bit 17) of the INTCTL0 selects either the reset or interrupt behavior for the WDT expiration.

Following **PWRGD** assertion, the WDT is disabled.

The software can enable the WDT by using coprocessor instructions (i.e, MCR or LDC) to write the value 1E1E 1E1EH followed by the value E1E1 E1E1H to the WDT Control register. When enabled, the WDT is initialized with FFFF FFFFH and begin to decrement towards 0000 0000H.

The software is required periodically to write the WDT initialization sequence (the value 1E1E 1E1EH followed by the value E1E1 E1E1H) to the WDT Control register in order to reset the timer value to FFFF FFFFH. For a 333 MHz Internal Bus frequency, this means that the sequence must be written approximately every fifteen seconds.

Note: The WDT always runs at Internal Bus frequency speed of 333MHz.

When the software fails to reinitialize the WDT prior to the timer value transitioning to zero, an Internal Bus Reset or Intel XScale<sup>®</sup> core Interrupt is generated. The Internal Bus Reset reinitializes all Internal Bus peripherals and the Intel XScale<sup>®</sup> core.

The WDT pending interrupt is cleared by software reinitializing the WDT as described above.

*Warning:* Once enabled, the WDT can **not** be disabled without resetting the Intel XScale<sup>®</sup> core.

# intel

# 16.1.3 Load/Store Access Latency for Timer Registers

As with all other load accesses from internal memory-mapped registers, a load instruction that accesses a timer register has a latency of one internal processor cycle. With one exception, a store access to a timer register completes and all state changes take effect before the next instruction begins execution. The exception to this is when disabling a timer. Latency associated with the disabling action is such that a timer interrupt may be posted immediately after the disabling instruction completes. This can occur when the timer is near zero as the store to TMRx occurs. In this case, the timer interrupt is posted immediately after the store to TMRx completes and before the next instruction can execute. Table 430 summarizes the timer access and response timings. Refer also to the individual register descriptions for details.

Note that the processor may delay the actual issuing of the load or store operation due to previous instruction activity and resource availability of processor functional units.

The processor ensures that the TMRx.tc bit is cleared within one internal bus clock after a load or store instruction accesses TMRx.

Name	Status	Action
(TMRx.tc) Terminal Count	READ	Timer clears this bit when user software accesses TMRx. This bit can be set 1 internal bus clock later. The timer sets this bit within 1 internal bus clock of TCRx reaching zero when TMRx.reload=0.
Bit 0	WRITE	Timer clears this bit within 1 internal bus clock after the software accesses TMRx. The timer ignores any value specified for TMRx.tc in a write request.
(TMRx.enable)	READ	Bit is available 1 internal bus clock after executing a read instruction from TMRx.
Bit 1	WRITE	Writing a '1' enables the internal bus clock to decrement TCRx within 1 internal bus clock after executing a store instruction to TMRx.
(TMRx.reload)	READ	Bit is available 1 internal bus clock after executing a read instruction from TMRx.
Enable Bit 2	WRITE	Writing a '1' enables the reload capability within 1 internal bus clock after the store instruction to TMRx has executed. The timer loads TRRx data into TCRx and decrements this value during the next internal bus clock cycle.
(TMRx.csel1:0)	READ	Bits are available 1 internal bus clock after executing a read instruction from TMRx.csel1:0 bit(s).
Select Bits 4-5	WRITE	The timer re-synchronizes the clock cycle used to decrement TCRx within one internal bus clock cycle after executing a store instruction to TMRx.csel1:0 bit(s).
(TCRx.d31:0) Timer Count	READ	The current TCRx count value is available within 1 internal bus clock cycle after executing a read instruction from TCRx. When the timer is running, the pre-decremented value is returned as the current value.
Register	WRITE	The value written to TCRx becomes the active value within 1 internal bus clock cycle. When the timer is running, the value written is decremented in the current clock cycle.
(TRRx.d31:0) Timer Reload	READ	The current TRRx count value is available within 1 internal bus clock after executing a read instruction from TRRx. When the timer is transferring the TRRx count into TCRx in the current count cycle, the timer returns the new TCRx count value to the executing read instruction.
Register	WRITE	The value written to TRRx becomes the active value stored in TRRx within 1 internal bus clock cycle. When the timer is transferring the TRRx value into the TCRx, data written to TRRx is also transferred into TCRx.

#### Table 430. Timer Responses to Register Bit Settings

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Timers



# **16.2** Timer Interrupts

Each timer is the source for one interrupt. When a timer detects a zero count in its TCRx, the timer generates an internal level-detected Timer Interrupt signal (TINTx) to the interrupt controller, and the interrupt source (INTSRC[1:0]) bit is set in the interrupt controller. Each timer interrupt can be selectively masked in the Interrupt Control (INTCTL[1:0]) registers. Refer to Section 17.5, "Intel® 80333 I/O Processor Interrupt Controller Unit" for a description of interrupt controller operation.

After servicing the timer interrupt, the interrupt service routine clears the pending request by writing a '1' to the appropriate bit of the Timer Interrupt Status Register (TISR).

When a timer generates a second interrupt request before the CPU services the first interrupt request, the second request may be lost.

When auto-reload is enabled for a timer, the timer continues to decrement the value in TCRx even after entry into the timer interrupt handler.

# intel

# 16.3 Timer State Diagram

Figure 143 shows the common states of the Timer Unit. For uncommon conditions see Section 16.5, "Uncommon TCRX and TRRX Conditions" on page 746.

Figure 143. Timer Unit State Diagram





# **16.4** Timer Registers

As shown in Table 431, each timer has three memory-mapped registers:

- Timer Mode Register programs the specific mode of operation or indicates the current programmed status of the timer. This register is described in Section 16.4.2, "Timer Mode Registers TMR0:1" on page 739.
- Timer Count Register contains the timer's current count. See Section 16.4.3, "Timer Count Register TCR0:1" on page 742.
- Timer Reload Register contains the timer's reload count. See Section 16.4.4, "Timer Reload Register TRR0:1" on page 743.

#### Table 431. Timer Registers

Timer Unit	Register Acronym	Register Name
	TMR0	Timer Mode Register 0
Timer 0	TCR0	Timer Count Register 0
	TRR0	Timer Reload Register 0
	TMR1	Timer Mode Register 1
Timer 1	TCR1	Timer Count Register 1
	TRR1	Timer Reload Register 1

#### 16.4.1 Power Up/Reset Initialization

Upon **RSTIN**# assertion, the timer registers are initialized to the values shown in Table 432.

#### Table 432. Timer Powerup Mode Settings

Mode/Control Bit	Notes
TMRx.tc = 0	No terminal count
TMRx.enable = 0	Prevents counting and assertion of TINTx
TMRx.reload = 0	Single terminal count mode
TMRx.pri = 0	Privileged Mode and User Mode Writes Allowed
TMRx.csel1:0 = 0	Timer Clock = internal bus clock
TCRx.d31:0 = 0	Undefined
TRRx.d31:0 = 0	Undefined
TINTx output	Deasserted



### **16.4.2** Timer Mode Registers – TMR0:1

The Timer Mode Register (TMRx) lets the user program the mode of operation and determine the current status of the timer. TMRx bits are described in the subsections following Table 433 and are summarized in Table 429.

Note: TMR0:1 registers are read-only from Memory-Mapped Register Address space.



#### 16.4.2.1 Bit 0 - Terminal Count Status Bit (TMRx.tc)

The TMRx.tc bit is set when the Timer Count Register (TCRx) decrements to 0 and bit 2 (TMRx.reload) is not set for a timer. The TMRx.tc bit allows applications to monitor timer status through software instead of interrupts. TMRx.tc remains set until software accesses (reads or writes) TMRx. The access clears TMRx.tc. The timer ignores any value specified for TMRx.tc in a write request.

When auto-reload is selected for a timer and the timer is enabled, the TMRx.tc bit status is unpredictable. Software should not rely on the value of the TMRx.tc bit when auto-reload is enabled.

The processor also clears the TMRx.tc bit upon hardware or software reset. Refer to Section 20.2, "Reset Overview" on page 870.



#### 16.4.2.2 Bit 1 - Timer Enable (TMRx.enable)

The TMRx.enable bit allows user software to control the timer's RUN/STOP status. When:

TMRx.enable = 1The Timer Count Register (TCRx) value decrements every Timer<br/>Clock (TCLOCK) cycle. TCLOCK is determined by the Timer Input<br/>Clock Select (TMRx.csel bits 0-1). See Section 16.4.2.5. When<br/>TMRx.reload=0, the timer automatically clears TMRx.enable when the<br/>count reaches zero. When TMRx.reload=1, the bit remains set. See<br/>Section 16.4.2.3.

TMRx.enable = 0 The timer is disabled and ignores all input transitions.

User software sets this bit. Once started, the timer continues to run, regardless of other processor activity. Three events can stop the timer:

- User software explicitly clearing this bit (i.e., TMRx.enable = 0).
- TCRx value decrements to 0, and the Timer Auto Reload Enable (TMRx.reload) bit = 0.
- Hardware or software reset. Refer to Section 20.2, "Reset Overview" on page 870.

#### 16.4.2.3 Bit 2 - Timer Auto Reload Enable (TMRx.reload)

The TMRx.reload bit determines whether the timer runs continuously or in single-shot mode. When TCRx = 0 and TMRx.enable = 1 and:

TMRx.reload = 1 The timer runs continuously. The processor:

- 1. Automatically loads TCRx with the value in the Timer Reload Register (TRRx), when TCRx value decrements to 0.
- 2. Decrements TCRx until it equals 0 again.

Steps 1 and 2 repeat until software clears TMRx bits 1 or 2.

TMRx.reload = 0 The timer runs until the Timer Count Register = 0. TRRx has no effect on the timer.

User software sets this bit. When TMRx.enable and TMRx.reload are set and TRRx does not equal 0, the timer continues to run in auto-reload mode, regardless of other processor activity. Two events can stop the timer:

- User software explicitly clearing either TMRx.enable or TMRx.reload.
- Hardware or software reset.

The processor clears this bit upon hardware or software reset.



#### 16.4.2.4 Bit 3 - Timer Register Privileged Read/Write Control (TMRx.pri)

The TMRx, bit[3], (where x = 0 or 1) enables or disables user-mode writes to the timer registers (TMRx, TCRx, TRRx). However, when TMR1[3] is set, the Watchdog timer register (WDTCR) is also disabled from user-mode writes..

*Note:* When TMR1.3 is set, the Watchdog timer register (WDTCR) is also disabled from user-mode writes.

When:

TMRx.pri = 1	The timer ignores the user mode write to the timer registers; however, writes from the privileged modes are allowed.
TMRx.pri = 0	The timer registers can be written from either the user mode or the privileged modes.

The processor clears TMRx.pri upon hardware or software reset.

#### 16.4.2.5 Bits 4, 5 - Timer Input Clock Select (TMRx.csel1:0)

User software programs the TMRx.csel bits to select the Timer Clock (TCLOCK) frequency. See Table 434. As shown in Figure 142, the internal bus clock is an input to the timer clock unit. These bits allow the application to specify whether TCLOCK runs at or slower than the internal bus clock frequency.

#### Table 434. Timer Input Clock (TCLOCK) Frequency Selection

Bit 5 TMRx.csel1	Bit 4 TMRx.csel0	Timer Clock (TCLOCK)
0	0	Timer Clock = internal bus clock
0	1	Timer Clock = internal bus clock / 4
1	0	Timer Clock = internal bus clock / 8
1	1	Timer Clock = internal bus clock / 16

The processor clears these bits upon hardware or software reset (TCLOCK = Core Clock).



# 16.4.3 Timer Count Register – TCR0:1

The Timer Count Register (TCRx) is a 32-bit register that contains the timer's current count. The register value decrements with each timer clock tick. When this register value decrements to zero (terminal count), a timer interrupt is generated. When TMRx.reload is not set for the timer, the status bit in the timer mode register (TMRx.tc) is set and remains set until the TMRx register is accessed. Table 435 shows the timer count register.

*Note:* TCR0:1 registers are read-only from Memory-Mapped Register Address space.





The valid programmable range is from 1H to FFFF FFFFH. Avoid programming TCRx to 0 as it has varying results as described in Section 16.5, "Uncommon TCRX and TRRX Conditions" on page 746.

User software can read or write TCRx whether the timer is running or stopped. Bit 3 of TMRx determines user read/write control (Section 16.4.2.5). The TCRx value is undefined after hardware or software reset.

# 16.4.4 Timer Reload Register – TRR0:1

The Timer Reload Register (TRRx; Table 436) is a 32-bit register that contains the timer's reload count. The timer loads the reload count value into TCRx when TMRx.reload is set (1), TMRx.enable is set (1) and TCRx equals zero.

As with TCRx, the valid programmable range is from 1H to FFFF FFFFH. Avoid programming a value of 0, as it may prevent TINTx from asserting continuously. (See Section 16.5, "Uncommon TCRX and TRRX Conditions" on page 746 for more information.)

User software can access TRRx whether the timer is running or stopped. Bit 3 of TMRx determines read/write control (Section 16.4.2.5, "Bits 4, 5 - Timer Input Clock Select (TMRx.csel1:0)" on page 741). TRRx value is undefined after hardware or software reset.

Note: TRR0:1 registers are read-only from Memory-Mapped Register Address space.

Table 436.Timer Reload Register – TRRx





# **16.4.5** Timer Interrupt Status Register – TISR

The Timer Interrupt Status Register (TISR; Table 437) is a two-bit register that contains the timer's pending interrupt status. The setting of these status bits represents the assertion of a "level-sensitive" interrupt request to the "Intel® 80333 I/O Processor Interrupt Controller Unit". After the interrupt service routine completes processing of the interrupt request, it needs to write a '1' to the appropriate bit in the TISR to clear the pending request.

TISR interrupt requests are cleared after hardware or software reset.

Note: TISR register is read-only from Memory-Mapped Register Address space.





# intel

# **16.4.6 Watch Dog Timer Control Register – WDTCR**

The Watch Dog Timer Control Register (WDTCR) is a 32-bit register that software can use to enable the WDT or read the current WDT count value. The register value decrements with each internal bus clock tick. When this register value decrements to zero (terminal count), an Internal Bus Reset or internal Interrupt is generated. The timer can be enabled and/or reinitialized by writing 1E1E 1E1EH immediately followed by E1E1 E1E1H to the WDTCR. The WDT pending interrupt is also cleared by writing this same sequence.

*Note:* WDTCR register is read-only from Memory-Mapped Register Address space.



#### Table 438. Watch Dog Timer Control Register -- WDTCR



# **16.5 Uncommon TCRx and TRRx Conditions**

Table 429 summarizes the most common settings for programming the timer registers. Under certain conditions, however, it may be useful to set the Timer Count Register or the Timer Reload Register to zero before enabling the timer. Table 439 details the conditions and results when these conditions are set.

#### Table 439. Uncommon TMRx Control Bit Settings

TRRx	TCRx	Bit 2 (TMRx.reload)	Bit 1 (TMRx.enable)	Action
Х	0	0	1	TMRx.tc and TINTx set, TMR.enable cleared
0	0	1	1	Timer and auto reload enabled, TINTx not generated and timer enable remains set.
0	Ν	1	1	Timer and auto reload enabled. TINTx set when TCRx=0. The timer remains enabled but further TINTx's are not generated.

**NOTE:** X = don't care

N = a number between 1H and FFFF FFFFH

# intel®

# Interrupt Controller Unit and IOAPIC 17

This chapter describes the Intel<sup>®</sup> 80333 I/O processor (80333) Interrupt Controller Unit and IOAPIC. The operation modes, setup, external memory interface, and implementation of the interrupts are described in this chapter.

# 17.1 Overview

The interrupt control unit manages the interrupt routing and interrupt sources to the Intel XScale<sup>®</sup> core. The IOAPICs manage the interrupts to be routed to the host processor through the upstream PCI Express interface.

An interrupt is an event that causes a temporary break in program execution so the processor can handle another task. Interrupts commonly request I/O services or synchronize the processor with some external hardware activity. For interrupt handler portability across the Intel XScale<sup>®</sup> microarchitecture family (ARM\* architecture compliant), the architecture defines a consistent exception handling mechanism. To manage exceptions which include interrupt requests in parallel with processor execution, the 80333 provides an on-chip programmable interrupt controller.

Requests for interrupt service come from many sources and are prioritized such that instruction execution is redirected only when an exception interrupt request is of higher priority than that of the executing task. On the 80333, interrupt requests may originate from external hardware sources, internal peripherals or software. The 80333 contains a number of integrated peripherals which may generate interrupts, including:

- DMA Channel 0
- DMA Channel 1
- ATU
- UART 0 and 1
- Timers 0 and 1
- Performance Monitoring Unit
- Watch Dog Timer

- I<sup>2</sup>C Bus Interface Units 0 and 1
- Application Accelerator Unit
- Messaging Unit
- Memory Controller Unit
- Peripheral Bus Interface Unit
- Core Bus Interface Unit

The interrupt controller can also intercept external processor interrupts and forward them to the integrated IOAPICs.

.

Interrupts are detected with the chip 8-bit interrupt port, an 12-bit GPIO port, and with a dedicated High-Priority Interrupt (**HPI#**) input in the Intel XScale<sup>®</sup> core interrupt controller. Interrupt requests originate from software by the **SWI** instruction.

Ultimately, all interrupt sources that are steered to the Intel XScale<sup>®</sup> core processor are combined into one of two internal interrupt exceptions: IRQ and FIQ.



# 17.1.1 **IOAPIC**

The 80333, intended for future PCI Express chipsets only supports the FSB interrupt delivery mode from its integrated IOAPICs. In the front-side bus interrupt delivery mechanism, interrupts are delivered as interrupt message transactions on the processor system bus. There are several ways a system interrupt can reach the front side bus, as shown in Figure 144.





These IOAPICs include the following features:

- Two IOAPIC controllers one per PCI bus segment.
- 24 interrupts per controller (internally).
- Up to 8 physical PCI interrupt pins per PCI bus.
- PCI virtual wire interrupt support via writing to Pin Assertion Register in the IOAPICs.
- Support for both IA-32 and Intel® Itanium front side bus messages.
- No support for serial bus delivery.
- Inband PCI Express boot interrupt message to be routed to ICH.
- No support for EOI special cycle propagation to PCI (No Software transparent IOAPIC support behind 80333).

# 17.2 Theory of Operation

# 17.2.1 Interrupt Controller Unit

The 80333 Interrupt Controller Unit (ICU) provides the ability to generate interrupts to both the Intel XScale<sup>®</sup> core and the integrated IOAPICs.

In addition to the internal peripherals, external devices may also generate interrupts to the Intel XScale<sup>®</sup> core. External devices can generate interrupts via the **XINT[7:0]**# pins, the **GPIO[7:0]/XINT[15:8]**# pins, and the **HPI**# pin. The Peripheral Interrupt Controller provides the ability to direct PCI interrupts. The routing logic enables, under software control, the ability to intercept external PCI interrupts and forward through the IOAPICs or to the Intel XScale<sup>®</sup> core.

The Interrupt Controller has two functions:

- Internal Peripheral Interrupt Control
- External Interrupt Routing

The internal peripheral interrupt control mechanism consolidates a number of interrupt sources for a given peripheral into a single interrupt driven to the Intel XScale<sup>®</sup> core. High performance data movement associated interrupts are fully demultiplexed into the ICU, however. In order to provide the executing software with the knowledge of interrupt source, coprocessor mapped status registers describe the source of the active interrupts and the vectors to interrupt handlers for the highest priority active sources. All of the peripheral interrupts are individually enabled from the respective peripheral control registers.

The Peripheral interrupt routing mechanism allows the host software (or Intel XScale<sup>®</sup> microarchitecture software) to route external interrupts to either the Intel XScale<sup>®</sup> core or either of the internal IOAPICs. This routing mechanism is controlled through a memory-mapped register accessible from the 80333.



# 17.2.2 **IOAPIC**

The IOAPIC supports three upstream interrupt mechanisms for forwarding interrupts to the system processors: IRQ# mechanism, MSI mechanism, and Virtual Wire mechanism.

#### 17.2.2.1 PCI IRQ# Mechanism

In this mode, a PCI device uses a PCI interrupt pin INTx# to signal an interrupt to a system IOAPIC controller like in a 80333. The IOAPIC controller would then convert the interrupt request into an upstream interrupt message (when the interrupt is unmasked) in the form of a memory write transaction with an address of 0FEEx\_xxxx and a 32-bit data giving information about the interrupt. X\_xxxx in the address carries the encoding of the destination processor. MCH converts this upstream memory write from PCI Express into a FSB interrupt message transaction which contains information about the processor for which the interrupt is meant and information about the interrupt itself – vector number, trigger mode etc. The address and data in the upstream interrupt message generated by the IOAPIC come out of the interrupt redirection table that software programs in the IOAPIC.

PCI IRQ# mechanism is level based. The IOAPIC receiving the interrupts would convert the level-signaled interrupts into the edge triggered semantics of upstream memory writes.

#### 17.2.2.2 PCI MSI Mechanism

PCI devices compliant with *PCI-to-PCI Bridge Specification*, Revision 2.2 and beyond optionally support a direct interrupt delivery mechanism to the front side bus through the PCI message signaled interrupt capability. All PCI-X devices are required to support this mechanism. In this mode, PCI or PCI-X devices do a direct memory write to the front side bus to cause an interrupt. These memory writes would flow through 80333 as any other memory write moving to PCI Express, bypassing the internal IOAPIC. The address of the message would be 0xFEEx\_xxxx and the 32-bit data would contain the interrupt vector information. Note that bit A2 may not be actually asserted on a 64-bit PCI bus. On a 64-bit bus, A2 may be '0', but the BEs are asserted to indicate the upper DW is valid. This must also be decoded as a redirected MSI. The address and data generated for the MSI memory write transactions are derived from the MSI message data and address fields defined as part of the MSI capability register model by the PCI specifications.

PCI MSI mechanism is inherently edge triggered.

#### 17.2.2.3 PCI Virtual Wire Mechanism

In this method, PCI devices are given a write path directly to the pin assertion register (memory offset 20h in the IOAPIC BAR range), PAR, in the IOAPIC, thus forming a virtual wire to the IOAPIC, that causes an interrupt. PCI devices generate a memory transaction on the PCI bus with an address equal to the IOAPIC\_MEM\_BAR + 40 and a 32-bit data equal to the interrupt vector number corresponding to the device. This information is stored in the device's MSI address and data registers that would be initialized by the system software which is PCI MSI aware. The PCI memory write transaction would be claimed by 80333 as any other memory write transaction and forwarded to the interrupt message in the form of an inbound memory write. The address for the upstream write would be 0xFEEx\_xxxh, where the x\_xxxx encodes the destination processor and the 32-bit data would carry information about the interrupt itself. Interrupts associated with the PCI virtual wire mechanism must be in edge-triggered mode, and the level setting (active high or active low) is irrelevant. Note that the standard hot-plug controller in 80333 cannot access the PAR (Pin Assertion Register) register in the IOAPIC. Instead the SHPC talks to the APIC via the interrupt input 23.

PCI virtual write interrupts are inherently edge triggered. PCI virtual wire mechanism is only implemented between the PCI bus and the corresponding IOAPIC i.e. PCI Segment A cannot write to the PAR register in IOAPIC B and vice-versa.



# 17.3 The Intel XScale<sup>®</sup> Core Exceptions Architecture

The Intel XScale<sup>®</sup> core supports five types of exceptions<sup>1</sup>, and a privileged processing mode for each type.

- IRQ and FIQ internal interrupt exceptions otherwise known as the normal and fast interrupts, respectively
- memory aborts (used to implement memory protection or virtual memory)
- attempted execution of an undefined instruction
- software interrupts (SWIs) (used to make a call to an Operating System)

When an exception occurs, some of the standard registers are replaced with registers specific to the exception mode. All exceptions have replacement (or banked) registers for R14 and R13, and one interrupt mode has more registers for fast interrupt processing.

After an exception, R14 holds the return address for exception processing, which is used both to return after the exception is processed and to address the instruction that caused the exception.

R13 is banked across exception modes to provide each exception handler with a private stack pointer (SP). The fast interrupt mode also banks R8 to R12, so that interrupt processing can begin without the need to save or restore these registers. There is a seventh processing mode, System Mode, that does not have any banked registers (it uses the User mode registers), which is used to run normal (non-exception) tasks that require a privileged processor mode.

#### 17.3.1 CPSR and SPSR

All other processor state is held in status registers. The current operating processor status is in the Current Program Status Register or CPSR. The CPSR holds:

- Four condition code flags (Negative, Zero, Carry and Overflow)
- Two interrupt disable bits (one for each type of interrupt)
- Five bits which encode the current processor mode

All five exception modes also have a Saved Program Status Register (SPSR) which holds the CPSR of the task immediately before the exception occurred. Both the CPSR and SPSR are accessed with special instructions.

#### 17.3.2 The Exception Process

When an exception occurs, the Intel XScale<sup>®</sup> core halts execution after the current instruction and begins execution at a fixed address in low memory, known as the exception vectors. There is a separate vector location for each exception (and two for memory aborts to distinguish between data and instruction accesses).

An operating system installs a handler on every exception at initialization. Privileged operating system tasks normally run in System mode to allow exceptions to occur within the operating system without state loss (exceptions overwrite their R14 when an exception occurs, and System mode is the only privileged mode that cannot be entered by an exception).

<sup>1.</sup> Exception Description from the ARM Architecture Reference Manual, p. 1-3, Copyright Advanced RISC Machines Ltd. (ARM) 1996



# 17.3.3 Exception Priorities and Vectors

It is important to note that fast interrupt (FIQ) is higher priority than the normal interrupt (IRQ). In addition, while an FIQ exception is executing, the IRQ exception is masked out.

When an exception is taken by the processor, the Program Counter (PC) is loaded with the vector associated with that exception as specified by Table 440.

Generally, the instruction at this location is required to be a branch instruction to the associated exception handler. However, in the case of an FIQ, this is not necessary since the vector location is at the very bottom of all the defined exception vectors, thus the entire FIQ exception handler can be placed at that vector location.

#### Table 440. Exception Priorities And Vectors

Exception	Priority	Vector <sup>a</sup>
Reset	1 (Highest)	0000 0000H
Data Abort	2	0000 0010H
FIQ	3	0000 001CH
IRQ	4	0000 0018H
Prefetch Abort	5	0000 000CH
Undefined Instructions	6 (Lowest)	0000 0004H
Software Interrupt (SWI) <sup>b</sup>	6 (Lowest)	0000 0008H

a. By enabling the Exception Vector Relocation mode (bit 13, CP15, Register 1), the Vectors (except Reset Vector) can be relocated to be based at FFFF 0000H rather than 0000 0000H. (i.e., FIQ Vector located at FFFF 001CH)

b. Undefined Instruction and SWI can not occur at the same time since SWI is a particular instruction decoding.

# 17.3.4 Software Requirements For Exception Handling

To use the processor's exception handling facilities, user software must provide the following items in memory:

- Exception Handler Routines
- Software handler to nest certain exceptions (i.e., FIQ and IRQ)

These items are established in memory as part of the initialization procedure.

#### 17.3.4.1 Nesting FIQ and IRQ Exceptions

Hardware does not provide support for nesting of any particular exception, including the FIQ and IRQ exceptions.

In order to provide support for nested interrupts, a software handler must be provided to save the Link Register (R14) and the SPSR (Saved Program Status Register) before reenabling the FIQ or IRQ exception.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Interrupt Controller Unit and IOAPIC

# 17.4 Intel<sup>®</sup> 80333 I/O Processor External Interrupt Interface

The interrupt controller attached to the Intel XScale<sup>®</sup> core, has the facilities necessary to handle all core processor and peripheral internal interrupts, as well as eight external interrupts (**XINT[7:0]**#/), eight GPIO inputs (**GPIO[7:0]**), and one High Priority Interrupt (**HPI**#).

80333 PCI Express interface defines two interrupt mechanisms which operate over the interface without dedicated interrupt signals. Interrupts from the downstream PCI bus segments, the Messaging Unit and Address Translation Unit are routed to the PCI Express interface through integrated IOAPICs.

# 17.4.1 Interrupt Inputs

The 17 external interrupt input pins of the 80333 have the following definitions:

- XINT[7:0]#/ External Interrupt (Input) These pins cause interrupts to be requested. Each pin is a level-detect input only. These pins are internally synchronized. These pins only act as interrupt inputs when they are unmasked in the INTCTL[1:0] registers.
- **GPIO**[7:0] **External Interrupt (Input)** These pins cause interrupts (**XINT**[15:8]#) to be requested. Each pin is a level-detect input only. These pins are internally synchronized. These pins only act as interrupt inputs when they are configured as general purpose inputs and are unmasked in the INTCTL[1:0] registers.
- HPI# <u>High-Priority Interrupt (Input)</u> Causes a high priority interrupt event to occur. The external HPI# input requires a level input and is maskable by the INTCTL[1:0] registers. This pin is internally synchronized.

The external interrupt input interface for the 80333 consists of the pins shown in Table 441.

#### Table 441.Interrupt Input Pin Descriptions (Sheet 1 of 2)

Signal	Description
XINT0#	Directed to the input 0 of the 80333 IOAPIC or the 80333 Interrupt Controller input <b>XINT0#</b> . When directed to the 80333 Interrupt Controller input <b>XINT0#</b> , the <b>XINT0#</b> input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.
XINT1#	Directed to the input 1 of the 80333 IOAPIC or the 80333 Interrupt Controller input <b>XINT1#</b> . When directed to the 80333 Interrupt Controller input <b>XINT1#</b> , the <b>XINT1#</b> input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.
XINT2#	Directed to the input 2 of the 80333 IOAPIC or the 80333 Interrupt Controller input <b>XINT2#</b> . When directed to the 80333 Interrupt Controller input <b>XINT2#</b> , the <b>XINT2#</b> input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.
XINT3#	Directed to the input 3 of the 80333 IOAPIC or the 80333 Interrupt Controller input <b>XINT3#</b> . When directed to the 80333 Interrupt Controller input <b>XINT3#</b> , the <b>XINT3#</b> input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.
XINT4#	Directed to the input 4 of the 80333 IOAPIC or the 80333 Interrupt Controller input XINT4#. When directed to the 80333 Interrupt Controller input XINT4#, the XINT4# input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.



#### Table 441. Interrupt Input Pin Descriptions (Sheet 2 of 2)

Signal	Description
XINT5#	Directed to the input 5 of the 80333 IOAPIC or the 80333 Interrupt Controller input <b>XINT5#</b> . When directed to the 80333 Interrupt Controller input <b>XINT5#</b> , the <b>XINT5#</b> input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.
XINT6#	Directed to the input 6 of the 80333 IOAPIC or the 80333 Interrupt Controller input <b>XINT6#</b> . When directed to the 80333 Interrupt Controller input <b>XINT6#</b> , the <b>XINT6#</b> input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.
XINT7#	Directed to the input 7 of the 80333 IOAPIC or the 80333 Interrupt Controller input <b>XINT7#</b> . When directed to the 80333 Interrupt Controller input <b>XINT7#</b> , the <b>XINT7#</b> input can be steered to either the FIQ or the IRQ internal interrupt input of the Intel XScale <sup>®</sup> core.
XINT[15:8]#	Interrupt Inputs dedicated to the Intel XScale <sup>®</sup> core. Functionary is muxed with <b>GPI0[7:0]</b> functionality. <b>XINT8#</b> corresponds to <b>GPI0[0]</b> , <b>XINT9#</b> corresponds to <b>GPI0[1]</b> , and so on. Interrupt functionality is available on these pins only when configured as General Purpose Inputs. To enable a given pin as an interrupt, it needs to be unmasked in the INTCTL1 register.
HPI#	<b>HPI#</b> is a dedicated interrupt input to the Intel XScale <sup>®</sup> core. This interrupt can be disabled by the INTCTL0 register, and can be steered to either the FIQ or the IRQ internal interrupt (not the IOAPICs).



# 17.4.2 Outbound Interrupts

The Messaging Unit (MU) has the capability of generating interrupts through the IOAPIC. The MU has four distinct messaging mechanisms. Each allows a host processor or external PCI agent and the 80333 to communicate through message passing and interrupt generation. The four mechanisms are:

- Message Registers allow the 80333 and external PCI agents to communicate by passing messages in one of four 32-bit Message Registers. In this context, a message is any 32-bit data value. Message registers combine aspects of mailbox registers and doorbell registers. Writes to the message registers may optionally cause interrupts.
- **Doorbell Registers** allow the 80333 to assert the PCI interrupt signals and allow external PCI agents to generate an interrupt to the Intel XScale<sup>®</sup> core.
- Circular Queues support a message passing scheme that uses four circular queues.
- Index Registers support a message passing scheme that uses a portion of the 80333 local memory to implement a large set of message registers.

All four mechanisms can result in Outbound Interrupts to a host processor.

The Messaging Unit (MU) is a part of the Address Translation Unit (ATU) and uses the ATU interrupt pin.

80333 does not have dedicated pins for interrupts. The PCI Express upstream interface of 80333 supports interrupt mechanism as defined in the *PCI Express Specification*, Revision 1.0a. The interrupt generated by the 80333 is from the ATU/MU. This interrupt is routed to the input **IRQ14**# of the A-segment IOAPIC.


## 17.4.3 Interrupt Routing

The 80333 has two levels of interrupt routing. First, eight external interrupts can be assigned to either the A-PCI or the B-PCI bus segment IOAPIC. In addition, each interrupt can be outed to either Intel XScale<sup>®</sup> core interrupt inputs or to the IOAPIC of the respective PCI bus segment. The A-PCI bus segment IOAPIC is Function 1 of Device 0 and the B-PCI bus segment IOAPIC is Function 3 of Device 0, and both are described in Chapter 2, "PCI Express-to-PCI Bridges".

Routing of interrupt inputs is controlled by the PCI Interrupt Routing Select Register (PIRSR). Table 442 summarizes the usage of the bits in the PIRSR.

 Table 442.
 Interrupt Routing Summary

PIRSR Select Bit	Bit Value	Description	
bit 0	1	XINT0# Input Pin routed to ICU XINT0# Input Pin	
	0	XINT0# Input Pin routed to IOAPIC input 0	
bit 1	1	XINT1# Input Pin routed to ICU XINT1# Input Pin	
Dit 1	0	XINT1# Input Pin routed to IOAPIC input 1	
bit 2	1	XINT2# Input Pin routed to ICU XINT2# Input Pin	
Dit 2	0	XINT2# Input Pin routed to IOAPIC input 2	
hit 3	1	XINT3# Input Pin routed to ICU XINT3# Input Pin	
bit 5	0	XINT3# Input Pin routed to IOAPIC input 3	
hit 4	1	XINT4# Input Pin routed to ICU XINT4# Input Pin	
Dit 4	0	XINT4# Input Pin routed to IOAPIC input 4	
bit 5	1	XINT5# Input Pin routed to ICU XINT5# Input Pin	
Dit 5	0	XINT5# Input Pin routed to IOAPIC input 5	
bit 6	1	XINT6# Input Pin routed to ICU XINT6# Input Pin	
Dit O	0	XINT6# Input Pin routed to IOAPIC input 6	
bit 7	1	XINT7# Input Pin routed to ICU XINT7# Input Pin	
Dit 1	0	XINT7# Input Pin routed to IOAPIC input 7	
bit 16	1	XINT0# Input Pin routed to the B Segment IOAPIC	
Dit TO	0	XINT0# Input Pin routed to the A Segment IOAPIC	
bit 17	1	XINT1# Input Pin routed to the B Segment IOAPIC	
Dit 17	0	XINT1# Input Pin routed to the A Segment IOAPIC	
bit 18	1	XINT2# Input Pin routed to the B Segment IOAPIC	
Dit 10	0	XINT2# Input Pin routed to the A Segment IOAPIC	
bit 19	1	XINT3# Input Pin routed to the B Segment IOAPIC	
bit 15	0	XINT3# Input Pin routed to the A Segment IOAPIC	
hit 20	1	XINT4# Input Pin routed to the B Segment IOAPIC	
bit 20	0	XINT4# Input Pin routed to the A Segment IOAPIC	
bit 21	1	XINT5# Input Pin routed to the B Segment IOAPIC	
	0	XINT5# Input Pin routed to the A Segment IOAPIC	
bit 22	1	XINT6# Input Pin routed to the B Segment IOAPIC	
	0	XINT6# Input Pin routed to the A Segment IOAPIC	
bit 23	1	XINT7# Input Pin routed to the B Segment IOAPIC	
511 25	0	XINT7# Input Pin routed to the A Segment IOAPIC	

*Note:* XINT0# through XINT7# of the 80333 Interrupt Controller only handles level sensitive inputs such as PCI interrupts. When any Select bit is cleared, the logic external to the ICU input must drive an inactive level ('1') to the corresponding interrupt input of the 80333 Interrupt Controller.



## 17.5 Intel<sup>®</sup> 80333 I/O Processor Interrupt Controller Unit

The 80333 Interrupt Controller Unit (ICU) provides a flexible, low-latency means for requesting interrupts and minimizing the core's interrupt handling burden.

All interrupt sources are combined into one of the two internal interrupt exceptions: IRQ and FIQ.

The interrupt controller provides the following features for managing hardware-requested interrupts:

- Flexibility to direct interrupt sources to either the FIQ or IRQ internal interrupt exception
- 17 external interrupt pins.
  - One high-priority maskable interrupt pin, HPI#.
  - Eight maskable Inputs, XINT[7:0]#.
  - Eight GPIO Pins (when configured as Inputs), GPIO[7:0]/XINT[15:8]#
- Two internal programmable timer and a Watch Dog Timer sources.
- Peripheral interrupt sources.

All interrupts are *level sensitive*: interrupt sources must keep asserting the interrupt signal until software causes the source to deassert it.

All interrupt sources are individually maskable with the ICUs Interrupt Control registers (

). Additionally, all interrupts may be quickly disabled by altering the F and I bits in the CPSR as specified in the *ARM Architecture Reference Manual* - ARM Limited. Order number: ARM DDI 0100E..

When software running on the 80333 is vectored to an Interrupt Service Routine (ISR), it reads the ICUs IRQ Interrupt Vector Register (IINTVEC) or FIQ Interrupt Vector register (FINTVEC) to quickly retrieve the address for the interrupt handler of the highest priority active interrupt source.



### 17.5.1 Programmer Model

Software has access to 15 registers in the ICU. These registers control, masking, prioritization, and vector generation for all interrupt sources.

#### 17.5.1.1 Active Interrupt Source Control and Status

The INTCTL[1:0] register is used to enable or disable (mask) individual interrupts. As mentioned, masking of all interrupts may still be accomplished via the CPSR register in the core. INTSTR[1:0] are used to direct internal interrupts to either FIQ or IRQ. IINTSRC[1:0] and FINTSRC[1:0] are read-only registers that record all currently active and unmasked interrupt sources; the architecture for the interrupt source registers and FIQ/IRQ generation is illustrated in Figure 145.

#### Figure 145. Interrupt Controller Block Diagram (Active Interrupt Source Registers)





#### 17.5.1.2 Prioritization and Vector Generation for Active Interrupt Sources

The INTPRI[3:0] registers reserve two bits for each source to assign one of four priority levels.

00 <sub>2</sub> - High Priority
01 <sub>2</sub> - Med/High Priority
10 <sub>2</sub> - Med/Low Priority
11 <sub>2</sub> - Low Priority

When interrupt vector generation is enabled and there are multiple requests pending either in the FINTSRC[1:0] or the IINTSRC[1:0] registers, the prioritizer selects a highest priority active source for each source register.

*Note:* When multiple interrupts at the same priority level are pending for either FIQ or IRQ, the highest priority active source is selected according to a fixed priority based on bit location. Highest order bit is first.

The INTBASE and INTSIZE registers are used to establish a contiguous Interrupt Service Routine (ISR) memory range for all of 64 possible sources. The architecture provides for an ISR ranging from 256 bytes to 64 Kbytes per source.

The actual vector value is a function of the INTBASE and the INTSIZE registers and is based on a fixed order of all 64 possible interrupt sources. The vectors begin at INTBASE with source 0 (i.e., IINTSRC0 bit 0), and end at INTBASE + INTSIZE (per source)\*63 with source 63 (i.e., IINTSRC1 bit 31).

#### Example 8. Determining the Location of the Interrupt Handler for Source 25

```
INTBASE = 0x81200000 ; 2 Mbyte Aligned Base Address
INTSIZE = 0xE ; 32 Kbytes per source (ISR Memory Range of 2 Mbytes)
ISR Address(25) = 0x81200000 + conv_hex(2^15*25) = 0x812C8000
```



Based on IINTSRC[1:0], FINTSRC[1:0], INTPRI[3:0], INTBASE, and INTSIZE, the interrupt controller generates the values provided by the IINTVEC and FINTVEC registers as illustrated in Figure 146. The IINTVEC and FINTVEC registers presents the vector for the active interrupt source with the highest priority to the IRQ and FIQ exception handlers, respectively.

#### Figure 146. Interrupt Controller Block Diagram (FIQ/IRQ Interrupt Vector Generation)



Note: The 80333 does not use all 64 possible sources.

ICU registers reside in Coprocessor 6 (CP6). They may be accessed/manipulated with the MCR, MRC, STC, and LDC instructions. The instruction *CRn* field denotes the accessed register number. The instruction *opcode\_1*, *opcode\_2*, and *CRm* fields should be zero. Most systems restricts access to CP6 to privileged processes. To control access to CP6, use Coprocessor Access Register.

An instruction that modifies an ICU register is guaranteed to take effect before the next instruction executes. For example, when an instruction masks an interrupt source, subsequent instructions execute in an environment in which the masked interrupt does not occur.



## 17.5.2 Operational Blocks

The ICU provides the connections to the Intel XScale<sup>®</sup> core. These connections are shown in Figure 147.





# intel

## 17.5.3 80333: Internal Peripheral Interrupt

The 80333 Interrupt Controller receives inputs from multiple internal interrupt sources. All pending interrupts required during normal operation of the various peripheral units are available in either the IINTSRC0, IINTSRC1, FINTSRC0 or FINTSRC1 registers depending on the value in INTSTR0 or INTSTR1. To provide the best latency for high performance event driven activities, the DMA channel and AAU interrupts are fully demultiplexed into the interrupt source registers for FIQ and IRQ so that software does not need to access these peripheral units to diagnose the exact source and cause of the interrupt. The IINTSRC[1:0] and the FINTSRC[1:0] registers also include pending interrupts that indicate that an error has occurred in one of the peripheral units. For the interrupts that indicate errors, more detail about the exact cause of the interrupt can be determined by reading the status register of the respective peripheral unit.

#### 17.5.3.1 Normal Interrupt Sources

The 80333 Interrupt Controller receives normal interrupts from the two DMA channels, Performance Monitoring Unit, the I<sup>2</sup>C Bus Interface Unit, the ATU, the Programmable Timers, the Messaging Unit, the Application Accelerator Unit (AAU) and the UARTs. The DMA/AAU channel interrupts for End of Transfer interrupt or End of Chain interrupt are demultiplexed into the interrupt controller. A Performance Monitoring Unit interrupt implies that at least one of the fourteen programmable event counters and/or the Global Time Stamp Counter has a pending interrupt condition.

A valid interrupt from any of these sources outputs a *level-sensitive* interrupt to the 80333 Interrupt Controller input. The corresponding IRQ or FIQ interrupt source register bit in the interrupt controller should remain active as long as the interrupt is pending in the peripheral unit. The appropriate interrupt source bit is cleared by clearing the source of the interrupt at the internal peripheral.

The normal interrupt sources which drive the inputs to the 80333 Interrupt Controller are detailed in Table 443.

*Note:* The UART and I<sup>2</sup>C Bus Interface Unit interrupt sources are combined as a single interrupt, and include both normal and error conditions within the respective units.

Table 443.Normal Interrupt Sources (Sheet 1 of 2)

Unit	Interrupt Condition	Register
DMA Channel 0	End of Chain	Section 6.14.2, "Channel Status Register x - CSRx" on page 400
DMA Channel 0	End of Transfer	Section 6.14.2, "Channel Status Register x - CSRx" on page 400
DMA Channel 1	End of Chain	Section 6.14.2, "Channel Status Register x - CSRx" on page 400
DMA Channel 1	End of Transfer	Section 6.14.2, "Channel Status Register x - CSRx" on page 400
Application Accelerator	End of Chain	Section 6.14.2, "Channel Status Register x - CSRx" on page 400
Application Accelerator	End of Transfer	Section 6.14.2, "Channel Status Register x - CSRx" on page 400



#### Table 443. Normal Interrupt Sources (Sheet 2 of 2)

Unit	Interrupt Condition	Register	
		Receive Buffer Full	
0		Transmit Buffer Empty	
l <sup>2</sup> C Bus	I <sup>2</sup> C Status	Slave Address Detect (General Call Address Detect)	
0	Register 0	STOP Detected	
		Bus Error Detected	
		Arbitration Lost Detected	
		Receive Buffer Full	
_		Transmit Buffer Empty	
l <sup>2</sup> C Bus	I <sup>2</sup> C Status	Slave Address Detect (General Call Address Detect)	
Interface Unit	Register 1	STOP Detected	
		Bus Error Detected	
		Arbitration Lost Detected	
		Index Register Interrupt	
	Inbound	Inbound Post Queue Interrupt	
Messaging	Interrupt	Inbound Doorbell Interrupt	
Offic	Register	Inbound Message 1 Interrupt	
		Inbound Message 0 Interrupt	
متانا	ATU Interrupt Status Register	ATU BIST Start	
Alo	Configure Register Write	Any of the ATU Configuration registers written by an inbound Configuration Write cycle or changes to the PCI Express-to-PCI Bridge bus number registers in the ATU. This includes dedicated interrupt status bits for configuration writes to the VPDAR, IABAR1 and IAUBAR1 registers.	
Timer 0	Timer Mode Register 0	Timer 0 has decremented to 0 interrupt.	
Timer 1	Timer Mode Register 1	Timer 1 has decremented to 0 interrupt.	
		Received Line Status	
	UART 0 Interrupt ID Register	Received Data is Available	
UART Unit 0		Character Timeout Indications	
		Transmit FIFO Data Request	
		Autobaud Lock Indication	
		Received Line Status	
	UART 1	Received Data is Available	
UART Unit 1	Interrupt ID Register	Character Timeout Indications	
		Transmit FIFO Data Request	
		Autobaud Lock Indication	



#### 17.5.3.2 Error Interrupt Sources

The 80333 Interrupt Controller receives error interrupts from the ATU, the Messaging Unit, the two DMA channels and the AAU. Each of these interrupts represent an error condition in the peripheral unit. Refer to the appropriate units for more details.

A valid interrupt from any of these sources outputs a *level-sensitive* interrupt to the 80333 Interrupt Controller input. The corresponding FIQ or IRQ interrupt source register bit in the interrupt controller should remain active as long as the interrupt is pending in the peripheral unit. The appropriate FIQ or IRQ interrupt source bit is cleared by clearing the source of the interrupt at the internal peripheral.

#### Table 444.Error Interrupt Sources

Unit	Register	Error Condition	
		Initiated Split Completion Error Message	
		Received Split Completion Error Message	
		Power State Transition	
		P_SERR# Asserted	
		PCI Detected Parity Error	
ATU	ATU Interrupt Status Register	IB Master Abort	
		P_SERR# Detected	
		PCI Master Abort	
		PCI Target Abort (master)	
		PCI Target Abort (target)	
		PCI Master Parity Error	
Messaging Unit	Inhound Interrupt Status Register	Outbound Free Queue Full Interrupt	
Messaging Onit	inbound interrupt Status Register	Error Doorbell Interrupt	
		IB Master Abort	
DMA Channel 0	Channel Status Register 0	PCI Master Abort	
DIVIA CHAIINEI U	Channel Status Register 0	PCI Target Abort (master)	
		Unexpected Split Completion	
		IB Master Abort	
DMA Channel 1	Channel Status Register 1	PCI Master Abort	
	Channel Glatus Register 1	PCI Target Abort (master)	
		Unexpected Split Completion	
Application Accelerator (AAU Error)	Accelerator Status Register	IB Master Abort	
Intel XScale <sup>®</sup> Core* Bus Interface Unit Bus Interface Unit Status Registe		Master Abort	
		ECC Error 0	
		ECC Error 1	
Memory Controller	Status Register	ECC Error N	
		Address Region Error	
		IB Discard Timer Expired	

The PCI Interrupt Routing Select Register and the Interrupt Source Register are described in Section 17.8.

Intel<sup>®</sup> 80333 I/O Processor Developer's Manual Interrupt Controller Unit and IOAPIC



## 17.5.4 High-Priority Interrupt (HPI#)

The HPI# pin generates an interrupt for implementation of critical interrupt routines.

## 17.5.5 Timer Interrupts

Each of the two timer units has an associated interrupt. Timer interrupts are connected directly to the 80333 interrupt controller and are posted in either the IINTSRC0 or FINTSRC0 registers. These interrupts are set up through the timer control registers described in Chapter 16, "Timers."

The Watch Dog Timer (WDT) can also be unmasked as an interrupt source. The WDT is connected directly to the interrupt controller and posted in either the IINTSRC0 or FINTSRC0 registers. This interrupt is set up through the INTCTL0 register and the WDTCR registers described in Chapter 16, "Timers."

### 17.5.6 Software Interrupts

The application program may use the SWI instruction to request interrupt service.

## 17.6 The Intel<sup>®</sup> 80333 I/O Processor IOAPIC

The 80333 contains two IOAPIC controllers, both of which reside on the primary bus. The intended use of these controllers is to have the interrupts from PCI bus A connected to the interrupt controller on device#0/function#1, and have the interrupts on PCI bus B connected to the interrupt controller on device#0/function#3.

80333 behaves as a normal PCI-to-PCI Bridge and can handle all the three system interrupt mechanisms detailed in the previous section. Figure 148 summarizes the 80333 interrupt architecture.

#### 17.6.1 PCI IRQ# Interrupts

80333 can handle 8 pin interrupts. 80333 has 8 pins (**XINT[7:0]**#) for these interrupts. Interrupts delivered by a pin can be either in level or edge mode, and may be either active high or active low through the "IOAPIC Redirection Table Low DWORD IRQxx - APIC\_RDLxx (Offset 10-3E)" on page 821. Since this IOAPIC is connected to a PCI bus, its most likely configuration will be as active low level, which will match the PCI pin polarity and functionality. Each pin is collected by the 80333, synchronized into the PCI clock domain, and scheduled for delivery when it is unmasked.

The 80333 supports up to 8 interrupt pins per PCI segment, which can be redirected to either IOAPIC. These pins are connected to IOAPIC redirection table entries 7 – 0 (of 24 entries). The standard Hot-Plug controller is hard-wired to redirection table entry 23 of the Segment-B IOAPIC. The Address Translation Unit (ATU) is hard-wired to redirection table entry 14 of the Segment-A IOAPIC. All other interrupts are only addressable through PCI virtual wire mechanism. The unused interrupt input pins must be connected to  $V_{CC}$  to ensure the boot interrupt works correctly.

#### Figure 148. Intel<sup>®</sup> 80333 I/O Processor IOAPIC Interrupt Support





#### 17.6.2 PCI MSI Interrupts

These interrupts which appear on PCI as inbound memory writes are decoded by 80333 in the PCI-to-PCI Bridge inverse decode window and passed upstream without any modifications. BIOS would setup the PCI-to-PCI Bridge decode register such that 0xFEEx\_xxxx falls in the inverse decode window of the 80333.

## 17.6.3 PCI Virtual Wire Interrupts

PCI devices that directly write to the Pin Assertion Register in the corresponding 80333 IOAPIC devices causing an upstream interrupt message. All 24 interrupt entries are addressable through this mechanism. Values above 23 are invalid in the pin assertion register and ignored by the 80333.

80333 decodes the memory write transactions to the pin assertion register in the PCI-to-PCI Bridge inverse decode window, just as any other upstream transaction.

## 17.6.4 PCI Express Legacy INTx Support and Boot Interrupt

The 80333 contains a capability to generate an in-band interrupt request on PCI Express when the IOAPIC is disabled. This in-band interrupt mechanism is necessary for systems that do not support the IOAPIC and for boot. PCI Express describes an in-band legacy wire-interrupt INTx mechanism for I/O devices to signal PCI-style level interrupts. 80333 generates a PCI Express INTx message as follows. Each interrupt pin input of the IOAPIC is compared with its mask. When the interrupt is masked in the 80333 IOAPIC, then that interrupt needs to cause a INTx message over PCI Express whenever asserted. When the interrupt is not masked, then that interrupt is being used by the 80333 IOAPIC and should not cause a INTx message on PCI Express.

In PCI Express, the legacy interrupts are virtualized using a pair of ASSERT and DEASSERT messages. This then gives a way to preserve the level-sensitive semantics of the PCI interrupts on PCI Express. PCI Express defines four virtual wire interrupts: INTA:INTD - corresponding to the four interrupt wires defined in the *PCI Local Bus Specification*, Revision 2.3. PCI Express routes its PIC interrupt pins and the internal interrupts, to PCI Express INTx interrupts according to Table 445:

#### Table 445. Exception Priorities and Vectors

IOAPIC Pin and Corresponding External Interrupt Input	Internal Interrupts	PCI Express INTx Interrupts
0 - XINT0#, 4 - XINT4#		INTA
1 - XINT1#, 5 - XINT5#	SHPC (PCI B)	INTB
2 - XINT2#, 6 - XINT6#	ATU	INTC
3 - XINT3#, 7 - XINT7#		INTD

Each row in the table indicates a logical ORing of interrupts in that row, internally in 80333. The ASSERT/DEASSERT message sent out on PCI Express captures the asserting/deasserting edge of the signal that represents the logical OR of the interrupts in that row. The pair of ASSERT and DEASSERT messages are then routed by the MCH to ICH using a DO\_PCI\_INIT message or by simply asserting a sideband interrupt pin to ICH.

The RequestorID used in the PCI Express INTx message that 80333 generates would be "80333 primary bus number: 80333 device number=0:0".

March 2005

*Note:* PCI Express INTx messages are not inhibited by the BME bit.



## 17.6.5 Buffer Flushing

80333 does not implement any buffer flushing feature (i.e., when 80333 receives an interrupt on its interrupt pin, it does not flush its posted write buffers in the inbound direction in the PCI interface). This is not required from 80333 because PCI device drivers ultimately have to guarantee that all posted writes from the device to the memory are all flushed before executing the interrupt service routine. It could do this by doing a dummy read to the device that originated the interrupt. Or the originating device could solve this problem when it did a read of the memory to guarantee that the write before it went through and then signal an interrupt.

## 17.6.6 EOI Special Cycles

80333 could get EOI special cycles over PCI Express in the IA32 FSB mode. This is the result of broadcasting the IA32 FSB EOI cycle. Both IOAPICs in 80333 would compare the vector number in the EOI data field with the vector field for each entry in the I/O Redirection Table. When a match is found, the Remote\_IRR bit for that I/O Redirection Entry in the IOAPIC will be cleared. Note, when multiple I/O Redirection entries, for any reason, assign the same vector for more than one interrupt input, each of those entries have the Remote\_IRR bit reset to '0'. EOI is not forwarded to PCI.

## 17.6.7 Upstream Interrupt Message Format

When an interrupt message needs to be sent over PCI Express (i.e., when the IRR bit is set for an interrupt, the 80333 performs a memory write on PCI Express, as seen in Table 446 and Table 447).

#### Table 446. Front Side Bus Delivery Address Format

Bit	Description
31:20	FEEh
19:12	Destination ID: Same as bits 63:56 of the I/O Redirection Table entry for the interrupt associated with this message.
11:4	Enhanced Destination ID: Same as bits 55:48 of I/O Redirection Table entry for interrupt associated with this message.
	Redirection Hint: This bit is used by the CPU host bridge to allow the interrupt message to be redirected.
3	<ul> <li>0 = The message will be delivered to the agent (CPU) listed in bits 19:4.</li> <li>1 = The message will be delivered to an agent with a lower interrupt priority</li> </ul>
	The Redirection Hint bit will be a 1 when bits 10:8 in the Delivery Mode field associated with corresponding interrupt are encoded as 001 (Lowest Priority). Otherwise, the Redirection Hint bit will be 0.
2	<b>Destination Mode:</b> This bit is used only the Redirection Hint bit is set to 1. When the Redirection Hint bit and the Destination Mode bit are both set to 1, then the logical destination mode is used, and the redirection is limited only to those processors that are part of the logical group as based on the logical ID.
1:0	00

#### Table 447. Front Side Bus Delivery Data Format

Bti	Description
31:16	0000H
15	<b>Trigger Mode:</b> 1 = Level, 0 = Edge. Same as the corresponding bit in the I/O Redirection Table for that interrupt.
14	<b>Delivery Status:</b> 1 = Assert, 0 = De-Assert. When using edge-triggered interrupts, then bit will always be 1, since only the assertion is sent. When using level-triggered interrupts, then this bit indicates the state of the interrupt input.
13:12	002
11	<b>Destination Mode:</b> 1 = Logical, 0 = Physical. Same as the corresponding bit in the Redirection Table
10:8	Delivery Mode: This is the same as the corresponding bits in the I/O Redirection Table for that interrupt.
7:0	Vector: This is the same as the corresponding bits in the I/O Redirection Table for that interrupt.



## 17.7 Default Status

The interrupt logic is reset by the PCI reset signal or through software. Table 448 shows the power-up and reset values.

#### Table 448. Default Interrupt Routing and Status Values

Register	Default Value	Description
INTCTL0	0000 0000H	All interrupts 31:0 masked.
INTCTL1	0000 0000H	All interrupts 63:32 masked.
INTSTR0	0000 0000H	All interrupts 31:0 steered to IRQ.
INTSTR1	0000 0000H	All interrupts 63:32 steered to IRQ.
IINTSRC0	0000 0000H	All IRQ interrupts 31:0 inactive.
IINTSRC1	0000 0000H	All IRQ interrupts 63:32 inactive.
FINTSRC0	0000 0000H	All FIQ interrupts 31:0 inactive.
FINTSRC1	0000 0000H	All FIQ interrupts 63:32 inactive.
PIRSR	00F0 0000H	XINT[3:0]# routed to the A PCI Segment IOAPIC (Function 1) IRQ inputs 3:0. XINT[7:4]# routed to the B PCI Segment IOAPIC (Function 3) IRQ inputs 7:4.

## 17.8 Interrupt Control Unit Registers

int

All Interrupt Controller registers are visible as 80333 memory mapped registers and can be accessed through the internal memory bus. Each is a 32-bit register and is memory-mapped in the Intel XScale<sup>®</sup> core memory space.

The PCI Interrupt Routing Select Register is accessible from the internal memory bus and also during PCI configuration cycles through the PCI configuration register space. See Chapter 3, "Address Translation Unit" for additional information regarding the PCI configuration cycles that access the PCI Interrupt Routing Select Register. The programmer's interface to the interrupt controller is through both coprocessor registers and memory-mapped control register. Table 449 describes these registers.

The coprocessor registers may be accessed/manipulated with the MCR, MRC, STC, and LDC instructions. The *CRn* field of the instruction denotes the register number to be accessed. The *opcode\_1*, *opcode\_2*, and *CRm* fields of the instruction should be zero. Most systems restricts access to CP6 to privileged processes. To control access to CP6, use the Coprocessor Access Register.

Register Name	Description	Coprocessor Register (CR <sub>m</sub> = 0) or MMR Address
INTCTL0	Interrupt Control Register 0	CP6, Register 0 (Read/Write) FFFF E790H (Read-Only)
INTCTL1	Interrupt Control Register 1	CP6, Register 1 (Read/Write) FFFF E794H (Read-Only)
INTSTR0	Interrupt Steer Register 0	CP6, Register 2 (Read/Write) FFFF E798H (Read-Only)
INTSTR1	Interrupt Steer Register 1	CP6, Register 3 (Read/Write) FFFF E79CH (Read-Only)
IINTSRC0	IRQ Interrupt Source Register 0	CP6, Register 4 (Read-Only) FFFF E7A0H (Read-Only)
IINTSRC1	IRQ Interrupt Source Register 1	CP6, Register 5 (Read-Only) FFFF E7A4H (Read-Only)
FINTSRC0	FIQ Interrupt Source Register 0	CP6, Register 6 (Read-Only) FFFF E7A8H (Read-Only)
FINTSRC1	FIQ Interrupt Source Register 1	CP6, Register 7 (Read-Only) FFFF E7ACH (Read-Only)
IPR0	Interrupt Priority Register 0	CP6, Register 8 (Read/Write) FFFF E7B0H (Read-Only)
IPR1	Interrupt Priority Register 1	CP6, Register 9 (Read/Write) FFFF E7B4H (Read-Only)
IPR2	Interrupt Priority Register 2	CP6, Register 10 (Read/Write) FFFF E7B8H (Read-Only)
IPR3	Interrupt Priority Register 3	CP6, Register 11 (Read/Write) FFFF E7BCH (Read-Only)
INTBASE	Interrupt Base Register	CP6, Register 12 (Read/Write) FFFF E7C0H (Read-Only)
INTSIZE	Interrupt Size Register	CP6, Register 13 (Read/Write) FFFF E7C4H (Read-Only)
IINTVEC	IRQ Interrupt Vector Register	CP6, Register 14 (Read/Write) FFFF E7C8H (Read-Only)
FINTVEC	FIQ Interrupt Vector Register	CP6, Register 15 (Read/Write) FFFF E7CCH (Read-Only)
PIRSR	PCI Interrupt Routing Select Register	FFFF E1ECH

#### Table 449. Interrupt Controller Register Addresses



#### 17.8.1 Interrupt Control Register 0 - INTCTL0

The Interrupt Control register 0 is a 32-bit Coprocessor 6 control register used to specify which of 32 interrupts are masked.

*Note:* INTCTL0 register is read-only from Memory-Mapped Register Address space.

 Table 450.
 Interrupt Control Register 0 - INTCTL0 (Sheet 1 of 3)





Table 450.         Interrupt Control Register 0 - INTCTL0 (Sheet 2 of 3)						
$\begin{array}{c} 31  28  24  20  16  12  8  4  0 \\ \hline \text{Coprocessor} \\ \text{Attributes} \end{array} \begin{bmatrix} 31  28  24  20  16  12  8  4  0 \\ \hline wrw rw $						
	CP6, Register 0       RV = Reserved       RC = Read Clear         Intel XScale® Core Local Bus Address       PR = Preserved       RC = Read Clear         FFFF E790H       RS = Read/Set       NA = Not Accessible					
Bit	Default	Description				
16	02	Intel XScale <sup>®</sup> Core PMU Interrupt Mask 0 = Masked 1 = Not Masked				
15	02	Peripheral Performance Monitor Interrupt Mask 0 = Masked 1 = Not Masked				
14	02	ATU/Start BIST Interrupt Mask 0 = Masked 1 = Not Masked				
13	02	Messaging Unit Inbound Post Queue Interrupt Mask 0 = Masked 1 = Not Masked				
12	02	Messaging Unit Interrupt Mask 0 = Masked 1 = Not Masked				
11	02	I <sup>2</sup> C Bus Interface 1 Interrupt Mask 0 = Masked 1 = Not Masked				
10	02	I <sup>2</sup> C Bus Interface 0 Interrupt Mask 0 = Masked 1 = Not Masked				
9	02	Timer 1 Interrupt Mask 0 = Masked 1 = Not Masked				
8	02	Timer 0 Interrupt Mask 0 = Masked 1 = Not Masked				
7	02	Application Accelerator End-Of-Chain Interrupt Mask 0 = Masked 1 = Not Masked				
6	02	Application Accelerator End-Of-Transfer Interrupt Mask 0 = Masked 1 = Not Masked				



 Table 450.
 Interrupt Control Register 0 - INTCTL0 (Sheet 3 of 3)



#### 17.8.2 Interrupt Control Register 1 - INTCTL1

The Interrupt Control register 1 is a 32-bit Coprocessor 6 control register used to specify which of 32 interrupts are masked.

*Note:* INTCTL1 register is read-only from Memory-Mapped Register Address space.

 Table 451.
 Interrupt Control Register 1 - INTCTL1 (Sheet 1 of 2)





Soprocessor       Attributes       31       28       24       20       16       12       8       4       0         Coprocessor       rw/rw/rw/rw/pr/pr/wr/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw/rw					
	CF	P6, Register 1	RV = Reserved	RC = Read Clear	
	Int FF	tel XScale <sup>®</sup> Core Local Bus Address FFF E794H	PR = Preserved RS = Read/Set	RO = Read Only NA = Not Accessible	
Bit	Default	De	scription		
21	02	Reserved.			
20	02	UART 1 Interrupt Mask 0 = Masked 1 = Not Masked			
19	02	UART 0 Interrupt Mask 0 = Masked 1 = Not Masked			
18:08	000H	Preserved			
7	02	XINT15# Interrupt Mask 0 = Masked 1 = Not Masked			
6	02	XINT14# Interrupt Mask 0 = Masked 1 = Not Masked			
5	02	XINT13# Interrupt Mask 0 = Masked 1 = Not Masked			
4	02	XINT12# Interrupt Mask 0 = Masked 1 = Not Masked			
3	02	XINT11# Interrupt Mask 0 = Masked 1 = Not Masked			
2	02	XINT10# Interrupt Mask 0 = Masked 1 = Not Masked			
1	02	XINT9# Interrupt Mask 0 = Masked 1 = Not Masked			
0	02	XINT8# Interrupt Mask 0 = Masked 1 = Not Masked			





## 17.8.3 Interrupt Steering Register 0 - INTSTR0

The Interrupt Steering Register 0 allows system designers to direct any of 32 internal or external interrupt sources to either one of the two internal interrupt exceptions, FIQ and IRQ.

When an interrupt is enabled with the INTCTL0 register, this register steers the interrupt to an internal interrupt exception.

*Note:* INTSTR0 register is read-only from Memory-Mapped Register Address space.

Table 452. Interrupt Steering Register 0 - INTSTR0 (Sheet 1 of 3)





31       28       24       20       16       12       8       4       0         Coprocessor Attributes						
	CP6, Re Intel XS	egister 2 cale <sup>®</sup> Core Local Bus Address	PR = Preserved RS = Read/Set	RO = Read Only NA = Not Accessible		
D:4						
DI	Derault	Watch Dog Timor Interrupt				
17	02	0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ				
16	02	Intel XScale <sup>®</sup> Core PMU Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ				
15	02	Peripheral Performance Monitor Interrupt Steer 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ	ing			
14	02	ATU/Start BIST Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ				
13	02	Messaging Unit Inbound Post Queue Interrupt 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ	Steering			
12	02	Messaging Unit Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ				
11	02	<ul> <li>I<sup>2</sup>C Bus Interface 1 Interrupt Steering</li> <li>0 = Interrupt Directed to Internal IRQ</li> <li>1 = Interrupt Directed to Internal FIQ</li> </ul>				
10	02	I <sup>2</sup> C Bus Interface 0 Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ				
9	02	Timer 1 Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ				
8	02	Timer 0 Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ				
7	02	Application Accelerator End-Of-Chain Interrupt 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ	Steering			
6	02	Application Accelerator End-Of-Transfer Interru 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ	upt Steering			









#### 17.8.4 Interrupt Steering Register 1 - INTSTR1

The Interrupt Steering Register 1 allows system designers to direct any of 32 internal or external interrupt sources to either one of the two internal interrupt exceptions, FIQ and IRQ.

When an interrupt is enabled with the INTCTL1 register, this register steers the interrupt to an internal interrupt exception.

*Note:* INTSTR1 register is read-only from Memory-Mapped Register Address space.

Table 453. Interrupt Steering Register 1 - INTSTR1 (Sheet 1 of 2)





Table 4	53. Inter	errupt Steering Register 1 - INTSTR1 (Sheet 2 of 2)			
31       28       24       20       16       12       8       4       0         Coprocessor Attributes					
	Intel XS CP6, Re Intel XS FFFF E	cale <sup>®</sup> Core Coprocessor address egister 3 cale <sup>®</sup> Core Local Bus Address 79CH	Attribute Legend: RV = Reserved PR = Preserved RS = Read/Set	RW = Read/Write RC = Read Clear RO = Read Only NA = Not Accessible	
Bit	Default		Description		
22	02	ATU Configuration Register Write Interrupt 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
21	02	Reserved.			
20	02	UART 1 Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
19	02	UART 0 Interrupt Steering 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
18:08	000H	Preserved			
7	02	XINT15# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
6	02	XINT14# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
5	02	XINT13# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
4	02	XINT12# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
3	02	XINT11# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
2	02	XINT10# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
1	02	XINT9# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			
0	02	XINT8# Interrupt Mask 0 = Interrupt Directed to Internal IRQ 1 = Interrupt Directed to Internal FIQ			



## 17.8.5 IRQ Interrupt Source Register 0 - IINTSRC0

The IRQ Interrupt Source register is a 32-bit Coprocessor 6 control register used to specify which of 32 interrupts that are steered to the internal IRQ exception are unmasked by the INTCTL0 register and active. The INTSTR0 control register is used to steer individual interrupts to the IRQ exception.

The IINTSRC0 register may be used by an Interrupt Service Routine (ISR) to determine quickly the source of an IRQ interrupt.

*Note:* IINTSRC0 register is read-only from Memory-Mapped Register Address space.



 Table 454.
 IRQ Interrupt Source Register 0 - IINTSRC0 (Sheet 1 of 3)



Table 454.         IRQ Interrupt Source Register 0 - IINTSRC0 (Sheet 2 of 3)				
31       28       24       20       16       12       8       4       0         Coprocessor Attributes				
Bit	Default	Description		
23:18	00 00002	Reserved		
17	02	Watch Dog Timer Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0		
		Intel XScale <sup>®</sup> Core PMU Interrupt		
16	02	0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0		
15	02	Peripheral Performance Monitor Interrupt - when set, at least one of the programmable event counters and/or the Global Time Stamp Counter contains an overflow condition. Application software identifies the counter by reading the Event Monitoring Interrupt Status register (EMISR). 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCT1 0		
14	02	<ul> <li>ATU/Start BIST Interrupt - when set, the host processor has set the start BIST request in the ATUBISTR register.</li> <li>0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0</li> </ul>		
		Messaging Unit Inbound Post Queue Interrupt		
13	02	<ul> <li>0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0</li> </ul>		
		Messaging Unit Interrupt		
12 0 <sub>2</sub> 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0		<ul> <li>0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0</li> </ul>		
		I <sup>2</sup> C Bus Interface 1 Interrupt		
11         02         0 = Not Interrupting or Not steered to internal IRQ exception or masked by INT           1 = Interrupting and steered to internal IRQ exception and unmasked by INTC		<ul> <li>0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0</li> </ul>		
		I <sup>2</sup> C Bus Interface 0 Interrupt		
10	02	<ul> <li>0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0</li> </ul>		
	02	Timer 1 Interrupt		
9		<ul> <li>0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0</li> </ul>		
	02	Timer 0 Interrupt		
8		0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0		

Copre Att At	31       28       24       20       16       12       8       4       0         Coprocessor Attributes				
	FFFF E7	Core Local bus Address     RS = Read/Set     NA = Not Accessible       7A0H			
Bit	Default	Description			
7	02	Application Accelerator End-Of-Chain Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0			
6	02	Application Accelerator End-Of-Transfer Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0			
5:4	002	Reserved			
3	02	DMA Channel 1 End-Of-Chain Interrupt         0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0         1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0			
2	02	DMA Channel 1 End-Of-Transfer Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0			
1	02	DMA Channel 0 End-Of-Chain Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0			
0	0 02 DMA Channel 0 End-Of-Transfer Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL0 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL0				

#### Table 454. IRQ Interrupt Source Register 0 - IINTSRC0 (Sheet 3 of 3)

intel



## 17.8.6 IRQ Interrupt Source Register 1 - IINTSRC1

The IRQ Interrupt Source register is a 32-bit Coprocessor 6 control register used to specify which of 32 interrupts that are steered to the internal IRQ exception are unmasked by the INTCTL1 register and active. The INTSTR1 control register is used to steer individual interrupts to the IRQ exception.

The IINTSRC1 register may be used by an Interrupt Service Routine (ISR) to determine quickly the source of an IRQ interrupt.

*Note:* IINTSRC1 register is read-only from Memory-Mapped Register Address space.

#### Table 455. IRQ Interrupt Source Register 1 - IINTSRC1 (Sheet 1 of 2)

31       28       24       20       16       12       8       4       0         Coprocessor Attributes					
	Intel XS FFFF E	cale <sup>®</sup> Core Local Bus Address PR = Preserved RO = Read Only 7A4H RS = Read/Set NA = Not Accessible			
Bit	Default	Description			
31	0 <sub>2</sub>	HPI# Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			
30	02	Messaging Unit Error Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			
29	02	Bus Interface Unit Master Abort Error Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			
28	02	Application Accelerator Unit Error Interrupt         0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1         1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			
27	02	Reserved			
26	02	DMA Channel 1 Error Interrupt         0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1         1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			
25	02	DMA Channel 0 Error Interrupt         0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1         1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			
24	<ul> <li>24</li> <li>02</li> <li>Memory Controller Unit Error Interrupt - when set, an error condition exists within the MCU. The bit indicates one of the following conditions:         <ul> <li>A single-bit correctable or uncorrectable ECC error.</li> <li>A multi-bit correctable or uncorrectable ECC error.</li> <li>0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1</li> <li>Interrupting and steered to internal IRQ exception and unmasked by INTCTL1</li> </ul> </li> </ul>				
23	23 0 <sub>2</sub> ATU Error Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1				

March 2005

intel	3
-------	---

31       28       24       20       16       12       8       4       0         Coprocessor Attributes       Image: Coprocessor attributes       Image: Coprocessor attribute attri				
Bit	Default	Description		
22	02	ATU Configuration Register Write Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
21	02	Reserved.		
20	02	UART 1 Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
19	02	UART 0 Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
18:08	02	Reserved		
7	02	XINT15# Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
6	02	XINT14# Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
5	02	XINT13# Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
4	02	XINT12# Interrupt         0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1         1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
3	02	XINT11# Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
2	02	XINT10# Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1		
1	02       0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1         1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			
0	0 02 XINT8# Interrupt 0 = Not Interrupting or Not steered to internal IRQ exception or masked by INTCTL1 1 = Interrupting and steered to internal IRQ exception and unmasked by INTCTL1			

#### Table 455. IRQ Interrupt Source Register 1 - IINTSRC1 (Sheet 2 of 2)



## 17.8.7 FIQ Interrupt Source Register 0 - FINTSRC0

The FIQ Interrupt Source register 0 is a 32-bit Coprocessor 6 control register used to specify which interrupts that are steered to the internal FIQ exception are unmasked by the INTCTL0 register and active. The INTSTR0 control register is used to steer individual interrupts to the FIQ exception.

The FINTSRC0 register may be used by an Interrupt Service Routine (ISR) to determine quickly the source of an FIQ interrupt.

Note: FINTSRC0 register is read-only from Memory-Mapped Register Address space.



Table 456. FIQ Interrupt Source Register 0 - FINTSRC0 (Sheet 1 of 3)



$\begin{array}{c ccccccccccccccccccccccccccccccccccc$					
	Intel XS FFFF E7	PR = Preserved     RO = Read Only       cale <sup>®</sup> Core Local Bus Address     RS = Read/Set     NA = Not Accessible       7A8H     RS = Read/Set     NA = Not Accessible			
Bit	Default	Description			
17	02	Watch Dog Timer Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
16	02	Intel XScale <sup>®</sup> Core PMU Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
15	02	Peripheral Performance Monitor Interrupt - when set, at least one of the programmable event counters and/or the Global Time Stamp Counter contains an overflow condition. Application software identifies the counter by reading the Event Monitoring Interrupt Status register (EMISR). 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
14	02	<ul> <li>ATU/Start BIST Interrupt - when set, the host processor has set the start BIST request in the ATUBISTR register.</li> <li>0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0</li> </ul>			
13	02	Messaging Unit Inbound Post Queue Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
12	02	Messaging Unit Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
11	02	I <sup>2</sup> C Bus Interface 1 Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
10	02	<ul> <li>I<sup>2</sup>C Bus Interface 0 Interrupt</li> <li>0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0</li> <li>1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0</li> </ul>			
9	02	Timer 1 Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
8	02	Timer 0 Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0			
7	7       02       Application Accelerator End-Of-Chain Interrupt         0       0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL0         1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL0				









## 17.8.8 FIQ Interrupt Source Register 1 - FINTSRC1

The FIQ Interrupt Source register 1 is a 32-bit Coprocessor 6 control register used to specify which interrupts that are steered to the internal FIQ exception are unmasked by the INTCTL1 register and active. The INTSTR1 control register is used to steer individual interrupts to the FIQ exception.

The FINTSRC1 register may be used by an Interrupt Service Routine (ISR) to determine quickly the source of an FIQ interrupt.

Note: FINTSRC1 register is read-only from Memory-Mapped Register Address space.

#### 28 24 20 16 12 4 Δ 31 Coprocessor Attributes Memory Attributes Intel XScale<sup>®</sup> Core Coprocessor address Attribute Legend: RW = Read/Write RV = Reserved RC = Read Clear CP6, Register 7 PR = Preserved RO = Read Only Intel XScale<sup>®</sup> Core Local Bus Address NA = Not Accessible RS = Read/Set FFFF E7ACH Bit Default Description HPI# Interrupt 31 02 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1 Messaging Unit Error Interrupt 30 02 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1 Bus Interface Unit Master Abort Error Interrupt 29 $0_{2}$ 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1 Application Accelerator Unit Error Interrupt 28 $0_{2}$ 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1 27 02 Reserved DMA Channel 1 Error Interrupt 26 02 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1 DMA Channel 0 Error Interrupt 25 $0_{2}$ 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1 Memory Controller Unit Error Interrupt - when set, an error condition exists within the MCU. The bit indicates one of the following conditions: A single-bit correctable or uncorrectable ECC error. 24 $0_{2}$ A multi-bit correctable or uncorrectable ECC error. ٠ 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1

 Table 457.
 FIQ Interrupt Source Register 1 - FINTSRC1 (Sheet 1 of 3)



Table 457.         FIQ Interrupt Source Register 1 - FINTSRC1 (Sheet 2 of 3)						
Coproces Attribu Men Attrib	31       28       24       20       16       12       8       4       0         Coprocessor Attributes					
Bit	Default	Description				
23	02	ATU Error Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1				
22	02	ATU Configuration Register Write Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1				
21	02	Reserved.				
20	02	UART 1 Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1				
19	02	UART 0 Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1				
18:08	000H	Reserved				
7	02	XINT15# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1				
6	6 02 XINT14# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1					
5	5 02 XINT13# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1					

31       28       24       20       16       12       8       4       0         Coprocessor Attributes					
Bit	Default	Description			
4	02	XINT12# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1			
3	02	XINT11# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1			
2	2 02 XINT10# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1				
1	1 0 <sub>2</sub> XINT9# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1			by INTCTL1 by INTCTL1	
0	0 02 XINT8# Interrupt 0 = Not Interrupting or Not steered to internal FIQ exception or masked by INTCTL1 1 = Interrupting and steered to internal FIQ exception and unmasked by INTCTL1			by INTCTL1 y INTCTL1	

#### Table 457. FIQ Interrupt Source Register 1 - FINTSRC1 (Sheet 3 of 3)

intel


## 17.8.9 Interrupt Priority Register 0 - IPR0

The Interrupt Priority Register 0 is a 32-bit Coprocessor 6 control register used to assign a priority level to interrupt sources 15 down to 0. The IPR0 control register is used to assign one of 4 priority levels to each interrupt source independent of the INTSTR[1:0] registers:

00 <sub>2</sub> - High Priority	
01 <sub>2</sub> - Med/High Priority	Ī
10 <sub>2</sub> - Med/Low Priority	Ī
11 <sub>2</sub> - Low Priority	

When interrupt vector generation is enabled and there are multiple requests pending either in the FINTSRC[1:0] or the IINTSRC[1:0] registers, the highest priority vectors pending for either FIQ or IRQ is presented in the FINTVEC or IINTVEC, respectively.

*Note:* When multiple interrupts at the same priority level are pending for either FIQ or IRQ, the vector is selected according to a fixed priority based on bit location. Highest order bit is first.

31 28 24 20 16 12 8 4 0 Coprocessor Attributes Memory Attribute Intel XScale<sup>®</sup> Core Coprocessor address RW = Read/Write Attribute Legend: RV = Reserved RC = Read Clear CP6, Register 8 PR = Preserved RO = Read Only Intel XScale<sup>®</sup> Core Local Bus Address RS = Read/Set NA = Not Accessible FFFF E7B0H Description Bit Default 31:30 002 Peripheral Performance Monitor Interrupt Priority 29:28  $00_{2}$ ATU/Start BIST Interrupt Priority 27:26 Messaging Unit Inbound Post Queue Interrupt Priority 002 002 25:24 Messaging Unit Interrupt Priority I<sup>2</sup>C Bus Interface 1 Interrupt Priority 23:22  $00_{2}$ 21:20 I<sup>2</sup>C Bus Interface 0 Interrupt Priority 002 19:18 002 Timer 1 Interrupt Priority 17:16 002 Timer 0 Interrupt Priority 15:14 002 Application Accelerator End-Of-Chain Interrupt Priority 13:12 002 Application Accelerator End-Of-Transfer Interrupt Priority 11:08 002 Preserved 002 07:06 DMA Channel 1 End-Of-Chain Interrupt Priority 05:04 DMA Channel 1 End-Of-Transfer Interrupt Priority 002 002 03:02 DMA Channel 0 End-Of-Chain Interrupt Priority DMA Channel 0 End-Of-Transfer Interrupt Priority 01:00  $00_{2}$ 

 Table 458.
 Interrupt Priority Register 0 - IPR0



## 17.8.10 Interrupt Priority Register 1 - IPR1

The Interrupt Priority Register 1 is a 32-bit Coprocessor 6 control register used to assign a priority level to interrupt sources 31 down to 15. The IPR1 control register is used to assign one of 4 priority levels to each interrupt source independent of the INTSTR[1:0] registers:

00 <sub>2</sub> - High Priority
01 <sub>2</sub> - Med/High Priority
10 <sub>2</sub> - Med/Low Priority
11 <sub>2</sub> - Low Priority

When interrupt vector generation is enabled and there are multiple requests pending either in the FINTSRC[1:0] or the IINTSRC[1:0] registers, the highest priority vectors pending for either FIQ or IRQ is presented in the FINTVEC or IINTVEC, respectively.

*Note:* When multiple interrupts at the same priority level are pending for either FIQ or IRQ, the vector is selected according to a fixed priority based on bit location. Highest order bit is first.



#### Table 459. Interrupt Priority Register 1 - IPR1



## 17.8.11 Interrupt Priority Register 2 - IPR2

The Interrupt Priority Register 2 is a 32-bit Coprocessor 6 control register used to assign a priority level to interrupt sources 47 down to 32. The IPR2 control register is used to assign one of 4 priority levels to each interrupt source independent of the INTSTR[1:0] registers:

00 <sub>2</sub> - High Priority
01 <sub>2</sub> - Med/High Priority
10 <sub>2</sub> - Med/Low Priority
11 <sub>2</sub> - Low Priority

When interrupt vector generation is enabled and there are multiple requests pending either in the FINTSRC[1:0] or the IINTSRC[1:0] registers, the highest priority vectors pending for either FIQ or IRQ is presented in the FINTVEC or IINTVEC, respectively.

*Note:* When multiple interrupts at the same priority level are pending for either FIQ or IRQ, the vector is selected according to a fixed priority based on bit location. Highest order bit is first.



#### Table 460. Interrupt Priority Register 2 - IPR2



## 17.8.12 Interrupt Priority Register 3 - IPR3

The Interrupt Priority Register 3 is a 32-bit Coprocessor 6 control register used to assign a priority level to interrupt sources 63 down to 48. The IPR3 control register is used to assign one of 4 priority levels to each interrupt source independent of the INTSTR[1:0] registers:

00 <sub>2</sub> - High Priority
01 <sub>2</sub> - Med/High Priority
10 <sub>2</sub> - Med/Low Priority
11 <sub>2</sub> - Low Priority

When interrupt vector generation is enabled and there are multiple requests pending either in the FINTSRC[1:0] or the IINTSRC[1:0] registers, the highest priority vectors pending for either FIQ or IRQ is presented in the FINTVEC or IINTVEC, respectively.

*Note:* When multiple interrupts at the same priority level are pending for either FIQ or IRQ, the vector is selected according to a fixed priority based on bit location. Highest order bit is first.

Copro Att M Att	31       28       24       20       16       12       8       4       0         Coprocessor Attributes       Implementation       Implementation <td< th=""></td<>					
Bit	Default	Description				
31:30	002	HPI# Interrupt Priority				
29:28	002	Messaging Unit Error Interrupt Priority	Aessaging Unit Error Interrupt Priority			
27:26	002	Bus Interface Unit Master Abort Error Interrupt Priority				
25:24	002	Application Accelerator Unit Error Interrupt Priority	Application Accelerator Unit Error Interrupt Priority			
23:22	002	Preserved				
21:20	002	DMA Channel 1 Error Interrupt Priority				
19:18	002	DMA Channel 0 Error Interrupt Priority				
17:16	002	Memory Controller Unit Error Interrupt Priority				
15:14	002	ATU Error Interrupt Priority				
13:12	002	ATU Configuration Register Write Interrupt Priority	ATU Configuration Register Write Interrupt Priority			
11:10	002	Reserved.				
9:8	002	UART 1 Interrupt Priority				
7:6	002	UART 0 Interrupt Priority	UART 0 Interrupt Priority			
5:0	00 00002 Preserved					

 Table 461.
 Interrupt Priority Register 3 - IPR3



## 17.8.13 Interrupt Base Register - INTBASE

The Interrupt Base Register indicates the beginning of the Interrupt Service Routine (ISR) memory range that contains the interrupt service routines for up to 64 sources. The starting address must be on a boundary equal to the granularity of the ISR memory range as specified by the INTSIZE register.

For instance, the upper 24 bits are used for a 256 byte range, upper 16 bits for a 64 Kbyte range.

*Note:* INTBASE register is read-only from Memory-Mapped Register Address space.



#### Table 462. Interrupt Base Register - INTBASE



## 17.8.14 Interrupt Size Register - INTSIZE

The Interrupt Size Register indicates the size of the Interrupt Service Routine (ISR) memory range that contains the interrupt service routines for up to 64 sources. The INTSIZE register can allocate from 4 bytes to 64 Kbytes of memory address space for the ISR per source. This means that the INTSIZE register can allocate a total ISR memory space that ranges in size from 256 bytes to 4 Mbytes.

Along with the starting address defined in the INTBASE register, the INTSIZE register fully specifies the ISR memory range.

*Note:* INTSIZE register is read-only from Memory-Mapped Register Address space.



Table 463. Interrupt Size Register - INTSIZE

## 17.8.15 IRQ Interrupt Vector Register - IINTVEC

The IRQ Interrupt Vector Register is a 32-bit Coprocessor 6 control register. Following an IRQ exception, the IRQ interrupt service routine reads the 32-bit vector to the ISR for the active IRQ source with the highest priority.

The actual vector value is a function of the INTBASE and the INTSIZE registers and is based on a fixed order of all 64 possible interrupt sources. The vectors begin at INTBASE with source 0 (i.e., IINTSRC0 bit 0), and end at INTBASE + INTSIZE (per source)\*63 with source 63 (i.e., IINTSRC1 bit 31).

Before returning to User Mode from Interrupt Mode, the software reads the IINTVEC register and process any lower priority IRQ sources that are active. When there are no longer any active IRQ sources, a read from the IINTVEC register returns FFFF FFFFH.



#### Table 464. IRQ Interrupt Vector Register- IINTVEC



## 17.8.16 FIQ Interrupt Vector Register - FINTVEC

The FIQ Interrupt Vector Register is a 32-bit Coprocessor 6 control register. Following an FIQ exception, the FIQ interrupt service routine reads the 32-bit vector to the ISR for the active FIQ source with the highest priority.

The actual vector value is a function of the INTBASE and the INTSIZE registers and is based on a fixed order of all 64 possible interrupt sources. The vectors begin at INTBASE with source 0 (i.e., FINTSRC0 bit 0), and end at INTBASE + INTSIZE (per source)\*63 with source 63 (i.e, FINTSRC1 bit 31).

Before returning to User Mode from Interrupt Mode, the software reads the FINTVEC register and process any lower priority FIQ sources that are active. When there are no longer any active FIQ sources, a read from the FINTVEC register returns FFFF FFFFH.







## 17.8.17 PCI Interrupt Routing Select Register - PIRSR

The PCI Interrupt Routing Select Register (PIRSR) determines the routing of the external input pins. The input pins consist of eight external interrupt inputs which are routed to either one of the IOAPICs or Intel XScale<sup>®</sup> core interrupts. The external interrupt pins are defined as "level sensitive," asserted low. The assertion and deassertion of the interrupt pins are synchronous to the PCI or processor clock.

All external interrupt inputs, **XINT[7:0]**#/, are required to be "level sensitive," asserted low. The interrupt assertion is cleared at the source by the interrupt handler.

Table 466 shows the bit definitions for programming the PCI Interrupt Routing Select Register.

#### Table 466. PCI Interrupt Routing Select Register - PIRSR (Sheet 1 of 2)





Au	IOP tributes PCI rv rv rv rv rv rv	28       24       20       16         16       16       16       16       16         17       17       17       17       17       17       17       17       17       17       17       17       17       17       17       16         17       17       16       16       16       16       16       16       17       16       16       16       16       16       16       16       16       16       17       16       16       16       16       16       16       16       16       17       17       17       17       17       17       17       17       17       17       17       17       17	12     8     4     0       v     rv     rv     rv     rv     rv     rv     rv       v     rv     rv     rv     rv     rv     rv     rv				
PCI Co EC - EF	nfiguration Offs <sup>-</sup> h	et Intel XScale <sup>®</sup> Core Local Bus Address FFFF E1ECH	Attribute Legend:RW = Read/WriteRV = ReservedRC = Read ClearPR = PreservedRO = Read OnlyRS = Read/SetNA= Not Accessible				
Bit	Default	Descri	iption				
		XINT0# Segment Routing Bit					
16	02	0 = XINT0# Input Pin routed to the A Segment IOA 1 = XINT0# Input Pin routed to the B Segment IOA	PIC PIC				
15:8	00H	Reserved (initialize to 0)					
		XINT7# Select Bit					
7	02	0 = Interrupt routed to IOAPIC input 7 1 = Interrupt routed to Intel XScale <sup>®</sup> core interrupt of	controller input ( <b>XINT7#</b> )				
		XINT6# Select Bit					
6	02	0 = Interrupt routed to IOAPIC input 6 1 = Interrupt routed to Intel XScale <sup>®</sup> core interrupt of	controller input (XINT6#)				
_		XINT5# Select Bit					
5	02	<ul> <li>0 = Interrupt routed to IOAPIC input 5</li> <li>1 = Interrupt routed to Intel XScale<sup>®</sup> core interrupt controller input (XINT5#)</li> </ul>					
		XINT4# Select Bit					
4	02	0 = Interrupt routed to IOAPIC input 4 1 = Interrupt routed to Intel XScale <sup>®</sup> core interrupt of	controller input (XINT4#)				
		XINT3# Select Bit					
3	02	0 = Interrupt routed to IOAPIC input 3 1 = Interrupt routed to Intel XScale <sup>®</sup> core interrupt of	controller input (XINT3#)				
		XINT2# Select Bit					
2	02	0 = Interrupt routed to IOAPIC input 2 1 = Interrupt routed to Intel XScale <sup>®</sup> core interrupt of	controller input (XINT2#)				
		XINT1# Select Bit					
1	02	0 = Interrupt routed to IOAPIC input 1 1 = Interrupt routed to Intel XScale <sup>®</sup> core interrupt of	controller input (XINT1#)				
		XINT0# Select Bit					
0	02	0 = Interrupt routed to IOAPIC input 0 1 = Interrupt routed to Intel XScale <sup>®</sup> core interrupt of	controller input (XINT0#)				

## Table 466. PCI Interrupt Routing Select Register - PIRSR (Sheet 2 of 2)



## 17.9 IOAPIC Registers

IOAPIC registers are accessible from downstream PCI Express configuration cycles and are addressable by the host processor through the PCI Express interface. Access to the IOAPIC register space is controlled by the APIC Configuration Space Disable bit in the Bridge Configuration register. This register space is disabled by default.

## 17.9.1 Configuration Space Registers

#### Table 467. IOAPIC Configuration Space Layout

				Offset			Offset
APIC_DID		D APIC_VID		00h		APIC_IDX Index Reg Alias	80h
APIC_STS		API	C_CMD	04h			84h
	APIC_CC		APIC_RID	08h			88h
APIC_BIST	APIC_HT	APIC_MLT	APIC_CLS	0Ch			8Ch
	APIC_	MBAR		10h	APIC_WIND Window Registe	er Alias	90h
				14h			94h
				18h			98h
				1Ch			9Ch
				20h		APIC_PAR Reg Alias	A0h
				24h			A4h
				28h			A8h
APIC_	_SSID	APIC	_SSVID	2Ch			Ach
				30h			B0h
			APIC_EXPC APPTR	34h			B4h
				38h			B8h
				3Ch			BCh
		APIC	C_ABAR	40h		APIC_EOI Reg Alias	C0h
APIC_E	XP_CAP	APIC_EXP NXTP	_ APIC_EXP_ CAPID	44h			C4h
	APIC_EX	P_DCAP		48h			C8h
APIC_EX	(P_DSTS	APIC_E	XP_DCTL	4Ch			CCh
	APIC_EX	(P_LCAP		50h			D0h
APIC_EX	(P_LSTS	APIC_E	EXP_LCTL	54h			D4h
				58h			D8h
				5Ch			DCh
				60h			E0h
				64h			E4h
				68h			E8h
APIC_P	M_CAP	APIC_PM _NXTP	APIC_PM _CAPID	6Ch			ECh
APIC_PM _DATA	APIC_PM _BSE	AP	IC_PM CSR	70h			F0h
							F4h
							F8h
							FCh
				AF	PIC_ADVCAP 100	)	-
					Reserved 104-F	FF	



## Table 468. IOAPIC Configuration Register Summary

Offset	Size (bits)	Symbol	Register Name
00	16	APIC_VID	IOAPIC Vendor ID
02	16	APIC_DID	IOAPIC Device ID
04	16	APIC_CMD	IOAPIC Device Command Register
06	16	APIC_STS	IOAPIC Device Status Register
08	8	APIC_RID	IOAPIC Revision ID
09	24	APIC_CC	IOAPIC Class Code
0C	8	APIC_CLS	IOAPIC Cache Line Size Register
0D	8	APIC_MLT	IOAPIC Master Latency Timer
0E	8	APIC_HDR	IOAPIC Header Type Register
10	32	APIC_MBAR	IOAPIC Memory Base Address Register
2C	16	APIC_SSVID	IOAPIC Subsystem Vendor ID
34	8	APIC_CAPPTR	IOAPIC Capability Pointer
2E	16	APIC_SSID	IOAPIC Subsystem ID
40	16	APIC_ABAR	IOAPIC Alternative Base Address Register
44	8	APIC_EXP_CAPID	IOAPIC PCI Express Capability ID Register
45	8	APIC_EXP_NXTP	IOAPIC PCI Express Next Capability Pointer Register
46	16	APIC_EXP_CAP	IOAPIC PCI Express Base Capability Information Register
48	32	APIC_EXP_DCAP	IOAPIC PCI Express Device Capability Register
4C	16	APIC_EXP_DCTL	IOAPIC PCI Express Device Control Register
4E	16	APIC_EXP_DSTS	IOAPIC PCI Express Device Status Register
50	32	APIC_EXP_LCAP	IOAPIC PCI Express Link Capability Register
54	16	APIC_EXP_LCTL	IOAPIC PCI Express Link Control Register
56	16	APIC_EXP_LSTS	IOAPIC PCI Express Link Status Register
80	8	APIC_IDX	IOAPIC Index Register (Alias of memory space register)
90	32	APIC_WIND	IOAPIC Window Register (Alias of memory space register)
A0	8	APIC_PAR	IOAPIC IRQ Pin Assertion Register (Alias of memory space register)
C0	8	APIC_EOI	IOAPIC EOI Register (Alias of memory space register)
100	32	APIC_EXP_ADVCAP	IOAPIC PCI Express Advanced Capability Register



## 17.9.1.1 IOAPIC Vendor ID - APIC\_VID (Offset 00)

This register contains the Vendor Identifier.

#### Table 469. IOAPIC Vendor ID - APIC\_VID

Bits	Туре	Reset	Description
15:00	RO	8086h	Vendor ID (VID): 16-bit field which indicates that Intel is the vendor.

## 17.9.1.2 IOAPIC Device ID - APIC\_DID (Offset 02)

This register contains the Device Identifier.

#### Table 470. IOAPIC Device ID - APIC\_DID

Bits	Туре	Default		Description
15:00	RO	A-Bus	0371H	Device ID (DID): Indicates what device number was assigned by the Intel. A unique ID
15.00	NO	B-Bus	0373H	exists for the IOAPIC of Bus Segment A (function 1) vs. bus Segment B (function 3).

## 17.9.1.3 IOAPIC Device Command Register - APIC\_CMD (Offset 04)

This register controls how the IOAPIC device behaves.

#### Table 471. IOAPIC Device Command Register - APIC\_CMD

Bits	Туре	Default	Description
15:11	RO	0 0000 <sub>2</sub>	Reserved
10	RO	02	INTx Mask: APIC does not generate interrupts on its own behalf. Reserved as 0.
09	RO	02	Fast Back-to-back enable (FBE): Reserved
08	RW	0 <sub>2</sub>	<b>SERR# Enable (SEE):</b> Controls the enable for the ERR_FATAL or ERR_NONFATAL messages on PCI Express for data parity errors to the APIC configure/memory space. Refer to section for details on how SERR enable bit is used in conjunction with ERR_FATAL or ERR_NONFATAL enable bits to generate the appropriate message.
07	RO	02	Wait Cycle Control (WCC): Reserved
06	RW	02	<b>Parity Error Response Enable (PERE):</b> Enables checking of parity. APIC function uses this bit to report data parity errors it receives on writes.
05	RO	02	VGA Palette Snoop Enable (VGA_PSE): Reserved
04	RO	02	Memory Write and Invalidate Enable (MWIE): Reserved
03	RO	02	Special Cycle Enable (SCE): Reserved
02	RW	02	<b>Bus Master Enable (BME):</b> Controls the IOAPICs ability to act as a master on PCI Express when forwarding front side bus interrupt messages.
01	RW	02	<b>Memory Space Enable (MSE):</b> Controls the IOAPICs response as a target to memory accesses that address the IOAPIC.
00	RO	02	I/O Space Enable (IOSE): Reserved

## 17.9.1.4 IOAPIC Status Register - APIC\_STS (Offset 06)

None of the bits in the IOAPIC status register are read/write.

#### Table 472. IOAPIC Status Register - APIC\_STS

Bits	Туре	Default	Description
15	RWC	02	<b>Detected Parity Error (DPE):</b> Indicates that a parity error was detected on cycles targeting the IOAPIC.
14	RWC	02	<b>Signaled System Error (SSE):</b> This bit is set to when ERR_FATAL or ERR_NONFATAL message is sent on PCI Express for data parity errors to APIC configuration or memory space and the SERR enable bit (bit 8 in APIC_CMD) is set. This bit is also set on error messages generated on PCI Express for errors not specified to a function.
13	RO	02	Received Master Abort (RMA): Reserved
12	RO	02	Received Target Abort (RTA): Reserved.
11	RO	02	Signaled Target Abort (STA): Reserved.
10:09	RO	002	<b>DEVSEL# Timing (DT):</b> Fast decode is performed by the IOAPIC.
08	RO	02	Master Data Parity Error (MDPE): Reserved.
07	RO	02	Fast Back-to-Back Capable (FBC): Reserved as not fast back-to-back capable.
06	RO	02	Reserved
05	RO	02	66 MHz Capable (C66): 66 MHz capable.
04	RO	1 <sub>2</sub>	<b>Capabilities List Enable (CAPE):</b> Indicates that the 80333 contains the capabilities pointer in the IOAPIC. Offset 34h indicates the offset for the first entry in the linked list of capabilities.
03	RO	02	Interrupt Status: Reserved 0 because IOAPIC does not generate interrupt on its own behalf.
02:00	RO	002	Reserved

## 17.9.1.5 IOAPIC Revision ID - APIC\_RID (Offset 08)

Revision ID Register.

#### Table 473. IOAPIC Revision ID - APIC\_RID

Bits	Туре	Default	Description
07:00	RO	00H	Revision ID (RID): Indicates the A-step of the IOAPIC in the 80333



## 17.9.1.6 IOAPIC Class Code - APIC\_CC (Offset 09)

This register contains the base class, sub-class and programming interface codes.

Bits	Туре	Default	Description
23:16	RO	08H	Base Class Code (BCC): The value of '08h' indicates that this is a generic system peripheral.
15:08	RO	00H	Sub Class Code (SCC): The value of '00h' indicates that this generic peripheral is an interrupt controller.
07:00	RO	20H	<b>Programming Interface (PIF):</b> The value of '20h' indicates that this interrupt peripheral is an IOAPIC.

#### Table 474. IOAPIC Class Code - APIC\_CC

## 17.9.1.7 IOAPIC Header Register - APIC\_BIST/HT/MLT/CLS (Offset 0C)

This is additional header information related to the IOAPIC

#### Table 475. IOAPIC Header Register - APIC\_BIST/HT/MLT/CLS

Bits	Туре	Default	Description
31:24	RO	00H	Built In Self Test (BIST): Reserved
23:16	RO	80H	<b>Header Type (HT):</b> This indicates that it is a type "00" header (normal PCI device) and that it is part of a multi-function device.
15:08	RO	00H	Latency Timer (MLT): Reserved
07:00	RO	00H	Cache Line Size (CLS): Reserved

## 17.9.1.8 IOAPIC Memory Base Register - APIC\_MBAR (Offset 10)

Contains the APIC Base Address for the APIC memory space.

#### Table 476. IOAPIC Memory Base Register - APIC\_MBAR

Bits	Туре	Default	Description
31:12	RW	0 0000H	Address (ADDR): These bits determine the base address of the IOAPIC
11:04	RO	00H	Reserved
03	RO	02	Prefetchable (PF): Indicates that the BAR is not prefetchable.
02:01	RO	002	<b>Location (LOC):</b> '00' indicates that the address can be located anywhere in the 32-bit address space.
00	RO	02	Space Indicator (SI): Indicates that the BAR is in memory space.



## 17.9.1.9 IOAPIC Subsystem Vendor ID - APIC\_SSVID (Offset 2C)

This register is initialized to logic 0 by the assertion of **PWRGD**. This register can be written only once after **PWRGD** de-assertion.

#### Table 477. IOAPIC Subsystem Vendor ID - APIC\_SSVID

Bits	Туре	Default	Description
15:00	RWOS	0000H	Subsystem Vendor ID (SSVID): Write once register for holding the subsystem vendor ID.

## 17.9.1.10 IOAPIC Subsystem ID - APIC\_SSID (Offset 2E)

This register is initialized to logic 0 by the assertion of **PWRGD**. This register can be written only once after **PWRGD** de-assertion.

#### Table 478. IOAPIC Subsystem ID - APIC\_SSID

Bits	Туре	Default	Description
15:00	RWOS	0000H	Subsystem ID (SSID): Write once register for sub-system ID.

## 17.9.1.11 IOAPIC Capability Pointer - APIC\_CAPPTR (Offset 34)

#### Table 479. IOAPIC Capability Pointer - APIC\_CAPPTR

Bits	Туре	Default	Description
7:0	RO	44H	CAPPTR: Indicates the presence of a capability list item – the PCI Express capability.

## 17.9.1.12 IOAPIC Alternate Base Address Register \_APIC\_ABAR (Offset 40)

This register contains an alternate base address in the legacy APIC range. This range can co-exist with the Memory BAR register range. This range is needed for OSes that support the APIC but do not yet support remapping the APIC anywhere in the 4 GB address space.

#### Table 480. IOAPIC Alternate Base Address Register - APIC\_ABAR

Bits	Туре	Default	Description
15	RW	02	<b>Enable (EN):</b> When set, the range FECX_YZ00 to FECX_YZFF is enabled as an alternate access method to the IOxAPIC registers. Bits 'XYZ' are defined below.
14:12	RO	0002	Reserved.
11:08	RW	ОН	<b>Base Address [19:16] (XBAD):</b> These bits determine the high order bits of the I/O APIC address map. When a memory address is recognized by the 80333 which matches FECX_YZ00 or FECX_YZ10, the 80333 will respond to the cycle and access the internal I/O APIC.
07:04	RW	ОН	<b>Base Address [15:12] (YBAD):</b> These bits determine the low order bits of the I/O APIC address map. When a memory address is recognized by the 80333 which matches FECX_YZ00 or FECX_YZ10, the 80333 will respond to the cycle and access the internal I/O APIC.
03:00	RW	ОН	<b>Base Address [11:8] (ZBAD):</b> These bits determine the low order bits of the I/O APIC address map. When a memory address is recognized by the 80333 which matches FECX_YZ00 or FECX_YZ10, the 80333 will respond to the cycle and access the internal I/O APIC.



#### 17.9.1.13 IOAPIC PCI Express Capability Identifier -APIC\_EXP\_CAPID (Offset 44)

This register stores the PCI Express capability ID value.

#### Table 481. IOAPIC PCI Express Capability Identifier - APIC\_EXP\_CAPID

Bits	Туре	Default	Description
7:0	RO	10H	PCI Express Capability ID: PCI Express capability identifier.

#### 17.9.1.14 IOAPIC PCI Express Next Item Pointer -APIC\_EXP\_NXTP (Offset 45)

This register stores the byte offset pointer to the next capability list item of the IOAPIC.

#### Table 482. IOAPIC PCI Express Next Item Pointer - APIC\_EXP\_NXTP

Bits	Туре	Default	Description
7:0	RO	6CH	Next Capability Pointer: The offset of the next capabilities list item, which is the Power Management capability.

#### 17.9.1.15 IOAPIC PCI Express Capability -APIC\_EXP\_CAP (Offset 46)

This register carries the version number of the capability item and other base information contained in the capability structure.

#### Table 483. IOAPIC PCI Express Capability - APIC\_EXP\_CAP

Bits	Туре	Default	Description
15:14	RV	002	Reserved
13:9	RO	0 0000 <sub>2</sub>	Interrupt Message Number: Not relevant for IOAPIC
8	RO	02	Slot Implemented: Not relevant for IOAPIC
7:4	RO	0H	Device/Port Number: Indicates PCI Express end-point device
3:0	RO	1H	Version Number: Indicates PCI Express capability structure version number



## 17.9.1.16 IOAPIC PCI Express Device Capabilities Register -APIC\_EXP\_DCAP (Offset 48)

This register carries information on the PCI Express link capabilities

#### Table 484. IOAPIC PCI Express Device Capabilities Register - APIC\_EXP\_DCAP

Bits	Туре	Default	Description
31:28	RsvdP	0	Reserved.
27:26	RO	00	Captured Slot Power Limit Scale: In combination with the Slot Power Limit, Limit value. Specifies the upper limit on power supplied by the slot. Power limit (in Watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Value field. This value is set by the Set Slot Power Limit message.
25:18	RO	000H	Captured Slot Power Limit Value: In combination with the Slot Power Limit, Scale value. Specifies the upper limit on power supplied by the slot. Power limit (in Watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. This value is set by the Set Slot Power Limit message.
17:15	RO	00	Reserved.
14	RO	02	Power Indicator Present: Not Supported
13	RO	02	Attention Indicator Present: Not Supported
12	RO	02	Attention Button Present: Not Supported
11:9	RO	0002	Endpoint L1 Acceptable Latency: 80333 does not support L1 ASPM
8:6	RO	0002	Endpoint L0s Acceptable Latency: 80333 wants the least latency possible out of L0s
5	RO	02	Extended Tag Field Supported: 80333 supports only a 5 bit tag
4:3	RO	002	Phantom Functions Supported: 80333 does not support phantom functions
2:0	RO	001 <sub>2</sub>	Supported Max Payload sizes: 80333 supports max 256 byte packets

March 2005

## intel

## 17.9.1.17 IOAPIC PCI Express Device Control Register -APIC\_EXP\_DCTL (Offset 4C)

This register carries command bits that control 80333 behavior on PCI Express.

#### Table 485. IOAPIC PCI Express Device Control Register - APIC\_EXP\_DCTL

Bits	Туре	Default	Description	
15	RsvdP	0	Reserved.	
14:12	RO	000	Max_Read_Request_Size: Does not apply to APIC.	
11	RO	0	Enable No Snoop: Is not relevant/used by APIC.	
10	RO	0	Auxiliary (AUX) Power PM Enable: 80333 ignores this since it does not support Auxiliary Power.	
9	RO	0	Phantom Function Enable: 80333 ignores this since it does not support Phantom functions.	
8	RO	0	Extended Tag Field Enable: 80333 ignores this since it supports only 5-bit tag.	
7:5	RW	001	Maximum Payload Size: Does not affect APIC since it does not do writes greater than a DWORD.	
4	RO	0	Enabled Relaxed Ordering: Not relevant to APIC.	
3	RW	0	Unsupported Request Reporting Enable: This bit is used by the APIC to enable reporting of ERR_FATAL/NONFATAL message for data parity errors on writes to the APIC (configuration or memory).	
2	RW	0	Report Fatal Errors: Used to gate the generation of the ERR_FATAL message on Fatal line errors. Refer to <i>PCI Express Specification</i> , Revision 1.0, on how this bit is used to report fatal errors in the context of Multi-Function devices like 80333.	
1	RW	0	Report Nonfatal Errors: Used by APIC to gate the generation of the ERR_NONFATAL message on data parity errors/cpmpletor-abort to it.	
0	RW	0	Report Correctable Errors: when set 80333 is enabled to generate ERR_CORR message on PCI-Express. Not used by APIC per se.	



## 17.9.1.18 IOAPIC PCI Express Device Status Register -APIC\_EXP\_DSTS (Offset 4E)

This register carries information on the PCI Express device status.

#### Table 486. IOAPIC PCI Express Device Status Register - APIC\_EXP\_DSTS

Bits	Туре	Default	Description
15:6	RO	000H	Reserved
5	RO	02	Transactions pending: Not relevant for APIC, since it does not generated non-posted requests
4	RO	02	Auxiliary Power Detected: 80333 does not support auxiliary power and hence this bit is reserved for 80333.
3	RWC	02	Unsupported Request Detected: APIC will set this bit whenever it receives a configuration or memory write with bad parity. It is also set on link unsupported request errors that are not specific to any function within 80333.
2	RWC	02	Detected Fatal Error: APIC does not set this bit on its own, but rather set on link fatal errors.
1	RWC	02	Detected Nonfatal Error: APIC sets this bit whenever it detects a write to APIC (configuration or memory) with bad data parity or it signals a completor abort for nonposted transactions greater than a DWORD. This bit is also set on link uncorrectable errors that are not specific to any functions within 80333.
0	RWC	02	Detected Correctable Error: APIC does not set this bit on its own, but rather set on link correctable errors

## 17.9.1.19 IOAPIC PCI Express Link Capabilities Register -APIC\_EXP\_LCAP (Offset 50)

#### Table 487. IOAPIC PCI Express Link Capabilities Register - APIC\_EXP\_LCAP

Bits	Туре	Default	Description
31:24	RO	02	Port Number: Not applicable to 80333
23:18	RO	02	Reserved
17:15	RO	111 <sub>2</sub>	L1 Exit Latency: L1 transition is not supported in 80333
14:12	RO	111 <sub>2</sub>	<ul> <li>L0s Exit Latency: The value in these bits is influenced by bit 6 in link control register. Note that software could write the bit 6 in link control register to either a 1 or 0 and these bits should change accordingly. The mapping is shown below</li> <li>Bit 6 in LCTLL0s Exit Latency</li> <li>0 DEM[5:3] (because currently L0s cannot work with different ref clocks)</li> <li>1 DEM[2:0]</li> </ul>
11:10	RO	01 <sub>2</sub>	ASPM Support: 80333 supports only L0s
9:4	RO	00 1000 <sub>2</sub>	Maximum Link Width: 80333 supports a X8 link maximum
3:0	RO	1H	Maximum Link Speed: 80333 supports 2.5 Gb/s

## intel

## 17.9.1.20 IOAPIC PCI Express Link Control Register -APIC\_EXP\_LCTL (Offset 54)

#### Table 488. IOAPIC PCI Express Link Control Register - APIC\_EXP\_LCTL

Bits	Туре	Default	Description			
15:7	RO	0 0000 0000 <sub>2</sub>	Reserved			
6	RW	0 <sub>2</sub>	<ul> <li>Common Clock Configuration: Indicates the relationship of the reference clock between 80333 and the component at the opposite end of 80333 Upstream PCI Express interface.</li> <li>0 = clock is asynchronous.</li> <li>1 = clock is common.</li> <li>This bit is used to reflect the proper L0s exit latency value in the APIC_EXP_LCAP register.</li> </ul>			
5	RO	02	Retrain Link: Not Applicable to 80333.			
4	RO	02	Disable Link: Not Applicable to 80333.			
3	RO	02	Read Request Return Parameter R Control: Not relevant for APIC.			
2	RO	02	Link Loopback Mode: 80333 enters the loopback mode when this bit is set in all the functions.			
1:0	RW	0 <sub>2</sub>	ASPM Control: Enables 80333 to enter L0s, not used by APIC per se: • 00 L0s entry disabled. • 01 80333 enters L0s per the Specification requirement for L0s entry. • 10 L0s entry disabled. • 11 80333 enters L0s per the Specification requirement for L0s entry.			



## 17.9.1.21 IOAPIC PCI Express Link Status Register -APIC\_EXP\_LSTS (Offset 56)

Bits	Туре	Default	Description			
15:13	RsvdZ	0002	Reserved.			
12	RWOS	1 <sub>2</sub>	Slot Clock Configuration – This bit indicates that when 80333 is on a PCI Express connector that it is using the same reference clock as is provided at the connector. A value of 0 indicates independent reference clock and a value of 1 indicates same reference clock. This bit is initialized by BIOS.			
11	RO	02	Link Training: This read-only bit indicates that Link Training is in progress. 80333 clears this bit once Link Training is complete.			
10	RO	02	Link Training Error: This read-only bit indicates that a link training error occurred. This bit is updated by hardware, with the result of link training, every time the link trains.			
9:4	RO	NA	Negotiated Link Width: This field indicates the negotiated width of PCI-Express Link.Defined encodings are: $00\ 0001_b$ $X1$ $00\ 0010_b$ $X2$ $00\ 0100_b$ $X4$ $00\ 1000_b$ $X8$ $00\ 1100_b$ $X12$ $01\ 0000_b$ $X12$ $01\ 0000_b$ $X12$ $01\ 0000_b$ $X12$ Valid values for 80333 are x1, x4, and x8 All other values are reserved.			
3:0	RO	1H	Link Speed: 80333 supports only 2.5 Gbps.			

#### Table 489. IOAPIC PCI Express Link Status Register - APIC\_EXP\_LSTS

## 17.9.1.22 IOAPIC Power Management Capability Identifier -APIC\_PM\_CAPID (Offset 6C)

#### Table 490. IOAPIC Power Management Capability Identifier - APIC\_PM\_CAPID (Offset 6C)

Bits	Туре	Default	Description
7:0	RO	01H	Capability ID: Capabilities ID indicates PCI compatible Power Management.

#### 17.9.1.23 IOAPIC Power Management Next Item Pointer -APIC\_PM\_NXTP (Offset 6D)

#### Table 491. IOAPIC Power Management Next Item Pointer - APIC\_PM\_NXTP (Offset 6D)

Bits	Туре	Default	Description
7:0	RO	00H	Next Pointer: This is the last capability.

## intel

## 17.9.1.24 IOAPIC Power Management Capabilities -APIC\_PM\_CAP – (Offset 6E)

#### Table 492. IOAPIC Power Management Capabilities - APIC\_PM\_CAP - (Offset 6E)

Bits	Туре	Default	Description
15:11	RO	00H	PME_Support – APIC does not support PME.
10	RO	0	D2 Support: APIC does not support D2 device state.
9	RO	0	D1 Support: APIC does not support D1 device state.
8:6	RO	0	Auxiliary Current: APIC does not support Auxiliary power.
5	RO	0	DSI – APIC does not require device-specific initialization when transitioned to D0 from D3hot state. So this bit is zero.
4	RO	0	Reserved.
3	RO	0	PME Clock: Not relevant to PCI Express and hence hardwired to 0.
2:0	RO	2	Version: APIC PM Implementation is compliant with <i>PCI Bus Power Management Interface Specification</i> , Revision 1.1.

## 17.9.1.25 IOAPIC Power Management Control/Status Register -APIC\_PM\_CSR (Offset 70)

#### Table 493. IOAPIC Power Management Control/Status Register - APIC\_PM\_CSR (Offset 70)

Bits	Туре	Default	Description
15	RO	0	PME Status: APIC never sets this bit.
14:13	RO	0	Data Scale: APIC does not implement Data register and hence these two bits are 0.
12:09	RO	0000	Data Select: Reserved since data register is not implemented.
08	RO	0	PME En: N/A to APIC.
07:02	RO	0	Reserved.
01:00	RO	0	PowerState – This 2 -bit field is used both to determine the current power state of a function and to set the function into a new power state. APIC supported field values are given below. 00b – D0 01b – Reserved 10b – Reserved 11b – D3 hot When software attempts to write an unsupported, optional state to this field, the write operation must complete normally on the bus; however, the data is discarded and no state change occurs. When in D3hot state, APIC responds to configuration transactions only and a transition from D3hot to D0 does not reset the APICs registers. Also, in D3hot state, APIC cannot generate any MSI. Virtual wire interrupts generated by APIC on behalf of PCI agents/SHPC are not masked by the D3hot state.



## 17.9.1.26 IOAPIC Power Management Bridge Support Extensions -APIC\_PM\_BSE (Offset 72)

#### Table 494. IOAPIC Power Management Bridge Support Extensions - APIC\_PM\_BSE (Offset 72)

Bits	Туре	Default	Description
07	RO	0	BPCC_En (Bus Power/Clock Control Enable): N/A to APIC.
06	RO	0	B2/B3#: N/A to APIC.
05:00	RO	0	Reserved.

#### 17.9.1.27 IOAPIC Power Management Data Field -APIC\_PM\_DATA (Offset 73)

#### Table 495. IOAPIC Power Management Data Field - APIC\_PM\_DATA (Offset 73)

Bits	Туре	Default	Description
7:0	RO	0	Data: APIC does not report the data register.



## **17.9.2 Direct Memory Space Registers**

This memory space is mapped by the Section 17.9.1.8, "IOAPIC Memory Base Register - APIC\_MBAR (Offset 10)" or Section 17.9.1.12, "IOAPIC Alternate Base Address Register \_APIC\_ABAR (Offset 40)" and accessed by memory transactions.

*Note:* The memory space offsets 00h, 10h, 20h, and 40h are aliased into configuration space starting at configuration offset 80h for the purpose of being accessed by the SMBus controller.

#### Table 496. IOAPIC Direct Memory Registers



*Note:* Accesses to portions of the IOAPIC BAR that are unimplemented will be aliased to lower addresses rather than returning 0s on reads and dropping writes.

#### 17.9.2.1 Access Disposition

Only DWORD reads and writes are allowed to the APIC memory space. Reads larger than a DWORD receive completor abort response.

#### 17.9.2.2 Register Summary

#### Table 497. IOAPIC Direct Memory Space Register Summary

Start	End	Symbol	Full Name	Default
00	00	IDX	Index Register	00h
10	13	WND	Window Register	00000000h
20	20	PAR	IRQ Pin Assertion Register	00h
40	40	EOI	EOI Register	00h



## 17.9.2.3 IOAPIC Index Register \_APIC\_IDX (Offset 00)

The Index Register will select which indirect register appears in the window register to be manipulated by software. Software will program this register to select the desired APIC indirect memory addressed internal register.

#### Table 498. IOAPIC Index Register - APIC\_IDX

Bit	Туре	Default	Description
07:00	RW	00H	Index (IDX): Indirect register to access.

## 17.9.2.4 IOAPIC Window Register - APIC\_WIND (Offset 10)

This is a 32-bit register specifying the data to be read or written to the register pointed to by the IOAPIC Index Register. This register can be accessed in byte quantities.

#### Table 499. IOAPIC Window Register - APIC\_WIND

Bit	Туре	Default	Description
31:00	RW	0000 0000H	<b>Window (WND):</b> Data to be written to the Index register on writes, and location of register data from the Index register on reads.

## 17.9.2.5 IOAPIC IRQ Pin Assertion Register - APIC\_PAR (Offset 20)

The IRQ Pin Assertion Register is present to provide a mechanism to scale the number of interrupt inputs into the I/O APIC without increasing the number of dedicated input pins. When a device that supports this interrupt assertion protocol requires interrupt service, that device will issue a write to this register. Bits 4:0 written to this register contain the IRQ number for this interrupt. The only valid values are 0-23. Bits 31:5 are ignored.

*Note:* The B-segmant IOAPIC IRQ Pin Assertion Register (APIC\_PAR) is not accessible from the A-segment.

Bit	Туре	Default	Description
31:05	RO	000 0000H	Reserved.
04:00	RW	X XXXX <sub>2</sub>	Assertion (PAR): Virtual pin to be asserted (active high). Writes to this register are treated with edge triggered semantics regardless of what is programmed in the RDL DWORD, though the interrupt message is generated directly from the contents of the RDL and RDH DWORDS.

#### Table 500. IOAPIC IRQ Pin Assertion Register - APIC\_PAR

## 17.9.2.6 IOAPIC EOI Register - APIC\_EOI (Offset 40)

The EOI register is present to provide a mechanism to maintain the level triggered semantics for level-triggered interrupts issued on the parallel bus. When a write is issued to this register, the IOAPIC will check the lower 8 bits written to this register, and compare it with the vector field for each entry in the I/O Redirection Table. When a match is found, the Remote\_IRR bit for that I/O Redirection Entry will be cleared. Note that when multiple I/O Redirection entries, for any reason, assign the same vector for more than one interrupt input, each of those entries will have the Remote\_IRR bit reset to '0'.

#### Table 501. IOAPIC EOI Register - APIC\_EOI

Bit	Туре	Default	Description
07:00	RW	XXH	End of Interrupt (EOI): Vector to be cleared by the EOI.



## 17.9.3 IOAPIC Indirect Memory Space Registers

This memory space is mapped by the Section 17.9.2.3, "IOAPIC Index Register \_APIC\_IDX (Offset 00)", and accessed by memory transactions.

## 17.9.3.1 Register Summary

#### Table 502. IOAPIC Indirect Memory Space Register Summary

Offset	Symbol	Full Name	Default
00	ID	APIC ID	00000000h
01	VS	Version	00170020h
02	ARBID	Arbitration ID	00000000h
03	BCFG	Boot Configuration	00000001h
10	RDL00	IRQ00 Redirection Table Low DWORD	00010000h
11	RDH00	IRQ00 Redirection Table High DWORD	00000000h
12	RDL01	IRQ01 Redirection Table Low DWORD	00010000h
13	RDH01	IRQ01 Redirection Table High DWORD	00000000h
14	RDL02	IRQ02 Redirection Table Low DWORD	00010000h
15	RDH02	IRQ02 Redirection Table High DWORD	00000000h
3E	RDL23	IRQ23 Redirection Table Low DWORD	00010000h
3F	RDH23	IRQ23 Redirection Table High DWORD	00000000h

## 17.9.3.2 IOAPIC ID - APIC\_ID (Offset 00)

The IOAPIC ID serves as a physical name of the IOAPIC.

#### Table 503.IOAPIC ID - APIC\_ID

Bit	Туре	Default	Description
31:28	RO	0H	Reserved.
27:24	RW	0H	<b>IOAPIC ID (APIC_ID):</b> Software must program this value before using the IOAPIC.
23:00	RO	00 0000H	Reserved.

## 17.9.3.3 IOAPIC Version - APIC\_VS (Offset 04)

Contains information related to this IOAPIC for driver / OS software.

Bit	Туре	Default	Description
31:24	RO	00H	Reserved.
23:16	RO	17H	<b>Maximum Redirection Entries (MAX):</b> This is the entry number of the highest entry in the redirection table. It is equal to the number of interrupt inputs minus one. This field is hardwired to 17h to indicate 24 interrupts.
15	RO	1 <sub>2</sub>	<b>IRQ Assertion Register Supported (PRQ):</b> This bit is set to 1 to indicate that this version of the IOAPIC implements the IRQ Assertion register and allows PCI devices to write to it to cause interrupts.
14:08	RO	00H	Reserved.
07:00	RO	20H	Version (VS): This identifies the implementation version. This field is hardwired to 20h indicate this is an IOAPIC.

#### Table 504.IOAPIC Version - APIC\_VS

## 17.9.3.4 IOAPIC Arbitration ID - APIC\_ARBID (Offset 08)

This register contains the APIC serial bus arbitration priority for the APIC, and is loaded whenever the APIC ID register is loaded. 80333 does not support APIC bus serial delivery and hence this register is never used.

#### Table 505. IOAPIC Arbitration ID - APIC\_ARBID

Bit	Туре	Reset	Description
31:28	RO	0H	Reserved.
27:24	RO	0H	Arbitration ID (ARBID): Reflects the APIC ID.
23:00	RO	00 0000H	Reserved.

## 17.9.3.5 IOAPIC Boot Configuration - APIC\_BCFG (Offset 0C)

The Boot Configuration contains information that is only supposed to be accessed by BIOS and is not for OS use. It contains bits that must be programmed before the OS takes control of interrupts.

#### Table 506. IOAPIC Boot Configuration - APIC\_BCFG

Bit	Туре	Reset	Description
31:01	RO	0000 0000H	Reserved.
0	RW	1 <sub>2</sub>	<b>Delivery Type (DT):</b> This bit is a default1 to indicate FSB delivery mode. A value of 0 has no effect. Its left as RW for software compatibility reasons.



## 17.9.3.6 IOAPIC Redirection Table Low DWORD IRQxx -APIC\_RDLxx (Offset 10-3E)

The information in this register is sent on the front side bus to address a local APIC. There is one of these registers for every interrupt. The first interrupt (input 0), has this entry at offset 10h. The second interrupt at 12h, third at 14h, etc. until the final interrupt (input 23) at 3Eh.

*Note:* The valid IRQ inputs to the IOAPIC are 0 through 7 both the A & B-segments, 12 for the A-Segment and 23 for the B-segment. All unused IRQs must have their Mask bit (bit 16) in the corresponding Redirection Table Low DWORD set to one to disable those sources.

#### Bit Туре Default **Description** 31:18 RO 0000H Reserved. Disable Flushing (DFLSH): This bit is maintained for any potential software compatibility, 17 RW $0_{2}$ but the 80333 will perform no flushing action, regardless of the setting of this bit. Mask (MSK): When cleared, an edge or level on this interrupt pin results in the delivery of 16 RW the interrupt to the destination. When set, Interrupts are not delivered nor held pending. 1<sub>2</sub> Setting this bit after the interrupt is accepted by a local APIC has no effect on that interrupt. Trigger Mode (TM): This field indicates the type of signal on the interrupt pin that triggers 15 RW 02 an interrupt. 0 indicates edge sensitive, 1 indicates level sensitive. Remote IRR (RIRR): This bit is used for level triggered interrupts; its meaning is undefined for edge triggered interrupts. For level triggered interrupts, this bit is set when Local APIC/s 14 RO $0_{2}$ accept the level interrupt sent by the IOAPIC. It is reset when an EOI message is received from a local APIC. Interrupt Input Pin Polarity (IP): This bit specifies the polarity of each interrupt signal 13 RW $0_{2}$ connected to the interrupt pins. A value of 0 means the signal is active high. A value of 1 means the signal is active low. Delivery Status (DS): This field contains the current status of the delivery of this interrupt. It is read only. Writes to this bit have no effect. 12 RO $0_{2}$ 0 - Idle - No activity for this interrupt 1 - Pending - Interrupt has been injected, but delivery is held up due to the APIC bus being busy or the inability of the receiving APIC unit to accept the interrupt at this time. Destination Mode (DSTM): This field determines the interpretation of the Destination field. 0 - Physical - Destination APIC ID is identified by RDH bits [59:56]. 11 RW $X_2$ 1 - Logical – Destination is identified by matching bits [63:56] with the Logical Destination in the Destination Format Register and Logical Destination Register in each local APIC. Delivery Mode (DELM): This field specifies how the APICs listed in the destination field should act upon reception of the interrupt. Certain Delivery Modes will only operate as intended when used in conjunction with a specific trigger mode. These encodings are described in more detail in each serial message. The encodings are: 000 - Fixed: Trigger Mode can be edge or level. 001 - Lowest Priority: Trigger Mode can be edge or level. 10:08 RW $XXX_2$ 010 - SMI/PMI: Not supported 011 - Reserved 100 - NMI: Not supported 101 - INIT: Not supported 110 - Reserved 111 - ExtINT: Not supported Vector (VCT): This field contains the interrupt vector for this interrupt. Values range 07:00 RW 00H between 10h and FEh.

#### Table 507. IOAPIC Redirection Table Low DWORD IRQxx - APIC\_RDLxx



## 17.9.3.7 IOAPIC Redirection Table High DWORD IRQxx -APIC\_RDHxx (Offset 11 - 3F)

The information in this register is sent on the APIC serial bus or front side bus to address a local APIC. There is one of these registers for every interrupt. The first interrupt (input 0), has this entry at offset 11h. The second interrupt at 13h, third at 15h, etc. until the final interrupt (input 23) at 3Fh.

#### Table 508. IOAPIC Redirection Table High DWORD IRQxx - APIC\_RDHxx

Bit	Туре	Default	Description
31:24	RW	ХХН	<b>Destination ID (DID):</b> This information is transferred in both serial mode and front side bus mode. In parallel mode, they are bits [19:12] of the address. In serial mode, they are sent at different times, depending on the type of message.
23:16	RW	XXH	<b>Extended Destination ID (EDID):</b> These bits are only sent to a local APIC when in front side bus mode. They become bits [11:4] of the address.
15:00	RO	0000H	Reserved.

# General Purpose I/O Unit

INTel®

This chapter describes the Intel<sup>®</sup> 80333 I/O processor (80333) General Purpose I/O Unit. The operation modes, setup, external memory interface, and implementation of the General Purpose I/Os (GPIOs) are described in this chapter.

## **18.1 General Purpose Input Output Support**

Eight pins are provided as General Purpose Input Output (GPIO) pins. The eight pins are **GPIO**[7:0]. These pins can be used by the Intel XScale<sup>®</sup> core to control or monitor external devices in the I/O subsystem.

The interface for the two serial UART ports are multiplexed on to XINT[15:8]#/GPIO[7:0].

*Warning:* To avoid serial port integrity problems, the user needs to ensure that the GPIO Output Data Register - GPOD bits associated with a UART port is cleared prior to setting the enable bit for that serial UART port. The user prepares to enable the serial UART ports by clearing GPOD bits associated with the respective shared GPIO pins. Refer to Section 18.1.4, "GPIO Pin Multiplexing" for the details on shared GPIO pins.

When a UART port is enabled through the UxIER Register bit 6 ("UART x Interrupt Enable Register" on page 663), the GPIO functionality described in the following sections is not available on the associated pins. UART ports are disabled following system reset.

## 18.1.1 General Purpose Inputs

The current state of the eight GPIO pins can be read in Section 18.2.2, "GPIO Input Data Register - GPID" on page 827).

*Note:* When configured as general purpose inputs, the eight GPIO pins can be used as up to eight additional external interrupt inputs dedicated to the Intel XScale<sup>®</sup> core. This feature is available on a per pin basis simply by programming the INTCTL[1:0] registers.

## **18.1.2 General Purpose Outputs**

The output function of the GPIO pins is controlled by two registers, as stated in Section 18.2.3, "GPIO Output Data Register - GPOD" on page 828) and Section 18.2.1, "GPIO Output Enable Register - GPOE" on page 826).

The output enables are mapped on a per bit basis to each of the data bits in the GPIO Output Data Register. When a bit of the GPIO Output Enable Register is cleared, the corresponding data bit value in the GPIO Output Data Register is actively driven on the appropriate GPIO pin.



## 18.1.3 Reset Initialization of General Purpose Input Output Function

The GPIO Input Data Register is initialized to the state of GPIO[7:0] upon assertion of PWRGD.

The GPIO Output Data Register is initialized to all zeros upon assertion of PWRGD.

The GPIO Output Enable Register is initialized to FFH upon assertion of **PWRGD**. This means that **GPIO**[7:0] initializes as inputs.

The **GPIO**[7:0] pins is tristated during **PWRGD** assertion.

## 18.1.4 GPIO Pin Multiplexing

The GPIO pins are multiplexed with the UART ports as specified in. The selection between GPIO and UART function for these multiplexed pins is controlled by the UART Unit Enable bit (UUE bit 6) of the Section 12.4.3, "UART x Interrupt Enable Register" on page 663.

#### Table 509.GPIO Pin Multiplexing

GPIO Pin	UART 0 Shared Port	UART 0 Port Signal	Default
GPIO[0]	UART0	U0_RXD	GPIO[0] - Input
GPIO[1]	UART0	U0_TXD	GPIO[1] - Input
GPIO[2]	UART0	U0_CTS#	GPIO[2] - Input
GPIO[3]	UART0	U0_RTS#	GPIO[3] - Input
GPIO Pin	UART 1 Shared Port	UART 1 Port Signal	Default
GPIO[4]	UART1	U1_RXD	GPIO[4] - Input
GPIO[5]	UART1	U1_TXD	GPIO[5] - Input
GPIO[6]	UART1	U1_CTS#	GPIO[6] - Input
GPIO[7]	UART1	U1_RTS#	GPIO[7] - Input



## **18.2 Register Definitions**

All GPIO registers are visible as 80333 memory mapped registers and can be accessed through the internal memory bus. Each is a 32-bit register and is memory-mapped in the Intel XScale<sup>®</sup> core memory space.

The programmer interface to the General Purpose I/O interface is through memory-mapped control registers. Table 510 describes these registers.

#### Table 510. General Purpose I/O Registers Addresses

Register Name	Description	MMR Address
GPOE	GPIO Output Enable Register	FFFF F780H
GPID	GPIO Input Data Register	FFFF F784H
GPOD	GPIO Output Data Register	FFFF F788H



## **18.2.1 GPIO Output Enable Register - GPOE**

The GPIO Output Enable Register enables on a per pin basis the output value contained in the GPIO Output Data Register onto the appropriate pin.

The GPIO Output Enable Register is initialized to FFH such that all of GPIO[7:0] are inputs.

In order to enable a particular GPIO pin to operate as an output following the deassertion of **PWRGD**, the user needs to write a '0' into the appropriate GPOE bit.



#### Table 511. GPIO Output Enable Register - GPOE



## 18.2.2 GPIO Input Data Register - GPID

The GPIO Input Data Register reflects the state of the appropriate **IRQ** bus pin following the deassertion of **PWRGD**.



Table 512.
 GPIO Input Data Register - GPID



## 18.2.3 **GPIO Output Data Register - GPOD**

The GPIO Output Data Register is driven on a per bit basis on the appropriate **IRQ** bus pin following the deassertion of **PWRGD** when the corresponding bit in the GPOE register is cleared.



IOP Attributes     31     28     24     20     16     12     8     4     0		
PCI Attributes		
		Intel XScale® Core Local Bus AddressAttribute Legend:RW = Read/WriteFFFF F788HRV = ReservedRC = Read ClearPR = PreservedRO = Read OnlyRS = Read/SetNA = Not Accessible
Bit	Default	Description
31:12	0	Reserved.
11	0	GPIO11 Output Data - Software can write to this bit to toggle the SCL[1] signal. When GPOD[11] is written high, SCL[1] is driven low.
10	00	GPIO10 Output Data - Software can write to this bit to toggle the SCL[0] signal. When GPOD[10] is written high, SCL[0] is driven low.
09:08	0	Preserved.
07	02	GPIO7 Output Data This bit value is driven on the <b>GPIO[7]</b> pin when bit 7 of the GPOE register is cleared.
06	02	GPIO6 Output Data This bit value is driven on the <b>GPIO[6]</b> pin when bit 6 of the GPOE register is cleared.
05	02	GPIO5 Output Data This bit value is driven on the <b>GPIO[5]</b> pin when bit 5 of the GPOE register is cleared.
04	02	GPIO4 Output Data This bit value is driven on the <b>GPIO[4]</b> pin when bit 4 of the GPOE register is cleared.
03	02	GPIO3 Output Data This bit value is driven on the <b>GPIO[3]</b> pin when bit 3 of the GPOE register is cleared.
02	02	GPIO2 Output Data This bit value is driven on the <b>GPIO[2]</b> pin when bit 2 of the GPOE register is cleared.
01	02	GPIO1 Output Data This bit value is driven on the <b>GPIO[1]</b> pin when bit 1 of the GPOE register is cleared.
00	02	GPIO0 Output Data This bit value is driven on the <b>GPIO[0]</b> pin when bit 0 of the GPOE register is cleared.
# Peripheral Registers

## 19

This chapter summarizes the registers for the integrated peripherals (MMR). Each register is defined in detail in the corresponding unit chapter.

## 19.1 Overview

The Peripheral Registers of the Intel<sup>®</sup> 80333 I/O processor (80333) can be accessed via three methods: the PCI Configuration Register interface, Peripheral Memory-Mapped Register interface and the high performance Intel XScale<sup>®</sup> core (ARM\* architecture compliant) Coprocessor Register interface.

The PCI Configuration Register interface is supported by the PCI Express interface and PCI Configuration Cycle transaction type. The 80333 supports access of this interface via upstream configuration cycle transactions as described in the PCI Express-to-PCI Bridge chapter.

The Peripheral Memory-Mapped Register (PMMR) interface gives software the ability to read and modify internal control registers. Each of these registers is accessed as a memory-mapped 32-bit register with a unique memory address. Access is accomplished through regular memory-format instructions from the Intel XScale<sup>®</sup> core.

The Intel XScale<sup>®</sup> core Coprocessor Register (CCR) interface gives software the ability to read and modify internal control registers at a very low latency as compared to the PMMR interface.

These registers are specific to the 80333 only. They support the:

- DMA Controller Unit
- Memory Controller
- Intel XScale<sup>®</sup> Core Bus Interface Unit
- Interrupt Controller Unit and IOAPIC
- Messaging Unit
- Arbitration Unit
- UARTs

- Address Translation Unit
- I<sup>2</sup>C Bus Interface Units
- Application Accelerator Unit
- General Purpose I/O Unit
- Peripheral Bus Interface Unit
- I/O Level Control

Each of these peripherals fully describe the independent functionality of the registers, control and usage.

Control and status registers for the Intel XScale<sup>®</sup> core use the CCR interface. Accesses to coprocessor registers do not generate external bus cycles. See the Intel<sup>®</sup> 80200 Processor based on Intel<sup>®</sup> XScale<sup>TM</sup> Microarchitecture Developer's Manual (Order Number: 273411) for a full description of the usage of the coprocessor register space in the Intel XScale<sup>®</sup> core. For completeness, these registers are included in the tables of registers for the CCR interface. These registers can only be accessed using the Intel XScale<sup>®</sup> core coprocessor instructions.

The PMMR interface provides full accessibility from the ATU, and the Intel XScale<sup>®</sup> core. Addresses FFFF E000H through FFFF FFFFH are allocated to the PMMR interface.



## **19.2** Accessing the PCI Configuration Registers

The 80333 PCI configuration space consists of the two bridges, two IOAPICs and the ATU. Table 514 summarizes the PCI programming model of the I/O processor. The configuration registers of the bridge and IOAPIC are not memory mapped, and are only accessible via PCI congruent space from the PCI Express port. The registers of the ATU are both memory mapped and mapped to PCI configuration space.

#### Table 514. Intel<sup>®</sup> 80333 I/O Processor PCI Programming Model

Bus	Device	Function	Unit
Upstream PCI Express Bus #	0	0	A-Segment PCI Express-to-PCI Bridge
Upstream PCI Express Bus #	0	1	A-Segment IOAPIC
Upstream PCI Express Bus #	0	2	B-Segment PCI Express-to-PCI Bridge
Upstream PCI Express Bus #	0	3	B-Segment IOAPIC
A-Segment PCI Bus #	Eh	0	Address Translation Unit

The registers of the ATU are both memory mapped and mapped to PCI configuration space. The ATU registers are listed in the Peripheral Memory Mapped Register section. The configuration registers of the bridges and IOAPICs are summarized in the PCI Configuration Register section, including address and access type.

For PCI Configuration Read transactions, the PMMR shall return a value of zero for registers declared as "reserved". For PCI Configuration Write transactions, the PMMR shall discard the data. For all other types of access, reading or writing a register declared as "reserved" is undefined.



## **19.3** Accessing Peripheral Memory-Mapped Registers

The PMMR interface is a slave device connected to the 80333 internal bus. This interface accepts data transactions which appear on the internal bus from the ATU and the Intel XScale<sup>®</sup> core.

The PMMR interface allows these devices to perform read, write, or read-modify-write transactions. The specific actions taken when modifying any value in the PMMR space is independently defined within each chapter which describes the functionality of the register.

Note: The PMMR interface does not support multi-DWORD burst accesses from any internal bus master.

All PMMR transactions are allowed from the Intel XScale<sup>®</sup> core operating in either user or supervisor mode. In addition, the PMMR does not provide any access exception to the Intel XScale<sup>®</sup> core.

## 19.4 Accessing Peripheral Registers Using the Core Coprocessor Register Interface

Registers may be accessed/manipulated through the CCR interface with the MCR, MRC, STC, and LDC instructions. The *CRn* field of the instruction denotes the register number to be accessed. The *opcode\_1*, and *opcode\_2* fields of the instruction should be zero. The *CRm* field must be set to 0 for the Interrupt Controller Unit and to 1 for the Programmable Timers. Most systems restricts access to coprocessor registers to privileged processes. To control access to a coprocessor register, use the Coprocessor Access Register as described in the *ARM Architecture Reference Manual* - ARM Limited. Order number: ARM DDI 0100E..

## 19.5 Architecturally Reserved Memory Space

The 80333 provides 4 Gbytes of address space. Portions of this address space is architecturally reserved and users are restricted as to their function. Figure 149 shows the reserved address space.

Addresses 0000 0000H through 0000 001FH are reserved for the exception vectors of the Intel XScale<sup>®</sup> Core.

Addresses FFFF 0000H through FFFF 001FH are reserved for the relocated exception vectors of the Intel XScale<sup>®</sup> Core.<sup>1</sup>

Addresses FFFF E000H through FFFF E8FFH are allocated to the PMMR interface. These registers are reserved for 80333 use and should not be written by the system designer.

Addresses FFFF E900H through FFFF FFFFH are allocated for future expansion of the PMMR interface. These registers are reserved for 80333 use and should not be written by the system designer.

<sup>1.</sup> By enabling the Exception Vector Relocation mode (bit 13, CP15, Register 1), the Exception Vectors (except Reset Vector at 0000 0000H) can be relocated to be based at FFFF 0000H rather than 0000 0000H. (i.e., FIQ Vector located at FFFF 001CH)



Figure 149 shows the Intel XScale<sup>®</sup> Core address space and addresses available to the applications.





For a full description of the address space for the Intel XScale<sup>®</sup> core Reset and Exception Vectors, refer to the *ARM Architecture Reference Manual*.



## **19.6 PCI Configuration Register Address Space**

The PCI Configuration space is accessible via configuration transactions from the PCI Express Primary PCI bus interface. The Intel XScale<sup>®</sup> core of 80333 can access the bridge configuration space as documented in Section 2.3.2, "Addressing" on page 59.

PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	A Bridge Vendor ID Register	16		00H
	A Bridge Device ID Register	16		00H
	A Bridge Command Register	16		01H
	A Bridge Primary Status Register	16		01H
	A Bridge Revision ID Register	8		02H
	A Bridge Class Code Register	24		02H
	A Bridge Cacheline Size Register	8		03H
	A Bridge Primary Master Latency Timer Register	8		03H
	A Bridge Header Type Register	8		03H
	Reserved	8		03H
	Reserved	32	PCI Standard Configuration Header	04H
3ridge	Reserved	32		05H
CIE	A Bridge Primary Bus Number	8		06H
<u>с</u>	A Bridge Secondary Bus Number	8		06H
iss t	A Bridge Subordinate Bus Number	8		06H
kpre	A Bridge Secondary Master Latency Timer Register	8		06H
Ê	A Bridge I/O Base Register	8	Device 0	07H
t PC	A Bridge I/O Limit Register	8	Function 0	07H
nen	A Bridge Secondary Status Register	16		07H
egn	A Bridge Memory Base Register	16		08H
8-A	A Bridge Memory Limit Register	16		08H
	A Bridge Prefetchable Memory Base Register	16		09H
	A Bridge Prefetchable Memory Limit Register	16		09H
	A Bridge Prefetchable Memory Upper Base Register	32		0AH
	A Bridge Prefetchable Memory Upper Limit Register	32		0BH
	A Bridge I/O High Base Register	16		0CH
	A Bridge I/O High Limit Register	16		0CH
	A Bridge Capabilities List Pointer	8		0DH
	Reserved	24		0DH
	Reserved	32	1	0EH
	A Bridge Interrupt Line Register	8	1	0FH
	A Bridge Interrupt Pin Register	8	1	0FH
	A Bridge Control Register	16		0FH

#### Table 515. PCI Configuration Register Locations (Sheet 1 of 12)



PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	A Bridge Configuration Register	16		10H
	A Bridge Multi Transaction Timer	8		10H
	A Bridge PCI Clock Control	8		10H
	A Bridge Express Capabilities Identifier	8		11H
	A Bridge Express Next Item Pointer	8		11H
	A Bridge Express Capability	16		11H
	A Bridge Express Device Capabilities Register	32		12H
	A Bridge Express Device Control Register	16		13H
	A Bridge Express Device Status Register	16		13H
	A Bridge Express Link Capabilities Register	32		14H
	A Bridge Express Link Control Register	16		15H
	A Bridge Express Link Status Register	16		15H
	Reserved	32		16H
	A Bridge MSI Capability Identifier	8	-	17H
ge	A Bridge MSI Next Item Pointer	8		17H
Bric	A Bridge MSI Message Control	16	PCI Extended Configuration Header Device 0 Function 0	17H
<u>c</u>	A Bridge MSI Message Address	64		18H - 19H
to F	A Bridge MSI Message Data	16		1AH
ess	Reserved	16		1AH
xpr	A Bridge Power Management Capability Identifier	8		1BH
G	A Bridge Power Management Next Item Pointer	8		1BH
Jt P	A Bridge Power Management Capabilities	16		1BH
mer	A Bridge Power Management Control/Status Register	16		1CH
Seg	A Bridge Power Management Bridge Support Extensions	8		1CH
A-	A Bridge Power Management Data Field	8		1CH
	Reserved	32		1DH
	A Bridge Standard Hot-Plug Controller Capability Identifier	8		1EH
	A Bridge Standard Hot-Plug Controller Next Item Pointer	8		1EH
	A Bridge Standard Hot-Plug Controller DWORD Select Register	8		1EH
	A Bridge Standard Hot-Plug Controller Status	8		1EH
	A Bridge Standard Hot-Plug Controller DWORD	32		1FH
	Reserved			20H through 35H
	A Bridge PCI-X Capabilities Identifier	8		36H
	A Bridge PCI-X Next Item Pointer	8		36H
	A Bridge PCI-X Secondary Status Register	16		36H
	A Bridge PCI-X Bridge Status Register	32		37H
	A Bridge PCI-S Upstream Split Transaction Control	32	]	38H
	A Bridge PCI-X Downstream Split Transaction Control	32	1	39H

## Table 515. PCI Configuration Register Locations (Sheet 2 of 12)

PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	A Bridge ECC Control and Status Register	32		3AH
	A Bridge ECC Error First Address Register	32		3BH
	A Bridge ECC Error Second Address Register	32		3CH
	A Bridge ECC Error Attribute Register	32		3DH
	Reserved	32		3EH
	A Bridge Bridge Initialization Register	32		3FH
	A Bridge Express Advanced Error Capability Identifier	32		40H
	A Bridge Express Uncorrectable Error Status	32		41H
	A Bridge Express Uncorrectable Error Mask	32		42H
ge	A Bridge Express Uncorrectable Error Severity	32		43H
Brid	A Bridge Express Uncorrectable Error Status	32		44H
<u>0</u>	A Bridge Express Correctable Error Mask	32		45H
to F	A Bridge Express Advanced Error Control and Capability Register	32	PCI Express Extended PCI Configuration Space Device 0 Function 0	46H
press	A Bridge Express Transaction Header Log	128		47H through 4AH
<u> </u>	A Bridge Uncorrectable PCI/X Error Status	16		4BH
ЪС	Reserved	16		4BH
lent	A Bridge Uncorrectable PCI/X Error Mask	16		4CH
egm	Reserved	16		4CH
S-A	A Bridge Uncorrectable PCI/X Error Severity	16		4DH
	Reserved	16		4DH
	A Bridge Uncorrectable PCI/X Error Pointer	16		4EH
	Reserved	16		4EH
	A Bride Uncorrectable PCI/X Error Transaction Header Log	128		4FH through 52H
	A Bridge Uncorrectable/Correctable PCI/X Data Error Log	64		53H through 54H
	A Bridge Other PCI/X Error Logs and Control	32		55H
	Reserved	x		56H through 5AH
	A Bridge PCI Compensation Control Register	32		5BH
	A Bridge Strap Status Register	16	Extended PCI	5CH
A-Segment PCI Express	Reserved		Configuration	
to PCI Bridge	A Bridge Additional Configuration Register	32	Space	5DH
	A Bridge Prefetch Control Registers	64	Function 0	5EH through 5FH

## Table 515. PCI Configuration Register Locations (Sheet 3 of 12)



PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	Reserved	x		60H through BFH
	A Bridge Express Power Budgeting Enhanced Capability Header	32		C0H
	A Bridge Express Power Budgeting Data Select Register	8		C1H
	Reserved	24		C1H
	A Bridge Express Power Budgeting Data Register	32		C2H
	A Bridge Express Power Budgeting Capability Register	8		СЗН
	Reserved	24		СЗН
	Reserved	32		C4H
	A Bridge Express Power Budgeting Information Register 0	32		C5H
	A Bridge Express Power Budgeting Information Register 1	32		C6H
	A Bridge Express Power Budgeting Information Register 2	32		C7H
ge	A Bridge Express Power Budgeting Information Register 3	32	PCI Express Extended PCI Configuration Space	C8H
Brid	A Bridge Express Power Budgeting Information Register 4	32		C9H
CI E	A Bridge Express Power Budgeting Information Register 5	32		CAH
р 10	A Bridge Express Power Budgeting Information Register 6	32		СВН
SSS	A Bridge Express Power Budgeting Information Register 7	32		ССН
xpre	A Bridge Express Power Budgeting Information Register 8	32		CDH
ш П	A Bridge Express Power Budgeting Information Register 9	32	Device 0	CEH
it P(	A Bridge Express Power Budgeting Information Register 10	32	Function 0	CFH
nen	A Bridge Express Power Budgeting Information Register 11	32		D0H
Segi	A Bridge Express Power Budgeting Information Register 12	32		D1H
A-9	A Bridge Express Power Budgeting Information Register 13	32		D2H
	A Bridge Express Power Budgeting Information Register 14	32		D3H
	A Bridge Express Power Budgeting Information Register 15	32		D4H
	A Bridge Express Power Budgeting Information Register 16	32		D5H
	A Bridge Express Power Budgeting Information Register 17	32		D6H
	A Bridge Express Power Budgeting Information Register 18	32		D7H
	A Bridge Express Power Budgeting Information Register 19	32		D8H
	A Bridge Express Power Budgeting Information Register 20	32		D9H
	A Bridge Express Power Budgeting Information Register 21	32		DAH
	A Bridge Express Power Budgeting Information Register 22	32		DBH
	A Bridge Express Power Budgeting Information Register 23	32		DCH
	Reserved	x		DDH through FFH

## Table 515. PCI Configuration Register Locations (Sheet 4 of 12)

PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	A IOAPIC Vendor ID	16		00H
	A IOAPIC Device ID	16		00H
	A IOAPIC Command Register	16		01H
	A IOAPIC Status Register	16		10H
	A IOAPIC Revision ID	8	1	02H
	A IOAPIC Class Code Register	24		02H
<u>0</u>	A IOAPIC Cache Line Size Register	8	-	03H
AP	A IOAPIC Master Latency Timer Register	8	PCI Standard	03H
t O	A IOAPIC Header Type Register	8	Configuration	03H
Jen	Reserved	8	Header	03H
ube	A IOAPIC Memory Base Address Register	32	Function 1	04H
A-Se	Reserved			05H through 0AH
	A IOAPIC Subsystem Vendor ID	16	-	0BH
	A IOAPIC Subsystem ID	16		0BH
	Reserved	32		0CH
	A IOAPIC Capabilities Pointer	8		0DH
	Reserved			0DH through 0FH
	A IOAPIC Alternate Base Address Register	16	-	10H
	A IOAPIC Express Capabilities Identifier	8		11H
	A IOAPIC Express Next Item Pointer	8		11H
	A IOAPIC Express Capability	16		11H
	A IOAPIC Express Device Capabilities Register	32		12H
	A IOAPIC Express Device Control Register	16		13H
~	A IOAPIC Express Device Status Register	16		13H
PIC	A IOAPIC Express Link Capabilities Register	32		14H
OA OA	A IOAPIC Express Link Control Register	16	PCI Extended	15H
int l	A IOAPIC Express Link Status Register	32	Header	15H
egme	Reserved	x	Device 0 Function 1	16H through 1AH
A-S	A IOAPIC Power Management Capability Identifier	8		1BH
	A IOAPIC Power Management Next Item Pointer	8	1	1BH
	A IOAPIC Power Management Capabilities	16	1	1BH
	A IOAPIC Power Management Control/Status Register	16	1	1CH
	A IOAPIC Power Management Bridge Support Extensions	8	1	1CH
	A IOAPIC Power Management Data Field	8	1	1CH
	Reserved			1DH through 1FH

## Table 515. PCI Configuration Register Locations (Sheet 5 of 12)



PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	A IOAPIC Index Register Alias	8	PCI Express Extended PCI Configuration Space Device 0 Function 1	20H
t IOAPIC	Reserved			20H through 23H
	A IOAPIC Window Register Alias	32		24H
	Reserved			24H through 27H
ner	A IOAPIC Pin Assertion Register Alias	8		28H
A-Segr	Reserved			28H through 2FH
	A IOAPIC End Of Interrupt Register Alias	8		30H
	Reserved			30H through 3FFH

## Table 515. PCI Configuration Register Locations (Sheet 6 of 12)

PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	B Bridge Vendor ID Register	16		00H
	B Bridge Device ID Register	16		00H
	B Bridge Command Register	16		01H
	B Bridge Primary Status Register	16		01H
	B Bridge Revision ID Register	8		02H
	B Bridge Class Code Register	24		02H
	B Bridge Cacheline Size Register	8		03H
	B Bridge Primary Master Latency Timer Register	8		03H
	B Bridge Header Type Register	8		03H
	Reserved	8		03H
	B Bridge Standard Hot-Plug Controller Base Address Register 0	32		04H
dge	B Bridge Standard Hot-Plug Controller Upper Base Address Register 0	32	PCI Standard Configuration Header	05H
B	B Bridge Primary Bus Number	8		06H
PCI	B Bridge Secondary Bus Number	8		06H
to to	B Bridge Subordinate Bus Number	8		06H
Less	B Bridge Secondary Master Latency Timer Register	8		06H
Exp	B Bridge I/O Base Register	8		07H
D D	B Bridge I/O Limit Register	8	Function 2	07H
ent F	B Bridge Secondary Status Register	16		07H
gme	B Bridge Memory Base Register	16		08H
, Ve	B Bridge Memory Limit Register	16		08H
<u>ن</u>	B Bridge Prefetchable Memory Base Register	16		09H
	B Bridge Prefetchable Memory Limit Register	16		09H
	B Bridge Prefetchable Memory Upper Base Register	32		0AH
	B Bridge Prefetchable Memory Upper Limit Register	32		0BH
	B Bridge I/O High Base Register	16		0CH
	B Bridge I/O High Limit Register	16		0CH
	B Bridge Capabilities List Pointer	8		0DH
	Reserved	24		0DH
	Reserved	32	1	0EH
	B Bridge Interrupt Line Register	8	1	0FH
	B Bridge Interrupt Pin Register	8	]	0FH
	B Bridge Control Register	16		0FH

## Table 515. PCI Configuration Register Locations (Sheet 7 of 12)



PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	B Bridge Configuration Register	16		10H
	B Bridge Multi Transaction Timer	8		10H
	B Bridge PCI Clock Control	8	-	10H
	B Bridge Express Capabilities Identifier	8		11H
	B Bridge Express Next Item Pointer	8		11H
	B Bridge Express Capability	16		11H
	B Bridge Express Device Capabilities Register	32		12H
	B Bridge Express Device Control Register	16		13H
	B Bridge Express Device Status Register	16		13H
	B Bridge Express Link Capabilities Register	32		14H
	B Bridge Express Link Control Register	16		15H
	B Bridge Express Link Status Register	16		15H
	Reserved	32		16H
	B Bridge MSI Capability Identifier	8	PCI Extended Configuration Header Device 0 Eurotion 2	17H
ge	B Bridge MSI Next Item Pointer	8		17H
Brid	B Bridge MSI Message Control	16		17H
- C	B Bridge MSI Message Address	64		18H - 19H
to F	B Bridge MSI Message Data	16		1AH
ess	Reserved	16		1AH
xpr	B Bridge Power Management Capability Identifier	8		1BH
U U U	B Bridge Power Management Next Item Pointer	8		1BH
Ē	B Bridge Power Management Capabilities	16	T dilotion 2	1BH
mer	B Bridge Power Management Control/Status Register	16		1CH
Seg	B Bridge Power Management Bridge Support Extensions	8		1CH
ů.	B Bridge Power Management Data Field	8		1CH
	Reserved	32		1DH
	B Bridge Standard Hot-Plug Controller Capability Identifier	8		1EH
	B Bridge Standard Hot-Plug Controller Next Item Pointer	8		1EH
	B Bridge Standard Hot-Plug Controller DWORD Select Register	8		1EH
	B Bridge Standard Hot-Plug Controller Status	8		1EH
	B Bridge Standard Hot-Plug Controller DWORD	32		1FH
	Reserved			20H through 35H
	B Bridge PCI-X Capabilities Identifier	8	]	36H
	B Bridge PCI-X Next Item Pointer	8	1	36H
	B Bridge PCI-X Secondary Status Register	16	]	36H
	B Bridge PCI-X Bridge Status Register	32		37H
	B Bridge PCI-S Upstream Split Transaction Control	32	1	38H
	B Bridge PCI-X Downstream Split Transaction Control	32		39H

## Table 515. PCI Configuration Register Locations (Sheet 8 of 12)

PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	B Bridge ECC Control and Status Register	32		3AH
	B Bridge ECC Error First Address Register	32		3BH
	B Bridge ECC Error Second Address Register	32		3CH
	B Bridge ECC Error Attribute Register	32	-	3DH
	Reserved	32		3EH
	B Bridge Bridge Initialization Register	32		3FH
	B Bridge Express Advanced Error Capability Identifier	32		40H
	B Bridge Express Uncorrectable Error Status	32		41H
	B Bridge Express Uncorrectable Error Mask	32		42H
ge	B Bridge Express Uncorrectable Error Severity	32		43H
Brid	B Bridge Express Uncorrectable Error Status	32		44H
D.	B Bridge Express Correctable Error Mask	32	PCI Express Extended PCI Configuration Space Device 0 Function 2	45H
to H	B Bridge Express Advanced Error Control and Capability Register	32		46H
kpress	B Bridge Express Transaction Header Log	128		47H through 4AH
Ê	B Bridge Uncorrectable PCI/X Error Status	16		4BH
РС	Reserved	16		4BH
lent	B Bridge Uncorrectable PCI/X Error Mask	16		4CH
egn	Reserved	16		4CH
о Ч	B Bridge Uncorrectable PCI/X Error Severity	16		4DH
	Reserved	16		4DH
	B Bridge Uncorrectable PCI/X Error Pointer	16		4EH
	Reserved	16		4EH
	B Bride Uncorrectable PCI/X Error Transaction Header Log	128		4FH through 52H
	B Bridge Uncorrectable/Correctable PCI/X Data Error Log	64		53H through 54H
	B Bridge Other PCI/X Error Logs and Control	32		55H
	Reserved	x		56H through 5AH
	B Bridge Strap Status Register	16	PCI Express	5CH
B-Segment	Reserved		Extended PCI	
PCI Express	B Bridge Additional Configuration Register	32	Space	5DH
to PCI Bridge	B Bridge Prefetch Control Registers	64	Device 0 Function 2	5EH through 5FH

## Table 515. PCI Configuration Register Locations (Sheet 9 of 12)



PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	Reserved	x		60H through BFH
	B Bridge Express Power Budgeting Enhanced Capability Header	32		C0H
	B Bridge Express Power Budgeting Data Select Register	8	-	C1H
	Reserved	24		C1H
	B Bridge Express Power Budgeting Data Register	32		C2H
	B Bridge Express Power Budgeting Capability Register	8		СЗН
	Reserved	24		СЗН
	Reserved	32		C4H
	B Bridge Express Power Budgeting Information Register 0	32		C5H
	B Bridge Express Power Budgeting Information Register 1	32		C6H
	B Bridge Express Power Budgeting Information Register 2	32	PCI Express Extended PCI Configuration Space	C7H
e	B Bridge Express Power Budgeting Information Register 3	32		C8H
Bride	B Bridge Express Power Budgeting Information Register 4	32		C9H
CI	B Bridge Express Power Budgeting Information Register 5	32		CAH
р С	B Bridge Express Power Budgeting Information Register 6	32		СВН
SSS	B Bridge Express Power Budgeting Information Register 7	32		ССН
xpre	B Bridge Express Power Budgeting Information Register 8	32		CDH
E E E E E E E E E E E E E E E E E E E	B Bridge Express Power Budgeting Information Register 9	32	Device 0	CEH
t PC	B Bridge Express Power Budgeting Information Register 10	32	Function 2	CFH
nen	B Bridge Express Power Budgeting Information Register 11	32		D0H
, egr	B Bridge Express Power Budgeting Information Register 12	32		D1H
5) (1)	B Bridge Express Power Budgeting Information Register 13	32		D2H
	B Bridge Express Power Budgeting Information Register 14	32		D3H
	B Bridge Express Power Budgeting Information Register 15	32		D4H
	B Bridge Express Power Budgeting Information Register 16	32		D5H
	B Bridge Express Power Budgeting Information Register 17	32		D6H
	B Bridge Express Power Budgeting Information Register 18	32		D7H
	B Bridge Express Power Budgeting Information Register 19	32		D8H
	B Bridge Express Power Budgeting Information Register 20	32		D9H
	B Bridge Express Power Budgeting Information Register 21	32		DAH
	B Bridge Express Power Budgeting Information Register 22	32		DBH
	B Bridge Express Power Budgeting Information Register 23	32	1	DCH
	Reserved	x		DDH through 3FFH

## Table 515. PCI Configuration Register Locations (Sheet 10 of 12)

PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	B IOAPIC Vendor ID	16		00H
	B IOAPIC Device ID	16		00H
	B IOAPIC Command Register	16		01H
	B IOAPIC Status Register	16		10H
	B IOAPIC Revision ID	8	1	02H
	B IOAPIC Class Code Register	24		02H
<u>ں</u>	B IOAPIC Cache Line Size Register	8		03H
AP	B IOAPIC Master Latency Timer Register	8	PCI Standard	03H
t O	B IOAPIC Header Type Register	8	Configuration	03H
Jen	Reserved	8	Header	03H
ube	B IOAPIC Memory Base Address Register	32	Function 3	04H
ъ. В	Reserved			05H through 0AH
	B IOAPIC Subsystem Vendor ID	16		0BH
	B IOAPIC Subsystem ID	16		0BH
	Reserved	32		0CH
	B IOAPIC Capabilities Pointer	8		0DH
	Reserved			0DH through 0FH
	B IOAPIC Alternate Base Address Register	16	-	10H
	B IOAPIC Express Capabilities Identifier	8		11H
	B IOAPIC Express Next Item Pointer	8		11H
	B IOAPIC Express Capability	16		11H
	B IOAPIC Express Device Capabilities Register	32		12H
	B IOAPIC Express Device Control Register	16		13H
~	B IOAPIC Express Device Status Register	16		13H
PIC	B IOAPIC Express Link Capabilities Register	32		14H
OA	B IOAPIC Express Link Control Register	16	PCI Extended	15H
int	B IOAPIC Express Link Status Register	32	Header	15H
egme	Reserved	x	Device 0 Function 1	16H through 1AH
က်	B IOAPIC Power Management Capability Identifier	8		1BH
	B IOAPIC Power Management Next Item Pointer	8	1	1BH
	B IOAPIC Power Management Capabilities	16	1	1BH
	B IOAPIC Power Management Control/Status Register	16	1	1CH
	B IOAPIC Power Management Bridge Support Extensions	8	1	1CH
	B IOAPIC Power Management Data Field	8	1	1CH
	Reserved			1DH through 1FH

### Table 515. PCI Configuration Register Locations (Sheet 11 of 12)



PCI Device	Register Description (Name)	Register Size in Bits	Register Access	PCI Configuration Space Register Number
	B IOAPIC Index Register Alias	8		20H
	Reserved		PCI Express Extended PCI Configuration Space Device 0 Function 3	20H through 23H
PIC	B IOAPIC Window Register Alias	32		24H
It IOA	Reserved			24H through 27H
ner	B IOAPIC Pin Assertion Register Alias	8		28H
B-Segr	Reserved			28H through 2FH
	B IOAPIC End Of Interrupt Register Alias	8	]	30H
	Reserved			30H through 3FFH

## Table 515. PCI Configuration Register Locations (Sheet 12 of 12)



## **19.7 Memory Mapped Register Address Space**

80333 has memory mapped register space for control of both the Standard Hot-Plug Controller of the PCI Express-to-PCI Bridge, and the IOAPIC function. Access to these registers is through memory read and write transactions. The addresses listed are the offsets from the Base Address Register of the respective unit.

PCI Device	Register Description (Name)	Register Size in Bits	Register Access	Register Memory Byte Offset
U	A IOAPIC ID	32	at et	00H
API	A IOAPIC Version	32	opet	04H
<u>Ô</u>	A IOAPIC Arbitration ID	32	Мар * + 0	08H
ent	A IOAPIC Boot Configuration	32	PCI Memory <sup>I</sup> IOAPIC BAR	0CH
шb	A IOAPIC Redirection Table Low	32		10H
A-Se	A IOAPIC Redirection Table High	32		14H
U U	B IOAPIC ID	32	ry Mapped at AR + Offset	00H
PIC	B IOAPIC Version	32		04H
10	B IOAPIC Arbitration ID	32		08H
ent	B IOAPIC Boot Configuration	32		0CH
gme	B IOAPIC Redirection Table Low	32	E E	10H
B-Se	B IOAPIC Redirection Table High	32	PCI Me IOAPI	14H

#### Table 516. PCI Configuration Register Locations (Sheet 1 of 2)



PCI Device	Register Description (Name)	Register Size in Bits	Register Access	Register Memory Byte Offset
	B Bridge SHPC Base Offset Register	32		00H
	B Bridge SHPC Slots Available Register 1	32		04H
	B Bridge SHPC Slots Available Register 2	32		08H
	B Bridge SHPC Slot Configure Register	32		0CH
	B Bridge SHPC Configuration Register	16		10H
	B Bridge SHPC MSI Control Register	8		12H
	B Bridge SHPC Programming Interface Register	8		13H
	B Bridge SHPC Command Register	16		14H
	B Bridge SHPC Command Status Register	16		16H
	B Bridge SHPC Interrupt Locator Register	32	S	18H
e	B Bridge SHPC System Interrupt Locator Register	32	dres	1CH
er	B Bridge SHPC Interrupt Control Register	32	PCI Memory Mapped at	20H
CI B troll	B Bridge SHPC Slot Status Register 1	16		24H
o Pr	B Bridge SHPC Slot Event Latch Register 1	8		26H
iss t ig C	B Bridge SHPC Slot Interrupt Mask Register 1	8		27H
plu.	B Bridge SHPC Slot Status Register 2	16		28H
pt E	B Bridge SHPC Slot Event Latch Register 2	8		2AH
L PC	B Bridge SHPC Slot Interrupt Mask Register 2	8		2BH
nent	B Bridge SHPC Slot Status Register 3	16		2CH
egn	B Bridge SHPC Slot Event Latch Register 3	8		2EH
S S S	B Bridge SHPC Slot Interrupt Mask Register 3	8		2FH
	B Bridge SHPC Slot Status Register 4	16	HS III	30H
	B Bridge SHPC Slot Event Latch Register 4	8		32H
	B Bridge SHPC Slot Interrupt Mask Register 4	8		33H
	B Bridge SHPC Slot Status Register 5	16		34H
	B Bridge SHPC Slot Event Latch Register 5	8		36H
	B Bridge SHPC Slot Interrupt Mask Register 5	8		37H
	Reserved			38H through C3H
	B Bridge SHPC Extended Control Register	32		C4
	Reserved			C7H through FFFH

## Table 516. PCI Configuration Register Locations (Sheet 2 of 2)



## **19.8 Peripheral Memory-Mapped Register Address Space**

The PMMR address space is divided to support the integrated peripherals on the 80333. Table 519 shows all of the 80333 integrated peripheral memory-mapped registers and their internal bus addresses.

#### Table 517. Intel XScale<sup>®</sup> Core Local Addresses Assigned to Integrated Peripherals

Integrated Peripheral	Internal Address Block
Address Translation Unit	FFFF E100H through FFFF E1FFH
Reserved	FFFF E200H through FFFF E2FFH
Messaging Unit	FFFF E300H through FFFF E3FFH
DMA Controller	FFFF E400H through FFFF E4FFH
Memory Controller	FFFF E500H through FFFF E5FFH
Intel XScale <sup>®</sup> Core Bus Interface Unit	FFFF E600H through FFFF E67FH
Peripheral Bus Interface Unit	FFFF E680H through FFFF E6FFH
Peripheral Performance Monitoring Unit	FFFF E700H through FFFF E77FH
Interrupt Controller Unit	FFFF E780H through FFFF E7EFH
Internal Arbitration Unit	FFFF E7F0H through FFFF E7FFH
Application Accelerator Unit	FFFF E800H through FFFF E8FFH
Reserved	FFFF E900H through FFFF F4FFH
DDR I/O Control	FFFF F500H through FFFF F5FFH
Reserved	FFFF F600H through FFFF F67FH
I <sup>2</sup> C Bus Interface Units	FFFF F680H through FFFF F6FFH
UART Units	FFFF F700H through FFFF F77FH
GPIO Unit	FFFF F780H through FFFF F7BFH
Reserved	FFFF F7C0H through FFFF F7FFH
I/O Pad Control	FFFF F800H through FFFF F8FFH
Reserved	FFFF F900H through FFFF FFFFH

The memory-mapped registers that are also accessible via PCI configuration transactions are:

• Address Translation Unit

The registers which must have the address translation logic configured to translate PCI addresses into the Intel XScale<sup>®</sup> Core address space, to access the memory-mapped registers from the PCI Express interface are:

- DMA Controllers
- Memory Controller
- I<sup>2</sup>C Bus Interface Unit
- Synchronous Serial Port UART Unit
- Messaging Unit
- Application Accelerator Unit
- Internal Arbitration Unit
- Peripheral Bus Interface Unit
- Performance Monitoring
- Interrupt Controller and General Purpose I/O Unit
- Interface Pad Control Registers



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Reserved	x	FFFF E000H through FFFF E0FFH	
	ATU Vendor ID Register	16	FFFF E100H	00H
	ATU Device ID Register	16	FFFF E102H	00H
	ATU Command Register	16	FFFF E104H	01H
	ATU Status Register	16	FFFF E106H	01H
	ATU Revision ID Register	8	FFFF E108H	02H
	ATU Class Code Register	24	FFFF E109H	02H
	ATU Cacheline Size Register	8	FFFF E10CH	03H
	ATU Latency Timer Register	8	FFFF E10DH	03H
	ATU Header Type Register	8	FFFF E10EH	03H
	BIST Register	8	FFFF E10FH	03H
nit	Inbound ATU Base Address Register 0	32	FFFF E110H	04H
U nu	Inbound ATU Upper Base Address Register 0	32	FFFF E114H	05H
latic	Inbound ATU Base Address Register 1	32	FFFF E118H	06H
ans	Inbound ATU Upper Base Address Register 1	32	FFFF E11CH	07H
s Tr	Inbound ATU Base Address Register 2	32	FFFF E120H	08H
dres	Inbound ATU Upper Base Address Register 2	32	FFFF E124H	09H
Adc	Reserved	32	FFFF E128H	0AH
	ATU Subsystem Vendor ID Register	16	FFFF E12CH	0BH
	ATU Subsystem ID Register	16	FFFF E12EH	0BH
	Expansion ROM Base Address Register	32	FFFF E130H	0CH
	ATU Capabilities Pointer Register	8	FFFF E134H	0DH
	Reserved	24	FFFF E135H	0DH
	Reserved	32	FFFF E138H	0EH
	ATU Interrupt Line Register	8	FFFF E13CH	0FH
	ATU Interrupt Pin Register	8	FFFF E13DH	0FH
	ATU Minimum Grant Register	8	FFFF E13EH	0FH
	ATU Maximum Latency Register	8	FFFF E13FH	0FH

### Table 518. Peripheral Memory-Mapped Register Locations (Sheet 1 of 14)

Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Inbound ATU Limit Register 0	32	FFFF E140H	10H
isters	Inbound ATU Translate Value Register 0	32	FFFF E144H	11H
	Expansion ROM Limit Register	32	FFFF E148H	12H
	Expansion ROM Translate Value Register	32	FFFF E14CH	13H
	Inbound ATU Limit Register 1	32	FFFF E150H	14H
	Inbound ATU Limit Register 2	32	FFFF E154H	15H
	Inbound ATU Translate Value Register 2	32	FFFF E158H	16H
	Outbound I/O Window Translate Value Register	32	FFFF E15CH	17H
	Outbound Memory Window Value Register 0	32	FFFF E160H	18H
	Outbound Upper 32-bit Memory Window Value Register 0	32	FFFF E164H	19H
	Outbound Memory Window Value Register 1	32	FFFF E168H	1AH
ers	Outbound Upper 32-bit Memory Window Value Register 1	32	FFFF E16CH	1BH
giste	Reserved	32	FFFF E170H	1CH
Re	Reserved	32	FFFF E174H	1DH
tion	Outbound Upper 32-bit Direct Window Value Register	32	FFFF E178H	1EH
ıfigura	PCI Express to PCI Bridge Secondary A-Segment Bus Number Register	8	FFFF E17CH	1FH
ed Cor	PCI Express to PCI Bridge Secondary B-Segment Bus Number Register	8	FFFF E17DH	1FH
pude	PCI Express to PCI Bridge Primary Bus Number Register	8	FFFF E17EH	1FH
Exte	PCI Express to PCI Bridge Device Number Register	8	FFFF E17FH	1FH
Jnit	ATU Configuration Register	32	FFFF E180H	20H
on L	PCI Configuration and Status Register	32	FFFF E184H	21H
latic	ATU Interrupt Status Register	32	FFFF E188H	22H
rans	ATU Interrupt Mask Register	32	FFFF E18CH	23H
s Tr	Inbound ATU Base Address Register 3	32	FFFF E190H	24H
dres	Inbound ATU Upper Base Address Register 3	32	FFFF E194H	25H
Ado	Inbound ATU Limit Register 3	32	FFFF E198H	26H
	Inbound ATU Translate Value Register 3	32	FFFF E19CH	27H
	Reserved	32	FFFF E1A0H	28H
	Outbound Configuration Cycle Address Register	32	FFFF E1A4H	29H
	Reserved	32	FFFF E1A8H	2AH
	Outbound Configuration Cycle Data Register	32	FFFF E1ACH	2BH
	Reserved	32	FFFF E1B0H	2CH
	Reserved	32	FFFF E1B4H	2DH
	VPD Capabilities Identifier Register	8	FFFF E1B8H	2EH
	VPD Next Item Pointer Register	8	FFFF F1B9H	2EH
	VPD Address Register	16	FFFF F1BAH	2EH
	VPD Data Register	32	FFFF F1BCH	2FH

Table 518.	Peripheral Memory-Mapped	Register	Locations	(Sheet 2 of	14)
------------	--------------------------	----------	-----------	-------------	-----



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Power Management Capability Identifier Register	8	FFFF E1C0H	30H
	Power Management Next Item Pointer Register	8	FFFF E1C1H	30H
	Power Management Capabilities Register	16	FFFF E1C2H	30H
	Power Management Control/Status Register	16	FFFF E1C4H	31H
	Reserved	16	FFFF E1C6H	31H
	Reserved	32	FFFF E1C8H	32H
	Reserved	32	FFFF E1CCH	33H
	MSI Capability Identifier Register	8	FFFF E1D0H	34H
	MSI Next Item Pointer Register	8	FFFF E1D1H	34H
Address Translation	MSI Message Control Register	16	FFFF E1D2H	34H
Unit	MSI Message Address Register	32	FFFF E1D4H	35H
Extended	MSI Message Upper Address Register	32	FFFF E1D8H	36H
ration	MSI Message Data Register	16	FFFF E1DCH	37H
Registers	Reserved	16	FFFF E1DEH	37H
	PCI-X Capability Identifier Register	8	FFFF E1E0H	38H
	PCI-X Next Item Pointer Register	8	FFFF E1E1H	38H
	PCI-X Command Register	16	FFFF E1E2H	38H
	PCI-X Status Register	32	FFFF E1E4H	39H
	Reserved	32	FFFF E1E8H	3AH
	PCI Interrupt Routing Select Register	32	FFFF E1ECH	3BH
	Reserved		FFFF E1F0H through FFFF E1FFH	
	Reserved	x	FFFF E200H through FFFF E2FFH	

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 3 of 14)

Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Reserved	x	FFFF E300H through FFFF E30CH	the
	Inbound Message Register 0	32	FFFF E310H	'inde s to ss
	Inbound Message Register 1	32	FFFF E314H	gh n W ore ddre
	Outbound Message Register 0	32	FFFF E318H	atio Cc d Ac
	Outbound Message Register 1	32	FFFF E31CH	e thi PCI ale
	Inbound Doorbell Register	32	FFFF E320H	lable ate XSc Mag
	Inbound Interrupt Status Register	32	FFFF E324H	vvail unc ansla ansla tel ) ory-
	Inbound Interrupt Mask Register	32	FFFF E328H	it tra lem
	Outbound Doorbell Register	32	FFFF E32CH	I U I M
	Outbound Interrupt Status Register	32	FFFF E330H	Or I
	Outbound Interrupt Mask Register	32	FFFF E334H	
g Unit	Reserved	x	FFFF E338H through FFFF E34FH	
agin	MU Configuration Register	32	FFFF E350H	
SS SS	Queue Base Address Register	32	FFFF E354H	
Ň	Reserved	32	FFFF E358H	the
	Reserved	32	FFFF E35CH	s to ess
	Inbound Free Head Pointer Register	32	FFFF E360H	ddr
	Inbound Free Tail Pointer Register	32	FFFF E364H	adc bd A C C
	Inbound Post Head Pointer Register	32	FFFF E368H	PCI
	Inbound Post Tail Pointer Register	32	FFFF E36CH	-Ma -Ma
	Outbound Free Head Pointer Register	32	FFFF E370H	ansk ntel nory
	Outbound Free Tail Pointer Register	32	FFFF E374H	t Tra I
	Outbound Post Head Pointer Register	32	FFFF E378H	1 I
	Outbound Post Tail Pointer Register	32	FFFF E37CH	~
	Index Address Register	32	FFFF E380H	
	Reserved	x	FFFF E384H through FFFF E3FFH	

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 4 of 14)



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Channel 0 Channel Control Register	32	FFFF E400H	
	Channel 0 Channel Status Register	32	FFFF E404H	
	Reserved	32	FFFF E408H	
	Channel 0 Descriptor Address Register	32	FFFF E40CH	
	Channel 0 Next Descriptor Address Register	32	FFFF E410H	
	Channel 0 PCI Address Register	32	FFFF E414H	
	Channel 0 PCI Upper Address Register	32	FFFF E418H	
	Channel 0 Internal Bus Address Register	32	FFFF E41CH	0
	Channel 0 Byte Count Register	32	FFFF E420H	the the
	Channel 0 Descriptor Control Register	32	FFFF E424H	ranslate PCI address to Intel XScale <sup>®</sup> Core emory-Mapped Address
introller	Reserved	x	FFFF E428H through FFFF E43FH	
O C	Channel 1 Channel Control Register	32	FFFF E440H	
/WC	Channel 1 Channel Status Register	32	FFFF E444H	
	Reserved	32	FFFF E448H	
	Channel 1 Descriptor Address Register	32	FFFF E44CH	Me Me
	Channel 1 Next Descriptor Address Register	32	FFFF E450H	Wr
	Channel 1 PCI Address Register	32	FFFF E454H	
	Channel 1 PCI Upper Address Register	32	FFFF E458H	
	Channel 1 Internal Bus Address Register	32	FFFF E45CH	
	Channel 1 Byte Count Register	32	FFFF E460H	
	Channel 1 Descriptor Control Register	32	FFFF E464H	
	Reserved	x	FFFF E468H through FFFF E4FFH	

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 5 of 14)

Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	SDRAM Initialization Register	32	FFFF E500H	
	SDRAM Control Register 0	32	FFFF E504H	
	SDRAM Control Register 1	32	FFFF E508H	
	SDRAM Base Register	32	FFFF E50CH	
	SDRAM Bank 0 Size Register	32	FFFF E510H	
	SDRAM Bank 1 Size Register	32	FFFF E514H	
	SDRAM 32-bit Region Size Register	32	FFFF E518H	
	ECC Control Register	32	FFFF E51CH	
	ECC Log 0 Register	32	FFFF E520H	
	ECC Log 1 Register	32	FFFF E524H	ω
	ECC Address 0 Register	32	FFFF E528H	s oth
<u> </u>	ECC Address 1 Register	32	FFFF E52CH	e Ires
olle	ECC Test Register	32	FFFF E530H	Add
ontr	Memory Controller Interrupt Status Register	32	FFFF E534H	ed ac
C A	Reserved	32	FFFF E538H	anslate PC Intel XSca mory-mapp
IOU	MCU Port Transaction Count Register	32	FFFF E53CH	
Me	MCU Preemption Control Register	32	FFFF E540H	
	Refresh Frequency Register	32	FFFF E548H	st Ti Me
	Reserved	32	FFFF E550H	Mu
	Reserved	32	FFFF E554H	
	Reserved	32	FFFF E558H	
	Reserved	32	FFFF E55CH	
	Reserved	32	FFFF E560H	
	Reserved	32	FFFF E564H	
	Reserved	32	FFFF E568H	
	Reserved	32	FFFF E56CH	
	Reserved	x	FFFF E58CH through FFFF E5FFH	
it e	BIU Status Register	32	FFFF E600H	
S	BIU Error Address Register	32	FFFF E604H	
ale <sup>®</sup> face	BIU Control Register	32	FFFF E608H	
Intel XScal Bus Interfa	Reserved	x	FFFF E60CH through FFFF E67FH	

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 6 of 14)



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	PBI Control Register	32	FFFF E680H	
	Reserved	32	FFFF E684H	
	PBI Base Address Register 0	32	FFFF E688H	
	PBI Limit Register 0	32	FFFF E68CH	Φ
Init	PBI Base Address Register 1	32	FFFF E690H	s oth
⊖ ∩	PBI Limit Register 1	32	FFFF E694H	e Ires
s Interfac	Reserved	x	FFFF E698H through FFFF E6BCH	PCI addre cale <sup>®</sup> Cor pped Add
IBu	PBI Memory-less Boot Register 0	32	FFFF E6C0H	XSc -ma
Periphera	Reserved	x	FFFF E6C4H through FFFF E6DFH	st Transla Intel Memory
	PBI Memory-less Boot Register 1	32	FFFF E6E0H	W
	PBI Memory-less Boot Register 2	32	FFFF E6E4H	
	Reserved	x	FFFF E6E8H through FFFF E6FFH	
	Global Timer Mode Register	32	FFFF E700H	
	Event Select Register	32	FFFF E704H	
	Event Monitoring Interrupt Status Register	32	FFFF E708H	
	Reserved	32	FFFF E70CH	
	Global Time Stamp Register	32	FFFF E710H	
nit	Programmable Event Counter Register 1	32	FFFF E714H	
D B	Programmable Event Counter Register 2	32	FFFF E718H	the
orin	Programmable Event Counter Register 3	32	FFFF E71CH	s to
onit	Programmable Event Counter Register 4	32	FFFF E720H	ore
e	Programmable Event Counter Register 5	32	FFFF E724H	add d A C
anc	Programmable Event Counter Register 6	32	FFFF E728H	2CI 2ale
orm	Programmable Event Counter Register 7	32	FFFF E72CH	xSc Ma
Perf	Programmable Event Counter Register 8	32	FFFF E730H	nsla ntel iory
ral F	Programmable Event Counter Register 9	32	FFFF E734H	Tra Ir Aem
phe	Programmable Event Counter Register 10	32	FFFF E738H	lust N
Perip	Programmable Event Counter Register 11	32	FFFF E73CH	2
	Programmable Event Counter Register 12	32	FFFF E740H	
	Programmable Event Counter Register 13	32	FFFF E744H	
	Programmable Event Counter Register 14	32	FFFF E748H	
	Reserved		FFFF E74CH through FFFF E77FH	

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 7 of 14)

Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Reserved	32	FFFF E780H	
	Reserved		FFFF E784H	
	Reserved		FFFF E788H	
	Reserved		FFFF E78CH	
	Interrupt Control Register 0	32	FFFF E790H	
	Interrupt Control Register 1	32	FFFF E794H	ē
	Interrupt Steer Register 0	32	FFFF E798H	ss to th
Ľ.	Interrupt Steer Register 1	32	FFFF E79CH	e e dres
olle	IRQ Interrupt Source Register 0	32	FFFF E7A0H	Addre
ontr	IRQ Interrupt Source Register 1	32	FFFF E7A4H	oed ®a Ded
pt C d Ti	FIQ Interrupt Source Register 0	32	FFFF E7A8H	e PC Sca
an	IRQ Interrupt Source Register 1	32	FFFF E7ACH	el X ry-n
Int	Interrupt Priority Register 0	32	FFFF E7B0H	Interio
	Interrupt Priority Register 1	32	FFFF E7B4H	St T Me
	Interrupt Priority Register 2	32	FFFF E7B8H	
	Interrupt Priority Register 3	32	FFFF E7BCH	
	Interrupt Base Register	32	FFFF E7C0H	
	Interrupt Size Register	32	FFFF E7C4H	
	IRQ Interrupt Vector Register	32	FFFF E7C8H	
	FIQ Interrupt Vector Register	32	FFFF E7CCH	
	Timer Mode Register 0	32	FFFF E7D0H	e
	Timer Mode Register 1	32	FFFF E7D4H	s to the
	Timer Count Register 0	32	FFFF E7D8H	ddress t Core Addres
÷	Timer Count Register 1	32	FFFF E7DCH	
- D	Timer Reload Register 0	32	FFFF E7E0H	ped ®a
mer	Timer Reload Register 1	32	FFFF E7E4H	e PC Sca
Ē	Timer Interrupt Status Register	32	FFFF E7E8H	slate slate ry-r
	Watch Dog Timer Control Register	32	FFFF E7ECH	Must Trans Inte Memoi
	Internal Arbitration Control Register	32	FFFF E7F0H	Ð
	Multi-Transaction Timer Register 1	32	FFFF E7F4H	s o th
Internal Arbitration Unit	Multi-Transaction Timer Register 2	32	FFFF E7F8H	e Ires:
	Reserved	x	FFFF E7FCH through FFFF E7FFH	Aust Translate PCI addres Intel XScale <sup>®</sup> Core Memory-mapped Add

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 8 of 14)



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Accelerator Control Register	32	FFFF E800H	
	Accelerator Status Register	32	FFFF E804H	
	Accelerator Descriptor Address Register	32	FFFF E808H	
	Accelerator Next Descriptor Address Register	32	FFFF E80CH	
	Accelerator Source Address 1 Register	32	FFFF E810H	
	Accelerator Source Address 2 Register	32	FFFF E814H	
	Accelerator Source Address 3 Register	32	FFFF E818H	
	Accelerator Source Address 4 Register	32	FFFF E81CH	
	Destination Address Register	32	FFFF E820H	
	Accelerator Byte Count Register	32	FFFF E824H	
	Accelerator Descriptor Control Register	32	FFFF E828H	
	Accelerator Source Address 5 Register	32	FFFF E82CH	Inslate PCI address to the net XScale <sup>®</sup> Core nory-mapped Address
tt.	Accelerator Source Address 6 Register	32	FFFF E830H	
Uni	Accelerator Source Address 7 Register	32	FFFF E834H	
ator	Accelerator Source Address 8 Register	32	FFFF E838H	
elera	Extended Descriptor Control Register 0	32	FFFF E83CH	
Acce	Accelerator Source Address 9 Register	32	FFFF E840H	
on 4	Accelerator Source Address 10 Register	32	FFFF E844H	
cati	Accelerator Source Address 11 Register	32	FFFF E848H	
ilqq	Accelerator Source Address 12 Register	32	FFFF E84CH	t Tra Nen
A	Accelerator Source Address 13 Register	32	FFFF E850H	Just I
	Accelerator Source Address 14 Register	32	FFFF E854H	2
	Accelerator Source Address 15 Register	32	FFFF E858H	
	Accelerator Source Address 16 Register	32	FFFF E85CH	
	Extended Descriptor Control Register 1	32	FFFF E860H	
	Accelerator Source Address 17 Register	32	FFFF E864H	
	Accelerator Source Address 18 Register	32	FFFF E868H	
	Accelerator Source Address 19 Register	32	FFFF E86CH	
	Accelerator Source Address 20 Register	32	FFFF E870H	
	Accelerator Source Address 21 Register	32	FFFF E874H	
	Accelerator Source Address 22 Register	32	FFFF E878H	1
	Accelerator Source Address 23 Register	32	FFFF E87CH	1
	Accelerator Source Address 24 Register	32	FFFF E880H	1

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 9 of 14)

Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	Extended Descriptor Control Register 2	32	FFFF E884H	
	Accelerator Source Address 25 Register	32	FFFF E888H	the
	Accelerator Source Address 26 Register	32	FFFF E88CH	s to
	Accelerator Source Address 27 Register	32	FFFF E890H	rese
Application	Accelerator Source Address 28 Register	32	FFFF E894H	Must Translate PCI add Intel XScale <sup>®</sup> Cc Memory-mapped Ac
Accelerator	Accelerator Source Address 29 Register	32	FFFF E898H	
Unit	Accelerator Source Address 30 Register	32	FFFF E89CH	
	Accelerator Source Address 31 Register	32	FFFF E8A0H	
	Accelerator Source Address 32 Register	32	FFFF E8A4H	
	Reserved	x	FFFF E8A8H through FFFF E8FFH	
	Reserved	x	FFFF E900H through FFFF EFFFH	
	Reserved	x	FFFF F000H through FFFF F4FFH	

### Table 518. Peripheral Memory-Mapped Register Locations (Sheet 10 of 14)



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	DCAL Control and Status Register	32	FFFF F500H	
	DCAL Address Register	32	FFFF F504H	
	DCAL Data Register 0	32	FFFF F508H	
	DCAL Data Register 1	32	FFFF F50CH	
	DCAL Data Register 2	32	FFFF F510H	
	DCAL Data Register 3	32	FFFF F514H	
	DCAL Data Register 4	32	FFFF F518H	
	DCAL Data Register 5	32	FFFF F51CH	
	DCAL Data Register 6	32	FFFF F520H	
	DCAL Data Register 7	32	FFFF F524H	
	DCAL Data Register 8	32	FFFF F528H	
	DCAL Data Register 9	32	FFFF F52CH	ē
	DCAL Data Register 10	32	FFFF F530H	ss to th
_	DCAL Data Register 11	32	FFFF F534H	st Translate PCI address t Intel XScale <sup>®</sup> Core Memory-mapped Addres
ntro	DCAL Data Register 12	32	FFFF F538H	
Ŝ	DCAL Data Register 13	32	FFFF F53CH	
<u>9</u>	DCAL Data Register 14	32	FFFF F540H	
DR	DCAL Data Register 15	32	FFFF F544H	
	DCAL Data Register 16	32	FFFF F548H	
	DCAL Data Register 17	32	FFFF F54CH	
	Receive Enable Delay Register	32	FFFF F550H	ML
	Slave Low Mix 0	32	FFFF F554H	
	Slave Low Mix 1	32	FFFF F558H	
	Slave High Mix 0	32	FFFF F55CH	
	Slave High Mix 1	32	FFFF F560H	
	Slave Length	32	FFFF F564H	
	Master Mix	32	FFFF F568H	
	Master Length	32	FFFF F56CH	
	DDR Drive Strength Status Register	32	FFFF F570H	
	DDR Drive Strength Control Register	32	FFFF F574H	
	DDR Miscellaneous Pad Control Register	32	FFFF F578H	-
	Reserved	32	FFFF F57CH	
	PBI Drive Strength Control Register	32	FFFF F580H	
	Reserved	x	FFFF F584H through FFFF F5BFH	

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 11 of 14)



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	PCI-A Drive Strength Control Register - PADSCR	32	FFFF F5C0H	
	Reserved	32	FFFF F5C4H	
itrol	PCI-A Drive Strength Value Register - PADSVR	32	FFFF F5C8H	
Cor	Reserved	32	FFFF F5CCH	
<u>9</u>	PCI-B Drive Strength Control Register - PBDSCR	32	FFFF.F5D0H	
- C	Reserved	32	FFFF.F5D4H	
-	PCI-B Drive Strength Value Register - PBDSVR	32	FFFF.F5D8H	
	Reserved	32	FFFF.F5DCH	
	Reserved	x	FFFF F5E0H through FFFF F67FH	
	I <sup>2</sup> C Control Register 0	32	FFFF F680H	
	I <sup>2</sup> C Status Register 0	32	FFFF F684H	
	I <sup>2</sup> C Data Buffer Register 0	32	FFFF F68CH	
	Reserved	32	FFFF F690H	the
	I <sup>2</sup> C Bus Monitor Register 0	32	FFFF F694H	s to
	Reserved	32	FFFF F698H	ore
I <sup>2</sup> C Ruc	Reserved	32	FFFF F69CH	add a A A
Interface	I <sup>2</sup> C Control Register 1	32	FFFF F6A0H	Ppe
Units	I <sup>2</sup> C Status Register 1	32	FFFF F6A4H	XSc -ma
	I <sup>2</sup> C Slave Address Register 1	32	FFFF F6A8H	nsla itel
	I <sup>2</sup> C Data Buffer Register 1	32	FFFF F6ACH	Tra len L
	Reserved	32	FFFF F6B0H	lust N
	I <sup>2</sup> C Bus Monitor Register 1	32	FFFF F6B4H	2
	Reserved	x	FFFF F6B8H through FFFF F6FFH	

### Table 518. Peripheral Memory-Mapped Register Locations (Sheet 12 of 14)



Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	UART 0 Receive Buffer Register (Read Only) (DLAB=0)	32	FFFF F700H	
	UART 0 Transmit Holding Register (Write Only) (DLAB=0)	32		
	UART 0 baud Divisor Latch Low byte (DLAB=1)	8		
	UART 0 Interrupt Enable Register (DLAB=0)	8	FFFF F704H	
	UART 0 baud Divisor Latch High byte (DLAB=1)	8		Φ
	UART 0 Interrupt ID Register (Read Only)	8	FFFF F708H	anslate PCI address to th Intel XScale <sup>®</sup> Core mory-mapped Address
	UART 0 FIFO Control Register (Write Only)	8		
el 0	UART 0 Line Control Register	8	FFFF F70CH	
ann	UART 0 Modem Control Register	8	FFFF F710H	
c	UART 0 Line Status Register	8	FFFF F714H	
ART	UART 0 Modem Status Register	8	FFFF F718H	
) )	UART 0 Scratch Pad Register	8	FFFF F71CH	
	Reserved	32	FFFF F720H	st Ti Me
	UART 0 FIFO Occupancy Register	8	FFFF F724H	m W
	UART 0 Autobaud Control Register	8	FFFF F728H	
	UART 0 Autobaud Count Register	16	FFFF F72CH	
	Reserved		FFFF F730H through FFFF F73FH	

## Table 518. Peripheral Memory-Mapped Register Locations (Sheet 13 of 14)

Peripheral	Register Description (Name)	Register Size in Bits	Internal Bus Address	PCI Configuration Space Register Number
	UART 1 Receive Buffer Register (Read Only) (DLAB=0)	32		
	UART 1 Transmit Holding Register (Write Only) (DLAB=0)	32	FFFF F740H	
	UART 1 baud Divisor Latch Low byte (DLAB=1)	8		
	UART 1 Interrupt Enable Register (DLAB=0)	8	FFFF F744H	
	UART 1 baud Divisor Latch High byte (DLAB=1)	8		
	UART 1 Interrupt ID Register (Read Only)	8		
	UART 1 FIFO Control Register (Write Only)	8	гггг г/40п	
el 1	UART 1 Line Control Register	8	FFFF F74CH	
ann	UART 1 Modem Control Register	8	FFFF F750H	the
ch	UART 1 Line Status Register	8	FFFF F754H	anslate PCI address to ntel XScale <sup>®</sup> Core nory-mapped Address
ART	UART 1 Modem Status Register	8	FFFF F758H	
n n	UART 1 Scratch Pad Register	8	FFFF F75CH	
	Reserved	32	FFFF F760H	
	UART 1 FIFO Occupancy Register	8	FFFF F764H	
	UART 1 Autobaud Control Register	8	FFFF F768H	
	UART 1 Autobaud Count Register	16	FFFF F76CH	t Tra
	Reserved	x	FFFF F770H through FFFF F77FH	Must
	GPIO Output Enable Register	32	FFFF F780H	
	GPIO Input Data Register	32	FFFF F784H	
GPIO	GPIO Output Data Register	32	FFFF F788H	
	SMBus Enable Register	32	FFFF F78CH	
	Reserved	x	FFFF F790H through FFFF F7FFH	
	Reserved	x	FFFF F800H through FFFF F9FFH	

### Table 518. Peripheral Memory-Mapped Register Locations (Sheet 14 of 14)

## 19.9 Coprocessor Register Space

The CCR address space is assigned to support the integrated peripherals on the 80333 that require low latency register access. Table 519 shows all of the 80333 integrated coprocessor registers and assigned coprocessor space. The *ARM Architecture Reference Manual* - ARM Limited. Order number: ARM DDI 0100E. provides for a total of 16 coprocessors each of which can contain up to 256 32 bit registers. For completeness, the coprocessor space reserved by the *ARM Architecture Reference Manual* - ARM Limited. Order number: ARM DDI 0100E. is shown.

All accesses to coprocessor registers that are not implemented have 'unpredictable' behavior and are designated as 'undefined' in the register table. Accesses to these registers results in an undefined instruction exception.

#### Table 519. Intel XScale<sup>®</sup> Core Coprocessor Registers Assigned to Integrated Peripherals

Integrated Peripheral	Coprocessor
Interrupt Control Unit	CP6
Programmable Timers	CP6
Core Performance Monitoring Unit	CP14
System Control <sup>a</sup>	CP15

a. Reserved by the ARM Architecture Reference Manual - ARM Limited. Order number: ARM DDI 0100E..

Peripheral	Register Description (Name)	Coprocessor	Field CR <sub>m</sub>	Coprocessor Register (Field <i>CR<sub>n</sub></i> )
	Interrupt Control Register 0			Register 0
	Interrupt Control Register 1			Register 1
	Interrupt Steering Register 0			Register 2
	Interrupt Steering Register 1			Register 3
	IRQ Interrupt Source Register 0			Register 4
Jnit	IRQ Interrupt Source Register 1			Register 5
	FIQ Interrupt Source Register 0			Register 6
ontr	FIQ Interrupt Source Register 1		0	Register 7
ot C	Interrupt Priority Register 0		0	Register 8
Inne	Interrupt Priority Register 1			Register 9
Inte	Interrupt Priority Register 2			Register 10
	Interrupt Priority Register 3			Register 11
	Interrupt Base Register			Register 12
	Interrupt Size Register	CP6		Register 13
	IRQ Interrupt Vector Register			Register 14
	FIQ Interrupt Vector Register			Register 15
	Timer Mode Register 0			Register 0
ij	Timer Mode Register 1			Register 1
, Ur	Timer Count Register 0			Register 2
ners	Timer Count Register 1			Register 3
μ	Timer Reload Register 0			Register 4
able	Timer Reload Register 1		1	Register 5
Ĕ	Timer Interrupt Status Register			Register 6
grai	Watch Dog Timer Control Register			Register 7
Pro	Undefined			Register 8 through Register 15

### Table 520. Coprocessor Register Locations (Sheet 1 of 2)



Peripheral	Register Description (Name)	Coprocessor	Field CR <sub>m</sub>	Coprocessor Register (Field <i>CR<sub>n</sub></i> )
	Performance Monitor Control Register			Register 0
	Clock Count Register			Register 1
	Performance Count Register 0			Register 2
	Performance Count Register 1			Register 3
ng Unit	Reserved			Register 4 and Register 5
litor	Core Clock Configuration Register (CCLKCFG)			Register 6
Mon	Reserved			Register 7
Ce	Access Transmit Debug Register (TX)	CP14	CP14 Function	Register 8
nan	Access Receive Debug Register (RX)			Register 9
Perforr	Access Debug Control and Status Register (DBGCSR)			Register 10
Dre	Access Trace Buffer Register (TBREG)			Register 11
ŭ	Access Checkpoint 0 Register (CHKPT0)			Register 12
	Access Checkpoint 1 Register (CHKPT1)			Register 13
	Access Transmit and Receive Debug Control Register			Register 14
	Debug Saved Program Status Register			Register 15
	ID and Cache Type Registers			Register 0
	Control and Auxiliary Control Registers			Register 1
	Translation Table Base Register			Register 2
	Domain Access Control Register			Register 3
	Undefined			Register 4
	Fault Status Register			Register 5
sic	Fault Address Register			Register 6
ontro	Cache Operations Register			Register 7
Ö	TLB Operations Register	CP15	CP 15 Function <sup>a</sup>	Register 8
sterr	Cache Lock Down			Register 9
Sys	TLB Lock Down			Register 10
	Reserved			Register 11 and Register 12
	Process ID Register	]		Register 13
	Breakpoint Registers	1		Register 14
	CP Access (PID) Coprocessor Access Control Register			Register 15

#### Table 520. Coprocessor Register Locations (Sheet 2 of 2)

a. Some of the CP15 registers are differentiated by the Opcode\_2 field. For CP15, the CR<sub>m</sub> is used in some cases to denote different control functions for a given coprocessor register rather than a distinct register decode. Please refer to the Intel<sup>®</sup> 80200 Processor based on Intel<sup>®</sup> XScale<sup>™</sup> Microarchitecture Developer's Manual (Order Number: 273411), for more details on the operation of CP15.
# int<sub>el®</sub> Clocking and Reset

20

This chapter describes the clocking and reset function. The intent of this chapter is to elaborate and clarify descriptions of the clocking and reset mechanisms.

#### 20.1 **Clocking Overview**

The Intel<sup>®</sup> 80333 I/O processor (80333) contains various clocking boundaries internally. The clocks for all of the units within the 80333 are generated from a single input clock. Clock regions are buffered or PLL versions of this clock input for the processor design. This input feeds the Phase Lock Loop (PLL) circuitry which generates the internal clocks for all clock regions. Figure 150 shows the 80333 block diagram and highlights the seven clocking regions for this processor.





Within each of the clocking regions identified exist various clock requirements for the 80333 units and for the output clocks pins provided for the external subsystem.

March 2005



# 20.1.1 Clocking Theory of Operation

Each region within the 80333 contains different clocking requirements. These requirements are summarized in the following sections.

# 20.1.2 Clocking Region 1

Region 1 is the Intel XScale<sup>®</sup> core (ARM\* architecture compliant). It supports clock frequencies up to a maximum of 800 MHz operation. Region 1 also incudes the core interface of the Intel XScale<sup>®</sup> core Bus Interface Unit. The region 1 clock is an integer multiple of one-half the frequency of the DDR SDRAM clocks of Region 2. The creates a synchronous architecture between the Bus Interface Unit and Memory Controller Unit for minimal latency between the Core processor and DDR SDRAM.

# 20.1.3 Clocking Region 2

Region 2 provides six DDR SDRAM output clocks, and the clocking for the Memory Controller Unit. The clocking unit contains three output clocks, called M\_CLK[2:0], and three complement clock outputs called M\_CLK[2:0]#. The M\_CLK[2:0] and M\_CLK[2:0]# outputs are used by the DDR SDRAM memory subsystem. The frequency of this region is based on the MEM\_TYPE reset strap input. It supports clock frequencies up to a maximum of 333 MHz operation for DDR333 SDRAM. When configured for DDR-II 400, the maximum frequency of the region is 400 MHz.

# 20.1.4 Clocking Region 3

Region 3 covers the 80333 internal bus. It supports clock frequencies up to 333 MHz. The IOP units interface to or reside in this region. Units which interface to this clock region include the Intel XScale<sup>®</sup> core Bus Interface Unit, the Memory Controller Unit, Interrupt Controller, Timers, ATU, MU, Performance Monitor Unit and the bridge interface to the AHB bus. Units which reside in this clock region, and are based entirely on clock region 3 clocking include DMA units, AAU unit, and the Peripheral Bus interface. The peripheral bus interface operates at one-quarter of the Internal Bus frequency or 66MHz.

# 20.1.5 Clocking Region 4

Region 4 obtains its input clock from the clocking unit specified in clocking region 3. This region is for use by low-speed peripheral units. Currently, these include the  $I^2C$  bus interface, the General Purpose I/O unit and the UART serial bus interface.

Region 4 contains an output clock (SCL) used for the  $I^2C$  bus interface (see Chapter 10, " $I^2C$  Bus Interface Units"). The SCL clock frequency for  $I^2C$  operation is 100 KHz or 400 KHz. SCL is generated from the internal bus clock. In order to use the  $I^2C$  interface, a clock divider value must be written into the  $I^2C$  Clock Count Register. The UART input clock is driven at 33.334 MHz, and divided within the unit for the serial interface baud rate.

# intel

# 20.1.6 Clocking Region 5

Region 5 covers the PCI bus segment "A", and obtains its input clock **A\_CLKIN** as the board level feedback of from the **A\_CLKOUT** output clock. Region 5 also generates four secondary PCI bus segment output clocks the 80333 based on a dedicated PLL. Eight of these output clocks, **A\_CLKO[3:0]** are used to drive the Primary PCI input clocks of secondary PCI devices attached to the 80333 downstream PCI interfaces and the, **A\_CLKOUT** used as the reference clock to this region's PCI interface. This region is the downstream PCI "A" interface for 80333 and meets clocking requirements as specified in *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a. This region supports clock frequencies up to 133 MHz.

This region also includes the PCI interface of the ATU.

# 20.1.7 Clocking Region 6

Region 6 covers the PCI bus segment "B", and obtains its input clock **B\_CLKIN** as the board-level feedback from the **B\_CLKOUT** output clock. Region 6 also generates five secondary PCI bus segment output clocks the 80333 based on a dedicated PLL. Eight of these output clocks, **B\_CLKO[4:0]** are used to drive the Primary PCI input clocks of secondary PCI devices attached to the 80333 downstream PCI interfaces and the, **A\_CLKOUT** used as the reference clock to this region's PCI interface. This region is the downstream PCI "B" interface for 80333 and meets clocking requirements as specified in *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0a. This region supports clock frequencies up to 133 MHz.

# 20.1.8 Clocking Region 7

Region 7 contains the main input clock for providing the 80333 with all of its clock sources. This input clock is provided by the system designer. This differential input clock, called the EXP\_CLK/EXP\_CLK# clock, is connected to the input pins. The 80333 supports an input frequency of 100 MHz per the *PCI Express Specification*, Revision 1.0 for normal operation on the PCI Express interface.



# 20.1.9 Clocking Region Summary

Table 521 summarizes all of the input clock pins, output clock pins, and clock strapping option pins used in the 80333.

#### Table 521.Clock Pin Summary

Pin	Input/Output	Description
EXP_CLK/EXP_CLK#	Input	Differential PCI Express Input 100MHz Clock
A_CLKO[3:0], B_CLKO[4:0]	Output	Secondary PCI Output Clocks
A_CLKOUT, B_CLKOUT	Output	Secondary PCI Output Reference Clock
A_CLKIN, B_CLKIN	Input	Secondary PCI Reference Clock Input
B_M66EN, A_M66EN	Input	PCI 66 MHz Enable
A_PCIX133EN, B_PCIX133EN	Input	Reset Strap to Limit Secondary Clock Frequency to 100MHz or 133MHz in PCI-X Mode
M_CLK[2:0]	Output	SDRAM Output Clocks
M_CLK[2:0]#	Output	Complementary SDRAM Output Clocks
В_НОРСК	Output	PCI Hot-Plug Controller Output Clock on PCI Bus Segment "B" for output serial stream
B_HIPCK	Output	PCI Hot-Plug Controller Output Clock on PCI Bus Segment "B" for input serial stream
SCL[1:0]	Input/Output	I <sup>2</sup> C Output Clocks
SCLK	Input/Output	SMBus Output Clock



Table 522 summarizes all of the clocks generated to the seven regions within the 80333. The clock units initializes appropriately based on the PCI-X initialization pattern driven during the rising edge of  $P_RST\#$ .

#### Table 522. 80333 Clock Region Summary

Region	Interface	Unit	Frequencies				
1		Intel XScale <sup>®</sup> core	500 MHz 667 MHz		800 MHz		
2	DDR SDRAM	MCU	333 MHz	333 MHz	N/A		
2	DDR-II SDRAM	Mee	400 MHz	N/A	400 MHz		
3	Internal Bus	DMA[1:0], AAU, ATU		333 MHz			
5	Peripheral Bus Interface	PBI	66 MHz				
		l <sup>2</sup> C	33 MHz				
4	4 AHB		33.334 MHz				
		GPIO	33 MHz				
5	PCI-A	PCI-A side of Bridge	33, 66, 100, 133MHz based on PC and initiation sequer		I reset strapping ce.		
6	PCI-B	PCI-B side of Bridge	33, 66, 100, 133MHz based on PCI reset s and initiation sequence.		I reset strapping ce.		
		Port	2.5 GHz				
7	PCI Express	Primary Side of A and B Bridge	250 MHz				

# 20.2 Reset Overview

There are seven levels of reset for the 80333. The main reset is controlled through the reset signal (**RSTIN**#). When this signal is asserted, the entire 80333 is placed in a reset state. In addition to the reset pin, the 80333 provides software control of internal units such as the PCI Express-to-PCI Bridges, Internal Bus and Intel XScale<sup>®</sup> core.

The 80333 seven levels of reset are:

- **PWRGD** this signal indicates stable power when high and causes a asynchronous reset of the entire chip when low. All reset straps are sampled at the rising edge of **PWRGD**.
- RSTIN# this is also an asynchronous reset to the PXH and can be used for resetting 80333.
- PCI Express Reset a message coming on the PCI Express interface and is not a physical signal.
- Software PCI Reset this reset is initiated by writing to bridge control register of the PCI configuration space. This reset is specified to the particular bridge that the software wished to reset. This is also commonly referred to as the SBR (secondary bus reset).
- Internal Bus Reset this reset is initiated by writing to ATU control register in PCI configuration space. This reset is specific to the integrated I/O processor and the associated peripheral units. This reset may also be generated by the expiration of the Watch Dog Timer as described in Section 16.1.2, "Watch Dog Timer Operation" on page 734
- Intel XScale<sup>®</sup> core reset this reset is initiated by two mechanisms. First is the **CORE\_RST#** reset strap, and the other is via software through ATU control register in PCI configuration space.

# intel®

Figure 151 shows the logical block diagram of the reset conditions.

#### Figure 151. Intel<sup>®</sup> 80333 I/O Processor Reset Block Diagram



# 20.2.1 PWRGD Reset Mechanism

All the voltage sources in the system are tracked by a system component that asserts the **PWRGD** signal only after 'all' the voltages have been stable for some predetermined time. The 80333 receives the **PWRGD** signal as an asynchronous input, meaning that there is no assumed relationship between the assertion or the de-assertion of **PWRGD** and the reference clock. While the **PWRGD** is de-asserted the 80333 will hold all logic in reset.

The **PWRGD** reset will clear all internal state machines and logic, and initialize all registers to their default states including 'sticky' error bits that are persistent through all other reset classes. To eliminate potential system reliability problems, all devices are also required to either tristate their outputs or to drive them to safe levels during such a power on reset.



# 20.2.2 RSTIN# Reset Mechanism

Once the system is up and running, a full system reset may be required to recover from system error conditions related to various device or subsystem failures. This hot reset mechanism is provided to accomplish this recovery without clearing the 'sticky' error status bits useful to track the cause of the device or subsystem error conditions.

A hot reset can be initiated by asserting the **RSTIN**# signal. This signal is treated as an asynchronous input to the 80333, meaning that there is no assumed relationship between the assertion or the de-assertion of **RSTIN**# and the host reference clock

When the reset signal (**RSTIN**#) is asserted, all configuration registers, internal control and enable signals, state machines, and output buffers are reset to their initialized state, including the bridge and I/O processor units. The affects on the units beyond the bridge are described in Section 20.2.6, "I/O Processor Reset" on page 873.

# 20.2.3 PCI Express Reset Mechanism

There is no reset signal on the PCI Express and all reset communication is in-band. The Root Complex communicates the fact that it is entering and coming out of a reset using messages. Downstream devices (80333) will respond by also going through a reset. This incoming message by nature of the PCI Express protocol is asynchronous to the reference clock. However when 80333 goes through a reset for its own reasons (**PWRGD**, **RSTIN#**) the PCI Express link goes down which will be inferred by the Upstream device and handled with a Hot-Plug reset (when Hot-Plug is enabled).

# 20.2.4 Software PCI Reset Mechanism

This reset is initiated by a write to the bridge control register and resets only the particular PCI segment. This reset can be used for various reasons in including recovering from error conditions on the secondary bus, to redo enumeration, to change the operating frequency of the bus (33/66/100/133 MHz), to change the operating mode of the bus (PCI or PCI-X) etc. This reset is synchronous to the PCI clock domain in which it is used. SBR is strictly restricted to the particular PCI segment and affects neither the other PCI segment nor the rest of the 80333 logic. Writes to the bridge control register with a new frequency etc, will have no effect until the SBR happens. The power up frequency of the PCI bus is shown in Table 2.5.1.1, "Initialization" on page 72. When hot-plug is enabled the bus always powers up in 33 MHz mode. With Hot-Plug disabled the frequency depends on the M66EN and the PCIXCAP pins.

The I/O processor units of 80333, as a device on the secondary PCI bus segment, is completely reset by a SBR. Refer to Section 20.2.6, "I/O Processor Reset" on page 873 for details of the I/O processor reset.



# 20.2.5 Hot-Plug Reset Mechanism

This reset is initiated by a write to the Hot-Plug command register with the change frequency command. Note that a write to this register might do any of the following:

- Change the frequency (33/66/100/133)
- Change the mode (PCI or PCI-X)
- Change nothing but rewrite the present settings

Any write (doing any of the three above) will cause a reset of the particular PCI segment to reset. This reset is asynchronous by nature to the PCI clock as the Hot-Plug logic runs off of an internal clock that 'may' be asynchronous to the PCI clock. This reset will cause the any newly written frequency or PCI mode information to take effect. 80333 will support all changes in mode because of Hot-Plug events including switching from PCI to PCI-X or changing the frequency anytime.

A Hot-Plug Reset event on PCI Bus segment "A" has the affect of resetting the integrated I/O processor of 80333 as described in Section 20.2.6, "I/O Processor Reset" on page 873.

# 20.2.6 I/O Processor Reset

Those reset mechanisms which assert secondary bus reset, results in a reset of the I/O processor in addition to the additional reset actions described in the respective reset mechanism descriptions above.

When the I/O processor is reset by the Secondary Bus Reset signal A\_RST#, the I/O processor:

- resets the internal bus, refer to Section 20.2.7, "Internal Bus Reset" on page 874.
- resets the Intel XScale<sup>®</sup> core, refer to Section 20.2.8, "Intel XScale<sup>®</sup> core Reset Mechanism" on page 877
- resets all internal units
- · resets all Memory Mapped Registers
- all configuration straps are not affected by A\_RST#, refer to Section 20.4.

The assertion and deassertion of the PCI reset signal (**A\_RST**#) is asynchronous with respect to **A\_CLKIN**. The rising edge of the **A\_RST**# signal must be monotonic through the input switching range and must meet the minimum slew rate. The *PCI Local Bus Specification*, Revision 2.3 defines the assertion of **A\_RST**# for a period of 1 ms after power is stable.

Upon the assertion of **A\_RST#**, all units within the 80333, excluding the bridge are reset. This reset resets all internal memory mapped registers (MMRs) to their default configuration state. The reset value for each register is defined within each register description.



# 20.2.7 Internal Bus Reset

The Reset Internal Bus bit in the PCI Configuration and Status Register (see Table 168, "PCI Configuration and Status Register - PCSR" on page 283) resets the Intel XScale<sup>®</sup> core and all units on the internal bus. Expiration of the Watch Dog Timer as described in Section 16.1.2, "Watch Dog Timer Operation" on page 734 also results in the same internal bus reset.

Before resetting, the ATU shall gracefully halt all PCI bus transactions when operating in PCI mode. Following an Internal Bus Reset when the PCI Bus is operating in PCI-X mode, the ATU may allow Split Completion transactions due to Outbound Split Requests to Master Abort. Also, the ATU may not initiate all of the Split Completion transactions required by any outstanding Inbound Split Requests at the time of the Internal Bus Reset.

While the *PCI-X Addendum to the PCI Local Bus Specification*, Revision 1.0 requires host software to gracefully handle both these cases without reinitializing the PCI bus, the user may want to avoid this scenario. This can be done by taking the following steps:

- 1. Disable the ATU from either claiming or initiating new transactions by clearing the *Bus Master Enable* and the *Memory Enable* in the ATU Command Register (see Section 3.10.3, "ATU Command Register - ATUCMD" on page 240).
- 2. Monitor the *Inbound Read Transaction Queue Status* and the *Outbound Read Transaction Queue Status* in the PCSR.
- 3. When both the Inbound Read Transaction queue and the Outbound Read Transaction queue are empty, software writes to the PCSR to initiate the Internal Bus Reset.

It is the responsibility of the software to ensure that the I<sup>2</sup>C bus and UART are idle before the reset occurs. The Intel XScale<sup>®</sup> core may or may not be held in reset when the Reset Internal Bus bit is cleared by software. This depends on the **CORE\_RST#** strap as described in section Section 20.2.8.

Expiration of the Watch Dog Timer does not result in a graceful Internal Bus Reset, as the queues, I<sup>2</sup>C bus and UART may be active when the reset occurs.

When the Internal Bus Reset is initiated there are sideband signals notifying the ATU that a reset is coming. The following describes the operation of each unit:

# intel

# Table 523. Internal Bus Reset Summary (Sheet 1 of 2)

Unit	Preparation for Reset	Reset Status
	The ATU and MU detect via a sideband signal that a internal bus reset is coming. Upon detecting this signal, the ATU and MU complete the following:	
	PCI Interface Outbound Transaction:	
	<ul> <li>When the ATU has already asserted its PCI request signal, and not yet started a transaction, the ATU deasserts its request and not start its transaction.</li> </ul>	
	<ul> <li>When the ATU has not yet requested the PCI bus, the ATU never asserts its request for the PCI bus.</li> </ul>	
	<ul> <li>When the ATU is in the middle of a transaction, the ATU performs the existing transaction. This means that an outbound write, the ATU transfers as much data as available in the queue. When an outbound read, the ATU reads the data until the transaction stops naturally (meaning that the target has ended the transaction or the ATU has read all of the data it has been requested to read). Once terminated by the ATU or the target, the ATU no longer requests the PCI bus.</li> </ul>	
	<ul> <li>In PCI-X mode, the ATU allows any outstanding split completions due to prior outstanding Split Requests to Master-Abort on the PCI Bus. Since, the IOP is only accessing Prefetchable Memory on the host, no error condition is created in the system.</li> </ul>	
⊇	PCI Interface Inbound Transaction:	Clear all ATU and MU interrupts
ATU/N	<ul> <li>Once the ATU and MU have detected that an internal bus reset is coming, they no longer claim any new transactions on the PCI bus. This results in a master abort to the initiating master.</li> </ul>	and HPIs. All ATU and MU MMRs reset to the default value.
	<ul> <li>In PCI-X mode, this may mean that data from an outstanding split request may not be returned to the host. It is the responsibility of the host software to handle this condition as a consequence of writing to the Internal Bus Reset bit.</li> </ul>	
	<ul> <li>The other requirement for the inbound ATU is to complete the configuration cycle which set the IB reset bit in the ATUs configuration space. In PCI-X Mode this means returning a Split Completion Message to the host indicating that the write has completed.</li> </ul>	
	Internal Bus Interface: Inbound Transaction:	
	<ul> <li>The ATU and MU go ahead and assert their IB request signals, and try to continue any pending transactions as normal. There are no special actions taken on the internal bus for inbound transactions.</li> </ul>	
	Internal Bus Interface: Outbound Transaction:	
	<ul> <li>For all ATU and MU outbound transactions, there are no special requirements since the BIU/Core/DMA is reset.</li> </ul>	
	Upon meeting both the outbound and inbound transaction requirements, the ATU and MU assert the sideband signal to the reset unit notifying ready-for-reset.	
DMA0/1/	No special requirements. DMA can be reset at any time.	Clear all interrupts and HPIs. All DMA MMRs reset to the default value.



# Table 523. Internal Bus Reset Summary (Sheet 2 of 2)

Unit	Preparation for Reset	Reset Status
BIU	No special requirements. BIU can be reset at any time.	Clear all interrupts and HPIs. The BIU behaves as the Core after reset. Meaning, the CORE_RST# reset strap determines when the core is to be held in reset. All BIU mmrs reset to the default value.
AA	No special requirements. AA can be reset at any time.	Clear all interrupts and HPIs. All AA MMRs reset to the default value.
IB-ARB	No special requirements. IB-ARB can be reset at any time.	Clear all interrupts and HPIs. All IB-ARB MMRs reset to the default value.
l <sup>2</sup> C	No special requirements. $I^2C$ can be reset at any time. When the $I^2C$ is in the middle of a transmission, the protocol may be violated when the $I^2C$ unit resets.	Clear all interrupts and HPIs. All I <sup>2</sup> C MMRs reset to the default value.
UART	No special requirements. UART can be reset at any time. When the UART is in the middle of a transmission, the protocol may be violated when the UART unit resets.	Clear all interrupts and HPIs. All UART MMRs reset to the default value.
MCU	No special requirements. MCU can be reset at any time.	Clear all interrupts and HPIs. All MCU MMRs reset to the default value. <b>M_RST#</b> is asserted.
PBI	No special requirements. PBI can be reset at any time.	Clear all interrupts and HPIs. All PBI MMRs reset to the default value.



# 20.2.8 Intel XScale<sup>®</sup> core Reset Mechanism

This reset is initiated the **CORE\_RST#** reset strap, and exited via software through ATU PCSR register in PCI configuration space. Upon the de-assertion of the reset signal, the **CORE\_RST#** is sampled as described in Section 20.4, "Reset Strapping Options" on page 879. When asserted, the Intel XScale<sup>®</sup> core processor is held in the reset state. Software is required to clear this bit to deassert Intel XScale<sup>®</sup> core reset. Software cannot set this bit.

# 20.2.9 Reset Summary

#### Table 524. Reset Summary

Reset Source	Affected Resets	Affected Units
PWRGD	All	Full Chip including sticky bits and Reset Straps
RSTIN#	All	Full Chip except Sticky bits and Reset Straps
PCI Express In Band	All	Full Chip except Sticky bits and Reset Straps
Secondary Bus Reset A-Segment	A_RST#, I_RST#	A segment bridge logic (not configuration register) and all IOP units.
Secondary Bus Reset B-Segment	B_RST#	B segment bridge logic (not configuration register)
Internal Bus Reset bit in ATU	I_RST#	All IOP units
Watch Dog Timer Expiration	I_RST#	All IOP units
CORE_RST#	Intel XScale <sup>®</sup> core reset	Intel XScale <sup>®</sup> core
Hot-Plug initiated Reset	B_HRST#	B Segment Hot-Plug slot adapter cards



# 20.3 Reset Sequencing

Figure 151 shows the reset sequencing for the 80333 reset in a motherboard implementation.

# Figure 152. Intel<sup>®</sup> 80333 I/O Processor System Reset



# intel

# 20.4 Reset Strapping Options

There are many initialization modes that can be selected when the processor is reset. Table 525 shows the configuration modes. All of the configuration modes defined are determined on the rising edge of **PWRGD**.

Upon the assertion of **PWRGD**, the 80333 samples a series of strapping pins to set configuration modes. One strap which alters the behavior of the 80333 on the assertion of **PWRGD** is the **CORE\_RST#** strap. When the **CORE\_RST#** pin is asserted on the rising edge of **PWRGD**, the 80333 continues to assert the individual reset to the Intel XScale<sup>®</sup> core. This mode, holds the Intel XScale<sup>®</sup> core in reset until the Core Processor Reset Bit in the PCI Configuration and Status Register (ATU) is cleared, thus allowing the Intel XScale<sup>®</sup> core to enter its initialization procedure.

The PCI interface of the 80333 ATU is integrated to the Secondary PCI bus segment "A" and therefore samples the **A\_REQ64**# signal to determine when this secondary bus segment is connected to a 64-bit external data path. The state of **A\_REQ64**# on the rising edge of the **PWRGD** signal notifies 80333 PCI Express-to-PCI Bridge PCI bus interface it is connected to a 64-bit or 32-bit PCI bus.

The remaining reset straps are sampled similarly and are described in Table 525 below.

NAME	DESCRIPTION						
	Bus Width: <b>P_BOOT16#</b> is latched at the rising edge of <b>PWRGD</b> and it indicates the default bus width for the PBI Memory Boot window.						
	<ul><li>0 = 16 bits wide (Requires pull-down resistor).</li><li>1 = 8 bits wide (Default mode).</li></ul>						
AD[6] / <b>RETRY</b>	Configuration Retry Mode: <b>RETRY</b> is latched at the rising edge of <b>PWRGD</b> and it determines when the PCI interface of the ATU disables PCI configuration cycles by signaling a Retry until the Configuration Cycle Retry bit is cleared in the PCI Configuration and Status Register. <b>RETRY</b> also controls the behavior of the Upstream PCI Express to configuration transactions.						
	<ul> <li>0 = Configuration Cycles enabled (Requires pull-down resistor).</li> <li>1 = Configuration Retry enabled in the ATU and Configuration Retry Status response enabled in the Upstream PCI Express interface. (Default mode).</li> </ul>						
AD[5] / CORE_RST#	RESET MODE is latched at the rising edge of <b>PWRGD</b> and it determines when the Intel XScale <sup>®</sup> core is held in reset until the Intel XScale <sup>®</sup> core Processor Reset bit is cleared in the PCI Configuration and Status Register.						
	0 = Hold in reset (Requires pull-down resistor). 1 = Don't hold in reset (Default mode).						
	Memory Frequency: <b>MEM_TYPE</b> is latched at the rising edge of <b>PWRGD</b> and it indicates the speed of the DDR SDRAM interface.						
	0 = DDR-II SDRAM @ 400MHz (Requires pull-down resistor). 1 = DDR SDRAM @ 333MHz (Default mode).						
AD[3] / A_PCIX133EN	PCI Bus Segment "A" 133MHz Enable: <b>A_PCIX133EN</b> is latched at the rising edge of <b>PWRGD</b> and it determines the maximum PCI-X mode operating frequency.						
	0 = 100MHz enabled (Requires pull-down resistor). 1 = 133MHz enabled (Default mode).						
AD[10] / <b>B_PCIX133EN</b>	PCI Bus Segment "B" 133MHz Enable: <b>B_PCIX133EN</b> is latched at the rising edge of <b>PWRGD</b> and it determines the maximum PCI-X mode operating frequency.						
	0 = 100MHz enabled (Requires pull-down resistor). 1 = 133MHz enabled (Default mode).						

#### Table 525. Reset Strap Signals (Sheet 1 of 2)



# Table 525.Reset Strap Signals (Sheet 2 of 2)

NAME	DESCRIPTION					
AD[20] / PCIODT_EN	<b>PCI Bus ODT ENABLE: PCIODT_EN</b> is latched on the rising edge of <b>PWRGD</b> and it determines when the PCI-X interface has On Die Termination enabled. PCI ODT enable is valid for 80333 A and B segments in both PCI-X Mode 1 and Mode 2 (B_MODE2 = 1).					
	<ul><li>0 = ODT disabled (Requires pull-down resistor).</li><li>1 = ODT enabled (Default mode).</li></ul>					
AD[15] / <b>B_HSLOT[3]</b>	Standard Hot-Plug Controller Enable: <b>B_HSLOT[3]</b> is latched at the rising edge of <b>PWRGD</b> and it indicates when the "B" Secondary PCI-X bus interface Standard Hot-Plug Controller is enable for operation					
	<ul><li>0 = SHPC Disabled (Requires pull-down resistor).</li><li>1 = SHPC Enable (Default mode).</li></ul>					
AD[14:12] /	Standard Hot-Plug Controller Slot Count: <b>B_HSLOT[2:0]</b> is latched at the rising edge of <b>PWRGD</b> and it indicates the number of "B" Secondary PCI-X bus slots and the mode of control for the Standard Hot-Plug Controller.					
	Refer to Chapter 14, "Standard Hot-Plug Controller" for details of the values for these reset straps.					
A[19] / SMB_MA5 A[18] / SMB_MA3 A[17] / SMB_MA2	Manageability Address (MA): is latched on the rising edge of <b>PWRGD</b> and maps to MA bit 5, 3, 2, and 1 where MA bits 7:0 represent the address the SMBus slave port will respond to when access is attempted. 0 = Requires pull-down resistor					
	1 = Derault mode					



# Test Logic Unit and Testability

# 21.1 Overview

This chapter summarizes the testability features (DFT) that are incorporated in the Intel<sup>®</sup> 80333 I/O processor (80333). The Test Logic Unit (TLU) is based on the IEEE 1149.1a *Standard Test Access Port and Boundary-Scan Architecture* and supports boundary scan.

# 21.2 Test Control/Observe Pins

The following table lists the complete set of test pins required.

### Table 526. Test Control/Observe Pins

Pin Name	I/O	Test Purpose
тск	Ι	Test Clock. Input clock for the test logic defined by IEEE1149.1Std. Maximum frequency: 200 MHz
TDI	Ι	Test Data Input. TAP controller defined by IEEE1149.1 Std receives serial test instructions and data via this pin.
TDO	0	Test Data Output. TAP controller defined by IEEE1149.1 Std sends the serial output from the instruction and data registers via this pin.
тмѕ	I	Test Mode Select. TAP controller defined by IEEE1149.1 decodes the signal received at this pin to <b>TMS</b> to control test operations.
TRST#	I	Test Reset. The optional <b>TRST#</b> input provides for asynchronous initialization of the TAP controller.



# 21.3 IEEE 1149.1 Standard Test Access Port (TAP)

The IOP contains test logic that is compatible with the IEEE Standard 1149.1-1990 Test Access Port (TAP) and Boundary Scan Architecture. Logic that conforms to this standard contains the following:

- 1. A Test Access Port (TAP) four inputs (**TDI**, **TMS**, **TRST#** and TCLK) and a single output (**TDO**).
- 2. A TAP controller.
- 3. An instruction register.
- 4. A group of test data registers.

Each of these is described in more detail below. Figure 153 shows a generic diagram for logic conforming to the IEEE 1149.1 test standard.

#### Figure 153. IEEE 1149.1 Standard. Block Diagram





# 21.3.1 TAP Pin Description

The internal test logic is accessed through the TAP pins. The following sections describe some of the rules and permissions of the IEEE 1149.1a Standard for the TAP pins.

# 21.3.1.1 Test Clock (TCK)

This is the clock input for the test logic defined by this standard, i.e. the TAP controller and associated registers. The TLU is a fully static design, thus all registers retain their states indefinitely when **TCK** is stopped at "0" or "1".

# 21.3.1.2 Test Mode Select (TMS)

This pin is used to control the operation of the TAP controller. The signal received at **TMS** is decoded by the TAP controller to control test operations. The state of **TMS** is sampled on the rising edge of **TCK**. Internally, there is a weak pull-up on this pin to provide a logic high when not driven, per standard definition.

# 21.3.1.3 Test Data Input (TDI)

This pin is used to provide serial input data to the instruction and test data registers. Data at **TDI** is sampled on the rising edge of **TCK**. Data shifted from **TDI** through a register to **TDO** appears non-inverted at **TDO** after a number of rising and falling edges of **TCK** determined by the length of the instruction or test data register selected. Internally, there is a weak pull-up on this pin to provide a logic high when not driven, per standard definition.

### 21.3.1.4 Test Data Output (TDO)

This is the serial data output pin. Changes in the state of **TDO** occur only following the falling edge of **TCK** while performing a shift operation. This pin is only driven while scanning (in SHDR or SHIR states) otherwise it is in inactive (high Z) state. The non-shift inactive state is provided to support parallel connection of **TDO** outputs at the board or module level.

### 21.3.1.5 Asynchronous Reset (TRST#)

The **TRST**# signal is used to asynchronously reset the TAP controller and boundary-scan registers. The TAP controller is not initialized by any other system input, including system reset. The TAP controller initializes asynchronously on the falling edge of **TRST**# to the Test-Logic-Reset (initial) state. The TAP controller is initialized at power-up by cirrostrati built into the test logic. Upon reset, the TAP instruction register initializes to the IDCODE instruction. Internally, there is a weak pull-up on this pin to provide a logic high when not driven.

# 21.3.2 TAP Controller

The TAP controller, shown in Figure 155, is a sixteen-state synchronous finite state machine that changes state on the rising edge of **TCK**. The controller next state is controlled by the state present at the **TMS** input. The TAP controller generates control signals, which together with **TCK** and control signals decoded from the instruction active in the instruction register, determine the operation of the test circuitry as defined by the IEEE Standard.

All state transitions occur based on the values of TMS on the rising edge of **TCK** Actions of the test logic (instruction register, data registers, etc.) occur on either the rising or falling edge of **TCK**, as show in Figure 154.. See the description of each state to learn which.

#### Figure 154. Timing of Actions in a TAP Controller State



For greater detail on the state machine and the public instructions, refer to IEEE 1149.1a Standard Test Access Port and Boundary-Scan Architecture Specification.

#### Figure 155. TAP Controller State Diagram



# 21.3.2.1 Test-Logic-Reset State

In this state, test logic is disabled to allow normal operation of the IOP. This is achieved by loading the instruction register with the IDCODE instruction. No matter what he state of the controller, it enters Test-Logic-Reset state when the **TMS** input is held high for at least five rising edges of **TCK**. The controller remains in this state while **TMS** is high. The TAP controller is also forced to enter this state by enabling **TRST**#.

When the controller exits the Test-Logic-Reset controller state as a result of an erroneous low signal on the **TMS** line at the time of a rising edge on **TCK** (for example, a glitch due to external interference), it returns to the Test-Logic-Reset state following three rising edges of **TCK** with the **TMS** line at the intended high logic level. Test logic operation is such that no disturbance is caused to on-chip system logic operation as the result of such an error.

Transition to next state: On the rising edge of **TCK**, when **TMS** is low move to Run-Test/Idle, else **TMS** remains high so stay in Reset.

### 21.3.2.2 Run-Test/Idle State

This state is a controller state between scan operations. The controller remains in this state as long as TMS is held low. In the Run-Test/Idle state, activity in selected test logic occurs only when certain instructions are present. Instructions that do not cause functions to execute generate no activity in the test logic while the controller is in this state.

The instruction register and all test data registers retain their current value in this state.

Transition to next state: When **TMS** is high on the rising edge of **TCK**, move to Select-DR-Scan, else remain in Idle.

### 21.3.2.3 Select-DR-Scan State

This is a temporary controller state. Here the decision is made to enter the Capture-DR column and initiate a scan sequence for the selected test data register.

All test data registers selected by the current instruction retain their previous value in this state.

Transition to next state: When **TMS** is low on the rising edge of **TCK**, move to Capture-DR, else move to Select-IR.

### 21.3.2.4 Capture-DR State

When the controller is in this state data is parallel-loaded into test data registers selected by the current instruction on the rising edge of **TCK**. Test data registers that do not have parallel inputs are not changed. Also when capturing is not required for the selected instruction, the register retains its previous state.

The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is low on the rising edge of **TCK**, move to Shift-DR, else move to Exit1-DR.



### 21.3.2.5 Shift-DR State

In this controller state, the test data register, which is connected between **TDI** and **TDO** as a result of the current instruction, shifts data one bit position nearer to its serial output on each rising edge of **TCK**. Test data registers that the current instruction select but do not place in the serial path, retain their previous value during this state.

The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is high on the rising edge of **TCK**, move to Exit1-DR, else remain at Shift-DR.

### 21.3.2.6 Exit1-DR State

This is a temporary controller state.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is held high on the rising edge of **TCK**, the controller enters the Update-DR state and the scanning process terminates. When **TMS** is held low on the rising edge of **TCK**, the controller enters the Pause-DR state.

#### 21.3.2.7 Pause-DR State

The Pause-DR state allows the test controller to temporarily halt the shifting of data through the test data register in the serial path between **TDI** and **TDO**.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is high on the rising edge of **TCK**, move to the Exit2-DR, else remain at Pause-DR.

### 21.3.2.8 Exit2-DR State

This is a temporary state.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is high on the rising edge of **TCK**, the controller enters the Update-DR state and the scanning process terminates. When TMS is held low on the rising edge of **TCK**, the controller re-enters the Shift-DR state



# 21.3.2.9 Update-DR State

Data is latched into the parallel output of shift registers from the shift register path, on the falling edge of **TCK**.

All of the test data register's shift-register bit positions selected by the current instruction retain their previous values. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** remains high on the rising edge of **TCK**, then the controller moves to the Select-DR state, else the controller moves to the Run-Test/Idle state.

### 21.3.2.10 Select-IR-Scan State

This is a temporary controller state. Here the decision is made to enter the Capture-IR column and initiate a scan sequence for the instruction register or to return to Test-Logic-Reset.

All test data registers selected by the current instruction retain their previous value during this state.

Transition to next state: When **TMS** is low on the rising edge of **TCK**, move to Capture-IR, else move to Test-Logic-Reset.

#### 21.3.2.11 Capture-IR State

In this state, the shift register contained in the instruction register loads the fixed value  $0000001_2$  on the rising edge of **TCK**.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is low on the rising edge of **TCK**, move to Shift-IR, else move to Exit1-IR.

### 21.3.2.12 Shift-IR State

In this state, the shift register contained in the instruction register is connected between **TDI** and **TDO** and shifts data one bit position nearer to its serial output on each rising edge of **TCK**.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is high on the rising edge of **TCK**, move to Exit1-IR, else remain at Shift-IR.

#### 21.3.2.13 Exit1-IR State

This is a temporary state.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is high during the next rising edge of **TCK**, the controller enters the Update-IR state and the scanning process terminates. When **TMS** is held low during the next rising edge of **TCK**, the controller enters the Pause-IR state.



# 21.3.2.14 Pause-IR State

This state allows the TAP controller to temporarily halt the shifting of data through the instruction register.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is high on the rising edge of **TCK**, move to Exit2-IR, else remain in the Pause-IR state.

### 21.3.2.15 Exit2-IR State

This is a temporary state.

All test data registers selected by the current instruction retain their previous value during this state. The instruction does not change while the TAP controller is in this state.

Transition to next state: When **TMS** is held high during the next rising edge of **TCK**, the controller enters the Update-IR state and the scanning process terminates. When **TMS** is held low during the next rising edge of **TCK**, the controller re-enters the Shift-IR state

### 21.3.2.16 Update-IR State

The instruction shifted into the instruction register is latched onto the parallel output from the shift-register path on the falling edge of **TCK**. Once latched, the new instruction becomes the current instruction.

All test data registers selected by the current instruction retain their previous values in this state.

Transition to next state: When **TMS** remains high on the rising edge of **TCK**, then the controller moves to the Select-DR state, else the controller moves to the Run-Test/Idle state.



# 21.3.3 TAP Controller Configuration

The Test Logic Unit (TLU) and the Intel XScale<sup>®</sup> core each contain a separate TAP controller. They are connected in series through the TDI-TDO stream as shown in Figure 156.







# 21.3.4 TAP Controller Registers

The IEEE 1149.1 architecture specifies an instruction register and a minimum of two test data registers: bypass and boundary scan. The IOP TAP controller meets this requirement.

### 21.3.4.1 Instruction Register

Each of the TAP controllers in the design contain an instruction register (IR). Each IR is a 7-bit, master/slave-configured, parallel-loadable, serial-shift register with latched outputs. They are used in each unit to select the test data register to be accessed.

When the TAP controller is in the Shift-IR state, instructions are loaded serially via **TDI** clocked by the rising edge of **TCK**. The shifted-in instruction becomes active upon latching from the serial-stages to the parallel-stages in the Update-IR state. At that time the IR outputs, along with the TAP finite state machine outputs, are decoded to select and control the test data register selected by that instruction. Upon latching, all actions caused by any previous instructions must terminate.

On activation of **TRST**#, the latched instruction asynchronously changes to the IDCODE instruction. Additionally, upon entering the Test-Logic-Reset state, the IDCODE instruction is latched and becomes active on the falling edge of **TCK**.

Because the TAP controllers are connected in series, each unit must receive its own 7-bit instruction during the Shift-IR state. The BYPASS instruction should be loaded into the unit that is not to be accessed.

For example, to load the IDCODE instruction into the TLUs IR, the data shifted into TDI during the Shift-IR state is:

```
1 1 1 1 1 1 1 1 1 1 1 1 1 0 -----> TDI
| BYPASS into | IDCODE |
| X-Scale core | into TLU |
```

As another example, to load the IDCODE instruction into the Intel XScale<sup>®</sup> microarchitecture instruction register, the data shifted into TDI during the Shift-IR state is:

1	1	1	1	1	1	0	1	1	1	1	1	1	1		 - >	TDI
	IDO	COI	ΣC	ir	nto	C			B	ΖP	ASS	5				
2	X – S	Sca	ale	e o	201	ce				ir	nto	5 C	ГLŪ	J		



### 21.3.4.2 Instructions

Since the instruction set for each TAP controllers is independent of one another, each is listed separately below.

#### Table 527. TLU TAP Controller Instruction Set

Instruction	Opcode	Description
		Intended for supporting the boundary-scan feature for testing device interconnects at the board/system level, the EXTEST instruction connects only the boundary-scan register between TDI and TDO in the Shift-DR state. When EXTEST is selected, all output signal pin values are driven by values shifted into the boundary-scan register and may change only on the falling-edge of <b>TCK</b> in the Update_DR state.
EXTEST	0000000 <sub>2</sub>	Also, when extest is selected, all system input pin states must be loaded into the boundary-scan register on the rising-edge of <b>TCK</b> in the Capture_DR state.
		Note: Data would typically be loaded onto the latched parallel outputs of boundary-scan shift registers using the SAMPLE/PRELOAD instruction prior to selection of the EXTEST instruction.
		SAMPLE/PRELOAD performs two functions:
SAMPLE/ PRELOAD	0000001 <sub>2</sub>	• When the TAP controller is in the Capture-DR state, the sample instruction occurs on the rising edge of <b>TCK</b> and provides a snapshot of the component's normal operation without interfering with that normal operation. The instruction causes Boundary-Scan register cells associated with outputs to sample the value being driven by or to the processor.
		• When the TAP controller is in the Update-DR state, the preload instruction occurs on the falling edge of <b>TCK</b> . This instruction causes the transfer of data held in the boundary-scan cells to the slave register cells. Typically the slave latched data is then applied to the system outputs by means of the extest instruction.
HIGHZ	0110001 <sub>2</sub>	HIGHZ puts all output pins into a tri-state mode. When this instruction is active, the bypass register is connected between <b>TDI</b> and <b>TDO</b> .
IDCODE	1111110 <sub>2</sub>	IDCODE is used in conjunction with the device identification register. When selected, IDCODE parallel-loads the hard-wired identification code (32 bits) into the identification register on the rising edge of <b>TCK</b> following entry into the Capture-DR state. The instruction selects only the identification register for connection between <b>TDI</b> and <b>TDO</b> in the Shift-DR state for serial access. NOTE: The device identification register is not altered by data being shifted in on <b>TDI</b> .
BYPASS	1111111 <sub>2</sub>	BYPASS selects the bypass register between TDI and TDO while in Shift-DR state, effectively bypassing the processor's test logic. '0' is captured in the Capture-DR state. While this instruction is in effect, all other test data registers have no effect on the operation of the system.

#### 21.3.4.2.1 High-Z

High-Z aids in board-level testing. A mounted device effectively removes itself electrically from a circuit board. This allows for system-level testing where a remote tester exercises the processor system.

"HIGH-Z" is an instruction defined by the IEEE 1149.1 Standard. The requirement for this instruction are 1) all system logic outputs are high impedance; 2) the TAP controller continues to operate with the bypass register connected between TDI and TDO. The part may have other settings, as long as they do not interfere with these two requirements.

*Note:* Following the use of High-Z the part may be in an indeterminate state. Therefore, the IOP requires a reset in order to continue normal operation.

#### Table 528. Intel XScale<sup>®</sup> Core TAP Controller Instruction Set

Instruction	Opcode	Description			
DBGRXx	0000010 <sub>2</sub>	Used by the Intel XScale <sup>®</sup> core debugger during software debug. See core EAS Ch 9 for more information.			
LDIC	0000111 <sub>2</sub>	Used by the Intel XScale <sup>®</sup> core to load the instruction cache. Used during core burn-in testing. See core EAS Ch 9 for more information.			
DCSR	0001001 <sub>2</sub>	Used by the Intel XScale <sup>®</sup> core core debugger during software debug. See core EAS Ch 9 for more information.			
DBGTXx	0010000 <sub>2</sub>	Used by the Intel XScale $^{\ensuremath{\mathbb{R}}}$ core debugger during software debug. See core EAS Ch 9 for more information.			
TRACE	0010001 <sub>2</sub>	Selects ETM registers.			
IDCODE	1111110 <sub>2</sub>	IDCODE is used in conjunction with the device identification register. When selected, IDCODE parallel-loads the hard-wired identification code (32 bits) into the identification register on the rising edge of <b>TCK</b> following entry into the Capture-DR state. The instruction selects only the identification register for connection between <b>TDI</b> and <b>TDO</b> in the Shift-DR state for serial access. NOTE: The device identification register is not altered by data being shifted in on <b>TDI</b> .			
BYPASS	1111111 <sub>2</sub>	BYPASS selects the bypass register between TDI and TDO while in Shift-DR state, effectively bypassing the processor's test logic.'0' is captured in the Capture-DR state. While this instruction is in effect, all other test data registers have no effect on the operation of the system.			



# 21.3.4.3 Boundary-Scan Register

The boundary-scan register is a set of serial-shiftable register cells, connected between each of the system pins and the on-chip system logic (power, ground and TAP pins excluded). This forms a single shift-register between TDI and TDO of all the system pins.

This is the most extensive, complex register in the test circuitry. This register allows testing of circuitry external to the component (e.g. board interconnect) in addition to device system logic. It permits the system signals (into and out of the system logic) to be sampled and examined without causing interference with the normal operation of the system logic. Further definition, rules, and specifics of the boundary-scan register can be found in the IEEE Std, 1149.1-1990, Chapter 10.

# 21.3.4.4 Bypass Register

The bypass register is a single-bit, serial-shift register that connects **TDI** and **TDO** when the BYPASS instruction is in effect. This allows rapid movement of test data to and from other components on the board, since this register provides the shortest path between **TDI** and **TDO**. This path can be selected when no test operation is being performed. While the bypass register is selected, data is transferred from **TDI** to **TDO** without inversion.

# 21.3.4.5 Device Identification Register

The device identification (ID) register is a 32-bit register used for storing the manufacturer identification, part number, and the version of the processor. It is a dedicated part of the test logic and is not usable in system functionality.

The identification register is selected only by the IDCODE instruction. When the Test-Logic-Reset state of the TAP controller is entered, IDCODE instruction is automatically loaded into the instruction register. The generic register format is discussed in IEEE 1149.1 Standard, chapter 11.

#### Figure 157. Device ID Register



 Table 529.
 Device ID Register Field Definitions

	Field		Definition						
	Version	Indicates stepping changes.							
	Part Number	part number	part number						
	Manufacturer ID	Manufacturer ID assigned by IEEE. (0000001001h indicates Intel)							
	1	IEEE 1149.1 requirement							
Table 530.	Intel XScale <sup>®</sup> Core Device ID Register Settings								
	80333 0x09278013 0000 1001 0010 0111 1000 0000 0001 0011								

# Table 531. TLU Device ID Register Settings

80333	0x09278013	0000 1001 0010 0111 1000 0000 0001 0011

The TLU and the Intel XScale<sup>®</sup> core both have the same device ID number.



**This Page Left Intentionally Blank** 



This Page Left Intentionally Blank



This Page Left Intentionally Blank

