



Intel XScale® Core

Developer's Manual

January, 2004



Information in this document is provided in connection with Intel® 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® internal code names are subject to change.

THIS SPECIFICATION, THE [Intel XScale® Core Developer's Manual](#) 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, 2004

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, IntelDX2, IntelDX4, IntelSX2, Intel ChatPad, Intel Create&Share, Intel Dot.Station, Intel GigaBlade, Intel InBusiness, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetStructure, Intel Play, Intel Play logo, Intel Pocket Concert, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel TeamStation, 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.

The ARM* and ARM Powered logo marks (the ARM marks) are trademarks of ARM, Ltd., and Intel uses these marks under license from ARM, Ltd.

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

Contents

1	Introduction.....	13
1.1	About This Document	13
1.1.1	How to Read This Document.....	13
1.1.2	Other Relevant Documents	14
1.2	High-Level Overview of the Intel XScale® Core.....	15
1.2.1	ARM Compatibility	15
1.2.2	Features.....	16
1.2.2.1	Multiply/Accumulate (MAC).....	16
1.2.2.2	Memory Management	17
1.2.2.3	Instruction Cache	17
1.2.2.4	Branch Target Buffer.....	17
1.2.2.5	Data Cache	17
1.2.2.6	Performance Monitoring.....	18
1.2.2.7	Power Management.....	18
1.2.2.8	Debug	18
1.2.2.9	JTAG.....	18
1.3	Terminology and Conventions	19
1.3.1	Number Representation.....	19
1.3.2	Terminology and Acronyms	19
2	Programming Model	21
2.1	ARM Architecture Compatibility	21
2.2	ARM Architecture Implementation Options.....	21
2.2.1	Big Endian versus Little Endian	21
2.2.2	26-Bit Architecture	21
2.2.3	Thumb.....	21
2.2.4	ARM DSP-Enhanced Instruction Set.....	22
2.2.5	Base Register Update.....	22
2.3	Extensions to ARM Architecture	23
2.3.1	DSP Coprocessor 0 (CP0).....	23
2.3.1.1	Multiply With Internal Accumulate Format	24
2.3.1.2	Internal Accumulator Access Format	27
2.3.2	New Page Attributes	29
2.3.3	Additions to CP15 Functionality	31
2.3.4	Event Architecture	32
2.3.4.1	Exception Summary.....	32
2.3.4.2	Event Priority.....	32
2.3.4.3	Prefetch Aborts	33
2.3.4.4	Data Aborts	34
2.3.4.5	Events from Preload Instructions	35
2.3.4.6	Debug Events	36
3	Memory Management.....	37
3.1	Overview.....	37
3.2	Architecture Model.....	38
3.2.1	Version 4 vs. Version 5.....	38
3.2.2	Memory Attributes.....	38
3.2.2.1	Page (P) Attribute Bit	38

3.2.2.2	Cacheable (C), Bufferable (B), and eXtension (X) Bits.....	38
3.2.2.3	Instruction Cache.....	38
3.2.2.4	Data Cache and Write Buffer.....	39
3.2.2.5	Details on Data Cache and Write Buffer Behavior.....	40
3.2.2.6	Memory Operation Ordering.....	40
3.2.3	Exceptions.....	40
3.3	Interaction of the MMU, Instruction Cache, and Data Cache.....	41
3.4	Control.....	42
3.4.1	Invalidate (Flush) Operation.....	42
3.4.2	Enabling/Disabling.....	42
3.4.3	Locking Entries.....	43
3.4.4	Round-Robin Replacement Algorithm.....	45
4	Instruction Cache.....	47
4.1	Overview.....	47
4.2	Operation.....	48
4.2.1	Operation When Instruction Cache is Enabled.....	48
4.2.2	Operation When The Instruction Cache Is Disabled.....	48
4.2.3	Fetch Policy.....	49
4.2.4	Round-Robin Replacement Algorithm.....	49
4.2.5	Parity Protection.....	50
4.2.6	Instruction Fetch Latency.....	51
4.2.7	Instruction Cache Coherency.....	51
4.3	Instruction Cache Control.....	52
4.3.1	Instruction Cache State at RESET.....	52
4.3.2	Enabling/Disabling.....	52
4.3.3	Invalidating the Instruction Cache.....	53
4.3.4	Locking Instructions in the Instruction Cache.....	54
4.3.5	Unlocking Instructions in the Instruction Cache.....	55
5	Branch Target Buffer.....	57
5.1	Branch Target Buffer (BTB) Operation.....	57
5.1.1	Reset.....	58
5.1.2	Update Policy.....	58
5.2	BTB Control.....	59
5.2.1	Disabling/Enabling.....	59
5.2.2	Invalidation.....	59
6	Data Cache.....	61
6.1	Overviews.....	61
6.1.1	Data Cache Overview.....	61
6.1.2	Mini-Data Cache Overview.....	63
6.1.3	Write Buffer and Fill Buffer Overview.....	64
6.2	Data Cache and Mini-Data Cache Operation.....	65
6.2.1	Operation When Caching is Enabled.....	65
6.2.2	Operation When Data Caching is Disabled.....	65
6.2.3	Cache Policies.....	65
6.2.3.1	Cacheability.....	65
6.2.3.2	Read Miss Policy.....	66
6.2.3.3	Write Miss Policy.....	67
6.2.3.4	Write-Back Versus Write-Through.....	67

6.2.4	Round-Robin Replacement Algorithm	68
6.2.5	Parity Protection	68
6.2.6	Atomic Accesses	68
6.3	Data Cache and Mini-Data Cache Control	69
6.3.1	Data Memory State After Reset	69
6.3.2	Enabling/Disabling	69
6.3.3	Invalidate and Clean Operations	69
6.3.3.1	Global Clean and Invalidate Operation	70
6.4	Re-configuring the Data Cache as Data RAM	71
6.5	Write Buffer/Fill Buffer Operation and Control	75
7	Configuration	77
7.1	Overview	77
7.2	CP15 Registers.....	80
7.2.1	Register 0: ID & Cache Type Registers	81
7.2.2	Register 1: Control & Auxiliary Control Registers	83
7.2.3	Register 2: Translation Table Base Register	85
7.2.4	Register 3: Domain Access Control Register.....	85
7.2.5	Register 4: Reserved	85
7.2.6	Register 5: Fault Status Register	86
7.2.7	Register 6: Fault address Register	86
7.2.8	Register 7: Cache Functions	87
7.2.9	Register 8: TLB Operations	89
7.2.10	Register 9: Cache Lock Down	90
7.2.11	Register 10: TLB Lock Down	91
7.2.12	Register 11-12: Reserved.....	91
7.2.13	Register 13: Process ID	91
7.2.13.1	The PID Register Affect On Addresses	92
7.2.14	Register 14: Breakpoint Registers	93
7.2.15	Register 15: Coprocessor Access Register	94
7.3	CP14 Registers.....	96
7.3.1	Performance Monitoring Registers	96
7.3.1.1	XSC1 Performance Monitoring Registers	96
7.3.1.2	XSC2 Performance Monitoring Registers	97
7.3.2	Clock and Power Management Registers.....	98
7.3.3	Software Debug Registers.....	99
8	Performance Monitoring	101
8.1	Overview	101
8.2	XSC1 Register Description (2 counter variant).....	102
8.2.1	Clock Counter (CCNT; CP14 - Register 1)	102
8.2.2	Performance Count Registers (PMN0 - PMN1; CP14 - Register 2 and 3, Respectively).....	103
8.2.3	Extending Count Duration Beyond 32 Bits	103
8.2.4	Performance Monitor Control Register (PMNC)	103
8.2.4.1	Managing PMNC.....	105
8.3	XSC2 Register Description (4 counter variant).....	106
8.3.1	Clock Counter (CCNT).....	106
8.3.2	Performance Count Registers (PMN0 - PMN3)	107
8.3.3	Performance Monitor Control Register (PMNC)	108
8.3.4	Interrupt Enable Register (INTEN).....	109

8.3.5	Overflow Flag Status Register (FLAG)	110
8.3.6	Event Select Register (EVTSEL)	111
8.3.7	Managing the Performance Monitor	112
8.4	Performance Monitoring Events	113
8.4.1	Instruction Cache Efficiency Mode	115
8.4.2	Data Cache Efficiency Mode	115
8.4.3	Instruction Fetch Latency Mode.....	115
8.4.4	Data/Bus Request Buffer Full Mode	116
8.4.5	Stall/Writeback Statistics	116
8.4.6	Instruction TLB Efficiency Mode	117
8.4.7	Data TLB Efficiency Mode	117
8.5	Multiple Performance Monitoring Run Statistics.....	118
8.6	Examples.....	119
8.6.1	XSC1 Example (2 counter variant)	119
8.6.2	XSC2 Example (4 counter variant)	120
9	Software Debug.....	121
9.1	Definitions.....	121
9.2	Debug Registers.....	121
9.3	Introduction.....	122
9.3.1	Halt Mode	122
9.3.2	Monitor Mode.....	122
9.4	Debug Control and Status Register (DCSR)	123
9.4.1	Global Enable Bit (GE)	124
9.4.2	Halt Mode Bit (H)	124
9.4.3	SOC Break (B).....	124
9.4.4	Vector Trap Bits (TF, TI, TD, TA, TS, TU, TR)	125
9.4.5	Sticky Abort Bit (SA)	125
9.4.6	Method of Entry Bits (MOE).....	125
9.4.7	Trace Buffer Mode Bit (M)	125
9.4.8	Trace Buffer Enable Bit (E).....	125
9.5	Debug Exceptions.....	126
9.5.1	Halt Mode	127
9.5.2	Monitor Mode.....	129
9.6	HW Breakpoint Resources	130
9.6.1	Instruction Breakpoints	130
9.6.2	Data Breakpoints	131
9.7	Software Breakpoints.....	133
9.8	Transmit/Receive Control Register (TXRXCTRL)	134
9.8.1	RX Register Ready Bit (RR)	135
9.8.2	Overflow Flag (OV).....	136
9.8.3	Download Flag (D).....	136
9.8.4	TX Register Ready Bit (TR)	137
9.8.5	Conditional Execution Using TXRXCTRL.....	137
9.9	Transmit Register (TX)	138
9.10	Receive Register (RX).....	138
9.11	Debug JTAG Access	139
9.11.1	SELDCSR JTAG Register	139
9.11.1.1	hold_reset	140
9.11.1.2	ext_dbg_break	140

9.11.1.3	DCSR (DBG_SR[34:3]).....	140
9.11.2	DBGTX JTAG Register.....	141
9.11.2.1	DBG_SR[0].....	141
9.11.2.2	TX (DBG_SR[34:3]).....	141
9.11.3	DBGRX JTAG Register.....	142
9.11.3.1	RX Write Logic.....	143
9.11.3.2	DBG_SR[0].....	143
9.11.3.3	flush_rr.....	143
9.11.3.4	hs_download.....	143
9.11.3.5	RX (DBG_SR[34:3]).....	143
9.11.3.6	rx_valid.....	144
9.12	Trace Buffer.....	145
9.12.1	Trace Buffer Registers.....	145
9.12.1.1	Checkpoint Registers.....	146
9.12.1.2	Trace Buffer Register (TBREG).....	147
9.13	Trace Buffer Entries.....	148
9.13.1	Message Byte.....	148
9.13.1.1	Exception Message Byte.....	149
9.13.1.2	Non-exception Message Byte.....	150
9.13.1.3	Address Bytes.....	151
9.13.2	Trace Buffer Usage.....	152
9.14	Downloading Code in the Instruction Cache.....	154
9.14.1	Mini Instruction Cache Overview.....	154
9.14.2	LDIC JTAG Command.....	155
9.14.3	LDIC JTAG Data Register.....	155
9.14.4	LDIC Cache Functions.....	156
9.14.5	Loading Instruction Cache During Reset.....	158
9.14.6	Dynamically Loading Instruction Cache After Reset.....	160
9.14.6.1	Dynamic Download Synchronization Code.....	162
10	Performance Considerations.....	163
10.1	Interrupt Latency.....	163
10.2	Branch Prediction.....	164
10.3	Addressing Modes.....	164
10.4	Instruction Latencies.....	165
10.4.1	Performance Terms.....	165
10.4.2	Branch Instruction Timings.....	167
10.4.3	Data Processing Instruction Timings.....	167
10.4.4	Multiply Instruction Timings.....	168
10.4.5	Saturated Arithmetic Instructions.....	170
10.4.6	Status Register Access Instructions.....	170
10.4.7	Load/Store Instructions.....	171
10.4.8	Semaphore Instructions.....	171
10.4.9	Coprocessor Instructions.....	172
10.4.10	Miscellaneous Instruction Timing.....	172
10.4.11	Thumb Instructions.....	173
A	Optimization Guide.....	175
A.1	Introduction.....	175
A.1.1	About This Guide.....	175
A.2	The Intel XScale® Core Pipeline.....	176

A.2.1	General Pipeline Characteristics	176
A.2.1.1.	Number of Pipeline Stages	176
A.2.1.2.	The Intel XScale® Core Pipeline Organization	177
A.2.1.3.	Out Of Order Completion	178
A.2.1.4.	Register Scoreboarding	178
A.2.1.5.	Use of Bypassing	178
A.2.2	Instruction Flow Through the Pipeline	179
A.2.2.1.	ARM* V5TE Instruction Execution	179
A.2.2.2.	Pipeline Stalls	179
A.2.3	Main Execution Pipeline	180
A.2.3.1.	F1 / F2 (Instruction Fetch) Pipestages	180
A.2.3.2.	ID (Instruction Decode) Pipestage	180
A.2.3.3.	RF (Register File / Shifter) Pipestage	181
A.2.3.4.	X1 (Execute) Pipestages	181
A.2.3.5.	X2 (Execute 2) Pipestage	181
A.2.3.6.	WB (write-back)	181
A.2.4	Memory Pipeline	182
A.2.4.1.	D1 and D2 Pipestage	182
A.2.5	Multiply/Multiply Accumulate (MAC) Pipeline	182
A.2.5.1.	Behavioral Description	182
A.3	Basic Optimizations	183
A.3.1	Conditional Instructions	183
A.3.1.1.	Optimizing Condition Checks	183
A.3.1.2.	Optimizing Branches	184
A.3.1.3.	Optimizing Complex Expressions	186
A.3.2	Bit Field Manipulation	187
A.3.3	Optimizing the Use of Immediate Values	188
A.3.4	Optimizing Integer Multiply and Divide	189
A.3.5	Effective Use of Addressing Modes	190
A.4	Cache and Prefetch Optimizations	191
A.4.1	Instruction Cache	191
A.4.1.1.	Cache Miss Cost	191
A.4.1.2.	Round-Robin Replacement Cache Policy	191
A.4.1.3.	Code Placement to Reduce Cache Misses	191
A.4.1.4.	Locking Code into the Instruction Cache	192
A.4.2	Data and Mini Cache	193
A.4.2.1.	Non Cacheable Regions	193
A.4.2.2.	Write-through and Write-back Cached Memory Regions	193
A.4.2.3.	Read Allocate and Read-write Allocate Memory Regions	194
A.4.2.4.	Creating On-chip RAM	194
A.4.2.5.	Mini-data Cache	195
A.4.2.6.	Data Alignment	196
A.4.2.7.	Literal Pools	197
A.4.3	Cache Considerations	198
A.4.3.1.	Cache Conflicts, Pollution and Pressure	198
A.4.3.2.	Memory Page Thrashing	198
A.4.4	Prefetch Considerations	199
A.4.4.1.	Prefetch Distances	199
A.4.4.2.	Prefetch Loop Scheduling	199
A.4.4.3.	Prefetch Loop Limitations	199
A.4.4.4.	Compute vs. Data Bus Bound	199
A.4.4.5.	Low Number of Iterations	200

	A.4.4.6. Bandwidth Limitations	200
	A.4.4.7. Cache Memory Considerations	201
	A.4.4.8. Cache Blocking	203
	A.4.4.9. Prefetch Unrolling	203
	A.4.4.10. Pointer Prefetch	204
	A.4.4.11. Loop Interchange	205
	A.4.4.12. Loop Fusion	205
	A.4.4.13. Prefetch to Reduce Register Pressure	206
A.5	Instruction Scheduling	207
	A.5.1 Scheduling Loads	207
	A.5.1.1. Scheduling Load and Store Double (LDRD/STRD)	210
	A.5.1.2. Scheduling Load and Store Multiple (LDM/STM)	211
	A.5.2 Scheduling Data Processing Instructions	212
	A.5.3 Scheduling Multiply Instructions	213
	A.5.4 Scheduling SWP and SWPB Instructions	214
	A.5.5 Scheduling the MRA and MAR Instructions (MRRC/MCRR)	215
	A.5.6 Scheduling the MIA and MIAPH Instructions	216
	A.5.7 Scheduling MRS and MSR Instructions	217
	A.5.8 Scheduling CP15 Coprocessor Instructions	217
A.6	Optimizing C Libraries	218
A.7	Optimizations for Size	218
	A.7.1 Space/Performance Trade Off	218
	A.7.1.1. Multiple Word Load and Store	218
	A.7.1.2. Use of Conditional Instructions	218
	A.7.1.3. Use of PLD Instructions	218
B	Test Features	219
	B.1 Overview	219

Figures

1-1	Architecture Features	16
3-1	Example of Locked Entries in TLB	45
4-1	Instruction Cache Organization	47
4-2	Locked Line Effect on Round Robin Replacement.....	54
5-1	BTB Entry	57
5-2	Branch History	58
6-1	Data Cache Organization	62
6-2	Mini-Data Cache Organization	63
6-3	Locked Line Effect on Round Robin Replacement.....	74
9-1	SELDCSR.....	139
9-2	DBGTX	141
9-3	DBGRX.....	142
9-4	Message Byte Formats.....	148
9-5	Indirect Branch Entry Address Byte Organization	151
9-6	High Level View of Trace Buffer	152
9-7	LDIC JTAG Data Register Hardware.....	155
9-8	Format of LDIC Cache Functions	157
9-9	Code Download During a Cold Reset For Debug.....	158
9-10	Downloading Code in IC During Program Execution.....	160
A-1	The Intel XScale® Core RISC Superpipeline.....	177

Tables

2-1	Multiply with Internal Accumulate Format	24
2-2	MIA{<cond>} acc0, Rm, Rs	25
2-3	MIAPH{<cond>} acc0, Rm, Rs	25
2-4	MIAXy{<cond>} acc0, Rm, Rs	26
2-5	Internal Accumulator Access Format	27
2-6	MAR{<cond>} acc0, RdLo, RdHi	28
2-7	MRA{<cond>} RdLo, RdHi, acc0	28
2-9	Second-level Descriptors for Coarse Page Table	30
2-10	Second-level Descriptors for Fine Page Table	30
2-8	First-level Descriptors	30
2-11	Exception Summary	32
2-12	Event Priority	32
2-13	Encoding of Fault Status for Prefetch Aborts	33
2-14	Encoding of Fault Status for Data Aborts	34
3-1	Data Cache and Buffer Behavior when X = 0	39
3-2	Data Cache and Buffer Behavior when X = 1	39
3-3	Memory Operations that Impose a Fence	40
3-4	Valid MMU & Data/mini-data Cache Combinations	41
7-1	MRC/MCR Format	78
7-2	LDC/STC Format when Accessing CP14	79
7-3	CP15 Registers	80
7-4	ID Register	81
7-5	Cache Type Register	82
7-6	ARM* Control Register	83
7-7	Auxiliary Control Register	84
7-8	Translation Table Base Register	85
7-9	Domain Access Control Register	85
7-10	Fault Status Register	86
7-11	Fault Address Register	86
7-12	Cache Functions	87
7-13	TLB Functions	89
7-14	Cache Lockdown Functions	90
7-15	Data Cache Lock Register	90
7-16	TLB Lockdown Functions	91
7-17	Accessing Process ID	91
7-18	Process ID Register	91
7-19	Accessing the Debug Registers	93
7-20	Coprocessor Access Register	95
7-21	Accessing the XSC1 Performance Monitoring Registers	96
7-22	Accessing the XSC2 Performance Monitoring Registers	97
7-23	PWRMODE Register	98
7-24	Clock and Power Management	98
7-25	CCCLKCFG Register	98
7-26	Accessing the Debug Registers	99
8-1	XSC1 Performance Monitoring Registers	102
8-2	Clock Count Register (CCNT)	102
8-3	Performance Monitor Count Register (PMN0 and PMN1)	103
8-4	Performance Monitor Control Register (CP14, register 0)	104
8-5	Performance Monitoring Registers	106

8-6	Clock Count Register (CCNT)	106
8-7	Performance Monitor Count Register (PMN0 - PMN3)	107
8-8	Performance Monitor Control Register	108
8-9	Interrupt Enable Register.....	109
8-10	Overflow Flag Status Register	110
8-11	Event Select Register	111
8-12	Performance Monitoring Events	113
8-13	Some Common Uses of the PMU	114
9-1	Debug Control and Status Register (DCSR)	123
9-2	Event Priority	126
9-3	Halt Mode R14_DBG Updating	127
9-4	Monitor Mode R14_DBG Updating.....	129
9-5	Instruction Breakpoint Address and Control Register (IBCRx).....	130
9-6	Data Breakpoint Register (DBRx).....	131
9-7	Data Breakpoint Controls Register (DBCON).....	131
9-8	TX RX Control Register (TXRXCTRL).....	134
9-9	Normal RX Handshaking	135
9-10	High-Speed Download Handshaking States	135
9-11	TX Handshaking	137
9-12	TXRXCTRL Mnemonic Extensions	137
9-13	TX Register.....	138
9-14	RX Register	138
9-15	CP 14 Trace Buffer Register Summary	145
9-16	Checkpoint Register (CHKPTx).....	146
9-17	TBREG Format.....	147
9-18	Message Byte Formats.....	148
9-19	LDIC Cache Functions	156
9-20	Steps For Loading Mini Instruction Cache During Reset.....	159
9-21	Steps For Dynamically Loading the Mini Instruction Cache	161
10-1	Branch Latency Penalty.....	164
10-2	Latency Example	166
10-3	Branch Instruction Timings (Those predicted by the BTB)	167
10-4	Branch Instruction Timings (Those not predicted by the BTB)	167
10-5	Data Processing Instruction Timings	167
10-6	Multiply Instruction Timings	168
10-7	Multiply Implicit Accumulate Instruction Timings	169
10-8	Implicit Accumulator Access Instruction Timings.....	169
10-9	Saturated Data Processing Instruction Timings	170
10-10	Status Register Access Instruction Timings.....	170
10-11	Load and Store Instruction Timings	171
10-12	Load and Store Multiple Instruction Timings.....	171
10-13	Semaphore Instruction Timings	171
10-14	CP15 Register Access Instruction Timings.....	172
10-15	CP14 Register Access Instruction Timings.....	172
10-16	Exception-Generating Instruction Timings	172
10-17	Count Leading Zeros Instruction Timings	172
A-1	Pipelines and Pipe stages	177

Introduction

1

1.1 About This Document

This document is the authoritative and definitive reference for the external architecture of the Intel XScale[®] core¹.

This document describes two variants of the Intel XScale[®] core that differ only in the performance monitoring and the size of the JTAG instruction register. Software can detect which variant it is running on by examining the CoreGen field of Coprocessor 15, ID Register (bits 15:13). (See [Table 7-4, “ID Register” on page 7-81](#) for more details.) A CoreGen value of 0x1 is referred to as XSC1 and a value of 0x2 is referred to as XSC2.

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, and pin-outs does not imply a commitment to implement them.

1.1.1 How to Read This Document

It is necessary to be familiar with the ARM Version 5TE Architecture in order to understand some aspects of this document.

Each chapter in this document focuses on a specific architectural feature of the Intel XScale[®] core.

- [Chapter 2, “Programming Model”](#)
- [Chapter 3, “Memory Management”](#)
- [Chapter 4, “Instruction Cache”](#)
- [Chapter 5, “Branch Target Buffer”](#)
- [Chapter 6, “Data Cache”](#)
- [Chapter 7, “Configuration”](#)
- [Chapter 8, “Performance Monitoring”](#)
- [Chapter 9, “Software Debug”](#)
- [Chapter 10, “Performance Considerations”](#)

Several appendices are also present:

- [Appendix A, “Optimization Guide”](#) covers instruction scheduling techniques.
- [Appendix B, “Test Features”](#) describes the JTAG unit.

Note: All the “buzz words” and acronyms found throughout this document are captured in [Section 1.3.2, “Terminology and Acronyms” on page 1-19](#), located at the end of this chapter.

1. ARM* architecture compliant.

1.1.2 Other Relevant Documents

- *ARM Architecture Version 5TE Specification* Document Number: ARM DDI 0100E
This document describes Version 5TE of the ARM Architecture which includes Thumb ISA and ARM DSP-Enhanced ISA. (ISBN 0 201 737191)
- *StrongARM SA-1100 Microprocessor Developer's Manual*, Intel Order # 278105
- *StrongARM SA-110 Microprocessor Technical Reference Manual*, Intel Order #278104

1.2 High-Level Overview of the Intel XScale[®] Core

The Intel XScale[®] core is an ARM V5TE compliant microprocessor. It has been designed for high performance and low-power; leading the industry in mW/MIPs. The core is not intended to be delivered as a stand alone product but as a building block for an ASSP (Application Specific Standard Product) with embedded markets such as handheld devices, networking, storage, remote access servers, etc.

The Intel XScale[®] core incorporates an extensive list of architecture features that allows it to achieve high performance. This rich feature set allows programmers to select the appropriate features that obtains the best performance for their application. Many of the architectural features added to the Intel XScale[®] 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 memory bus external to the core.

The Intel XScale[®] core has been equipped to efficiently handle audio processing through the support of 16-bit data types and 16-bit operations. These audio coding enhancements center around multiply and accumulate operations which accelerate many of the audio filter operations.

1.2.1 ARM Compatibility

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

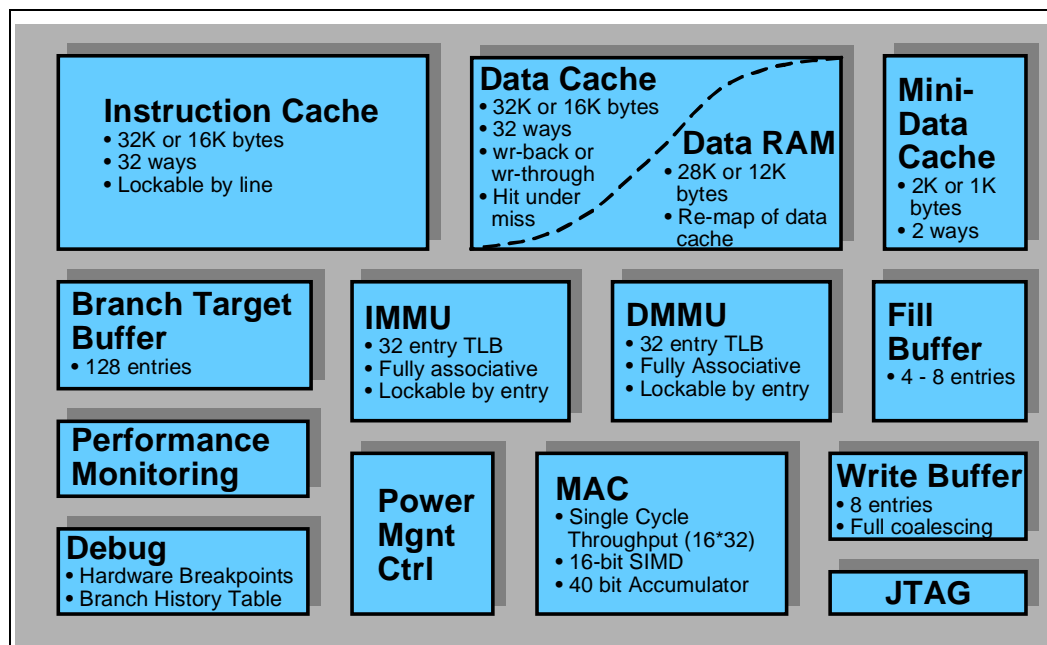
The Intel XScale[®] core provides the Thumb instruction set (ARM V5T) and the ARM V5E DSP extensions.

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

1.2.2 Features

Figure 1-1 shows the major functional blocks of the Intel XScale® core. The following sections give a brief, high-level overview of these blocks.

Figure 1-1. Architecture Features



1.2.2.1 Multiply/Accumulate (MAC)

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

See [Section 2.3, “Extensions to ARM Architecture”](#) on page 2-23 for more details.

1.2.2.2 Memory Management

The Intel XScale® 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

Chapter 3, “Memory Management” discusses this in more detail.

1.2.2.3 Instruction Cache

The Intel XScale® core comes with either a 16 K or 32 K byte instruction cache. The size is determined by the ASSP. The instruction cache is 32-way set associative and has 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.

Chapter 4, “Instruction Cache” discusses this in more detail.

1.2.2.4 Branch Target Buffer

The Intel XScale® 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. See Chapter 5, “Branch Target Buffer” for more details.

1.2.2.5 Data Cache

The Intel XScale® core comes with either a 16 K or 32 K byte data cache. The size is determined by the ASSP. Besides the main data cache, a mini-data cache is provided whose size is 1/16th the main data cache. So a 32 K, 16 K byte main data cache would have a 2 K, 1 K byte mini-data cache respectively. The main data cache is 32-way set associative and the mini-data cache is 2-way set associative. 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.

Chapter 6, “Data Cache” discusses all this in more detail.

The Intel XScale® 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. See Section 6.4, “Re-configuring the Data Cache as Data RAM” on page 6-71 for more information on this.

1.2.2.6 Performance Monitoring

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

Chapter 8, “Performance Monitoring” discusses this in more detail.

1.2.2.7 Power Management

The Intel XScale® core incorporates a power and clock management unit that can assist ASSPs in controlling their clocking and managing their power. These features are described in [Section 7.3, “CP14 Registers”](#) on page 7-96.

1.2.2.8 Debug

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

Chapter 9, “Software Debug” discusses this in more detail.

1.2.2.9 JTAG

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

[Appendix B](#) discusses this in more detail.

1.3 Terminology and Conventions

1.3.1 Number Representation

All numbers in this document can be assumed to be base 10 unless designated otherwise. In text and pseudo code descriptions, hexadecimal numbers have a prefix of 0x and binary numbers have a prefix of 0b. For example, 107 would be represented as 0x6B in hexadecimal and 0b1101011 in binary.

1.3.2 Terminology and Acronyms

ASSP	Application Specific Standard Product
Assert	This term refers to the logically active value of a signal or bit.
BTB	Branch Target Buffer
Clean	<p>A clean operation updates external memory with the contents of the specified line in the data/mini-data cache if any of the dirty bits are set and the line is valid. There are two dirty bits associated with each line in the cache so only the portion that is dirty will get written back to external memory.</p> <p>After this operation, the line is still valid and both dirty bits are deasserted.</p>
Coalescing	Coalescing means bringing together a new store operation with an existing store operation already resident in the write buffer. The new store is placed in the same write buffer entry as an existing store when the address of the new store falls in the 4 word aligned address of the existing entry. This includes, in PCI terminology, write merging, write collapsing, and write combining.
Deassert	This term refers to the logically inactive value of a signal or bit.
Flush	A flush operation invalidates the location(s) in the cache by deasserting the valid bit. Individual entries (lines) may be flushed or the entire cache may be flushed with one command. Once an entry is flushed in the cache it can no longer be used by the program.
XSC1	XSC1 refers to a variant of the Intel XScale [®] core denoted by a CoreGen (Coprocesor 15, ID Register) value of 0x1. This variant has a 2 counter performance monitor and a 5-bit JTAG instruction register. See Table 7-4, "ID Register" on page 7-81 for more details.
XSC2	XSC2 refers to a variant of the Intel XScale [®] core denoted by a CoreGen (Coprocesor 15, ID Register) value of 0x2. This variant has a 4 counter performance monitor and a 7-bit JTAG instruction register. See Table 7-4, "ID Register" on page 7-81 for more details.
Reserved	A <i>reserved</i> field is a field that may be used by an implementation. If the initial value of a reserved field is supplied by software, this value must be zero. Software should not modify reserved fields or depend on any values in reserved fields.



This Page Intentionally Left Blank

Programming Model

2

This chapter describes the programming model of the Intel XScale[®] core, namely the implementation options and extensions to the ARM Version 5TE architecture.

2.1 ARM Architecture Compatibility

The Intel XScale[®] core implements the integer instruction set architecture specified in ARM V5TE. T refers to the Thumb instruction set and E refers to the DSP-Enhanced instruction set.

ARM V5TE introduces a few more architecture features over ARM V4, specifically the addition of tiny pages (1 Kbyte), a new instruction (**CLZ**) that counts the leading zeroes in a data value, enhanced ARM-Thumb transfer instructions and a modification of the system control coprocessor, CP15.

2.2 ARM Architecture Implementation Options

2.2.1 Big Endian versus Little Endian

The Intel XScale[®] core supports both big and little endian data representation. The B-bit of the Control Register (Coprocessor 15, register 1, bit 7) selects big and little endian mode. To run in big endian mode, the B bit must be set before attempting any sub-word accesses to memory, or undefined results will occur. Note that this bit takes effect even if the MMU is disabled.

2.2.2 26-Bit Architecture

The Intel XScale[®] core does not support 26-bit architecture.

2.2.3 Thumb

The Intel XScale[®] core supports the Thumb instruction set.

2.2.4 ARM DSP-Enhanced Instruction Set

The Intel XScale® core implements the ARM DSP-enhanced instruction set which is a set of instructions that boost the performance of signal processing applications. There are new multiply instructions that operate on 16-bit data values and new saturation instructions. Some of the new instructions are:

- SMLAxy $32 \leq 16 \times 16 + 32$
- SMLAWy $32 \leq 32 \times 16 + 32$
- SMLALxy $64 \leq 16 \times 16 + 64$
- SMULxy $32 \leq 16 \times 16$
- SMULWy $32 \leq 32 \times 16$
- QADD adds two registers and saturates the result if an overflow occurred
- QDADD doubles and saturates one of the input registers then add and saturate
- QSUB subtracts two registers and saturates the result if an overflow occurred
- QDSUB doubles and saturates one of the input registers then subtract and saturate

The Intel XScale® core also implements LDRD, STRD and PLD instructions with the following implementation notes:

- PLD is interpreted as a read operation by the MMU and is ignored by the data breakpoint unit (i.e., PLD will never generate data breakpoint events).
- PLD to a non-cacheable page performs no action. Also, if the targeted cache line is already resident, this instruction has no affect.
- Both LDRD and STRD instructions will generate an alignment exception when the address bits [2:0] = 0b100.

MCRR and MRRC are only supported on the Intel XScale® core when directed to coprocessor 0 and are used to access the internal accumulator. See [Section 2.3.1.2](#) for more information. Access to coprocessors 15 and 14 generate an undefined instruction exception. Refer to the Intel XScale® core implementation option section of the ASSP architecture specification for the behavior when accessing all other coprocessors.

2.2.5 Base Register Update

If a data abort is signalled on a memory instruction that specifies writeback, the contents of the base register will not be updated. This holds for all load and store instructions. This behavior matches that of the first generation StrongARM processor and is referred to in the ARM V5TE architecture as the *Base Restored Abort Model*.

2.3 Extensions to ARM Architecture

The Intel XScale[®] core made a few extensions to the ARM Version 5TE architecture to meet the needs of various markets and design requirements. The following is a list of the extensions which are discussed in the next sections.

- A DSP coprocessor (CP0) has been added that contains a 40-bit accumulator and eight new instructions.
- New page attributes were added to the page table descriptors. The C and B page attribute encoding was extended by one more bit to allow for more encodings: write allocate and mini-data cache. An ASSP definable attribute (P bit) was also added.
- Additional functionality has been added to coprocessor 15. Coprocessor 14 was also created.
- Enhancements were made to the Event Architecture, which include instruction cache and data cache parity error exceptions, breakpoint events, and imprecise external data aborts.

2.3.1 DSP Coprocessor 0 (CP0)

The Intel XScale[®] core adds a DSP coprocessor to the architecture for the purpose of increasing the performance and the precision of audio processing algorithms. This coprocessor contains a 40-bit accumulator and 8 new instructions.

Note: Products using the Intel XScale[®] core may extend the definition of CP0. If this is the case, a complete definition can be found in the Intel XScale[®] core implementation option section of the ASSP architecture specification. For this very reason, software should not rely on behavior that is specific to the 40-bit length of the accumulator, since the length may be extended.

The 40-bit accumulator is referenced by several new instructions that were added to the architecture; **MIA**, **MIAPH** and **MIAxy** are multiply/accumulate instructions that reference the 40-bit accumulator instead of a register specified accumulator. **MAR** and **MRA** provide the ability to read and write the 40-bit accumulator.

Access to CP0 is always allowed in all processor modes when bit 0 of the Coprocessor Access Register is set. Any access to CP0 when this bit is clear will cause an undefined exception. (See [Section 7.2.15, "Register 15: Coprocessor Access Register" on page 7-94](#) for more details).

Note: Only privileged software can set this bit in the Coprocessor Access Register.

The 40-bit accumulator will need to be saved on a context switch if multiple processes are using it.

Two new instruction formats were added for coprocessor 0: Multiply with Internal Accumulate Format and Internal Accumulate Access Format. The formats and instructions are described next.

2.3.1.1 Multiply With Internal Accumulate Format

A new multiply format has been created to define operations on 40-bit accumulators. Table 2-1, “Multiply with Internal Accumulate Format” on page 2-24 shows the layout of the new format. The opcode for this format lies within the coprocessor register transfer instruction type. These instructions have their own syntax.

Table 2-1. Multiply with Internal Accumulate Format

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																											
cond				1	1	1	0	0	0	1	0	opcode_3				Rs		0	0	0	0	acc			1	Rm	
Bits	Description															Notes											
31:28	cond - ARM condition codes															-											
19:16	opcode_3 - specifies the type of multiply with internal accumulate															The Intel XScale® core defines the following: 0b0000 = MIA 0b1000 = MIAPH 0b1100 = MIABB 0b1101 = MIABT 0b1110 = MIATB 0b1111 = MIATT The effect of all other encodings are unpredictable.											
15:12	Rs - Multiplier																										
7:5	acc - select 1 of 8 accumulators															The Intel XScale® core only implements acc0; access to any other acc has an unpredictable effect.											
3:0	Rm - Multiplicand															-											

Two new fields were created for this format, *acc* and *opcode_3*. The *acc* field specifies 1 of 8 internal accumulators to operate on and *opcode_3* defines the operation for this format. The Intel XScale® core defines a single 40-bit accumulator referred to as *acc0*; future implementations may define multiple internal accumulators. The Intel XScale® core uses *opcode_3* to define six instructions, **MIA**, **MIAPH**, **MIABB**, **MIABT**, **MIATB** and **MIATT**.

Table 2-2. MIA{<cond>} acc0, Rm, Rs

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	
cond			1	1	1	0	0	0	1	0	Rs			0	0	0	0	0	0	0	1	Rm										
<p>Operation: if ConditionPassed(<cond>) then $acc0 = (Rm[31:0] * Rs[31:0])[39:0] + acc0[39:0]$</p> <p>Exceptions: none</p> <p>Qualifiers Condition Code No condition code flags are updated</p> <p>Notes: Early termination is supported. Instruction timings can be found in Section 10.4.4, "Multiply Instruction Timings" on page 10-168. Specifying R15 for register Rs or Rm has unpredictable results. acc0 is defined to be 0b000 on the core.</p>																																

The **MIA** instruction operates similarly to **MLA** except that the 40-bit accumulator is used. **MIA** multiplies the signed value in register Rs (multiplier) by the signed value in register Rm (multiplicand) and then adds the result to the 40-bit accumulator (acc0).

MIA does not support unsigned multiplication; all values in Rs and Rm will be interpreted as signed data values. **MIA** is useful for operating on signed 16-bit data that was loaded into a general purpose register by **LDRSH**.

The instruction is only executed if the condition specified in the instruction matches the condition code status.

Table 2-3. MIAPH{<cond>} acc0, Rm, Rs

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	
cond			1	1	1	0	0	0	1	0	Rs			0	0	0	0	0	0	0	1	Rm										
<p>Operation: if ConditionPassed(<cond>) then $acc0 = sign_extend(Rm[31:16] * Rs[31:16]) + sign_extend(Rm[15:0] * Rs[15:0]) + acc0[39:0]$</p> <p>Exceptions: none</p> <p>Qualifiers Condition Code S bit is always cleared; no condition code flags are updated</p> <p>Notes: Instruction timings can be found in Section 10.4.4, "Multiply Instruction Timings" on page 10-168. Specifying R15 for register Rs or Rm has unpredictable results. acc0 is defined to be 0b000 on the core.</p>																																

The **MIAPH** instruction performs two 16-bit signed multiplies on packed half word data and accumulates these to a single 40-bit accumulator. The first signed multiplication is performed on the lower 16 bits of the value in register Rs with the lower 16 bits of the value in register Rm. The second signed multiplication is performed on the upper 16 bits of the value in register Rs with the upper 16 bits of the value in register Rm. Both signed 32-bit products are sign extended and then added to the value in the 40-bit accumulator (acc0).

The instruction is only executed if the condition specified in the instruction matches the condition code status.

Table 2-4. **MIAXy{<cond>} acc0, Rm, Rs**

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
cond		1	1	1	0	0	0	1	0	1	1	x	y	Rs				0	0	0	0	0	0	0	1	Rm					
<p>Operation: if ConditionPassed(<cond>) then if (bit[17] == 0) <operand1> = Rm[15:0] else <operand1> = Rm[31:16] if (bit[16] == 0) <operand2> = Rs[15:0] else <operand2> = Rs[31:16] acc0[39:0] = sign_extend(<operand1> * <operand2>) + acc0[39:0]</p> <p>Exceptions: none</p> <p>Qualifiers Condition Code S bit is always cleared; no condition code flags are updated</p> <p>Notes: Instruction timings can be found in Section 10.4.4, "Multiply Instruction Timings" on page 10-168. Specifying R15 for register Rs or Rm has unpredictable results. acc0 is defined to be 0b000 on the core.</p>																															

The **MIAXy** instruction performs one 16-bit signed multiply and accumulates these to a single 40-bit accumulator. **x** refers to either the upper half or lower half of register Rm (multiplicand) and **y** refers to the upper or lower half of Rs (multiplier). A value of 0x1 will select bits [31:16] of the register which is specified in the mnemonic as T (for top). A value of 0x0 will select bits [15:0] of the register which is specified in the mnemonic as B (for bottom).

MIAXy does not support unsigned multiplication; all values in Rs and Rm will be interpreted as signed data values.

The instruction is only executed if the condition specified in the instruction matches the condition code status.

2.3.1.2 Internal Accumulator Access Format

The Intel XScale® core defines a new instruction format for accessing internal accumulators in CP0. Table 2-5, “Internal Accumulator Access Format” on page 2-27 shows that the opcode falls into the coprocessor register transfer space.

The *RdHi* and *RdLo* fields allow up to 64 bits of data transfer between StrongARM registers and an internal accumulator. The *acc* field specifies 1 of 8 internal accumulators to transfer data to/from. The core implements a single 40-bit accumulator referred to as *acc0*; future implementations can specify multiple internal accumulators of varying sizes, up to 64 bits.

Access to the internal accumulator is allowed in all processor modes (user and privileged) as long bit 0 of the Coprocessor Access Register is set. (See Section 7.2.15, “Register 15: Coprocessor Access Register” on page 7-94 for more details).

The Intel XScale® core implements two instructions **MAR** and **MRA** that move two ARM registers to *acc0* and move *acc0* to two ARM registers, respectively.

Table 2-5. Internal Accumulator Access Format

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																																										
cond								1	1	0	0	0	1	0	L	RdHi								RdLo								0	0	0	0	0	0	0	0	0	0	acc
Bits	Description		Notes																																							
31:28	cond - ARM condition codes		-																																							
20	L - move to/from internal accumulator 0= move to internal accumulator (MAR) 1= move from internal accumulator (MRA)		-																																							
19:16	RdHi - specifies the high order eight (39:32) bits of the internal accumulator.		On a read of the acc, this 8-bit high order field will be sign extended. On a write to the acc, the lower 8 bits of this register will be written to acc[39:32]																																							
15:12	RdLo - specifies the low order 32 bits of the internal accumulator		-																																							
7:4	Should be zero		This field could be used in future implementations to specify the type of saturation to perform on the read of an internal accumulator. (e.g., a signed saturation to 16-bits may be useful for some filter algorithms.)																																							
3	Should be zero		-																																							
2:0	acc - specifies 1 of 8 internal accumulators		The core only implements <i>acc0</i> ; access to any other <i>acc</i> is unpredictable																																							

Note: **MAR** has the same encoding as **MCRR** (to coprocessor 0) and **MRA** has the same encoding as **MRRC** (to coprocessor 0). These instructions move 64-bits of data to/from ARM registers from/to coprocessor registers. **MCRR** and **MRRC** are defined in ARM's DSP instruction set.

Disassemblers not aware of **MAR** and **MRA** will produce the following syntax:

```
MCRR{<cond>} p0, 0x0, RdLo, RdHi, c0
MRRC{<cond>} p0, 0x0, RdLo, RdHi, c0
```

Table 2-6. MAR{<cond>} acc0, RdLo, RdHi

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	
cond		1	1	0	0	0	1	0	0	RdHi				RdLo				0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Operation: if ConditionPassed(<cond>) then
 acc0[39:32] = RdHi[7:0]
 acc0[31:0] = RdLo[31:0]

Exceptions: none

Qualifiers Condition Code
 No condition code flags are updated

Notes: Instruction timings can be found in
 [Section 10.4.4, "Multiply Instruction Timings" on page 10-168](#)
 Specifying R15 as either RdHi or RdLo has unpredictable results.

The MAR instruction moves the value in register RdLo to bits[31:0] of the 40-bit accumulator (acc0) and moves bits[7:0] of the value in register RdHi into bits[39:32] of acc0.

The instruction is only executed if the condition specified in the instruction matches the condition code status.

This instruction executes in any processor mode.

Table 2-7. MRA{<cond>} RdLo, RdHi, acc0

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	
cond		1	1	0	0	0	1	0	1	RdHi				RdLo				0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Operation: if ConditionPassed(<cond>) then
 RdHi[31:0] = sign_extend(acc0[39:32])
 RdLo[31:0] = acc0[31:0]

Exceptions: none

Qualifiers Condition Code
 No condition code flags are updated

Notes: Instruction timings can be found in
 [Section 10.4.4, "Multiply Instruction Timings" on page 10-168](#)
 Specifying the same register for RdHi and RdLo has unpredictable results.
 Specifying R15 as either RdHi or RdLo has unpredictable results.

The MRA instruction moves the 40-bit accumulator value (acc0) into two registers. Bits[31:0] of the value in acc0 are moved into the register RdLo. Bits[39:32] of the value in acc0 are sign extended to 32 bits and moved into the register RdHi.

The instruction is only executed if the condition specified in the instruction matches the condition code status.

This instruction executes in any processor mode.

2.3.2 New Page Attributes

The Intel XScale[®] core extends the page attributes defined by the C and B bits in the page descriptors with an additional X bit. This bit allows four more attributes to be encoded when X=1. These new encodings include allocating data for the mini-data cache and write-allocate caching. A full description of the encodings can be found in [Section 3.2.2, “Memory Attributes” on page 3-38](#).

The Intel XScale[®] core retains ARM definitions of the C and B encoding when X = 0, which is different than the StrongARM products. The memory attribute for the mini-data cache has been moved and replaced with the write-through caching attribute.

When write-allocate is enabled, a store operation that misses the data cache (cacheable data only) will generate a line fill. If disabled, a line fill only occurs when a load operation misses the data cache (cacheable data only).

Write-through caching causes all store operations to be written to memory, whether they are cacheable or not cacheable. This feature is useful for maintaining data cache coherency.

The Intel XScale[®] core also adds a P bit in the first level descriptors to allow an ASSP to identify a new memory attribute. Refer to the Intel XScale[®] core implementation option section of the ASSP architecture specification to find out how the P bit has been defined. Bit 1 in the Control Register (coprocessor 15, register 1, opcode=1) is used to assigned the P bit memory attribute for memory accesses made during page table walks.

These attributes are programmed in the translation table descriptors, which are highlighted in [Table 2-8, “First-level Descriptors” on page 2-30](#), [Table 2-9, “Second-level Descriptors for Coarse Page Table” on page 2-30](#) and [Table 2-10, “Second-level Descriptors for Fine Page Table” on page 2-30](#). Two second-level descriptor formats have been defined for the core, one is used for the coarse page table and the other is used for the fine page table.

Table 2-8. First-level Descriptors

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																			
SBZ															0	0			
Coarse page table base address										P	Domain			SBZ		0	1		
Section base address					SBZ		TEX		AP	P	Domain			0	C	B	1	0	
Fine page table base address										SBZ		P	Domain			SBZ		1	1

Table 2-9. Second-level Descriptors for Coarse Page Table

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																
SBZ															0	0
Large page base address					TEX		AP3	AP2	AP1	AP0	C	B	0	1		
Small page base address					AP3		AP2	AP1	AP0	C	B	1	0			
Extended small page base address					SBZ		TEX			AP	C	B	1	1		

Table 2-10. Second-level Descriptors for Fine Page Table

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																
SBZ															0	0
Large page base address					TEX		AP3	AP2	AP1	AP0	C	B	0	1		
Small page base address					AP3		AP2	AP1	AP0	C	B	1	0			
Tiny Page Base Address					TEX			AP	C	B	1	1				

The TEX (Type Extension) field is present in several of the descriptor types. In the core, only the LSB of this field is defined; this is called the X bit. The remaining bits should be programmed as zero (SBZ).

A Small Page descriptor does not have a TEX field. For these descriptors, TEX is implicitly zero; that is, they operate as if the X bit had a '0' value.

The X bit, when set, modifies the meaning of the C and B bits. Description of page attributes and their encoding can be found in [Chapter 3, "Memory Management"](#).

2.3.3 Additions to CP15 Functionality

To accommodate the functionality in the Intel XScale® core, registers in CP15 and CP14 have been added or augmented. See [Chapter 7, “Configuration”](#) for details.

At times it is necessary to be able to guarantee exactly when a CP15 update takes effect. For example, when enabling memory address translation (turning on the MMU), it is vital to know when the MMU is actually guaranteed to be in operation. To address this need, a processor-specific code sequence is defined for the core. The sequence -- called CPWAIT -- is shown in [Example 2-1 on page 2-31](#).

Example 2-1. CPWAIT: Canonical method to wait for CP15 update

```
;; The following macro should be used when software needs to be
;; assured that a CP15 update has taken effect.
;; It may only be used while in a privileged mode, because it
;; accesses CP15.

MACRO CPWAIT

    MRC P15, 0, R0, C2, C0, 0          ; arbitrary read of CP15
    MOV R0, R0                        ; wait for it
    SUB PC, PC, #4                    ; branch to next instruction

    ; At this point, any previous CP15 writes are
    ; guaranteed to have taken effect.

ENDM
```

When setting multiple CP15 registers, system software may opt to delay the assurance of their update. This is accomplished by executing CPWAIT only after the sequence of MCR instructions.

Note: The CPWAIT sequence guarantees that CP15 side-effects are complete by the time the CPWAIT is complete. It is possible, however, that the CP15 side-effect will take place before CPWAIT completes or is issued. Programmers should take care that this does not affect the correctness of their code.

2.3.4 Event Architecture

2.3.4.1 Exception Summary

Table 2-11 shows all the exceptions that the core may generate, and the attributes of each. Subsequent sections give details on each exception.

Table 2-11. Exception Summary

Exception Description	Exception Type ^a	Precise?	Updates FAR?
Reset	Reset	N	N
FIQ	FIQ	N	N
IRQ	IRQ	N	N
External Instruction	Prefetch	Y	N
Instruction MMU	Prefetch	Y	N
Instruction Cache Parity	Prefetch	Y	N
Lock Abort	Data	Y	N
MMU Data	Data	Y	Y
External Data	Data	N	N
Data Cache Parity	Data	N	N
Software Interrupt	Software Interrupt	Y	N
Undefined Instruction	Undefined Instruction	Y	N
Debug Events ^b	varies	varies	N

- a. Exception types are those described in the ARM, section 2.5.
b. Refer to [Chapter 9, "Software Debug"](#) for more details

2.3.4.2 Event Priority

The Intel XScale® core follows the exception priority specified in the *ARM Architecture Reference Manual*. The processor has additional exceptions that might be generated while debugging. For information on these debug exceptions, see [Chapter 9, "Software Debug"](#).

Table 2-12. Event Priority

Exception	Priority
Reset	1 (Highest)
Data Abort (Precise & Imprecise)	2
FIQ	3
IRQ	4
Prefetch Abort	5
Undefined Instruction, SWI	6 (Lowest)

2.3.4.3 Prefetch Aborts

The Intel XScale® core detects three types of prefetch aborts: Instruction MMU abort, external abort on an instruction access, and an instruction cache parity error. These aborts are described in [Table 2-13](#).

When a prefetch abort occurs, hardware reports the highest priority one in the extended Status field of the Fault Status Register. The value placed in R14_ABORT (the link register in abort mode) is the address of the aborted instruction + 4.

Table 2-13. Encoding of Fault Status for Prefetch Aborts

Priority	Sources	FS[10,3:0] ^a	Domain	FAR
Highest	Instruction MMU Exception Several exceptions can generate this encoding: - translation faults - domain faults, and - permission faults It is up to software to figure out which one occurred.	0b10000	invalid	invalid
	External Instruction Error Exception This exception occurs when the external memory system reports an error on an instruction cache fetch.	0b10110	invalid	invalid
Lowest	Instruction Cache Parity Error Exception	0b11000	invalid	invalid

a. All other encodings not listed in the table are reserved.

2.3.4.4 Data Aborts

Two types of data aborts exist in the Intel XScale® core: precise and imprecise. A precise data abort is defined as one where R14_ABORT always contains the PC (+8) of the instruction that caused the exception. An imprecise abort is one where R14_ABORT contains the PC (+4) of the next instruction to execute and not the address of the instruction that caused the abort. In other words, instruction execution will have advanced beyond the instruction that caused the data abort.

On the core, precise data aborts are recoverable and imprecise data aborts are not recoverable.

Precise Data Aborts

- A lock abort is a precise data abort; the extended Status field of the Fault Status Register is set to 0xb10100. This abort occurs when a lock operation directed to the MMU (instruction or data) or instruction cache causes an exception, due to either a translation fault, access permission fault or external bus fault.

The Fault Address Register is undefined and R14_ABORT is the address of the aborted instruction + 8.

- A data MMU abort is precise. These are due to an alignment fault, translation fault, domain fault, permission fault or external data abort on an MMU translation. The status field is set to a predetermined ARM definition which is shown in [Table 2-14, “Encoding of Fault Status for Data Aborts”](#) on page 2-34.

The Fault Address Register is set to the effective data address of the instruction and R14_ABORT is the address of the aborted instruction + 8.

Table 2-14. Encoding of Fault Status for Data Aborts

Priority	Sources		FS[10,3:0] ^a	Domain	FAR
Highest	Alignment		0b000x1	invalid	valid
	External Abort on Translation	First level	0b01100	invalid	valid
		Second level	0b01110	valid	valid
	Translation	Section	0b00101	invalid	valid
		Page	0b00111	valid	valid
	Domain	Section	0b01001	valid	valid
		Page	0b01011	valid	valid
	Permission	Section	0b01101	valid	valid
		Page	0b01111	valid	valid
	Lock Abort This data abort occurs on an MMU lock operation (data or instruction TLB) or on an Instruction Cache lock operation.		0b10100	invalid	invalid
	Imprecise External Data Abort		0b10110	invalid	invalid
Lowest	Data Cache Parity Error Exception		0b11000	invalid	invalid

a. All other encodings not listed in the table are reserved.

Imprecise data aborts

- A data cache parity error is imprecise; the extended Status field of the Fault Status Register is set to 0xb11000.
- All external data aborts except for those generated on a data MMU translation are imprecise.

The Fault Address Register for all imprecise data aborts is undefined and R14_ABORT is the address of the next instruction to execute + 4, which is the same for both ARM and Thumb mode.

Although the core guarantees the *Base Restored Abort Model* for precise aborts, it cannot do so in the case of imprecise aborts. A Data Abort handler may encounter an updated base register if it is invoked because of an imprecise abort.

Imprecise data aborts may create scenarios difficult for an abort handler to recover. Both external data aborts and data cache parity errors may result in corrupted targeted register data. Because these faults are imprecise, it is possible corrupted data will have been used before the Data Abort fault handler is invoked. Because of this, software should treat imprecise data aborts as unrecoverable. Even memory accesses marked as “stall until complete” (see [Section 3.2.2.4](#)) can result in imprecise data aborts. For these types of accesses, the fault is somewhat less imprecise than the general case: it is guaranteed to be raised within three instructions of the instruction that caused it. In other words, if a “stall until complete” LD or ST instruction triggers an imprecise fault, then that fault will be seen by the program within three instructions.

With this knowledge, it is possible to write code that accesses “stall until complete” memory with impunity. Simply place several NOP instructions after such an access. If an imprecise fault occurs, it will do so during the NOPs; the data abort handler will see identical register and memory state as it would with a precise exception, and so should be able to recover. An example of this is shown in [Example 2-2 on page 2-35](#).

Example 2-2. Shielding Code from Potential Imprecise Aborts

```

; Example of code that maintains architectural state through the
; window where an imprecise fault might occur.

        LD      R0, [R1]                ; R1 points to stall-until-complete
                                           ; region of memory

        NOP
        NOP
        NOP

        ; Code beyond this point is guaranteed not to see any aborts
        ; from the LD.

```

If a system design precludes events that could cause external aborts, then such precautions are not necessary.

Multiple Data Aborts

Multiple data aborts may be detected by hardware but only the highest priority one will be reported. If the reported data abort is precise, software can correct the cause of the abort and re-execute the aborted instruction. If the lower priority abort still exists, it will be reported. Software can handle each abort separately until the instruction successfully executes.

If the reported data abort is imprecise, software needs to check the SPSR to see if the previous context was executing in abort mode. If this is the case, the link back to the current process has been lost and the data abort is unrecoverable.

2.3.4.5 Events from Preload Instructions

A **PLD** instruction will never cause the Data MMU to fault for any of the following reasons:

- Domain Fault
- Permission Fault
- Translation Fault

If execution of the **PLD** would cause one of the above faults, then the **PLD** causes no effect.

This feature allows software to issue **PLDs** speculatively. For example, [Example 2-3 on page 2-36](#) places a **PLD** instruction early in the loop. This **PLD** is used to fetch data for the next loop iteration. In this example, the list is terminated with a node that has a null pointer. When execution reaches the end of the list, the **PLD** on address 0x0 will not cause a fault. Rather, it will be ignored and the loop will terminate normally.

Example 2-3. Speculatively issuing PLD

```
;; R0 points to a node in a linked list. A node has the following layout:
;; Offset  Contents
;;-----
;;      0  data
;;      4  pointer to next node
;; This code computes the sum of all nodes in a list. The sum is placed into R9.
;;
        MOV R9, #0      ; Clear accumulator
sumList:
        LDR R1, [R0, #4] ; R1 gets pointer to next node
        LDR R3, [R0]     ; R3 gets data from current node
        PLD [R1]         ; Speculatively start load of next node
        ADD R9, R9, R3   ; Add into accumulator
        MOVS R0, R1     ; Advance to next node. At end of list?
        BNE sumList     ; If not then loop
```

2.3.4.6 Debug Events

Debug events are covered in [Section 9.5, “Debug Exceptions” on page 9-126](#).

Memory Management

3

This chapter describes the memory management unit implemented in the Intel XScale[®] core.

3.1 Overview

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

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

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

- cacheable
- bufferable
- line allocate policy
- write policy
- I/O
- mini Data Cache
- Coalescing
- an ASSP definable attribute - P bit (Refer to the Intel XScale[®] core implementation section of the ASSP architecture specification for more information.)

See [Section 3.2.2, “Memory Attributes” on page 3-38](#) for a description of page attributes and [Section 2.3.2, “New Page Attributes” on page 2-29](#) to find out where these attributes have been mapped in the MMU descriptors.

Note: The virtual address with which the TLBs are accessed may be remapped by the PID register. See [Section 7.2.13, “Register 13: Process ID” on page 7-91](#) for a description of the PID register.

3.2 Architecture Model

3.2.1 Version 4 vs. Version 5

ARM* MMU Version 5 Architecture introduces the support of tiny pages, which are 1 KByte in size. The reserved field in the first-level descriptor (encoding 0b11) is used as the fine page table base address. The exact bit fields and the format of the first and second-level descriptors can be found in [Section 2.3.2, “New Page Attributes” on page 2-29](#).

3.2.2 Memory Attributes

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

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

3.2.2.1 Page (P) Attribute Bit

The P bit allows an ASSP to assign its own page attribute to a memory region. This bit is only present in the first level descriptors. Refer to the Intel XScale® core implementation section of the ASSP architecture specification to find out how this has been defined. Accesses to memory for page table walks do not use the MMU. The core provides ASSP definable memory attributes for these accesses in the Auxiliary Control Register. See [Table 7-7, “Auxiliary Control Register” on page 7-84](#).

3.2.2.2 Cacheable (C), Bufferable (B), and eXtension (X) Bits

3.2.2.3 Instruction Cache

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

3.2.2.4 Data Cache and Write Buffer

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

If the X bit for a descriptor is zero, the C and B bits operate as mandated by the ARM architecture. This behavior is detailed in [Table 3-1](#).

If the X bit for a descriptor is one, the C and B bits' meaning is extended, as detailed in [Table 3-2](#).

Table 3-1. Data Cache and Buffer Behavior when X = 0

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

- a. Normally, the processor will continue executing after a data access if no dependency on that access is encountered. With this setting, the processor will stall execution until the data access completes. This guarantees to software that the data access has taken effect by the time execution of the data access instruction completes. External data aborts from such accesses will be imprecise (but see [Section 2.3.4.4](#) for a method to shield code from this imprecision).

Table 3-2. Data Cache and Buffer Behavior when X = 1

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

- a. Normally, bufferable writes can coalesce with previously buffered data in the same address range
 b. See [Section 7.2.2](#) for a description of this register

3.2.2.5 Details on Data Cache and Write Buffer Behavior

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

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

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

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

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

More details on cache policies may be gleaned from [Section 6.2.3, “Cache Policies” on page 6-65](#).

3.2.2.6 Memory Operation Ordering

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

[Table 3-3 on page 3-40](#) shows the circumstances in which memops act as fences.

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

Table 3-3. Memory Operations that Impose a Fence

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

3.2.3 Exceptions

The MMU may generate prefetch aborts for instruction accesses and data aborts for data memory accesses. The types and priorities of these exceptions are described in [Section 2.3.4, “Event Architecture” on page 2-32](#).

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

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

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

Table 3-4. Valid MMU & Data/mini-data Cache Combinations

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

3.4 Control

3.4.1 Invalidate (Flush) Operation

The entire instruction and data TLB can be invalidated at the same time with one command or they can be invalidated separately. An individual entry in the data or instruction TLB can also be invalidated. See [Table 7-13, "TLB Functions"](#) on [page 7-89](#) for a listing of commands supported by the core.

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

3.4.2 Enabling/Disabling

The MMU is enabled by setting bit 0 in coprocessor 15, register 1 (Control Register).

When the MMU is disabled, accesses to the instruction cache default to cacheable and all accesses to data memory are made non-cacheable.

A recommended code sequence for enabling the MMU is shown in [Example 3-1 on page 3-42](#).

Example 3-1. Enabling the MMU

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

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

; For a description of CPWAIT, see
; Section 2.3.3, "Additions to CP15 Functionality" on page 2-31
CPWAIT
; The MMU is guaranteed to be enabled at this point; the next instruction or
; data address will be translated.
```

3.4.3 Locking Entries

Individual entries can be locked into the instruction and data TLBs. See [Table 7-14, “Cache Lockdown Functions” on page 7-90](#) for the exact commands. If a lock operation finds the virtual address translation already resident in the TLB, the results are unpredictable. An invalidate by entry command before the lock command will ensure proper operation. Software can also accomplish this by invalidating all entries, as shown in [Example 3-2 on page 3-43](#).

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

A procedure for locking entries into the instruction TLB is shown in [Example 3-2 on page 3-43](#).

If a MMU abort is generated during an instruction or data TLB lock operation, the Fault Status Register is updated to indicate a Lock Abort (see [Section 2.3.4.4, “Data Aborts” on page 2-34](#)), and the exception is reported as a data abort.

Example 3-2. Locking Entries into the Instruction TLB

```

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

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

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

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

CPWAIT

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

```

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

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

The proper procedure for locking entries into the data TLB is shown in [Example 3-3 on page 3-44](#).

Example 3-3. Locking Entries into the Data TLB

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

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

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

CPWAIT                    ; wait for locks to complete

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

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

3.4.4 Round-Robin Replacement Algorithm


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

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

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

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

Figure 3-1. Example of Locked Entries in TLB



This Page Intentionally Left Blank

Instruction Cache

4

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

4.1 Overview

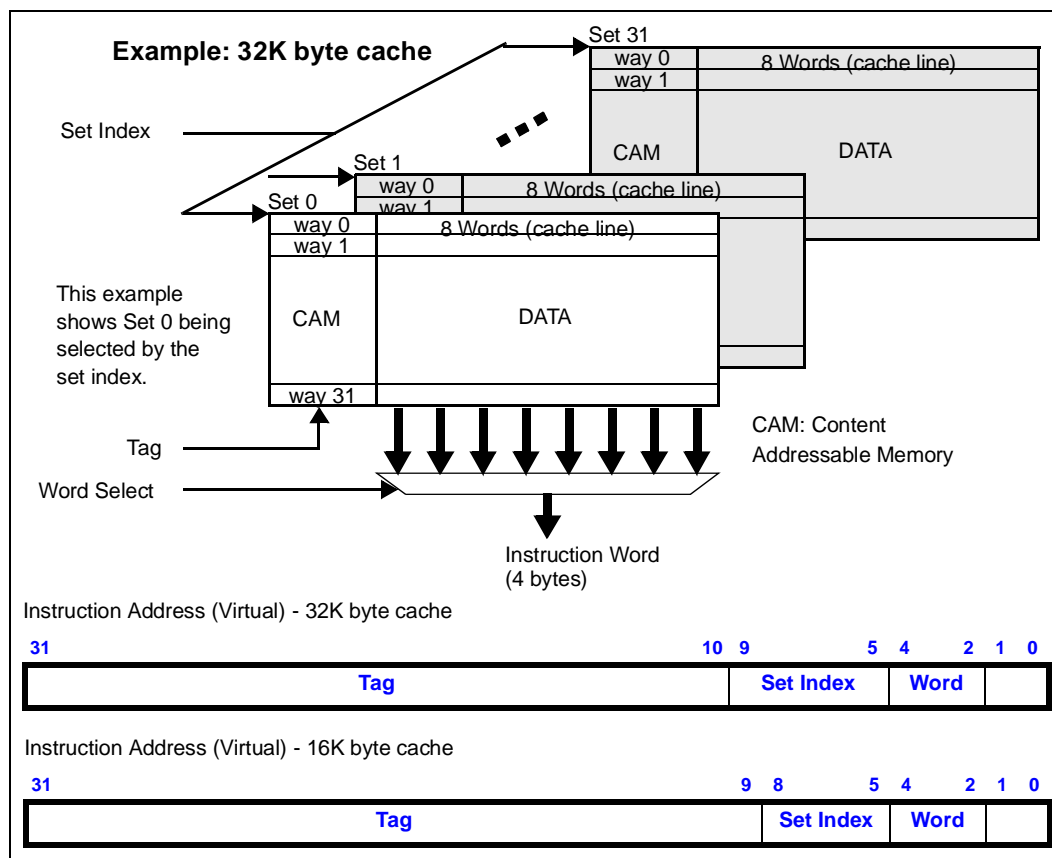
Figure 4-1 shows the cache organization and how the instruction address is used to access the cache.

The instruction cache is available as a **32K or 16K byte, 32-way set associative cache**. The size determines the number of sets; a 32K byte cache has 32 sets and the 16K byte cache has 16 sets. Each set, irrespective of size, contains 32 ways. Each way of a set contains eight 32-bit words and one valid bit, which is referred to as a line. The replacement policy is a round-robin algorithm and the cache also supports the ability to lock code in at a line granularity.

The instruction cache is virtually addressed and virtually tagged.

Note: The virtual address presented to the instruction cache may be remapped by the PID register. See Section 7.2.13, "Register 13: Process ID" on page 7-91 for a description of the PID register.

Figure 4-1. Instruction Cache Organization



4.2 Operation

4.2.1 Operation When Instruction Cache is Enabled

When the cache is enabled, it compares every instruction request address against the addresses of instructions that it is currently holding. If the cache contains the requested instruction, the access “hits” the cache, and the cache returns the requested instruction. If the cache does not contain the requested instruction, the access “misses” the cache, and the cache requests a fetch from external memory of the 8-word line (32 bytes) that contains the requested instruction using the fetch policy described in [Section 4.2.3](#). As the fetch returns instructions to the cache, they are placed in one of two fetch buffers and the requested instruction is delivered to the instruction decoder.

A fetched line will be written into the cache if it is cacheable. Code is designated as cacheable when the Memory Management Unit (MMU) is disabled or when the MMU is enable and the cacheable (C) bit is set to 1 in its corresponding page. See [Chapter 3, “Memory Management”](#) for a discussion on page attributes.

Note that an instruction fetch may “miss” the cache but “hit” one of the fetch buffers. When this happens, the requested instruction will be delivered to the instruction decoder in the same manner as a cache “hit.”

4.2.2 Operation When The Instruction Cache Is Disabled

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

Disabling the instruction cache *does not* disable instruction buffering that may occur within the instruction fetch buffers. Two 8-word instruction fetch buffers will always be enabled in the cache disabled mode. So long as instruction fetches continue to “hit” within either buffer (even in the presence of forward and backward branches), no external fetches for instructions are generated. A miss causes one or the other buffer to be filled from external memory using the fill policy described in [Section 4.2.3](#).

4.2.3 Fetch Policy

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

A miss causes the following:

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

4.2.4 Round-Robin Replacement Algorithm

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

After reset, way 31 is pointed to by the round-robin pointer for all the sets. Once a line is written into way 31, the round-robin pointer points to the first available way of a set, beginning with way 0 if no lines have been locked into that particular set. Locking lines into the instruction cache effectively reduces the available lines for cache updating. For example, if the first three lines of a set were locked down, the round-robin pointer would point to the line at way 3 after it rolled over from way 31. Refer to [Section 4.3.4, “Locking Instructions in the Instruction Cache”](#) on page 4-54 for more details on cache locking.

4.2.5 Parity Protection

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

A software exception handler can recover from an instruction cache parity error. This can be accomplished by invalidating the instruction cache and the branch target buffer and then returning to the instruction that caused the prefetch abort exception. A simplified code example is shown in [Example 4-1 on page 4-50](#). A more complex handler might choose to invalidate the specific line that caused the exception and then invalidate the BTB.

Example 4-1. Recovering from an Instruction Cache Parity Error

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

CPWAIT                ; wait for effect (see Section 2.3.3 for a
                        ; description of CPWAIT)

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

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

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

4.2.6 Instruction Fetch Latency

The instruction fetch latency is dependent on the core to memory frequency ratio, system bus bandwidth, system memory, etc., which are all particular to each ASSP. So, refer to the Intel XScale[®] core implementation option section of the ASSP architecture specification for exact details on instruction fetch latency.

4.2.7 Instruction Cache Coherency

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

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

To achieve cache coherence, instruction cache contents can be invalidated after code modification in external memory is complete. Refer to [Section 4.3.3, “Invalidating the Instruction Cache”](#) on [page 4-53](#) for the proper procedure in invalidating the instruction cache.

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

Naturally, when writing code as data, care must be taken to force it completely out of the processor into external memory before attempting to execute it. If writing into a non-cacheable region, flushing the write buffers is sufficient precaution (see [Section 7.2.8](#) for a description of this operation). If writing to a cacheable region, then the data cache should be submitted to a Clean/Invalidate operation (see [Section 6.3.3.1](#)) to ensure coherency.

4.3 Instruction Cache Control

4.3.1 Instruction Cache State at RESET

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

4.3.2 Enabling/Disabling

The instruction cache is enabled by setting bit 12 in coprocessor 15, register 1 (Control Register). This process is illustrated in [Example 4-2, Enabling the Instruction Cache](#).

Example 4-2. Enabling the Instruction Cache

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

CPWAIT
```

4.3.3 Invalidating the Instruction Cache

The entire instruction cache along with the fetch buffers are invalidated by writing to coprocessor 15, register 7. (See [Table 7-12, “Cache Functions” on page 7-87](#) for the exact command.) This command does not unlock any lines that were locked in the instruction cache nor does it invalidate those locked lines. To invalidate the entire cache including locked lines, the unlock instruction cache command needs to be executed before the invalidate command. This unlock command can also be found in [Table 7-14, “Cache Lockdown Functions” on page 7-90](#).

There is an inherent delay from the execution of the instruction cache invalidate command to where the next instruction will see the result of the invalidate. The following routine can be used to guarantee proper synchronization.

Example 4-3. Invalidating the Instruction Cache

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

CPWAIT

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

The Intel XScale® core also supports invalidating an individual line from the instruction cache. See [Table 7-12, “Cache Functions” on page 7-87](#) for the exact command.

4.3.4 Locking Instructions in the Instruction Cache

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

Lines can be locked into the instruction cache by initiating a write to coprocessor 15. (See Table 7-14, "Cache Lockdown Functions" on page 7-90 for the exact command.) Register *Rd* contains the virtual address of the line to be locked into the cache.

There are several requirements for locking down code:

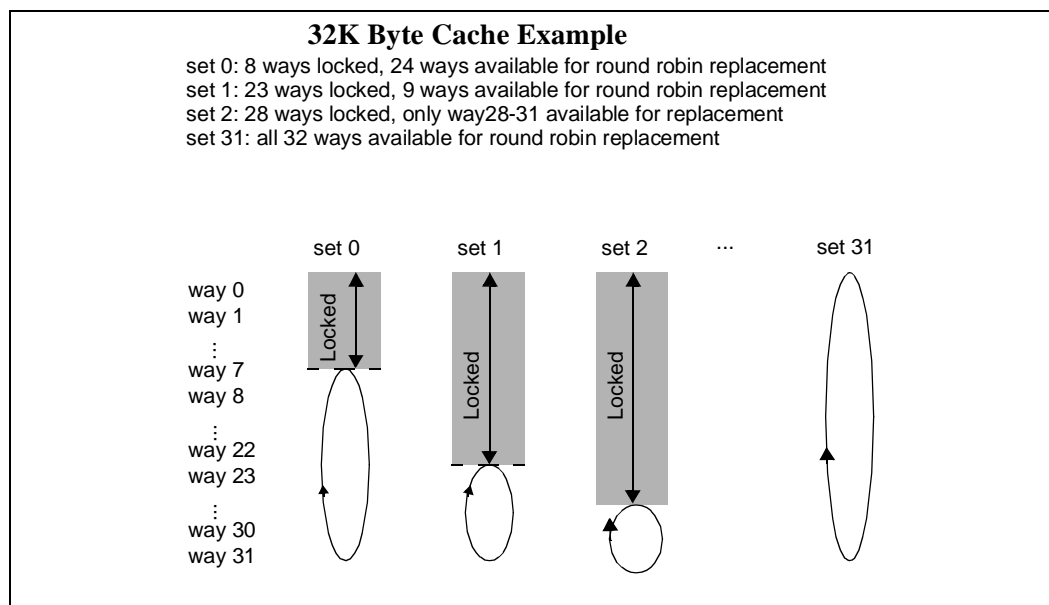
7. the routine used to lock lines down in the cache must be placed in non-cacheable memory, which means the MMU is enabled. As a corollary: no fetches of cacheable code should occur while locking instructions into the cache. the code being locked into the cache must be cacheable
8. the instruction cache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

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

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

Lines are locked into a set starting at way 0 and may progress up to way 27; which set a line gets locked into depends on the set index of the virtual address. Figure 4-2 is an example (32K byte cache) of where lines of code may be locked into the cache along with how the round-robin pointer is affected.

Figure 4-2. Locked Line Effect on Round Robin Replacement



Software can lock down several different routines located at different memory locations. This may cause some sets to have more locked lines than others as shown in [Figure 4-2](#).

[Example 4-4 on page 4-55](#) shows how a routine, called “lockMe” in this example, might be locked into the instruction cache. Note that it is possible to receive an exception while locking code (see [Section 2.3.4, “Event Architecture” on page 2-32](#)).

Example 4-4. Locking Code into the Cache

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

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

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

4.3.5 Unlocking Instructions in the Instruction Cache

The Intel XScale® core provides a global unlock command for the instruction cache. Writing to coprocessor 15, register 9 unlocks all the locked lines in the instruction cache and leaves them valid. These lines then become available for the round-robin replacement algorithm. (See [Table 7-14, “Cache Lockdown Functions” on page 7-90](#) for the exact command.)

This Page Intentionally Left Blank

Branch Target Buffer

5


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

This chapter is primarily for those optimizing their code for performance. An understanding of the branch target buffer is needed in this case so that code can be scheduled to best utilize the performance benefits of the branch target buffer.

5.1 Branch Target Buffer (BTB) Operation

The BTB stores the history of branches that have executed along with their targets. [Figure 5-1](#) shows an entry in the BTB, where the tag is the instruction address of a previously executed branch and the data contains the target address of the previously executed branch along with two bits of history information.

Figure 5-1. BTB Entry



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

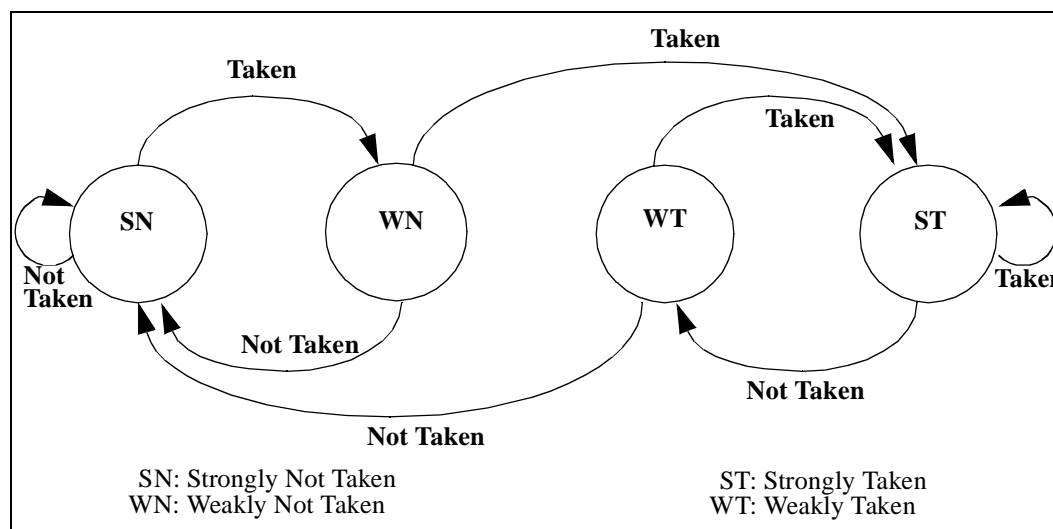
Bit[1] of the instruction address is included in the tag comparison in order to support Thumb execution. This organization means that two consecutive Thumb branch (B) instructions, with instruction address bits[8:2] the same, will contend for the same BTB entry. Thumb also requires 31 bits for the branch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

The history bits represent four possible prediction states for a branch entry in the BTB. Figure 5-2, “Branch History” on page 5-58 shows these states along with the possible transitions. The initial state for branches stored in the BTB is Weakly-Taken (WT). Every time a branch that exists in the BTB is executed, the history bits are updated to reflect the latest outcome of the branch, either taken or not-taken.

Chapter 10, “Performance Considerations” describes which instructions are dynamically predicted by the BTB and the performance penalty for mispredicting a branch.

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

Figure 5-2. Branch History



5.1.1 Reset

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

5.1.2 Update Policy

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

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

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

Once a branch is stored in the BTB, the history bits are updated upon every execution of the branch as shown in Figure 5-2.

5.2 BTB Control

5.2.1 Disabling/Enabling

The BTB is always disabled with Reset. Software can enable the BTB through a bit in a coprocessor register (see [Section 7.2.2](#)).

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

5.2.2 Invalidation

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

1. Reset
2. Software can directly invalidate the BTB via a CP15, register 7 function. Refer to [Section 7.2.8, "Register 7: Cache Functions" on page 7-87](#).
3. The BTB is invalidated when the Process ID Register is written.
4. The BTB is invalidated when the instruction cache is invalidated via CP15, register 7 functions.

This Page Intentionally Left Blank

Data Cache

6

The Intel XScale® core data cache enhances performance by reducing the number of data accesses to and from external memory. There are two data cache structures in the core, a data cache with two size options (32 K or 16 Kbytes) and a mini-data cache that is 1/16th the size of the main data cache. An eight entry write buffer and a four entry fill buffer are also implemented to decouple the core instruction execution from external memory accesses, which increases overall system performance.

6.1 Overviews

6.1.1 Data Cache Overview

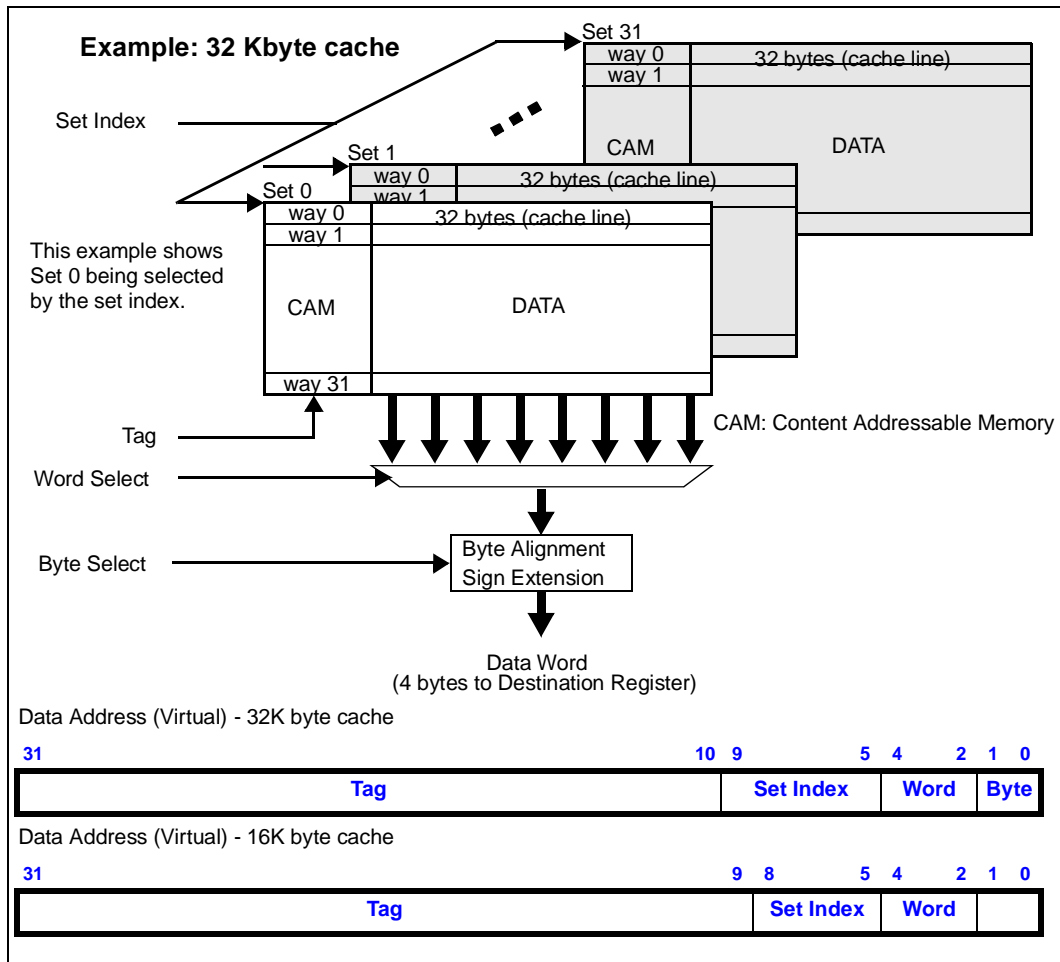
The data cache is available as a **32 K or 16 Kbyte, 32-way set associative cache**. The size determines the number of sets; a 32 Kbyte cache has 32 sets and the 16 Kbyte cache has 16 sets. Each set, irrespective of size, contains 32 ways. Each way of a set contains 32 bytes (one cache line) and one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

Figure 6-1, “Data Cache Organization” on page 6-62 shows the cache organization and how the data address is used to access the cache.

Cache policies may be adjusted for particular regions of memory by altering page attribute bits in the MMU descriptor that controls that memory. See [Section 3.2.2](#) for a description of these bits.

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

Figure 6-1. Data Cache Organization



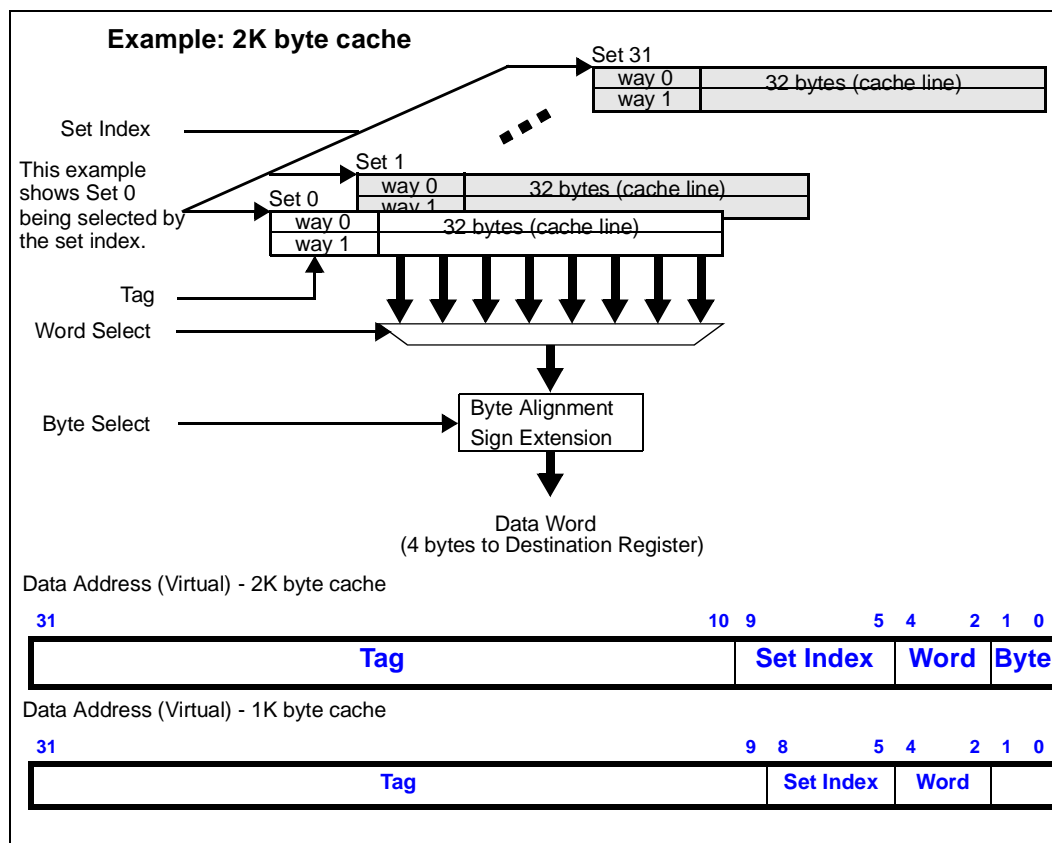
6.1.2 Mini-Data Cache Overview

The mini-data cache is $1/16^{\text{th}}$ the size of the data cache, so depending on the data cache size selected the available sizes are 2 K or 1 Kbytes. The 2 Kbyte version has 32 sets and the 1 Kbyte version has 16 sets; both versions are 2-way set associative. Each way of a set contains 32 bytes (one cache line) and one valid bit. There also exist 2 dirty bits for every line, one for the lower 16 bytes and the other one for the upper 16 bytes. When a store hits the cache the dirty bit associated with it is set. The replacement policy is a round-robin algorithm.

Figure 6-2, “Mini-Data Cache Organization” on page 6-63 shows the cache organization and how the data address is used to access the cache.

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

Figure 6-2. Mini-Data Cache Organization



6.1.3 Write Buffer and Fill Buffer Overview

The Intel XScale® core employs an eight entry write buffer, each entry containing 16 bytes. Stores to external memory are first placed in the write buffer and subsequently taken out when the bus is available.

The write buffer supports the coalescing of multiple store requests to external memory. An incoming store may coalesce with any of the eight entries.

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

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

Pended operations complete in program order.

6.2 Data Cache and Mini-Data Cache Operation

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

6.2.1 Operation When Caching is Enabled

When the data/mini-data cache is enabled for an access, the data/mini-data cache compares the address of the request against the addresses of data that it is currently holding. If the line containing the address of the request is resident in the cache, the access “hits” the cache. For a load operation the cache returns the requested data to the destination register and for a store operation the data is stored into the cache. The data associated with the store may also be written to external memory if write-through caching is specified for that area of memory. If the cache does not contain the requested data, the access ‘misses’ the cache, and the sequence of events that follows depends on the configuration of the cache, the configuration of the MMU and the page attributes, which are described in [Section 6.2.3.2, “Read Miss Policy” on page 6-66](#) and [Section 6.2.3.3, “Write Miss Policy” on page 6-67](#) for a load “miss” and store “miss” respectively.

6.2.2 Operation When Data Caching is Disabled

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

6.2.3 Cache Policies

6.2.3.1 Cacheability

Data at a specified address is cacheable given the following:

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

6.2.3.2 Read Miss Policy

The following sequence of events occurs when a cacheable (see [Section 6.2.3.1, “Cacheability” on page 6-65](#)) load operation misses the cache:

1. The fill buffer is checked to see if an outstanding fill request already exists for that line.
If so, the current request is placed in the pending buffer and waits until the previously requested fill completes, after which it accesses the cache again, to obtain the request data and returns it to the destination register.
If there is no outstanding fill request for that line, the current load request is placed in the fill buffer and a 32-byte external memory read request is made. If the pending buffer or fill buffer is full, the core will stall until an entry is available.
2. A line is allocated in the cache to receive the 32 bytes of fill data. The line selected is determined by the round-robin pointer (see [Section 6.2.4, “Round-Robin Replacement Algorithm” on page 6-68](#)). The line chosen may contain a valid line previously allocated in the cache. In this case both dirty bits are examined and if set, the four words associated with a dirty bit that's asserted will be written back to external memory as a four word burst operation.
3. When the data requested by the load is returned from external memory, it is immediately sent to the destination register specified by the load. A system that returns the requested data back first, with respect to the other bytes of the line, will obtain the best performance.
4. As data returns from external memory it is written into the cache in the previously allocated line.

A load operation that misses the cache and is NOT cacheable makes a request from external memory for the exact data size of the original load request. For example, **LDRH** requests exactly two bytes from external memory, **LDR** requests 4 bytes from external memory, etc. This request is placed in the fill buffer until, the data is returned from external memory, which is then forwarded back to the destination register(s).

6.2.3.3 Write Miss Policy

A write operation that misses the cache will request a 32-byte cache line from external memory if the access is cacheable and write allocation is specified in the page. In this case the following sequence of events occur:

1. The fill buffer is checked to see if an outstanding fill request already exists for that line.
If so, the current request is placed in the pending buffer and waits until the previously requested fill completes, after which it writes its data into the recently allocated cache line.
If there is no outstanding fill request for that line, the current store request is placed in the fill buffer and a 32-byte external memory read request is made. If the pending buffer or fill buffer is full, the core will stall until an entry is available.
2. The 32 bytes of data can be returned back to the core in any word order, i.e, the eight words in the line can be returned in any order. Note that it does not matter, for performance reasons, which order the data is returned to the core since the store operation has to wait until the entire line is written into the cache before it can complete.
3. When the entire 32-byte line has returned from external memory, a line is allocated in the cache, selected by the round-robin pointer (see [Section 6.2.4, “Round-Robin Replacement Algorithm” on page 6-68](#)). The line to be written into the cache may replace a valid line previously allocated in the cache. In this case both dirty bits are examined and if any are set, the four words associated with a dirty bit that's asserted will be written back to external memory as a 4 word burst operation. This write operation will be placed in the write buffer.
4. The line is written into the cache along with the data associated with the store operation.

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

6.2.3.4 Write-Back Versus Write-Through

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

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

6.2.4 Round-Robin Replacement Algorithm

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

After reset, way 31 is pointed to by the round-robin pointer for all the sets. Once a line is written into way 31, the round-robin pointer points to the first available way of a set, beginning with way 0 if no lines have been re-configured as data RAM in that particular set. Re-configuring lines as data RAM effectively reduces the available lines for cache updating. For example, if the first three lines of a set were re-configured, the round-robin pointer would point to the line at way 3 after it rolled over from way 31. Refer to [Section 6.4, “Re-configuring the Data Cache as Data RAM”](#) on [page 6-71](#) for more details on data RAM.

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

6.2.5 Parity Protection

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

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

A data abort due to a data/mini-data cache parity error may not be recoverable if the data address that caused the abort occurred on a line in the cache that has a write-back caching policy. Prior updates to this line may be lost; in this case the software exception handler should perform a “clean and clear” operation on the data cache, ignoring subsequent parity errors, and restart the offending process. This operation is shown in [Section 6.3.3.1](#).

6.2.6 Atomic Accesses

The **SWP** and **SWPB** instructions generate an atomic load and store operation allowing a memory semaphore to be loaded and altered without interruption. These accesses may hit or miss the data/mini-data cache depending on configuration of the cache, configuration of the MMU, and the page attributes. Refer to the ASSP architecture specification for a product specific definition.

6.3 Data Cache and Mini-Data Cache Control

6.3.1 Data Memory State After Reset

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

6.3.2 Enabling/Disabling

The data cache and mini-data cache are enabled by setting bit 2 in coprocessor 15, register 1 (Control Register). See [Chapter 7, “Configuration”](#), for a description of this register and others.

[Equation 6-1](#) shows code that enables the data and mini-data caches. Note that the MMU must be enabled to use the data cache.

Example 6-1. Enabling the Data Cache

```
enableDCache:

    MCR p15, 0, r0, c7, c10, 4; Drain pending data operations...
                                ; (see Section 7.2.8, “Register 7: Cache Functions”)
    MRC p15, 0, r0, c1, c0, 0; Get current control register
    ORR r0, r0, #4           ; Enable DCache by setting 'C' (bit 2)
    MCR p15, 0, r0, c1, c0, 0; And update the Control register
```

6.3.3 Invalidate and Clean Operations

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

This same register also provides the command to invalidate the entire data cache and mini-data cache. Refer to [Table 7-12, “Cache Functions” on page 7-87](#) for a listing of the commands. These global invalidate commands have no effect on lines locked in the data cache. Locked lines must be unlocked before they can be invalidated. This is accomplished by the Unlock Data Cache command found in [Table 7-14, “Cache Lockdown Functions” on page 7-90](#).

6.3.3.1 Global Clean and Invalidate Operation

A simple software routine is used to globally clean the data cache. It takes advantage of the line-allocate data cache operation, which allocates a line into the data cache. This allocation evicts any cache dirty data back to external memory. [Example 6-2](#) shows how data cache can be cleaned.

Example 6-2. Global Clean Operation

```

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

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

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

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

```

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

The line-allocate command will not operate on the mini Data Cache, so system software must clean this cache by reading 2 Kbytes of contiguous unused data into it. This data must be unused and reserved for this purpose so that it will not already be in the cache. It must reside in a page that is marked as mini Data Cache cacheable (see [Section 2.3.2](#)).

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

6.4 Re-configuring the Data Cache as Data RAM

Software has the ability to lock tags associated with 32-byte lines in the data cache, thus creating the appearance of data RAM. Any subsequent access to this line will always hit the cache unless it is invalidated. Once a line is locked into the data cache it is no longer available for cache allocation on a line fill. Up to 28 lines in each set can be reconfigured as data RAM, such that the maximum data RAM size is 28 Kbytes for the 32 Kbytes cache and 12 Kbytes for the 16 Kbytes cache.

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

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

Code examples for these two applications are shown in [Example 6-3 on page 6-72](#) and [Example 6-4 on page 6-73](#). The difference between these two routines is that [Example 6-3 on page 6-72](#) actually requests the entire line of data from external memory and [Example 6-4 on page 6-73](#) uses the line-allocate operation to lock the tag into the cache. No external memory request is made, which means software can map any unallocated area of memory as data RAM. However, the line-allocate operation does validate the target address with the MMU, so system software must ensure that the memory has a valid descriptor in the page table.

Another item to note in [Example 6-4 on page 6-73](#) is that the 32 bytes of data located in a newly allocated line in the cache must be initialized by software before it can be read. The line allocate operation does not initialize the 32 bytes and therefore reading from that line will produce unpredictable results.

In both examples, the code drains the pending loads before and after locking data. This step ensures that outstanding loads do not end up in the wrong place -- either unintentionally locked into the cache or mistakenly left out in the proverbial cold (not locked into the nice warm cache with their brethren). Note also that a drain operation has been placed after the operation that locks the tag into the cache. This drains ensures predictable results if a programmer tries to lock more than 28 lines in a set; the tag will get allocated in this case but not locked into the cache.

Example 6-3. Locking Data into the Data Cache

```
; R1 contains the virtual address of a region of memory to lock,  
; configured with C=1 and B=1  
; R0 is the number of 32-byte lines to lock into the data cache. In this  
; example 16 lines of data are locked into the cache.  
; MMU and data cache are enabled prior to this code.  
  
MACRO DRAIN  
    MCR P15, 0, R0, C7, C10, 4    ; drain pending loads and stores  
ENDM  
  
DRAIN  
  
MOV R2, #0x1  
MCR P15,0,R2,C9,C2,0    ; Put the data cache in lock mode  
CPWAIT  
MOV R0, #16  
LOOP1:  
MCR P15,0,R1,C7,C10,1    ; Write back the line if it's dirty in the cache  
MCR P15,0,R1, C7,C6,1    ; Flush/Invalidate the line from the cache  
LDR R2, [R1], #32        ; Load and lock 32 bytes of data located at [R1]  
                            ; into the data cache. Post-increment the address  
                            ; in R1 to the next cache line.  
SUBS R0, R0, #1; Decrement loop count  
BNE LOOP1  
  
    ; Turn off data cache locking  
MOV R2, #0x0  
MCR P15,0,R2,C9,C2,0    ; Take the data cache out of lock mode.  
CPWAIT
```


Example 6-4. Creating Data RAM

```

; R1 contains the virtual address of a region of memory to configure as data RAM,
; which is aligned on a 32-byte boundary.
; MMU is configured so that the memory region is cacheable.
; R0 is the number of 32-byte lines to designate as data RAM. In this example 16
; lines of the data cache are re-configured as data RAM.
; The inner loop is used to initialize the newly allocated lines
; MMU and data cache are enabled prior to this code.

MACRO ALLOCATE Rx
    MCR P15, 0, Rx, C7, C2, 5
ENDM

MACRO DRAIN
    MCR P15, 0, R0, C7, C10, 4    ; drain pending loads and stores
ENDM

DRAIN
MOV R4, #0x0
MOV R5, #0x0
MOV R2, #0x1
MCR P15,0,R2,C9,C2,0            ; Put the data cache in lock mode
CPWAIT

MOV R0, #16
LOOP1:
ALLOCATE R1                    ; Allocate and lock a tag into the data cache at
                               ; address [R1].
; initialize 32 bytes of newly allocated line
DRAIN
STRD R4, [R1],#8    ;
STRD R4, [R1],#8    ;
STRD R4, [R1],#8    ;
STRD R4, [R1],#8    ;

SUBS R0, R0, #1      ; Decrement loop count
BNE LOOP1
; Turn off data cache locking

DRAIN                ; Finish all pending operations

MOV R2, #0x0
MCR P15,0,R2,C9,C2,0; Take the data cache out of lock mode.
CPWAIT


```

Tags can be locked into the data cache by enabling the data cache lock mode bit located in coprocessor 15, register 9. (See [Table 7-14, “Cache Lockdown Functions”](#) on page 7-90 for the exact command.) Once enabled, any new lines allocated into the data cache will be locked down.

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

Lines are locked into a set starting at way0 and may progress up to way 27; which set a line gets locked into depends on the set index of the virtual address of the request. [Figure 6-3, “Locked Line Effect on Round Robin Replacement”](#) on page 6-74 is an example of where lines of code may be locked into the cache along with how the round-robin pointer is affected.

Figure 6-3. Locked Line Effect on Round Robin Replacement



Software can lock down data located at different memory locations. This may cause some sets to have more locked lines than others as shown in [Figure 6-3](#).

Lines are unlocked in the data cache by performing an unlock operation. See [Section 7.2.10, “Register 9: Cache Lock Down”](#) on page 7-90 for more information about locking and unlocking the data cache.

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

6.5 Write Buffer/Fill Buffer Operation and Control

See [Section 1.3.2, “Terminology and Acronyms”](#) on page 1-19 for a definition of coalescing.

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

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

The write buffer and fill buffer support a drain operation, such that before the next instruction executes, all the core data requests to external memory have completed. Note that an ASSP may also include operations external to the core in the drain operation. (Refer to the Intel XScale[®] core implementation option section in the ASSP architecture specification for more details.) See [Table 7-12, “Cache Functions”](#) on page 7-87 for the exact command.

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

If software is running in a privileged mode, it can explicitly drain all buffered writes. For details on this operation, see the description of Drain Write Buffer in [Section 7.2.8, “Register 7: Cache Functions”](#) on page 7-87.

This Page Intentionally Left Blank

Configuration

7

This chapter describes the System Control Coprocessor (CP15) and coprocessor 14 (CP14). CP15 configures the MMU, caches, buffers and other system attributes. CP14 contains the performance monitor registers, clock and power management registers and the debug registers.

7.1 Overview

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

All CP14 registers can be accessed through **MRC** and **MCR** coprocessor instructions. **LDC** and **STC** coprocessor instructions can only access the clock and power management registers, and the debug registers. The performance monitoring registers can't be accessed by **LDC** and **STC** because CRm != 0x0, which can't be expressed by **LDC** or **STC**. Access to all registers is allowed only in privileged mode. Any access to CP14 in user mode will cause an undefined instruction exception.

Coprocessors, CP15 and CP14, on the Intel XScale® core do not support access via **CDP**, **MRRC**, or **MCRR** instructions. An attempt to access these coprocessors with these instructions will result in an undefined instruction exception.

Many of the MCR commands available in CP15 modify hardware state sometime after execution. A software sequence is available for those wishing to determine when this update occurs and can be found in [Section 2.3.3, "Additions to CP15 Functionality" on page 2-31](#).

The Intel XScale® core includes an extra level of virtual address translation in the form of a PID (Process ID) register and associated logic. For a detailed description of this facility, see [Section 7.2.13, "Register 13: Process ID" on page 7-91](#). Privileged code needs to be aware of this facility because, when interacting with CP15, some addresses are modified by the PID and others are not. An address that has yet to be modified by the PID ("PIDified") is known as a *virtual address (VA)*. An address that has been through the PID logic, but not translated into a physical address, is a *modified virtual address (MVA)*.

The format of **MRC** and **MCR** is shown in [Table 7-1](#).

The Intel XScale® core implements CP15, CP14 and CP0 coprocessors, which is specified by *cp_num*. CP0 supports instructions specific for DSP and is described in [Chapter 2, “Programming Model.”](#) Refer to the Intel XScale® core implementation option section of the ASSP architecture specification to find out what other coprocessors, if any, are supported in the ASSP.

Unless otherwise noted, unused bits in coprocessor registers have unpredictable values when read. For compatibility with future implementations, software should not rely on the values in those bits.

Table 7-1. MRC/MCR Format

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
cond	1	1	1	0	opcode_1	n	CRn	Rd	cp_num	opcode_2	1	CRm																			
Bits	Description																Notes														
31:28	cond - ARM* condition codes																-														
23:21	opcode_1 - Reserved																Should be programmed to zero for future compatibility														
20	n - Read or write coprocessor register 0 = MCR 1 = MRC																-														
19:16	CRn - specifies which coprocessor register																-														
15:12	Rd - General Purpose Register, R0..R15																-														
11:8	cp_num - coprocessor number																The Intel XScale® core defines three coprocessors: 0b1111 = CP15 0b1110 = CP14 0x0000 = CP0 NOTE: Refer to the Intel XScale® core implementation option section of the ASSP architecture specification to see if there are any other coprocessors defined by the ASSP.														
7:5	opcode_2 - Function bits																This field should be programmed to zero for future compatibility unless a value has been specified in the command.														
3:0	CRm - Function bits																This field should be programmed to zero for future compatibility unless a value has been specified in the command.														

The format of **LDC** and **STC** for CP14 is shown in Table 7-2. **LDC** and **STC** follow the programming notes in the *ARM Architecture Reference Manual*. Note that access to CP15 with **LDC** and **STC** will cause an undefined exception and accesses to all other coprocessors is defined in the Intel XScale® core implementation option section of the ASSP architecture specification.

LDC and **STC** transfer a single 32-bit word between a coprocessor register and memory. These instructions do not allow the programmer to specify values for **opcode_1**, **opcode_2**, or **Rm**; those fields implicitly contain zero, which means the performance monitoring registers are not accessible.

Table 7-2. LDC/STC Format when Accessing CP14

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																															
cond				1	1	0	P	U	N	W	L	Rn				CRd				cp_num				8_bit_word_offset							
Bits	Description		Notes																												
31:28	cond - ARM* condition codes		-																												
24:23,21	P, U, W - specifies 1 of 3 addressing modes identified by addressing mode 5 in the <i>ARM Architecture Reference Manual</i> .		-																												
22	N - should be 0 for CP14 coprocessors. Setting this bit to 1 has will have an undefined effect.		-																												
20	L - Load or Store 0 = STC 1 = LDC		-																												
19:16	Rn - specifies the base register		-																												
15:12	CRd - specifies the coprocessor register		-																												
11:8	cp_num - coprocessor number		The Intel XScale® core defines the following: 0b1111 = Undefined Exception 0b1110 = CP14 NOTE: Refer to the Intel XScale® core implementation option section of the ASSP architecture specification to find out the meaning of the other encodings.																												
7:0	8-bit word offset		-																												

7.2 CP15 Registers

Table 7-3 lists the CP15 registers implemented in the Intel XScale® core.

Table 7-3. CP15 Registers

Register (CRn)	Opc_1	CRm	Opc_2	Access	Description
0	0	0	0	Read / Write-Ignored	ID
0	0	0	1	Read / Write-Ignored	Cache Type
1	0	0	0	Read / Write	Control
1	0	0	1	Read / Write	Auxiliary Control
2	0	0	0	Read / Write	Translation Table Base
3	0	0	0	Read / Write	Domain Access Control
4	-	-	-	Unpredictable	Reserved
5	0	0	0	Read / Write	Fault Status
6	0	0	0	Read / Write	Fault Address
7	0	Varies ^a	Varies ^a	Read-unpredictable / Write	Cache Operations
8	0	Varies ^a	Varies ^a	Read-unpredictable / Write	TLB Operations
9	0	Varies ^a	Varies ^a	Varies ^a	Cache Lock Down
10	0	Varies ^a	Varies ^a	Read-unpredictable / Write	TLB Lock Down
11 - 12	-	-	-	Unpredictable	Reserved
13	0	0	0	Read / Write	Process ID (PID)
14	0	Varies ^a	0	Read / Write	Breakpoint Registers
15	0	1	0	Read / Write	Coprocessor Access

a. The value varies depending on the specified function. Refer to the register description for a list of values.

7.2.1 Register 0: ID & Cache Type Registers

Register 0 houses two read-only register 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 ASSP, where a portion of it is defined by the ASSP. Refer to the Intel XScale® core implementation option section of the ASSP architecture specification for the exact encoding.

Table 7-4. ID 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																															
0 1 1 0 1 0 0 1								0 0 0 0 0 1 0 1								Core Gen				Core Revision				Product Number				Product Revision			
reset value: As Shown																															
Bits	Access	Description																													
31:24	Read / Write Ignored	Implementation trademark (0x69 = 'i' = Intel Corporation)																													
23:16	Read / Write Ignored	Architecture version = ARM* Version 5TE																													
15:13	Read / Write Ignored	Intel XScale® core Generation 0b001 = XSC1 0b010 = XSC2 This field reflects a specific set of architecture features supported by the core. If new features are added/deleted/modified this field will change. This allows software, that is not dependent on ASSP features, to target code at a specific core generation. The difference between XSC1 and XSC2 is: <ul style="list-style-type: none"> the performance monitoring facility (Chapter 8, "Performance Monitoring") size of the JTAG instruction register (Appendix B, "Test Features") 																													
12:10	Read / Write Ignored	Core Revision: This field reflects revisions of core generations. Differences may include errata that dictate different operating conditions, software work-around, etc.																													
9:4	Read / Write Ignored	Product Number (Defined by the ASSP)																													
3:0	Read / Write Ignored	Product Revision (Defined by the ASSP)																													

The Cache Type Register is selected when *opcode_2*=1 and describes the cache configuration of the core.

Table 7-5. Cache Type 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
0	0	0	0	1	0	1	1	0	0	0	Dsize	1	0	1	0	1	0	0	0	0	Isize	1	0	1	0	1	0	1	0	0	
reset value: As Shown																															
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 (Dsize) 0b101 = 16 KB 0b110 = 32 KB																												
17:15	Read / Write Ignored		Data cache associativity = 0b101 = 32-way																												
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 (Isize) 0b101 = 16KB 0b110 = 32 KB																												
5:3	Read / Write Ignored		Instruction cache associativity = 0b101 = 32-way																												
2	Read-as-Zero / Write Ignored		Reserved																												
1:0	Read / Write Ignored		Instruction cache line length = 0b10 = 8 words/line																												

7.2.2 Register 1: Control & Auxiliary Control Registers

Register 1 is made up of two registers, one that is compliant with ARM Version 5TE and referred by *opcode_2* = 0x0, and the other which is specific to the core is referred by *opcode_2* = 0x1. The latter is known as the Auxiliary Control Register.

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. If the MMU is enabled, the exception vectors will be accessed via the usual translation method involving the PID register (see [Section 7.2.13, "Register 13: Process ID" on page 7-91](#)) and the TLBs. To avoid automatic application of the PID to exception vector accesses, software may relocate the exceptions to high memory.

Table 7-6. ARM* 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																																																									
																												V	I	Z	0	R	S	B	1	1	1	1	C	A	M																
reset value: writable bits set to 0																																																									
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 Target Buffer Enable (Z) 0 = Disabled 1 = Enabled																																																							
10	Read-as-Zero / Write-as-Zero	Reserved																																																							
9	Read / Write	ROM Protection (R) This selects the access checks performed by the memory management unit. See the <i>ARM Architecture Reference Manual</i> for more information.																																																							
8	Read / Write	System Protection (S) This selects the access checks performed by the memory management unit. See the <i>ARM Architecture Reference Manual</i> for more information.																																																							
7	Read / Write	Big/Little Endian (B) 0 = Little-endian operation 1 = Big-endian operation																																																							
6:3	Read-as-One / Write-as-One	= 0b1111																																																							
2	Read / Write	Data cache enable/disable (C) 0 = Disabled 1 = Enabled																																																							
1	Read / Write	Alignment fault enable/disable (A) 0 = Disabled 1 = Enabled																																																							
0	Read / Write	Memory management unit enable/disable (M) 0 = Disabled 1 = Enabled																																																							

The mini-data cache attribute bits, in the Auxiliary Control Register, are used to control the allocation policy for the mini-data cache and whether it will use write-back caching or write-through caching.

Note: The configuration of the mini-data cache should be 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 7-7. 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		P	K
reset value: writable bits set to 0					
Bits	Access	Description			
31:6	Read-Unpredictable / Write-as-Zero	Reserved			
5:4	Read / Write	Mini Data Cache Attributes (MD) All configurations of the Mini-data cache are cacheable, stores are buffered in the write buffer and stores will be coalesced in the write buffer as long as coalescing is globally enable (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			
1	Read / Write	Page Table Memory Attribute (P) This field is defined by the ASSP. Refer to the Intel XScale® core implementation option section of the ASSP architecture specification for more information.			
0	Read / Write	Write Buffer Coalescing Disable (K) 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			

7.2.3 Register 2: Translation Table Base Register

Table 7-8. Translation Table Base 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																															
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	Reserved																													

7.2.4 Register 3: Domain Access Control Register

Table 7-9. Domain Access 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															
D15	D14	D13	D12	D11	D10	D9	D8	D7	D6	D5	D4	D3	D2	D1	D0
reset value: unpredictable															
Bits	Access	Description													
31:0	Read / Write	Access permissions for all 16 domains - The meaning of each field can be found in the <i>ARM Architecture Reference Manual</i> .													

7.2.5 Register 4: Reserved

Register 4 is reserved. Reading and writing this register yields unpredictable results.

7.2.6 Register 5: Fault Status Register

The Fault Status Register (FSR) indicates which fault has occurred, which could be either a prefetch abort or a data abort. Bit 10 extends the encoding of the status field for prefetch aborts and data aborts. The definition of the extended status field is found in [Section 2.3.4, “Event Architecture” on page 2-32](#). 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 will update this register with the source of the exception. Software is not required to clear these fields.

Table 7-10. Fault Status 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																																
																												X	D	0	Domain	Status
reset value: unpredictable																																
Bits	Access	Description																														
31:11	Read-unpredictable / Write-as-Zero	Reserved																														
10	Read / Write	Status Field Extension (X) This bit is used to extend the encoding of the Status field, when there is a prefetch abort and when there is a data abort. The definition of this field can be found in Section 2.3.4, “Event Architecture” on page 2-32																														
9	Read / Write	Debug Event (D) This flag indicates a debug event has occurred and that the cause of the debug event is found in the MOE field of the debug control register (CP14, register 10)																														
8	Read-as-zero / Write-as-Zero	= 0																														
7:4	Read / Write	Domain - 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																														

7.2.7 Register 6: Fault address Register

Table 7-11. Fault Address 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																															
Fault Virtual Address																															
reset value: unpredictable																															
Bits	Access	Description																													
31:0	Read / Write	Fault Virtual Address - Contains the MVA of the data access that caused the memory abort																													

7.2.8 Register 7: Cache Functions

This register should be 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 core 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 will not generate any of the following:

- translation faults
- domain faults
- permission faults

The invalidate instruction cache line command does not invalidate the BTB. If 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 7-12](#) work regardless of whether the cache is enabled or disabled.

Since the Clean DCache Line function reads from the data cache, it is capable of generating a parity fault. The other operations will not generate parity faults.

Table 7-12. Cache Functions

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

The line-allocate command allocates a tag into the data cache specified by bits [31:5] of Rd. If a valid dirty line (with a different MVA) already exists at this location it will be evicted. The 32 bytes of data associated with the newly allocated line are not initialized and therefore will generate unpredictable results if 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. Specific uses of these commands can be found in [Chapter 6, "Data Cache"](#).

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.
- If 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 will never get evicted. However, if a valid store is made to that line it will be marked as “dirty” and will get written back to external memory if another line is allocated to the same cache location. This eviction will produce unpredictable results.

To avoid this situation, the line-allocate operation should only be used if one of the following can be guaranteed:

- The virtual address associated with this command is not one that will be 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.

7.2.9 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 7-13](#) work regardless of whether the TLB is enabled or disabled.

This register should be accessed as write-only. Reads from this register, as with an MRC, have an undefined effect.

Table 7-13. 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

7.2.10 Register 9: Cache Lock Down

Register 9 is used for locking down entries into the instruction cache and data cache. (The protocol for locking down entries can be found in [Chapter 6, “Data Cache”](#).)

[Table 7-14](#) shows the command for locking down entries in the instruction and data cache. The entry to lock in the instruction cache is specified by the virtual address in Rd. The data cache locking mechanism follows a different procedure than the instruction cache. 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 7-15](#).

Lock/unlock operations on a disabled cache have an undefined effect.

Read and write access is allowed to the data cache lock register bit[0]. All other accesses to register 9 should be write-only; reads, as with an MRC, have an undefined effect.

Table 7-14. 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 7-15. Data Cache Lock 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		
		L
reset value: writable bits set to 0		
Bits	Access	Description
31:1	Read-unpredictable / Write-as-Zero	Reserved
0	Read / Write	Data Cache Lock Mode (L) 0 = No locking occurs 1 = Any fill into the data cache while this bit is set gets locked in

7.2.11 Register 10: TLB Lock Down

Register 10 is used for locking down entries into the instruction TLB, and data TLB. (The protocol for locking down entries can be found in [Chapter 3, “Memory Management”](#).) Lock/unlock operations on a TLB when the MMU is disabled have an undefined effect.

This register should be accessed as write-only. Reads from this register, as with an MRC, have an undefined effect.

Table 7-16 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 7-16. 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

7.2.12 Register 11-12: Reserved

These registers are reserved. Reading and writing them yields unpredictable results.

7.2.13 Register 13: Process ID

The Intel XScale® core supports 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 replaces 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. If 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 core would not require invalidating on a process switch.

Table 7-17. 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 7-18. Process ID 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		
Process ID		
reset value: 0x0000,0000		
Bits	Access	Description
31:25	Read / Write	Process ID - This field is used for remapping the virtual address when bits 31-25 of the virtual address are zero.
24:0	Read-as-Zero / Write-as-Zero	Reserved - Should be programmed to zero for future compatibility

7.2.13.1 The PID Register Affect On Addresses

All addresses generated and used by User Mode code are eligible for being “PIDified” 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 (see [Section 7.2.14, “Register 14: Breakpoint Registers” on page 7-93](#)) 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 (see [Chapter 3, “Memory Management”](#)) are MVAs.

7.2.14 Register 14: Breakpoint Registers

The Intel XScale® core 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).

Refer to [Chapter 9, “Software Debug”](#) for more information on these features of the Intel XScale® core.

Table 7-19. Accessing the Debug Registers

Function	opcode_2	CRm	Instruction
Access Instruction Breakpoint Control Register 0 (IBCR0)	0b000	0b1000	MRC p15, 0, Rd, c14, c8, 0 ; read MCR p15, 0, Rd, c14, c8, 0 ; write
Access Instruction Breakpoint Control Register 1 (IBCR1)	0b000	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

7.2.15 Register 15: Coprocessor Access Register

This register is selected when *opcode_2* = 0 and *CRm* = 1.

This register controls access rights to all the coprocessors in the system except for CP15 and CP14. Both CP15 and CP14 can only be accessed in privilege mode. This register is accessed with an MCR or MRC with the *CRm* field set to 1.

This register controls access to CP0 and other coprocessors (CP1 through CP13) that may exist in an ASSP. (See the Intel XScale® core implementation option section of the ASSP architecture specification for a list of coprocessors that may have been implemented.) A typical use for this register is for an operating system to control resource sharing among applications. Initially, all applications are denied access to shared resources by clearing the appropriate coprocessor bit in the Coprocessor Access Register. An application may request the use of a shared resource (e.g., the accumulator in CP0) by issuing an access to the resource, which will result in an undefined exception. The operating system may grant access to this coprocessor by setting the appropriate bit in the Coprocessor Access Register and return to the application where the access is retried.

Sharing resources among different applications requires a state saving mechanism. Two possibilities are:

- The operating system, during a context switch, could save the state of the coprocessor if 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 7-1. Disallowing access to CP0

```
;; The following code clears bit 0 of the CPAR.  
;; This will cause 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 7-20. Coprocessor Access 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
																0	0	CP13	CP12	CP11	CP10	CP9	CP8	CP7	CP6	CP5	CP4	CP3	CP2	CP1	CP0
reset value: 0x0000,0000																															
Bits	Access	Description																													
31:16	Read-unpredictable / Write-as-Zero	Reserved - Should be programmed to zero for future compatibility																													
15:14	Read-as-Zero/Write-as-Zero	Reserved - Should be programmed to zero for future compatibility																													
13:1	Read / Write	Coprocessor Access Rights- Each bit in this field corresponds to the access rights for each coprocessor. Refer to the Intel XScale® core implementation option section of the ASSP architecture specification to find out which, if any, coprocessors exist and for the definition of these bits.																													
0	Read / Write	Coprocessor Access Rights- This bit corresponds to the access rights for CP0. 0 = Access denied. Any attempt to access the corresponding coprocessor will generate an undefined exception. 1 = Access allowed. Includes read and write accesses.																													

7.3 CP14 Registers

CP14 contains software debug registers, clock and power management registers and the performance monitor registers.

All other registers are reserved in CP14. Reading and writing them yields unpredictable results.

7.3.1 Performance Monitoring Registers

There are two variants of the performance monitoring facility; the number, location and definition of the registers are different between them. Software can determine which variant it is running on by examining the CoreGen field of Coprocessor 15, ID Register (bits 15:13). (See [Table 7-4, “ID Register”](#) on page 7-81 for more details.) A CoreGen value of 0x1 is referred to as XSC1 and a value of 0x2 is referred to as XSC2. The main difference between the two is that XSC1 has two 32-bit performance counters while XSC2 has four 32-bit performance counters.

7.3.1.1 XSC1 Performance Monitoring Registers

The performance monitoring unit in XSC1 contains a control register (PMNC), a clock counter (CCNT) and two event counters (PMN0 and PMN1). The format of these registers can be found in [Chapter 8, “Performance Monitoring”](#), along with a description on how to use the performance monitoring facility.

Opcode_2 and CRm should be zero.

Table 7-21. Accessing the XSC1 Performance Monitoring Registers

Description	CRn Register#	CRm Register#	Instruction
(PMNC) Performance Monitor Control Register	0b0000	0b0000	Read: MRC p14, 0, Rd, c0, c0, 0 Write: MCR p14, 0, Rd, c0, c0, 0
(CCNT) Clock Counter Register	0b0001	0b0000	Read: MRC p14, 0, Rd, c1, c0, 0 Write: MCR p14, 0, Rd, c1, c0, 0
(PMN0) Performance Count Register 0	0b0010	0b0000	Read: MRC p14, 0, Rd, c2, c0, 0 Write: MCR p14, 0, Rd, c2, c0, 0
(PMN1) Performance Count Register 1	0b0011	0b0000	Read: MRC p14, 0, Rd, c3, c0, 0 Write: MCR p14, 0, Rd, c3, c0, 0

7.3.1.2 XSC2 Performance Monitoring Registers

The performance monitoring unit in XSC2 contains a control register (PMNC), a clock counter (CCNT), interrupt enable register (INTEN), overflow flag register (FLAG), event selection register (EVTSEL) and four event counters (PMN0 through PMN3). The format of these registers can be found in [Chapter 8, "Performance Monitoring"](#), along with a description on how to use the performance monitoring facility.

Opcode_2 should be zero on all accesses.

These registers can't be accessed by **LDC** and **STC** coprocessor instructions.

Table 7-22. Accessing the XSC2 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

7.3.2 Clock and Power Management Registers

These registers contain functions for managing the core clock and power.

Power management modes are supported through the PWRMODE Register (CRn = 0x7, CRm = 0x0). The function and definition of these modes is defined by the ASSP. The user should refer to the Intel XScale® core implementation option section of the ASSP architecture specification for specifics on the use of these registers.

To enter any of these modes, write the appropriate data to the PWRMODE register. Software may read this register, but since software only runs during ACTIVE mode, it will always read zeroes from the **M** field.

Table 7-23. PWRMODE 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																															
																															M
reset value: writable bits set to 0																															
Bits	Access	Description																													
31:4	Read-unpredictable / Write-as-Zero	Reserved																													
3:0	Read / Write	Mode (M) 0 = ACTIVE All other values are defined by the ASSP																													

Software can change core clock frequency by writing to the CCLKCFG register (CRn = 0x6, CRm = 0x0). This function informs the clocking unit (located external to the core) to change core clock frequency. Software can read CCLKCFG to determine current operating frequency. Exact definition of this register can be found in the Intel XScale® core implementation option section of the ASSP architecture specification.

Table 7-24. Clock and Power Management

Function	Data	Instruction
Power Mode Function (Defined by ASSP)	Defined by ASSP	MCR p14, 0, Rd, c7, c0, 0
Read CCLKCFG	ignored	MRC p14, 0, Rd, c6, c0, 0
Write CCLKCFG	CCLKCFG value	MCR p14, 0, Rd, c6, c0, 0

Table 7-25. CCLKCFG 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																															
																															CCLKCFG
reset value: unpredictable																															
Bits	Access	Description																													
31:4	Read-unpredictable / Write-as-Zero	Reserved																													
3:0	Read / Write	Core Clock Configuration (CCLKCFG) This field is used to configure the core clock frequency and is defined by the ASSP.																													

7.3.3 Software Debug Registers

Software debug is supported by address breakpoint registers (Coprocessor 15, register 14), serial communication over the JTAG interface and a trace buffer. Registers 8, 9 and 14 are used for the serial interface, register 10 is for general control and registers 11 through 13 support a 256 entry trace buffer. These registers are explained in more detail in [Chapter 9, “Software Debug”](#).

Opcode_2 and CRm should be zero.

Table 7-26. Accessing the Debug Registers

Function	CRn (Register #)	Instruction
Transmit Debug Register (TX)	0b1000	MCR p14, 0, Rd, c8, c0, 0
Receive Debug Register (RX)	0b1001	MRC p14, 0, Rd, c9, c0, 0
Debug Control and Status Register (DBGCSR)	0b1010	MCR p14, 0, Rd, c10, c0, 0 MRC p14, 0, Rd, c10, c0, 0
Trace Buffer Register (TBREG)	0b1011	MRC p14, 0, Rd, c11, c0, 0
Checkpoint 0 Register (CHKPT0)	0b1100	MCR p14, 0, Rd, c12, c0, 0 MRC p14, 0, Rd, c12, c0, 0
Checkpoint 1 Register (CHKPT1)	0b1101	MCR p14, 0, Rd, c13, c0, 0 MRC p14, 0, Rd, c13, c0, 0
Transmit and Receive Debug Control Register	0b1110	MCR p14, 0, Rd, c14, c0, 0 MRC p14, 0, Rd, c14, c0, 0

This Page Intentionally Left Blank

Performance Monitoring

8

This chapter describes the performance monitoring facility of the Intel XScale® core. The events that are monitored can provide performance information for compiler writers, system application developers and software programmers.

There are two variants of the performance monitoring facility; the number, location and definition of the registers are different between them. Software can determine which variant it is running on by examining the CoreGen field of Coprocessor 15, ID Register (bits 15:13). (See [Table 7-4, “ID Register”](#) on page 7-81 for more details.) A CoreGen value of 0x1 is referred to as XSC1 and a value of 0x2 is referred to as XSC2. The main difference between the two is that XSC1 has two 32-bit performance counters while XSC2 has four 32-bit performance counters.

8.1 Overview

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

The Intel XScale® core can monitor either occurrence events or duration events. When counting occurrence events, a counter is incremented each time a specified event takes place and when measuring duration, a counter counts the number of processor clocks that occur while a specified condition is true. If any of the counters overflow, an interrupt request will occur if it's enabled. (What happens to the interrupt request is definable by the ASSP, which typically contains an interrupt controller that handles priority, masking, steering to FIQ or IRQ, etc. Refer to the Intel XScale® core implementation option section of the ASSP architecture specification for more details.) Each counter has its own interrupt request enable. The counters continue to monitor events even after an overflow occurs, until disabled by software.

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

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

All of the performance monitoring registers are accessible through Coprocessor 14 (CP14). Access is allowed in privileged mode only. Note that these registers can't be accessed with **LDC** or **STC** coprocessor instructions.

8.2 XSC1 Register Description (2 counter variant)

Table 8-1 contains details on accessing these registers with **MRC** and **MCR** coprocessor instructions.

Table 8-1. XSC1 Performance Monitoring Registers

Description	CRn Register#	CRm Register#	Instruction
(PMNC) Performance Monitor Control Register	0b0000	0b0000	Read: MRC p14, 0, Rd, c0, c0, 0 Write: MCR p14, 0, Rd, c0, c0, 0
(CCNT) Clock Counter Register	0b0001	0b0000	Read: MRC p14, 0, Rd, c1, c0, 0 Write: MCR p14, 0, Rd, c1, c0, 0
(PMN0) Performance Count Register 0	0b0010	0b0000	Read: MRC p14, 0, Rd, c2, c0, 0 Write: MCR p14, 0, Rd, c2, c0, 0
(PMN1) Performance Count Register 1	0b0011	0b0000	Read: MRC p14, 0, Rd, c3, c0, 0 Write: MCR p14, 0, Rd, c3, c0, 0

8.2.1 Clock Counter (CCNT; CP14 - Register 1)

The format of CCNT is shown in Table 8-6. The clock counter is reset to '0' by Performance Monitor Control Register (PMNC) or can be set to a predetermined value by directly writing to it. It counts core clock cycles. When CCNT reaches its maximum value 0xFFFF,FFFF, the next clock cycle will cause it to roll over to zero and set the overflow flag (bit 6) in PMNC. An IRQ or FIQ will be reported if it is enabled via bit 6 in the PMNC register.

Table 8-2. Clock Count Register (CCNT)

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		
Clock Counter		
reset value: unpredictable		
Bits	Access	Description
31:0	Read / Write	32-bit clock counter - Reset to '0' by PMNC register. When the clock counter reaches its maximum value 0xFFFF,FFFF, the next cycle will cause it to roll over to zero and generate an IRQ or FIQ if enabled.

8.2.2 Performance Count Registers (PMN0 - PMN1; CP14 - Register 2 and 3, Respectively)

There are two 32-bit event counters; their format is shown in Table 8-7. The event counters are reset to '0' by the PMNC register or can be set to a predetermined value by directly writing to them. When an event counter reaches its maximum value 0xFFFF,FFFF, the next event it needs to count will cause it to roll over to zero and set the overflow flag (bit 8 or 9) in PMNC. An IRQ or FIQ interrupt will be reported if it is enabled via bit 4 or 5 in the PMNC register.

Table 8-3. Performance Monitor Count Register (PMN0 and PMN1)

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		
Event Counter		
reset value: unpredictable		
Bits	Access	Description
31:0	Read / Write	32-bit event counter - Reset to '0' by PMNC register. When an event counter reaches its maximum value 0xFFFF,FFFF, the next event it needs to count will cause it to roll over to zero and generate an IRQ interrupt if enabled.

8.2.3 Extending Count Duration Beyond 32 Bits

To increase the monitoring duration, software can extend the count duration beyond 32 bits by counting the number of overflow interrupts each 32-bit counter generates. This can be done in the interrupt service routine (ISR) where an increment to some memory location every time the interrupt occurs will enable longer durations of performance monitoring. This does intrude upon program execution but is negligible, since the ISR execution time is in the order of tens of cycles compared to the number of cycles it took to generate an overflow interrupt (2^{32}).

8.2.4 Performance Monitor Control Register (PMNC)

The performance monitor control register (PMNC) is a coprocessor register that:

- controls which events PMN0 and PMN1 will monitor
- detects which counter overflowed
- enables/disables interrupt reporting
- extends CCNT counting by six more bits (cycles between counter rollover = 2^{38})
- resets all counters to zero
- and enables the entire mechanism

Table 8-8 shows the format of the PMNC register.

Table 8-4. Performance Monitor Control Register (CP14, register 0)

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																															
evtCount1								evtCount0								flag				inten			D	C	P	E					
reset value: E and inten are 0, others unpredictable																															
Bits	Access	Description																													
31:28	Read-unpredictable / Write-as-0	Reserved																													
27:20	Read / Write	Event Count1 - identifies the source of events that PMN1 counts. See Table 8-12 for a description of the values this field may contain.																													
19:12	Read / Write	Event Count0 - identifies the source of events that PMN0 counts. See Table 8-12 for a description of the values this field may contain.																													
11	Read-unpredictable / Write-as-0	Reserved																													
10:8	Read / Write	Overflow/Interrupt Flag - identifies which counter overflowed Bit 10 = clock counter overflow flag Bit 9 = performance counter 1 overflow flag Bit 8 = performance counter 0 overflow flag Read Values: 0 = no overflow 1 = overflow has occurred Write Values: 0 = no change 1 = clear this bit																													
7	Read-unpredictable / Write-as-0	Reserved																													
6:4	Read / Write	Interrupt Enable - used to enable/disable interrupt reporting for each counter Bit 6 = clock counter interrupt enable 0 = disable interrupt 1 = enable interrupt Bit 5 = performance counter 1 interrupt enable 0 = disable interrupt 1 = enable interrupt Bit 4 = performance counter 0 interrupt enable 0 = disable interrupt 1 = enable interrupt																													
3	Read / Write	Clock Counter Divider (D) - 0 = CCNT counts every processor clock cycle 1 = CCNT counts every 64 th processor clock cycle																													
2	Read-unpredictable / Write	Clock Counter Reset (C) - 0 = no action 1 = reset the clock counter to '0x0'																													
1	Read-unpredictable / Write	Performance Counter Reset (P) - 0 = no action 1 = reset both performance counters to '0x0'																													
0	Read / Write	Enable (E) - 0 = all 3 counters are disabled 1 = all 3 counters are enabled																													

8.2.4.1 Managing PMNC

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

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

8.3 XSC2 Register Description (4 counter variant)

Table 8-5 contains details on accessing these registers with **MRC** and **MCR** coprocessor instructions.

Table 8-5. 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

8.3.1 Clock Counter (CCNT)

The format of CCNT is shown in Table 8-6. The clock counter is reset to '0' by setting bit 2 in the Performance Monitor Control Register (PMNC) or can be set to a predetermined value by directly writing to it. It counts core clock cycles. When CCNT reaches its maximum value 0xFFFF,FFFF, the next clock cycle will cause it to roll over to zero and set the overflow flag (bit 0) in FLAG. An interrupt request will occur if it is enabled via bit 0 in INTEN.

Table 8-6. Clock Count Register (CCNT)

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																															
Clock Counter																															
reset value: unpredictable																															
Bits	Access	Description																													
31:0	Read / Write	32-bit clock counter - Reset to '0' by PMNC register. When the clock counter reaches its maximum value 0xFFFF,FFFF, the next cycle will cause it to roll over to zero and generate an interrupt request if enabled.																													

8.3.2 Performance Count Registers (PMN0 - PMN3)

There are four 32-bit event counters; their format is shown in Table 8-7. The event counters are reset to '0' by setting bit 1 in the PMNC register or can be set to a predetermined value by directly writing to them. When an event counter reaches its maximum value 0xFFFF,FFFF, the next event it needs to count will cause it to roll over to zero and set its corresponding overflow flag (bit 1,2,3 or 4) in FLAG. An interrupt request will be generated if its corresponding interrupt enable (bit 1,2,3 or 4) is set in INTEN.

Table 8-7. Performance Monitor Count Register (PMN0 - PMN3)

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		
Event Counter		
reset value: unpredictable		
Bits	Access	Description
31:0	Read / Write	32-bit event counter - Reset to '0' by PMNC register. When an event counter reaches its maximum value 0xFFFF,FFFF, the next event it needs to count will cause it to roll over to zero and generate an interrupt request if enabled.

8.3.3 Performance Monitor Control Register (PMNC)

The performance monitor control register (PMNC) is a coprocessor register that:

- contains the PMU ID
- extends CCNT counting by six more bits (cycles between counter rollover = 2^{38})
- resets all counters to zero
- and enables the entire mechanism

Table 8-8 shows the format of the PMNC register.

Table 8-8. Performance Monitor 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																															
ID																								D	C	P	E				
reset value: E = 0, ID = 0x14, others unpredictable																															
Bits	Access	Description																													
31:24	Read / Write Ignored	Performance Monitor Identification (ID) - XSC2 = 0x14																													
23:4	Read-unpredictable / Write-as-0	Reserved																													
3	Read / Write	Clock Counter Divider (D) - 0 = CCNT counts every processor clock cycle 1 = CCNT counts every 64 th processor clock cycle																													
2	Read-unpredictable / Write	Clock Counter Reset (C) - 0 = no action 1 = reset the clock counter to '0x0'																													
1	Read-unpredictable / Write	Performance Counter Reset (P) - 0 = no action 1 = reset all performance counters to '0x0'																													
0	Read / Write	Enable (E) - 0 = all counters are disabled 1 = all counters are enabled																													

8.3.4 Interrupt Enable Register (INTEN)

Each counter can generate an interrupt request when it overflows. INTEN enables interrupt requesting for each counter.

Table 8-9. Interrupt Enable 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																																
																												P	P	P	P	C
																												3	2	1	0	
reset value: [4:0] = 0b00000, others unpredictable																																
Bits	Access	Description																														
31:5	Read-unpredictable / Write-as-0	Reserved																														
4	Read / Write	PMN3 Interrupt Enable (P3) - 0 = disable interrupt 1 = enable interrupt																														
3	Read / Write	PMN2 Interrupt Enable (P2) - 0 = disable interrupt 1 = enable interrupt																														
2	Read / Write	PMN1 Interrupt Enable (P1) - 0 = disable interrupt 1 = enable interrupt																														
1	Read / Write	PMN0 Interrupt Enable (P0) - 0 = disable interrupt 1 = enable interrupt																														
0	Read / Write	CCNT Interrupt Enable (C) - 0 = disable interrupt 1 = enable interrupt																														

8.3.5 Overflow Flag Status Register (FLAG)

FLAG identifies which counter has overflowed and also indicates an interrupt has been requested if the overflowing counter's corresponding interrupt enable bit (contained within INTEN) is asserted. An overflow is cleared by writing a '1' to the overflow bit.

Table 8-10. Overflow Flag Status 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																																
																												P	P	P	P	C
reset value: [4:0] = 0b00000, others unpredictable																																
Bits	Access	Description																														
31:5	Read-unpredictable / Write-as-0	Reserved																														
4	Read / Write	PMN3 Overflow Flag (P3) - Read Values: 0 = no overflow 1 = overflow has occurred Write Values: 0 = no change 1 = clear this bit																														
3	Read / Write	PMN2 Overflow Flag (P2) - Read Values: 0 = no overflow 1 = overflow has occurred Write Values: 0 = no change 1 = clear this bit																														
2	Read / Write	PMN1 Overflow Flag (P1) - Read Values: 0 = no overflow 1 = overflow has occurred Write Values: 0 = no change 1 = clear this bit																														
1	Read / Write	PMN0 Overflow Flag (P0) - Read Values: 0 = no overflow 1 = overflow has occurred Write Values: 0 = no change 1 = clear this bit																														
0	Read / Write	CCNT Overflow Flag (C) - Read Values: 0 = no overflow 1 = overflow has occurred Write Values: 0 = no change 1 = clear this bit																														

8.3.6 Event Select Register (EVTSEL)

EVTSEL is used to select events for PMN0, PMN1, PMN2 and PMN3. Refer to [Table 8-12](#), "Performance Monitoring Events" on page 8-113 for a list of possible events.

Table 8-11. Event Select 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															
evtCount3				evtCount2				evtCount1				evtCount0			
reset value: unpredictable															
Bits	Access	Description													
31:24	Read / Write	Event Count 3 (evtCount3) - Identifies the source of events that PMN3 counts. See Table 8-12 for a description of the values this field may contain.													
23:16	Read / Write	Event Count 2 (evtCount2) - Identifies the source of events that PMN2 counts. See Table 8-12 for a description of the values this field may contain.													
15:8	Read / Write	Event Count 1 (evtCount1) - Identifies the source of events that PMN1 counts. See Table 8-12 for a description of the values this field may contain.													
7:0	Read / Write	Event Count 0 (evtCount0) - Identifies the source of events that PMN0 counts. See Table 8-12 for a description of the values this field may contain.													

8.3.7 Managing the Performance Monitor

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

- An interrupt request will be generated when a counter's overflow flag is set and its associated interrupt enable bit is set in INTEN. The interrupt request will remain asserted until software clears the overflow flag by writing a one to the flag that is set. (Note that the product specific interrupt unit and the CPSR must have enabled the interrupt in order for software to receive it.) The interrupt request can also be deasserted by clearing the corresponding interrupt enable bit. Disabling the facility (PMNC.E) doesn't deassert the interrupt request.
- The counters continue to record events even after they overflow.
- To change an event for a performance counter, first disable the facility (PMNC.E) and then modify EVTSEL. Not doing so will cause unpredictable results.
- Simultaneously resetting and disabling the counter will cause unpredictable results. To disable an event for a performance counter and reset the event counter, first disable the facility (PMNC.E) and then reset the counter.
- To increase the monitoring duration, software can extend the count duration beyond 32 bits by counting the number of overflow interrupts each 32-bit counter generates. This can be done in the interrupt service routine (ISR) where an increment to some memory location every time the interrupt occurs will enable longer durations of performance monitoring. This does intrude upon program execution but is negligible, since the ISR execution time is in the order of tens of cycles compared to the number of cycles it took to generate an overflow interrupt (2^{32}).
- Power can be saved by selecting event 0xFF for any unused event counter. This only applies when other event counters are in use. When the performance monitor is not used at all (PMNC.E = 0x0), hardware ensures minimal power consumption.

8.4 Performance Monitoring Events

Table 8-12 lists events that may be monitored. Each of the Performance Monitor Count Registers can count any listed event. Software selects which event is counted by each PMNx register by programming the evtCountx fields.

Table 8-12. Performance Monitoring Events

Event Number (evtCountx)	Event Definition
0x0	Instruction cache miss requires fetch from external memory.
0x1	Instruction cache cannot deliver an instruction. This could indicate an ICache miss or an ITLB miss. This event will occur every cycle in which the condition is present.
0x2	Stall due to a data dependency. This event will occur every cycle in which the condition is present.
0x3	Instruction TLB miss.
0x4	Data TLB miss.
0x5	Branch instruction executed, branch may or may not have changed program flow. (Counts only B and BL instructions, in both ARM and Thumb mode.)
0x6	Branch mispredicted. (Counts only B and BL instructions, in both ARM and Thumb mode.)
0x7	Instruction executed.
0x8	Stall because the data cache buffers are full. This event will occur every cycle in which the condition is present.
0x9	Stall because the data cache buffers are full. This event will occur once for each contiguous sequence of this type of stall.
0xA	Data cache access, not including Cache Operations (defined in Section 7.2.8)
0xB	Data cache miss, not including Cache Operations (defined in Section 7.2.8)
0xC	Data cache write-back. This event occurs once for each 1/2 line (four words) that are written back from the cache.
0xD	Software changed the PC. All 'b', 'bl', 'blx', 'mov[s] pc, Rm', 'ldm Rn, {Rx, pc}', 'ldr pc, [Rm]', 'pop {pc}' will be counted. An 'mcr p<cp>, 0,pc, ...', will not. The count also does not increment when an event occurs and the PC changes to the event address, e.g., IRQ, FIQ, SWI, etc.
0x10 through 0x17	Defined by ASSP. See the Intel XScale® core implementation option section of the ASSP architecture specification for more details.
0xFF	Power saving event. This event deactivates the corresponding PMU event counter
all others	Reserved, unpredictable results

Some typical combinations of counted events are listed in this section and summarized in Table 8-13. In this section, we call such an event combination a *mode*.

Table 8-13. Some Common Uses of the PMU

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

Note: PMN0 and PMN1 were used for illustration purposes only. Given there are four event counters, more elaborate combination of events could be constructed. For example, one performance run could select 0xA, 0xB, 0xC, 0x9 events from which data cache performance statistics could be gathered (like hit rates, number of writebacks per data cache miss, and number of times the data cache buffers fill up per request).

8.4.1 Instruction Cache Efficiency Mode

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

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

Statistics derived from these two events:

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

8.4.2 Data Cache Efficiency Mode

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

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

PMN1 counts the number of data cache and mini-data cache misses. Cache operations do not contribute to this count. See [Section 7.2.8](#) for a description of these operations.

The statistic derived from these two events is:

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

8.4.3 Instruction Fetch Latency Mode

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

PMN1 counts the number of instruction fetch requests to external memory. Each of these requests loads 32 bytes at a time. This is the same event as measured in instruction cache efficiency mode.

Statistics derived from these two events:

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

8.4.4 Data/Bus Request Buffer Full Mode

The Data Cache has buffers available to service cache misses or uncachable accesses. For every memory request that the Data Cache receives from the processor core a buffer is speculatively allocated in case an external memory request is required or temporary storage is needed for an unaligned access. If no buffers are available, the Data Cache will stall the processor core. How often the Data Cache stalls depends on the performance of the bus external to the core and what the memory access latency is for Data Cache miss requests to external memory. If the core memory access latency is high, possibly due to starvation, these Data Cache buffers will become full. This performance monitoring mode is provided to see if the core is being starved of the external bus, which will effect the performance of the application running on the core.

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

Statistics derived from these two events:

- *The average number of cycles the processor stalled on a data-cache access that may overflow the data-cache buffers.* This is calculated by dividing PMN0 by PMN1. This statistic lets you know if the duration event cycles are due to many requests or are attributed to just a few requests. If the average is high then the Intel XScale® core may be starved of the external bus.
- *The percentage of total execution cycles the processor stalled because a Data Cache request buffer was not available.* This is calculated by dividing PMN0 by CCNT, which was used to measure total execution time.

8.4.5 Stall/Writeback Statistics

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

- **Load-use penalty:** attempting to use the result of a load before the load completes. To avoid the penalty, software should delay using the result of a load until it's available. This penalty shows the latency effect of data-cache access.
- **Multiply/Accumulate-use penalty:** attempting to use the result of a multiply or multiply-accumulate operation before the operation completes. Again, to avoid the penalty, software should delay using the result until it's available.
- **ALU use penalty:** there are a few isolated cases where back to back ALU operations may result in one cycle delay in the execution. These cases are defined in [Chapter 10, "Performance Considerations"](#).

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

Statistics derived from these two events:

- *The percentage of total execution cycles the processor stalled because of a data dependency.* This is calculated by dividing PMN0 by CCNT, which was used to measure total execution time. Often a compiler can reschedule code to avoid these penalties when given the right optimization switches.
- Total number of data writeback requests to external memory can be derived solely with PMN1.

8.4.6 Instruction TLB Efficiency Mode

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

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

Statistics derived from these two events:

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

8.4.7 Data TLB Efficiency Mode

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

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

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

The statistic derived from these two events is:

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

8.5 Multiple Performance Monitoring Run Statistics

There may be times when the number of events to be monitored exceed the number of counters. In this case, multiple performance monitoring runs can be done, capturing different events from each run. For example, the first run could monitor the events associated with instruction cache performance and the second run could monitor the events associated with data cache performance. By combining the results, statistics like total number of memory requests could be derived.

8.6 Examples

The same example is shown below for both variants (XSC1 and XSC2).

8.6.1 XSC1 Example (2 counter variant)

In this example, the events selected with the Instruction Cache Efficiency mode are monitored and CCNT is used to measure total execution time. Sampling time ends when PMN0 overflows which will generate an IRQ interrupt.

Example 8-1. Configuring the Performance Monitor

```

; Configure PMNC with the following values:
;   evtCount0 = 7, evtCount1 = 0 instruction cache efficiency
;   inten = 0x7set all counters to trigger an interrupt on
;         overflow
;   C = 1   reset CCNT register
;   P = 1   reset PMN0 and PMN1 registers
;   E = 1   enable counting
MOV  R0,#0x7777
MCR  P14,0,R0,C0,c0,0; write R0 to PMNC
; Counting begins

```

Counter overflow can be dealt with in the IRQ interrupt service routine as shown below:

Example 8-2. Interrupt Handling

```

IRQ_INTERRUPT_SERVICE_ROUTINE:
; Assume that performance counting interrupts are the only IRQ in the system
MRC  P14,0,R1,C0,c0,0; read the PMNC register
BIC  R2,R1,#1       ; clear the enable bit
MCR  P14,0,R2,C0,c0,0; clear interrupt flag and disable counting
MRC  P14,0,R3,C1,c0,0; read CCNT register
MRC  P14,0,R4,C2,c0,0; read PMN0 register
MRC  P14,0,R5,C3,c0,0; read PMN1 register

<process the results>
SUBS PC,R14,#4      ; return from interrupt

```

As an example, assume the following values in CCNT, PMN0, PMN1 and PMNC:

Example 8-3. Computing the Results

```

; Assume CCNT overflowed

CCNT = 0x0000,0020 ;Overflowed and continued counting
Number of instructions executed = PMN0 = 0x6AAA,AAAA
Number of instruction cache miss requests = PMN1 = 0x0555,5555
Instruction Cache miss-rate = 100 * PMN1/PMN0 = 5%
CPI = (CCNT + 2^32)/Number of instructions executed = 2.4 cycles/instruction

```

In the contrived example above, the instruction cache had a miss-rate of 5% and CPI was 2.4.

8.6.2 XSC2 Example (4 counter variant)

In this example, the events selected with the Instruction Cache Efficiency mode are monitored and CCNT is used to measure total execution time. Sampling time ends when PMN0 overflows which will generate an IRQ interrupt.

Example 8-4. Configuring the Performance Monitor

```

; Configure the performance monitor with the following values:
; EVTSEL.evtCount0 = 7, EVTSEL.evtCount1 = 0 instruction cache efficiency
; INTEN.inten = 0x7 set all counters to trigger an interrupt on overflow
; PMNC.C = 1 reset CCNT register
; PMNC.P = 1 reset PMN0 and PMN1 registers
; PMNC.E = 1 enable counting
MOV R0,#0x700
MCR P14,0,R0,C8,c1,0 ; setup EVTSEL
MOV R0,#0x7
MCR P14,0,R0,C4,c1,0 ; setup INTEN
MCR P14,0,R0,C0,c1,0 ; setup PMNC, reset counters & enable
; Counting begins

```

Counter overflow can be dealt with in the IRQ interrupt service routine as shown below:

Example 8-5. Interrupt Handling

```

IRQ_INTERRUPT_SERVICE_ROUTINE:
; Assume that performance counting interrupts are the only IRQ in the system
MRC P14,0,R1,C0,c1,0 ; read the PMNC register
BIC R2,R1,#1 ; clear the enable bit, preserve other bits in PMNC
MCR P14,0,R2,C0,c1,0 ; disable counting
MRC P14,0,R3,C1,c1,0 ; read CCNT register
MRC P14,0,R4,C0,c2,0 ; read PMN0 register
MRC P14,0,R5,C1,c2,0 ; read PMN1 register

<process the results>
SUBS PC,R14,#4 ; return from interrupt

```

As an example, assume the following values in CCNT, PMN0, PMN1 and PMNC:

Example 8-6. Computing the Results

```

; Assume CCNT overflowed

CCNT = 0x0000,0020 ;Overflowed and continued counting
Number of instructions executed = PMN0 = 0x6AAA,AAAA
Number of instruction cache miss requests = PMN1 = 0x0555,5555
Instruction Cache miss-rate = 100 * PMN1/PMN0 = 5%
CPI = (CCNT + 2^32)/Number of instructions executed = 2.4 cycles/instruction

```

In the contrived example above, the instruction cache had a miss-rate of 5% and CPI was 2.4.

Software Debug

9

This chapter describes the software debug and related features implemented in Elkhart, namely:

- debug modes, registers and exceptions.
- a serial debug communication link via the JTAG interface.
- a trace buffer.
- a mechanism and process for loading the instruction cache through JTAG.

9.1 Definitions

debug handler: Debug handler is the event handler that runs on Elkhart, when a debug event occurs.

debugger: The debugger is software that runs on a host system outside of Elkhart.

9.2 Debug Registers

CP15 Registers

CRn = 14; CRm = 8: instruction breakpoint register 0 (IBCR0)
CRn = 14; CRm = 9: instruction breakpoint register 1 (IBCR1)
CRn = 14; CRm = 0: data breakpoint register 0 (DBR0)
CRn = 14; CRm = 3: data breakpoint register 1 (DBR1)
CRn = 14; CRm = 4: data breakpoint control register (DBCON)

CP15 registers are accessible using MRC and MCR. CRn and CRm specify the register to access. The opcode_1 and opcode_2 fields are not used and should be set to 0.

CP14 Registers

CRn = 8; CRm = 0: TX Register (TX)
CRn = 9; CRm = 0: RX Register (RX)
CRn = 10; CRm = 0: Debug Control and Status Register (DCSR)
CRn = 11; CRm = 0: Trace Buffer Register (TBREG)
CRn = 12; CRm = 0: Checkpoint Register 0 (CHKPT0)
CRn = 13; CRm = 0: Checkpoint Register 1 (CHKPT1)
CRn = 14; CRm = 0: TXRX Control Register (TXRXCTRL)

CP14 registers are accessible using MRC, MCR, LDC and STC (CDP to any CP14 registers will cause an undefined instruction trap). CRn and CRm specify the register to access. The opcode_1 and opcode_2 fields are not used and should be set to 0.

Software access to all debug registers must be done from a privileged mode. User mode access will generate an undefined instruction exception. Specifying registers which do not exist has unpredictable results.

The TX and RX registers, certain bits in the TXRXCTRL register, and certain bits in the DCSR can be accessed by a debugger through the JTAG interface.

9.3 Introduction

The Elkhart debug unit, when used with a debugger application, allows software running on an Elkhart target to be debugged. The debug unit allows the debugger to stop program execution and re-direct execution to a debug handling routine. Once program execution has stopped, the debugger can examine or modify processor state, co-processor state, or memory. The debugger can then restart execution of the application.

On Elkhart, one of two debug modes can be used:

- Halt Mode
- Monitor Mode

9.3.1 Halt Mode

When the debug unit is configured for Halt Mode, the reset vector is overloaded to serve as the debug vector. A new processor mode, DEBUG mode (CPSR[4:0] = 0x15), is added to allow debug exceptions to be handled similarly to other types of ARM* exceptions.

When a debug exception occurs, the processor switches to debug mode and redirects execution to a debug handler, via the reset vector. After the debug handler begins execution, the debugger can communicate with the debug handler to examine or alter processor state or memory through the JTAG interface.

The debug handler can be downloaded and locked directly into the instruction cache through JTAG so external memory is not required to contain debug handler code.

9.3.2 Monitor Mode

In Monitor Mode, debug exceptions are handled like ARM prefetch aborts or ARM data aborts, depending on the cause of the exception.

When a debug exception occurs, the processor switches to abort mode and branches to a debug handler using the pre-fetch abort vector or data abort vector. The debugger then communicates with the debug handler to access processor state or memory contents.

9.4 Debug Control and Status Register (DCSR)

The DCSR register is the main control register for the debug unit. Table 9-1 shows the format of the register. The DCSR register can be accessed in privileged modes by software running on the core or by a debugger through the JTAG interface. Refer to Section 9.11.1, “SELDCSR JTAG Register” for details about accessing DCSR through JTAG.

Table 9-1. Debug Control and Status Register (DCSR) (Sheet 1 of 2)

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																															
GE H B						TF TI			TD TA TS TU TR						SA			MOE			M E										
Bits	Access	Description	Reset Value	TRST Value																											
31	SW Read / Write JTAG Read-Only	Global Enable (GE) 0: disables all debug functionality 1: enables all debug functionality	0	unchanged																											
30	SW Read Only JTAG Read / Write	Halt Mode (H) 0: Monitor Mode 1: Halt Mode	unchanged	0																											
29	SW Read-Only JTAG Read-Only	SOC Break (B) Value of SOC break core input	undefined	undefined																											
28:24	Read-undefined / Write-As-Zero	Reserved	undefined	undefined																											
23	SW Read Only JTAG Read / Write	Trap FIQ (TF)	unchanged	0																											
22	SW Read Only JTAG Read / Write	Trap IRQ (TI)	unchanged	0																											
21	Read-undefined / Write-As-Zero	Reserved	undefined	undefined																											
20	SW Read Only JTAG Read / Write	Trap Data Abort (TD)	unchanged	0																											
19	SW Read Only JTAG Read / Write	Trap Prefetch Abort (TA)	unchanged	0																											
18	SW Read Only JTAG Read / Write	Trap Software Interrupt (TS)	unchanged	0																											
17	SW Read Only JTAG Read / Write	Trap Undefined Instruction (TU)	unchanged	0																											
16	SW Read Only JTAG Read / Write	Trap Reset (TR)	unchanged	0																											

Table 9-1. Debug Control and Status Register (DCSR) (Sheet 2 of 2)

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																															
GE H B			TF TI			TD TA TS TU TR			SA			MOE			M E																
Bits	Access	Description	Reset Value	TRST Value																											
15:6	Read-undefined / Write-As-Zero	Reserved	undefined	undefined																											
5	SW Read / Write JTAG Read-Only	Sticky Abort (SA)	0	unchanged																											
4:2	SW Read / Write JTAG Read-Only	Method Of Entry (MOE) 000: Processor Reset 001: Instruction Breakpoint Hit 010: Data Breakpoint Hit 011: BKPT Instruction Executed 100: External Debug Event (JTAG Debug Break or SOC Debug Break) 101: Vector Trap Occurred 110: Trace Buffer Full Break 111: Reserved	0b000	unchanged																											
1	SW Read / Write JTAG Read-Only	Trace Buffer Mode (M) 0: wrap-around mode 1: fill-once mode	0	unchanged																											
0	SW Read / Write JTAG Read-Only	Trace Buffer Enable (E) 0: Disabled 1: Enabled	0	unchanged																											

9.4.1 Global Enable Bit (GE)

The Global Enable bit disables and enables all debug functionality (except the reset vector trap). Following a processor reset, this bit is clear so all debug functionality is disabled. When debug functionality is disabled, the BKPT instruction becomes a noop and external debug breaks, hardware breakpoints, and non-reset vector traps are ignored.

9.4.2 Halt Mode Bit (H)

The Halt Mode bit configures the debug unit for either Halt Mode or Monitor Mode.

9.4.3 SOC Break (B)

Reading the SOC Break bit returns the value of the SOC break input into the Intel XScale® core¹.

1. Use of the SOC break input to the core (used to generate SOC debug breaks) is product specific and is targeted towards chips that need system-on-a-chip debug capabilities. Refer to the ASSP architecture specification for more information.

9.4.4 Vector Trap Bits (TF, TI, TD, TA, TS, TU, TR)

The Vector Trap bits allow instruction breakpoints to be set on exception vectors without using up any of the breakpoint registers. When a bit is set, it acts as if an instruction breakpoint was set up on the corresponding exception vector. A debug exception is generated before the instruction in the exception vector executes.

Software running on Elkhart must set the Global Enable bit and the debugger must set the Halt Mode bit and the appropriate vector trap bit through JTAG to set up a non-reset vector trap.

To set up a reset vector trap, the debugger sets the Halt Mode bit and reset vector trap bit through JTAG. The Global Enable bit does not effect the reset vector trap. A reset vector trap can be set up before or during a processor reset. When processor reset is de-asserted, a debug exception occurs before the instruction in the reset vector executes.

9.4.5 Sticky Abort Bit (SA)

The Sticky Abort bit is only valid in Halt Mode. It indicates a data abort occurred within the Special Debug State (see [Section 9.5.1, "Halt Mode"](#)). Since Special Debug State disables all exceptions, a data abort exception does not occur. However, the processor sets the Sticky Abort bit to indicate a data abort was detected. The debugger can use this bit to determine if a data abort was detected during the Special Debug State. The sticky abort bit must be cleared by the debug handler before exiting the debug handler.

9.4.6 Method of Entry Bits (MOE)

The Method of Entry bits specify the cause of the most recent debug exception. When multiple exceptions occur in parallel, the processor places the highest priority exception (based on the priorities in [Table 9-2](#)) in the MOE field.

9.4.7 Trace Buffer Mode Bit (M)

The Trace Buffer Mode bit selects one of two trace buffer modes:

- Wrap-around mode - Trace buffer fills up and wraps around until a debug exception occurs.
- Fill-once mode - The trace buffer automatically generates a debug exception (trace buffer full break) when it becomes full.

9.4.8 Trace Buffer Enable Bit (E)

The Trace Buffer Enable bit enables and disables the trace buffer. Both DCSR.e and DCSR.ge must be set to enable the trace buffer. The processor automatically clears this bit to disable the trace buffer when a debug exception occurs. For more details on the trace buffer refer to [Section 9.12, "Trace Buffer"](#).

9.5 Debug Exceptions

A debug exception causes the processor to re-direct execution to a debug event handling routine. The Elkhart debug architecture defines the following debug exceptions:

- instruction breakpoint
- data breakpoint
- software breakpoint
- external debug break
- exception vector trap
- trace-buffer full break
- SOC debug break

When a debug exception occurs, the processor's actions depend on whether the debug unit is configured for Halt Mode or Monitor Mode.

Table 9-2 shows the priority of debug exceptions relative to other processor exceptions.

Table 9-2. Event Priority

Event	Priority
Reset	1 (highest)
Vector Trap	2
data abort (precise)	3
data bkpt	4
data abort (imprecise)	5
external debug break, trace-buffer full, SOC debug break	6
FIQ	7
IRQ	8
instruction breakpoint	9
pre-fetch abort	10
undef, SWI, software Bkpt	11

9.5.1 Halt Mode

The debugger turns on Halt Mode through the JTAG interface by scanning in a value that sets the bit in DCSR. The debugger turns off Halt Mode through JTAG, either by scanning in a new DCSR value or by a TRST. Processor reset does not effect the value of the Halt Mode bit.

When Halt Mode is active, the processor uses the reset vector as the debug vector. The debug handler and exception vectors can be downloaded directly into the instruction cache, to intercept the default vectors and reset handler, or they can be resident in external memory. Downloading into the instruction cache allows a system with memory problems, or no external memory, to be debugged. Refer to [Section 9.14, "Downloading Code in the Instruction Cache"](#) on page 9-154 for details about downloading code into the instruction cache.

During Halt Mode, software running on Elkhart cannot access DCSR, or any of hardware breakpoint registers, unless the processor is in Special Debug State (SDS), described below.

When a debug exception occurs during Halt Mode, or an SOC debug break occurs in Monitor Mode, the processor takes the following actions:

- disables the trace buffer
- sets DCSR.moe encoding
- processor enters a Special Debug State (SDS)
- R14_DBG is updated as follows:

Table 9-3. Halt Mode R14_DBG Updating

Debug Exception Type	DBG_r14 Value	
	ARM Mode	Thumb Mode
Data Breakpoint	PC of breakpointed memory instruction + 8	PC of breakpointed memory instruction + 6
Instruction Breakpoint, SW Breakpoint	PC of breakpointed instruction + 4	PC of breakpointed instruction + 4
Vector Trap	PC of trapped exception vector + 4	NA
Trace Buffer Full Break, SOC Debug Break, External Debug Break	PC of next instruction to execute + 4	PC of next instruction to execute + 4

- SPSR_dbg = CPSR
- CPSR[4:0] = 0b10101 (DEBUG mode)
- CPSR[5] = 0
- CPSR[6] = 1
- CPSR[7] = 1
- PC = 0x0 or 0xFFFF0000

The FSR.D bit, which is set for all Monitor Mode debug exceptions (including SOC debug breaks), is unaffected by debug exceptions during Halt Mode.

Following a debug exception, the processor switches to debug mode and enters SDS, which allows the following special functionality:

- All events are disabled. SWI or undefined instructions have unpredictable results. The processor ignores pre-fetch aborts, FIQ and IRQ (SDS disables FIQ and IRQ regardless of the enable values in the CPSR). The processor reports data aborts detected during SDS by setting the Sticky Abort bit in the DCSR, but does not generate an exception (processor also sets up FSR and FAR as it normally would for a data abort).
- Normally, during Halt Mode, software cannot write the hardware breakpoint registers or the DCSR. However, during the SDS, software has write access to the breakpoint registers (see [Section 9.6, “HW Breakpoint Resources”](#)) and the DCSR (see [Table 9-1, “Debug Control and Status Register \(DCSR\)”](#) on page 9-123).
- The IMMU is disabled. In Halt Mode, since the debug handler would typically be downloaded directly into the IC, it would not be appropriate to do TLB accesses or translation walks, since there may not be any external memory or if there is, the translation table or TLB may not contain a valid mapping for the debug handler code. To avoid these problems, the processor internally disables the IMMU during SDS.
- The PID is disabled for instruction fetches. This prevents fetches of the debug handler code from being remapped to a different address than where the code was downloaded.

The SDS remains in effect regardless of the processor mode. This allows the debug handler to switch to other modes, maintaining SDS functionality. Entering user mode may cause unpredictable behavior. The processor exits SDS following a CPSR restore operation.

When exiting, the debug handler should use:

```
subs pc, lr, #4
```

This restores CPSR, turns off all of SDS functionality, and branches to the target instruction.

9.5.2 Monitor Mode

In Monitor Mode, the processor handles debug exceptions like normal ARM exceptions, except for SOC debug breaks, which are handled like Halt Mode exceptions. If debug functionality is enabled and the processor is in Monitor Mode, debug exceptions cause either a data abort or a pre-fetch abort.

The following debug exceptions cause data aborts:

- data breakpoint
- external debug break
- trace-buffer full break

The following debug exceptions cause pre-fetch aborts:

- instruction breakpoint
- BKPT instruction

The processor ignores vector traps during Monitor Mode.

When an exception occurs in Monitor Mode, the processor takes the following actions:

- disables the trace buffer
- sets DCSR.moe encoding
- sets FSR[9]
- R14_DBG is updated as follows:

Table 9-4. Monitor Mode R14_DBG Updating

Debug Exception Type	DBG_r14 Value	
	ARM mode	Thumb mode
Data Breakpoint	PC of breakpointed memory instruction + 8	PC of breakpointed memory instruction + 6
Instruction Breakpoint, SW Breakpoint	PC of breakpointed instruction + 4	PC of breakpointed instruction + 4
Trace Buffer Full Break, External Debug Break	PC of next instruction to execute + 4	PC of next instruction to execute + 4

- SPSR_abt = CPSR
- CPSR[4:0] = 0b10111 (ABORT mode)
- CPSR[5] = 0
- CPSR[6] = unchanged
- CPSR[7] = 1
- PC = 0xC or 0xFFFF000C (for Prefetch Aborts),
PC = 0x10 or 0xFFFF0010 (for Data Aborts)

During abort mode, external debug breaks and trace buffer full breaks are internally pended. When the processor exits abort mode, either through a CPSR restore or a write directly to the CPSR, the pended debug breaks will immediately generate a debug exception. Any pending debug breaks are cleared out when any type of debug exception occurs. Note that SOC debug breaks are not pended during abort mode; they are handled immediately when detected.

When exiting, the debug handler should do a CPSR restore operation that branches to the next instruction to be executed in the program under debug.

9.6 HW Breakpoint Resources

The Elkhart debug architecture defines two instruction and two data breakpoint registers, denoted IBCR0, IBCR1, DBR0, and DBR1.

The instruction and data address breakpoint registers are 32-bit registers. The instruction breakpoint causes a break before execution of the target instruction. The data breakpoint causes a break after the memory access has been issued.

In this section Modified Virtual Address (MVA) refers to the virtual address ORed with the PID. Refer to [Section 7.2.13, “Register 13: Process ID” on page 7-91](#) for more details on the PID. The processor does not OR the PID with the specified breakpoint address prior to doing address comparison. This must be done by the programmer and written to the breakpoint register as the MVA. This applies to data and instruction breakpoints.

9.6.1 Instruction Breakpoints

The Debug architecture defines two instruction breakpoint registers (IBCR0 and IBCR1). The format of these registers is shown in [Table 9-5, “Instruction Breakpoint Address and Control Register \(IBCRx\)”](#). In ARM mode, the upper 30 bits contain a word aligned MVA to break on. In Thumb mode, the upper 31 bits contain a half-word aligned MVA to break on. In both modes, bit 0 enables and disables that instruction breakpoint register. Enabling instruction breakpoints while debug is globally disabled (DCSR.ge=0) may result in unpredictable behavior.

Table 9-5. Instruction Breakpoint Address and Control Register (IBCRx)

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		
IBCRx		
reset value: unpredictable address, disabled		
Bits	Access	Description
31:1	Read / Write	Instruction Breakpoint MVA in ARM* mode, IBCRx[1] is ignored
0	Read / Write	IBCRx Enable (E) - 0 = Breakpoint disabled 1 = Breakpoint enabled

An instruction breakpoint will generate a debug exception before the instruction at the address specified in the ICBR executes. When an instruction breakpoint occurs, the processor sets the DBCR.moe bits to 0b001.

Software must disable the breakpoint before exiting the handler. This allows the breakpointed instruction to execute after the exception is handled.

Single step execution is accomplished using the instruction breakpoint registers and must be completely handled in software (either on the host or by the debug handler).

9.6.2 Data Breakpoints

The Elkhart debug architecture defines two data breakpoint registers (DBR0, DBR1). The format of the registers is shown in Table 9-6.

Table 9-6. Data Breakpoint Register (DBRx)

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																															
DBRx																															
reset value: unpredictable																															
Bits	Access	Description																													
31:0	Read / Write	DBR0: Data Breakpoint MVA DBR1: Data Address Mask OR Data Breakpoint MVA																													

DBR0 is a dedicated data address breakpoint register. DBR1 can be programmed for 1 of 2 operations:

- data address mask
- second data address breakpoint

The DBCON register controls the functionality of DBR1, as well as the enables for both DBRs. DBCON also controls what type of memory access to break on.

Table 9-7. Data Breakpoint Controls Register (DBCON)

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																																
																												M			E1	E0
reset value: 0x00000000																																
Bits	Access	Description																														
31:9	Read-as-Zero / Write-ignored	Reserved																														
8	Read / Write	DBR1 Mode (M) - 0: DBR1 = Data Address Breakpoint 1: DBR1 = Data Address Mask																														
7:4	Read-as-Zero / Write-ignored	Reserved																														
3:2	Read / Write	DBR1 Enable (E1) - When DBR1 = Data Address Breakpoint 0b00: DBR1 disabled 0b01: DBR1 enabled, Store only 0b10: DBR1 enabled, Any data access, load or store 0b11: DBR1 enabled, Load only When DBR1 = Data Address Mask this field has no effect																														
1:0	Read / Write	DBR0 Enable (E0) - 0b00: DBR0 disabled 0b01: DBR0 enabled, Store only 0b10: DBR0 enabled, Any data access, load or store 0b11: DBR0 enabled, Load only																														

When DBR1 is programmed as a data address mask, it is used in conjunction with the address in DBR0. The bits set in DBR1 are ignored by the processor when comparing the address of a memory access with the address in DBR0. Using DBR1 as a data address mask allows a range of addresses to generate a data breakpoint. When DBR1 is selected as a data address mask, it is unaffected by the E1 field of DBCON. The mask is used only when DBR0 is enabled.

When DBR1 is programmed as a second data address breakpoint, it functions independently of DBR0. In this case, the DBCON.E1 controls DBR1.

A data breakpoint is triggered if the memory access matches the access type and the address of any byte within the memory access matches the address in DBRx. For example, LDR triggers a breakpoint if DBCON.E0 is 0b10 or 0b11, and the address of any of the 4 bytes accessed by the load matches the address in DBR0.

The processor does not trigger data breakpoints for the PLD instruction or any CP15, register 7,8,9, or 10 functions. Any other type of memory access can trigger a data breakpoint. For data breakpoint purposes the SWP and SWPB instructions are treated as stores - they will not cause a data breakpoint if the breakpoint is set up to break on loads only and an address match occurs.

On unaligned memory accesses, breakpoint address comparison is done on a word-aligned address (aligned down to word boundary).

When a memory access triggers a data breakpoint, the breakpoint is reported after the access is issued. The memory access will not be aborted by the processor. The actual timing of when the access completes with respect to the start of the debug handler depends on the memory configuration.

On a data breakpoint, the processor generates a debug exception and re-directs execution to the debug handler before the next instruction executes. The processor reports the data breakpoint by setting the DCSR.MOE to 0b010. The link register of a data breakpoint is always PC (of the next instruction to execute) + 4, regardless of whether the processor is configured for Monitor Mode or Halt Mode.

When setting a data breakpoint, the DBR registers should only be programmed while that data breakpoint register is disabled. Programming the DBR registers while they are enabled, may result in unpredictable behavior.

9.7 Software Breakpoints

Mnemonics: BKPT (See ARM Architecture Reference Manual, ARMv5T)

Operation: If DCSR[31] = 0, BKPT is a nop;
If DCSR[31] = 1, BKPT causes a debug exception

The processor handles the software breakpoint as described in [Section 9.5, “Debug Exceptions”](#) on [page 9-126](#).

9.8 Transmit/Receive Control Register (TXRXCTRL)

Communications between the debug handler and debugger are controlled through handshaking bits that ensures the debugger and debug handler make synchronized accesses to TX and RX. The debugger side of the handshaking is accessed through the DBGTX (Section 9.11.2, “DBGTX JTAG Register”) and DBGRX (Section 9.11.3, “DBGRX JTAG Register”) JTAG Data Registers, depending on the direction of the data transfer. The debug handler uses separate handshaking bits in TXRXCTRL register for accessing TX and RX.

The TXRXCTRL register also contains two other bits that support high-speed download. One bit indicates an overflow condition that occurs when the debugger attempts to write the RX register before the debug handler has read the previous data written to RX. The other bit is used by the debug handler as a branch flag during high-speed download.

All of the bits in the TXRXCTRL register are placed such that they can be read directly into the CC flags in the CPSR with an MRC (with Rd = PC). The subsequent instruction can then conditionally execute based on the updated CC value

Table 9-8. TX RX Control Register (TXRXCTRL)

		31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
R	O	D	T																															
R	V	D	R																															
Bits	Access	Description	Reset Value	TRST Value																														
31	SW Read-only / Write-ignored JTAG Write-only	RR RX Register Ready	0	0																														
30	SW Read / Write	OV RX overflow sticky flag	0	unchanged																														
29	SW Read-only/ Write-ignored JTAG Write-only	D High-speed download flag	unchanged	0																														
28	SW Read-only/ Write-ignored JTAG Write-only	TR TX Register Ready	0	unchanged																														
27:0	Read-as-Zero / Write-ignored	Reserved	undefined	undefined																														

9.8.1 RX Register Ready Bit (RR)

The debugger and debug handler use the RR bit to synchronize accesses to RX. Normally, the debugger and debug handler use a handshaking scheme that requires both sides to poll the RR bit. To support higher download performance for large amounts of data, a high-speed download handshaking scheme can be used in which only the debug handler polls the RR bit before accessing the RX register, while the debugger continuously downloads data.

Table 9-9 shows the normal handshaking used to access the RX register.

Table 9-9. Normal RX Handshaking

Debugger Actions
<p>Debugger wants to send data to debug handler.</p> <p>Before writing new data to the RX register, the debugger polls RR through JTAG until the bit is cleared.</p> <p>After the debugger reads a '0' from the RR bit, it scans data into JTAG to write to the RX register and sets the valid bit. The write to the RX register automatically sets the RR bit.</p>
Debug Handler Actions
<p>Debug handler is expecting data from the debugger.</p> <p>The debug handler polls the RR bit until it is set, indicating data in the RX register is valid.</p> <p>Once the RR bit is set, the debug handler reads the new data from the RX register. The read operation automatically clears the RR bit.</p>

When data is being downloaded by the debugger, part of the normal handshaking can be bypassed to allow the download rate to be increased. Table 9-10 shows the handshaking used when the debugger is doing a high-speed download. Note that before the high-speed download can start, both the debugger and debug handler must be synchronized, such that the debug handler is executing a routine that supports the high-speed download.

Although it is similar to the normal handshaking, the debugger polling of RR is bypassed with the assumption that the debug handler can read the previous data from RX before the debugger can scan in the new data.

Table 9-10. High-Speed Download Handshaking States

Debugger Actions
<p>Debugger wants to transfer code into the Elkhart system memory.</p> <p>Prior to starting download, the debugger must poll RR bit until it is clear. Once the RR bit is clear, indicating the debug handler is ready, the debugger starts the download.</p> <p>The debugger scans data into JTAG to write to the RX register with the download bit and the valid bit set. Following the write to RX, the RR bit and D bit are automatically set in TXRXCTRL.</p> <p>Without polling of RR to see whether the debug handler has read the data just scanned in, the debugger continues scanning in new data into JTAG for RX, with the download bit and the valid bit set.</p> <p>An overflow condition occurs if the debug handler does not read the previous data before the debugger completes scanning in the new data, (see Section 9.8.2, "Overflow Flag (OV)" for more details on the overflow condition).</p> <p>After completing the download, the debugger clears the D bit allowing the debug handler to exit the download loop.</p>
Debug Handler Actions
<p>Debug handler is in a routine waiting to write data out to memory. The routine loops based on the D bit in TXRXCTRL.</p> <p>The debug handler polls the RR bit until it is set. It then reads the Rx register, and writes it out to memory. The handler loops, repeating these operations until the debugger clears the D bit.</p>

9.8.2 Overflow Flag (OV)

The Overflow flag is a sticky flag that is set when the debugger writes to the RX register while the RR bit is set.

The flag is used during high-speed download to indicate that some data was lost. The assumption during high-speed download is that the time it takes for the debugger to shift in the next data word is greater than the time necessary for the debug handler to process the previous data word. So, before the debugger shifts in the next data word, the handler will be polling for that data.

However, if the handler incurs stalls that are long enough such that the handler is still processing the previous data when the debugger completes shifting in the next data word, an overflow condition occurs and the OV bit is set.

Once set, the overflow flag will remain set, until cleared by a write to TXXCTRL with an MCR. After the debugger completes the download, it can examine the OV bit to determine if an overflow occurred. The debug handler software is responsible for saving the address of the last valid store before the overflow occurred.

9.8.3 Download Flag (D)

The value of the download flag is set by the debugger through JTAG. This flag is used during high-speed download to replace a loop counter.

The download flag becomes especially useful when an overflow occurs. If a loop counter is used, and an overflow occurs, the debug handler cannot determine how many data words overflowed. Therefore the debug handler counter may get out of sync with the debugger - the debugger may finish downloading the data, but the debug handler counter may indicate there is more data to be downloaded - this may result in unpredictable behavior of the debug handler.

Using the download flag, the debug handler loops until the debugger clears the flag. Therefore, when doing a high-speed download, for each data word downloaded, the debugger should set the D bit.

9.8.4 TX Register Ready Bit (TR)

The debugger and debug handler use the TR bit to synchronize accesses to the TX register. The debugger and debug handler must poll the TR bit before accessing the TX register. [Table 9-11](#) shows the handshaking used to access the TX register.

Table 9-11. TX Handshaking

Debugger Actions
Debugger is expecting data from the debug handler. Before reading data from the TX register, the debugger polls the TR bit through JTAG until the bit is set. NOTE: while polling TR, the debugger must scan out the TR bit and the TX register data. Reading a '1' from the TR bit, indicates that the TX data scanned out is valid The action of scanning out data when the TR bit is set, automatically clears TR.
Debug Handler Actions
Debug handler wants to send data to the debugger (in response to a previous request). The debug handler polls the TR bit to determine when the TX register is empty (any previous data has been read out by the debugger). The handler polls the TR bit until it is clear. Once the TR bit is clear, the debug handler writes new data to the TX register. The write operation automatically sets the TR bit.

9.8.5 Conditional Execution Using TXRXCTRL

All of the bits in TXRXCTRL are placed such that they can be read directly into the CC flags using an MCR instruction. To simplify the debug handler, the TXRXCTRL register should be read using the following instruction:

```
mrc p14, 0, r15, C14, C0, 0
```

This instruction will directly update the condition codes in the CPSR. The debug handler can then conditionally execute based on each CC bit. [Table 9-12](#) shows the mnemonic extension to conditionally execute based on whether the TXRXCTRL bit is set or clear.

Table 9-12. TXRXCTRL Mnemonic Extensions

TXRXCTRL bit	mnemonic extension to execute if bit set	mnemonic extension to execute if bit clear
31 (to N flag)	MI	PL
30 (to Z flag)	EQ	NE
29 (to C flag)	CS	CC
28 (to V flag)	VS	VC

The following example is a code sequence in which the debug handler polls the TXRXCTRL handshaking bit to determine when the debugger has completed its write to RX and the data is ready for the debug handler to read.

```
loop: mrc    p14, 0, r15, c14, c0, 0 # read the handshaking bit in TXRXCTRL
      mrcmi p14, 0, r0, c9, c0, 0 # if RX is valid, read it
      bpl   loop                # if RX is not valid, loop
```

9.9 Transmit Register (TX)

The TX register is the debug handler transmit buffer. The debug handler sends data to the debugger through this register.

Table 9-13. TX 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																															
TX																															
reset value: unpredictable																TRST value: unchanged															
Bits	Access	Description																													
31:0	SW Read / Write JTAG Read-only	Debug handler writes data to send to debugger																													

Since the TX register is accessed by the debug handler (using MCR/MRC) and the debugger (through JTAG), handshaking is required to prevent the debug handler from writing new data before the debugger reads the previous data.

The TX register handshaking is described in [Table 9-11, “TX Handshaking” on page 9-137](#).

9.10 Receive Register (RX)

The RX register is the receive buffer used by the debug handler to get data sent by the debugger through the JTAG interface.

Table 9-14. RX 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																															
RX																															
reset value: unpredictable																TRST value: unpredictable															
Bits	Access	Description																													
31:0	SW Read-only JTAG Write-only	Software reads to receives data/commands from debugger																													

Since the RX register is accessed by the debug handler (using MRC) and the debugger (through JTAG), handshaking is required to prevent the debugger from writing new data to the register before the debug handler reads the previous data out. The handshaking is described in [Section 9.8.1, “RX Register Ready Bit \(RR\)”](#).

9.11 Debug JTAG Access

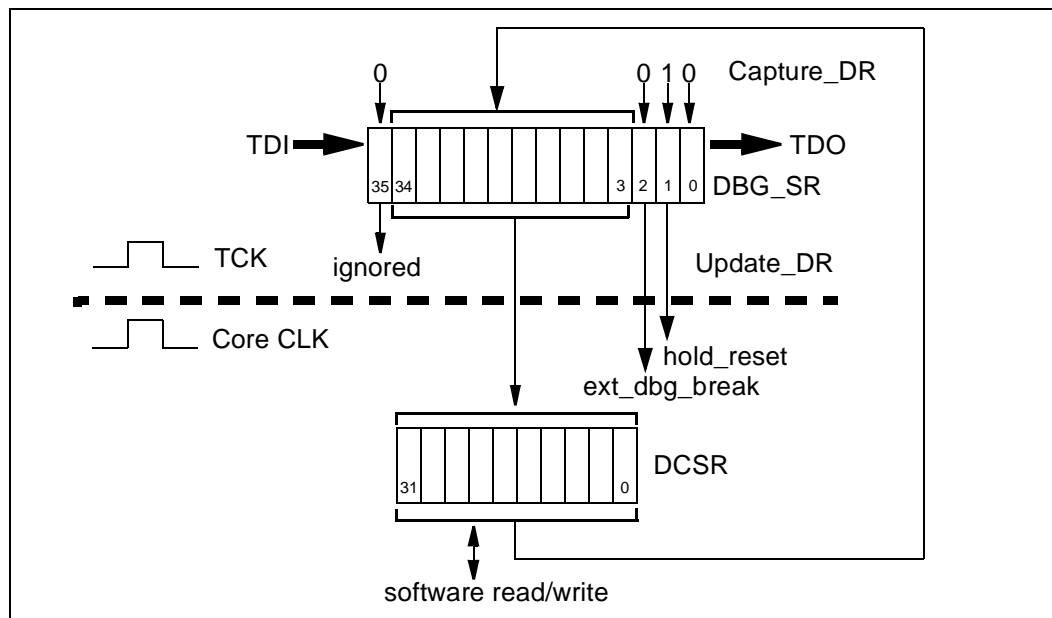
There are four JTAG instructions used by the debugger during software debug: LDIC, SELDCSR, DBGTX and DBGRX. LDIC is described in [Section 9.14, “Downloading Code in the Instruction Cache”](#). The other three JTAG instructions are described in this section. SELDCSR, DBGTX and DBGRX each use a 36-bit shift register to scan in new data and scan out captured data.

9.11.1 SELDCSR JTAG Register

The ‘SELDCSR’ JTAG instruction selects the DCSR JTAG data register. The JTAG opcode is ‘0b0001001’. When the SELDCSR JTAG instruction is in the JTAG instruction register, the debugger can directly access the Debug Control and Status Register (DCSR). The debugger can only modify certain bits through JTAG, but can read the entire register.

The SELDCSR instruction also allows the debugger to generate an external debug break and set the hold_reset signal, which is used when downloading code into the mini instruction cache during reset.

Figure 9-1. SELDCSR



A **Capture_DR** loads the current **DCSR** value into **DBG_SR[34:3]**. The other bits in **DBG_SR** are loaded as shown in [Figure 9-1](#).

A new **DCSR** value can be scanned into **DBG_SR**, and the previous value out, during the **Shift_DR** state. When scanning in a new **DCSR** value into the **DBG_SR**, care must be taken to also set up **DBG_SR[2:1]** to prevent undesirable behavior.

Update_DR parallel loads the new **DCSR** value into the **DCSR**. All bits defined as JTAG writable in [Table 9-1, “Debug Control and Status Register \(DCSR\)”](#) on page 9-123 are updated.

A debugger and the debug handler running on Elkhart must synchronize access the **DCSR**. If one side writes the **DCSR** at the same side the other side reads the **DCSR**, the results are unpredictable.

9.11.1.1 hold_reset

The debugger uses hold_reset when loading code into the instruction cache during a processor reset. Details about loading code into the instruction cache are in [Section 9.14, “Downloading Code in the Instruction Cache”](#).

The debugger must set hold_reset before or during assertion of the reset pin. Once hold_reset is set, the reset pin can be de-asserted, and the processor will internally remain in reset. The debugger can then load debug handler code into the instruction cache before the processor begins executing any code.

Once the code download is complete, the debugger must clear hold_reset. This takes the processor out of reset, and execution begins at the reset vector.

A debugger can set hold_reset in one of 2 ways:

- Either by taking the JTAG state machine into the Capture_DR state, which automatically loads DBG_SR[1] with '1', then the Exit2 state, followed by the Update_Dr state. This will set the hold_reset, clear ext_dbg_break, and leave the DCSR unchanged (the DCSR bits captured in DBG_SR[34:3] are written back to the DCSR on the Update_DR).
- Alternatively, a '1' can be scanned into DBG_SR[1], with the appropriate value scanned in for the DCSR and ext_dbg_break.

The hold_reset bit can only be cleared by scanning in a '0' to DBG_SR[1] and scanning in the appropriate values for the DCSR and ext_dbg_break.

9.11.1.2 ext_dbg_break

The ext_dbg_break allows the debugger to asynchronously re-direct execution on the core to a debug handling routine.

A debugger sets an external debug break by scanning a '1' into DBG_SR[2] (and scanning in the desired value for the DCSR JTAG writable bits in DBG_SR[34:3]).

Once ext_dbg_break is set, it remains set internally until a debug exception occurs. In Monitor Mode, external debug breaks detected during abort mode are pended until the processor exits abort mode. In Halt Mode, external debug breaks detected during SDS are pended until the processor exits SDS. When an external debug break is detected outside of these two cases, the processor ceases executing instructions as quickly as possible, clears the internal ext_dbg_break bit, and branches to the debug handler (Halt Mode) or abort handler (Monitor Mode).

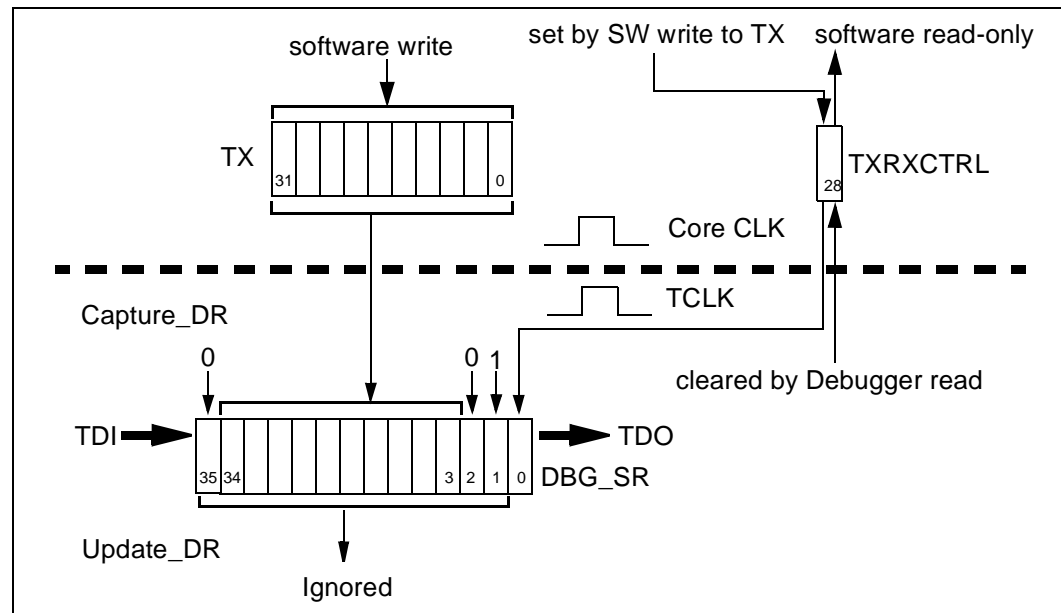
9.11.1.3 DCSR (DBG_SR[34:3])

The JTAG writable bits in the DCSR are updated with the value loaded into DBG_SR[34:3] following an Update_DR.

9.11.2 DBGTX JTAG Register

The 'DBGTX' JTAG instruction selects the DBGTX JTAG data register. The JTAG opcode for this instruction is '0b0010000'. The debug handler uses the DBGTX data register to send data to the debugger. A protocol can be setup between the debugger and debug handler to allow the debug handler to signal an entry into debug mode, and once in debug mode to transmit data requested by the debugger.

Figure 9-2. DBGTX



A Capture_DR loads the TX register value into DBG_SR[34:3] and TXRXCTRL.TR into DBG_SR[0]. The other bits in DBG_SR are loaded as shown in Figure 9-1.

The captured TX value is scanned out during the Shift_DR state. Transitioning from Shift_DR immediately to Capture_DR after capturing a '1' in DBG_SR[0] automatically clears TXRXCTRL.TR.

Data scanned in is ignored on an Update_DR.

9.11.2.1 DBG_SR[0]

DBG_SR[0] is used for part of the synchronization that occurs between the debugger and debug handler for accessing TX. The debugger polls DBG_SR[0] to determine when the TX register contains valid data from the debug handler.

A '1' captured in DBG_SR[0] indicates the captured TX data is valid. After capturing valid data, the debugger must place the JTAG state machine in the Shift_DR state to guarantee that a debugger read clears TXRXCTRL.TR. A '0' indicates there is no new data from the debug handler in the TX register.

9.11.2.2 TX (DBG_SR[34:3])

DBG_SR[34:3] is updated with the contents of the TX register following an Update_DR.

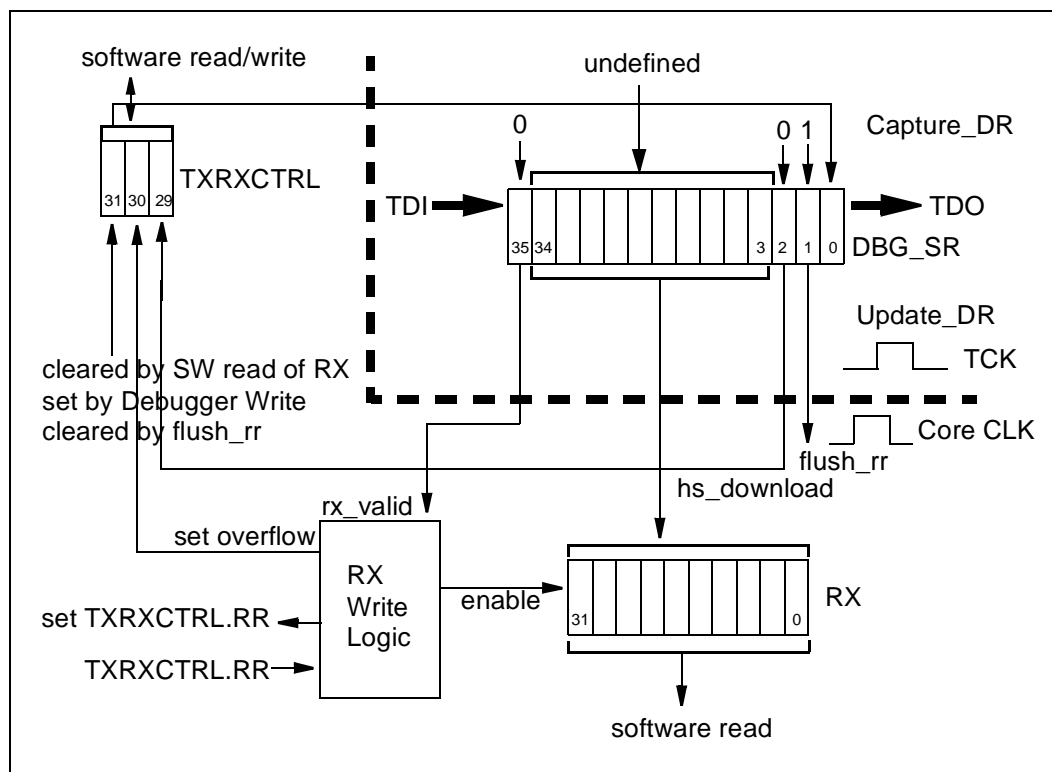
Note: If DBG_SR[0] is '0' following an Update_DR, the contents of DBG_SR[34:3] are unpredictable.

9.11.3 DBGRX JTAG Register

The 'DBGRX' JTAG instruction selects the DBGRX JTAG data register. The JTAG opcode for this instruction is '0b0000010'. The debug handler uses the DBGRX data register to receive information from the debugger. A protocol can be setup between the debugger and debug handler to allow the handler to identify data values and commands.

The DBGRX data register also contain bits to support high-speed download and to "invalidate" the contents of the RX register.

Figure 9-3. DBGRX



A Capture_DR loads the value of TXRXCTRL.RR into DBG_SR[0]. The other bits in DBG_SR are loaded as shown in Figure 9-3.

The captured data is scanned out during the Shift_DR state. Care must be taken while scanning in data. While polling TXRXCTRL.RR, incorrectly setting rx_valid or flush_rr may cause unpredictable behavior following an Update_DR.

Following an Update_DR the scanned in data takes effect.

9.11.3.1 RX Write Logic

The RX write logic (Figure 9-3) serves the following functions:

- 1) RX Write Enable: The RX register only gets updated when rx_valid is set and is unaffected if rx_valid is clear or an overflow occurs. In particular, when the debugger is polling DBG_SR[0], as long as rx_valid is 0, Update_DR does not modify RX.
- 2) Set TXRXCTRL.RR: When the debugger writes new data to RX, TXRXCTRL.RR is automatically set signalling to the debug handler that the RX register contains valid data.
- 3) Set TXRXCTRL.OV: When the debugger scans in a value with rx_valid set and TXRXCTRL.RR is already set, the TXRXCTRL.OV is automatically set. For instance, during high-speed download, the debugger does not poll to see if the handler has read the previous data. If the debug handler stalls long enough, the debugger may try to write a new data to RX before the handler has read the previous data. When this condition is occurs, the RX write logic sets TXRXCTRL.OV and blocks the write to the RX register.

9.11.3.2 DBG_SR[0]

DBG_SR[0] is used for part of the synchronization that occurs between the debugger and debug handler for accessing RX. The debugger polls DBG_SR[0] to determine when the handler has read the previous data from RX, and it is safe to write new data.

A '1' read in DBG_SR[0] indicates that the RX register contains valid data which has not yet been read by the debug handler. A '0' indicates it is safe for the debugger to write new data to the RX register.

9.11.3.3 flush_rr

The flush_rr bit allows the debugger to flush any previous data written to RX. Setting flush_rr clears TXRXCTRL.RR.

9.11.3.4 hs_download

The hs_download bit is provided for use during high speed download. This bit is written directly to TXRXCTRL.D. The debugger can use this bit to improve performance when downloading a block of code or data to the Elkhart system memory.

A protocol can be setup between the debugger and debug handler using this bit. For example, while this bit is set, the debugger can continuously download new data without polling TXRXCTRL.RR. The debug handler uses TXRXCTRL.D as a branch flag to loop while there is more data to come. The debugger clears this bit to indicate the end of the block and allow the debug handler to exit its loop.

Using hs_download as a branch flags eliminates the need for a loop counter in the debug handler code. This avoids the problem were the debugger's loop counter is out of synchronization with the debug handler's counter because of overflow conditions that may have occurred.

9.11.3.5 RX (DBG_SR[34:3])

DBG_SR[34:3] is written to RX following an Update_DR when the RX Write Logic enables the RX register to be updated.

9.11.3.6 rx_valid

The debugger sets the rx_valid bit to indicate the data scanned into DBG_SR[34:3] is valid data to be written to RX. When this bit is set, the data scanned into the DBG_SR will be written to RX following an Update_DR. If rx_valid is not set Update_DR does not affect RX.

This bit does not affect the actions of the flush_rr or hs_download bits.

9.12 Trace Buffer

The 256 entry trace buffer provides the ability to capture control flow information to be used for debugging an application. Two modes are supported:

1. The buffer fills up completely and generates a debug exception. Then SW empties the buffer.
2. The buffer fills up and wraps around until it is disabled. Then SW empties the buffer.

9.12.1 Trace Buffer Registers

CP14 contains three registers (see [Table 9-15](#)) for use with the trace buffer. These CP14 registers are accessible using MRC, MCR, LDC and STC (CDP to any CP14 registers will cause an undefined instruction trap). The CRn and CRm fields specify the register to access. The opcode_1 and opcode_2 fields are not used and should be set to 0.

Table 9-15. CP 14 Trace Buffer Register Summary

CRn	CRm	Register Name
11	0	Trace Buffer Register (TBREG)
12	0	Checkpoint 0 Register (CHKPT0)
13	0	Checkpoint 1 Register (CHKPT1)

Any access to the trace buffer registers in User mode will cause an undefined instruction exception. Specifying registers which do not exist has unpredictable results.

9.12.1.1 Checkpoint Registers

When the debugger reconstructs a trace history, it is required to start at the oldest trace buffer entry and construct a trace going forward. In fill-once mode and wrap-around mode when the buffer does not wrap around, the trace can be reconstructed by starting from the point in the code where the trace buffer was first enabled.

The difficulty occurs in wrap-around mode when the trace buffer wraps around at least once. In this case the debugger gets a snapshot of the last N control flow changes in the program, where $N \leq$ size of buffer. The debugger does not know the starting address of the oldest entry read from the trace buffer. The checkpoint registers provide reference addresses to help reduce this problem.

Table 9-16. Checkpoint Register (CHKPTx)

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																															
CHKPTx																															
reset value: Unpredictable																															
Bits	Access	Description																													
31:0	Read/Write	CHKPTx: target address for corresponding entry in trace buffer																													

The two checkpoint registers (CHKPT0, CHKPT1) on Elkhart provide the debugger with two reference addresses to use for re-constructing the trace history.

When the trace buffer is enabled, reading and writing to either checkpoint register has unpredictable results. When the trace buffer is disabled, writing to a checkpoint register sets the register to the value written. Reading the checkpoint registers returns the value of the register.

In normal usage, the checkpoint registers are used to hold target addresses of specific entries in the trace buffer. Only direct and indirect entries get checkpointed. Exception and roll-over messages are never checkpointed. When an entry is checkpointed, the processor sets bit 6 of the message byte to indicate this (refer to [Table 9-18, "Message Byte Formats"](#))

When the trace buffer contains only one checkpointed entry, the corresponding checkpoint register is CHKPT0. When the trace buffer wraps around, two entries will typically be checkpointed, usually about half a buffers length apart. In this case, the first (oldest) checkpointed entry read from the trace buffer corresponds to CHKPT1, the second checkpointed entry corresponds to CHKPT0.

Although the checkpoint registers are provided for wrap-around mode, they are still valid in fill-once mode.

9.12.1.2 Trace Buffer Register (TBREG)

The trace buffer is read through TBREG, using MRC and MCR. Software should only read the trace buffer when it is disabled. Reading the trace buffer while it is enabled, may cause unpredictable behavior of the trace buffer. Writes to the trace buffer have unpredictable results. Reading the trace buffer returns the oldest byte in the trace buffer in the least significant byte of TBREG. The byte is either a message byte or one byte of the 32 bit address associated with an indirect branch message. Table 9-17 shows the format of the trace buffer register.

Table 9-17. TBREG Format

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
Data																															
reset value: unpredictable																															
Bits	Access		Description																												
31:8	Read-as-Zero/Write-ignored		Reserved																												
7:0	Read / Write-unpredictable		Message Byte or Address Byte																												

9.13 Trace Buffer Entries

Trace buffer entries consist of either one or five bytes. Most entries are one byte messages indicating the type of control flow change. The target address of the control flow change represented by the message byte is either encoded in the message byte (like for exceptions) or can be determined by looking at the instruction word (like for direct branches). Indirect branches require five bytes per entry. One byte is the message byte identifying it as an indirect branch. The other four bytes make up the target address of the indirect branch. The following sections describe the trace buffer entries in detail.

9.13.1 Message Byte

There are two message formats, (exception and non-exception) as shown in Figure 9-4.

Figure 9-4. Message Byte Formats

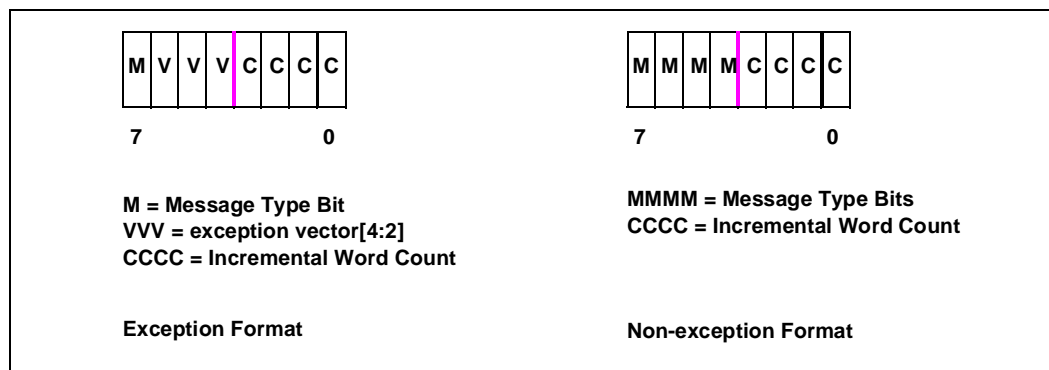


Table 9-18 shows all of the possible trace messages.

Table 9-18. Message Byte Formats

Message Name	Message Byte Type	Message Byte format	# address bytes
Exception	exception	0b0VVV CCCC	0
Direct Branch ^a	non-exception	0b1000 CCCC	0
Checkpointed Direct Branch ^a	non-exception	0b1100 CCCC	0
Indirect Branch ^b	non-exception	0b1001 CCCC	4
Checkpointed Indirect Branch ^b	non-exception	0b1101 CCCC	4
Roll-over	non-exception	0b1111 1111	0

- a. Direct branches include ARM and THUMB bl, b
- b. Indirect branches include ARM ldm, ldr, and dproc to PC; ARM and THUMB bx, blx(1) and blx(2); and THUMB pop.

9.13.1.1 Exception Message Byte

When any kind of exception occurs, an exception message is placed in the trace buffer. In an exception message byte, the message type bit (M) is always 0.

The vector exception (VVV) field is used to specify bits[4:2] of the vector address (offset from the base of default or relocated vector table). The vector allows the debugger to identify which exception occurred.

The incremental word count (CCCC) is the instruction count since the last control flow change (not including the current instruction for undef, SWI, and pre-fetch abort). The instruction count includes instructions that were executed and conditional instructions that were not executed due to the condition of the instruction not matching the CC flags.

A count value of 0 indicates that 0 instructions executed since the last control flow change and the current exception. For example, if a branch is immediate followed by a SWI, a direct branch exception message (for the branch) is followed by an exception message (for the SWI) in the trace buffer. The count value in the exception message will be 0, meaning that 0 instructions executed after the last control flow change (the branch) and before the current control flow change (the SWI). Instead of the SWI, if an IRQ was handled immediately after the branch (before any other instructions executed), the count would still be 0, since no instructions executed after the branch and before the interrupt was handled.

A count of 0b1111 indicates that 15 instructions executed between the last branch and the exception. In this case, an exception was either caused by the 16th instruction (if it is an undefined instruction exception, pre-fetch abort, or SWI) or handled before the 16th instruction executed (for FIQ, IRQ, or data abort).

Note: there is a special case for the count field related to precise data aborts. For a precise data abort on a load to the PC (LDR or LDM), the count is consistent with the above description (i.e. aborting instruction is not counted). For all other precise data aborts, the instruction that causes the data abort is included in the count value of the exception message.

9.13.1.2 Non-exception Message Byte

Non-exception message bytes are used for direct branches, indirect branches, and rollovers.

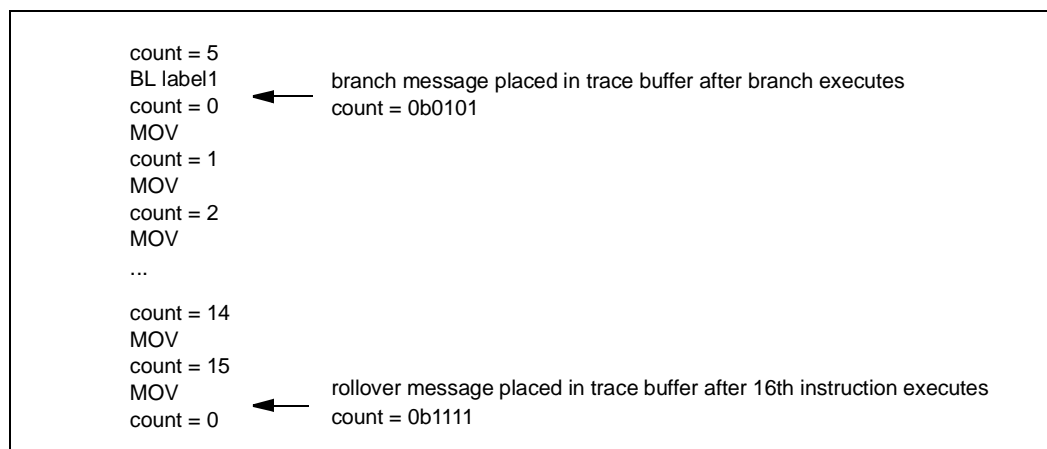
In a non-exception message byte, the 4-bit message type field (MMMM) specifies the type of message (refer to [Table 9-18](#)).

The incremental word count (CCCC) is the instruction count since the last control flow change (excluding the current branch). The instruction count includes instructions that were executed and conditional instructions that were not executed due to the condition of the instruction not matching the CC flags. In the case of back-to-back branches the word count would be 0 indicating that no instructions executed after the last branch and before the current one.

A rollover message is used to keep track of long traces of code that do not have control flow changes. The rollover message means that 16 instructions have executed since the last message byte was written to the trace buffer.

If the incremental counter reaches its maximum value of 15, a rollover message is written to the trace buffer following the next instruction (which will be the 16th instruction to execute). This is shown in [Example 9-1](#). The count in the rollover message is 0b1111, indicating that 15 instructions have executed after the last branch and before the current non-branch instruction that caused the rollover message.

Example 9-1. Rollover Messages Examples

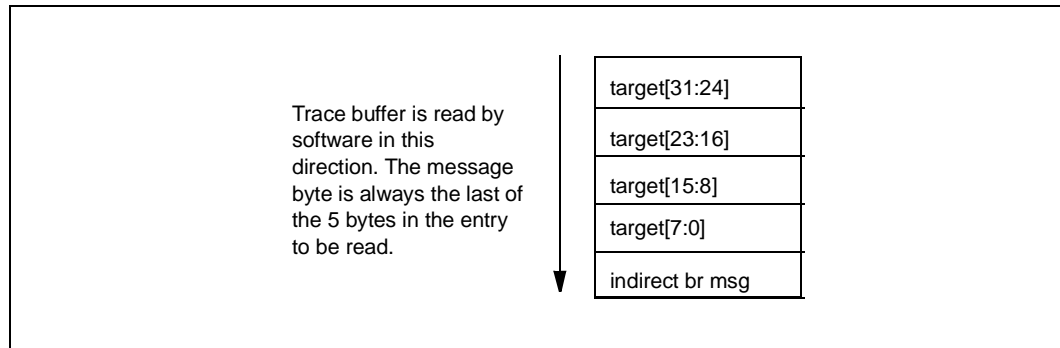


If the 16th instruction is a branch (direct or indirect), the appropriate branch message is placed in the trace buffer instead of the roll-over message. The incremental counter is still set to 0b1111, meaning 15 instructions executed between the last branch and the current branch.

9.13.1.3 Address Bytes

Only indirect branch entries contain address bytes in addition to the message byte. Indirect branch entries always have four address bytes indicating the target of that indirect branch. When reading the trace buffer the MSB of the target address is read out first; the LSB is the fourth byte read out; and the indirect branch message byte is the fifth byte read out. The byte organization of the indirect branch message is shown in Figure 9-5.

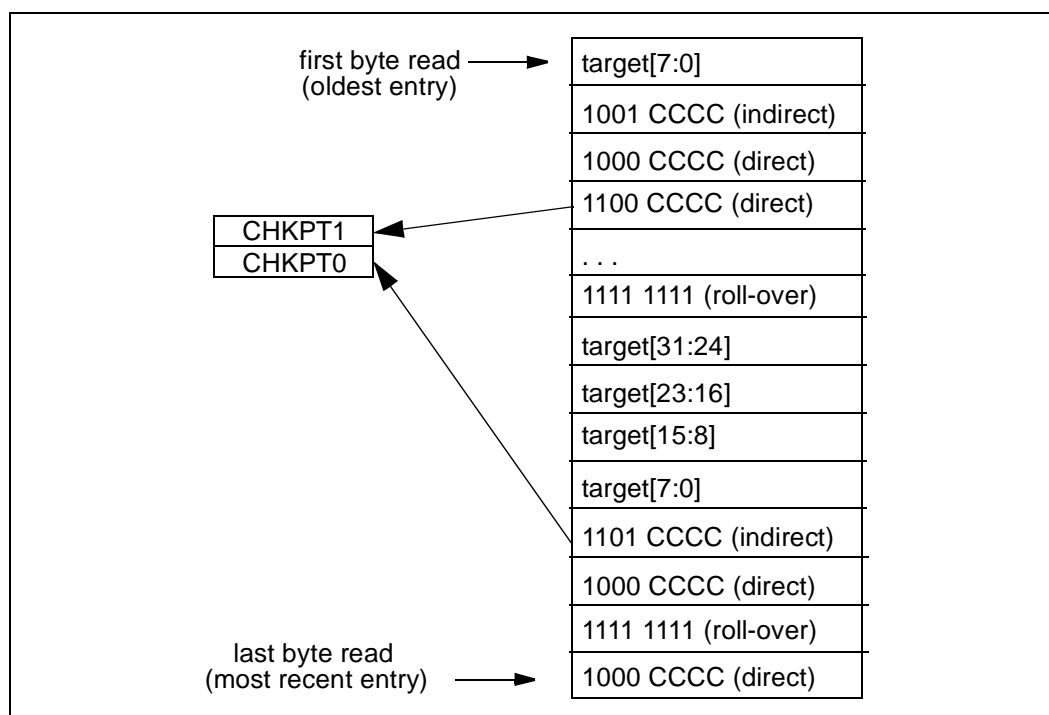
Figure 9-5. Indirect Branch Entry Address Byte Organization



9.13.2 Trace Buffer Usage

The Elkhart trace buffer is 256 bytes in length. The first byte read from the buffer represents the oldest trace history information in the buffer. The last (256th) byte read represents the most recent entry in the buffer. The last byte read from the buffer will always be a message byte. This provides the debugger with a starting point for parsing the entries out of the buffer. Because the debugger needs the last byte as a starting point when parsing the buffer, the entire trace buffer must be read (256 bytes on Elkhart) before the buffer can be parsed. Figure 9-6 is a high level view of the trace buffer.

Figure 9-6. High Level View of Trace Buffer



The trace buffer must be initialized prior to its initial usage, then again prior to each subsequent usage. Initialization is done by reading the entire trace buffer. The process of reading the trace buffer also clears it out (all entries are set to 0b0000 0000), so when the trace buffer has been used to capture a trace, the process of reading the captured trace data also re-initializes the trace buffer for its next usage.

The trace buffer can be used to capture a trace up to a processor reset. A processor reset disables the trace buffer, but the contents are unaffected. The trace buffer captures a trace up to the processor reset.

The trace buffer does not capture reset events or debug exceptions.

Since the trace buffer is cleared out before it is used, all entries are initially 0b0000 0000. In fill-once mode, these 0's can be used to identify the first valid entry in the trace buffer. In wrap around mode, in addition to identifying the first valid entry, these 0 entries can be used to determine whether a wrap around occurred.

As the trace buffer is read, the oldest entries are read first. Reading a series of 5 (or more) consecutive "0b0000 0000" entries in the oldest entries indicates that the trace buffer has not wrapped around and the first valid entry will be the first non-zero entry read out.

Reading 4 or less consecutive "0b0000 0000" entries requires a bit more intelligence in the debugger. The debugger must determine whether these 0's are part of the address of an indirect branch message, or whether they are part of the "0b0000 0000" that the trace buffer was initialized with. If the first non-zero message byte is an indirect branch message, then these 0's are part of the address since the address is always read before the indirect branch message (see [Section 9.13.1.3, "Address Bytes"](#)). If the first non-zero entry is any other type of message byte, then these 0's indicate that the trace buffer has not wrapped around and that first non-zero entry is the start of the trace.

If the oldest entry from the trace buffer is non-zero, then the trace buffer has either wrapped around or just filled up.

Once the trace buffer has been read and parsed, the debugger should re-create the trace history from oldest trace buffer entry to latest. Trying to re-create the trace going backwards from the latest trace buffer entry may not work in most cases, because once a branch message is encountered, it may not be possible to determine the source of the branch.

In fill-once mode, the return from the debug handler to the application should generate an indirect branch message. The address placed in the trace buffer will be that of the target application instruction. Using this as a starting point, re-creating a trace going forward in time should be straightforward.

In wrap around mode, the debugger should use the checkpoint registers and address bytes from indirect branch entries to re-create the trace going forward. The drawback is that some of the oldest entries in the trace buffer may be untraceable, depending on where the earliest checkpoint (or indirect branch entry) is located. The best case is when the oldest entry in the trace buffer was checkpointed, so the entire trace buffer can be used to re-create the trace. The worst case is when the first checkpoint is in the middle of the trace buffer and no indirect branch messages exist before this checkpoint. In this case, the debugger would have to start at its known address (the first checkpoint) which is half way through the buffer and work forward from there.

9.14 Downloading Code in the Instruction Cache

On Elkhart, a mini instruction cache, physically separate¹ from the main instruction cache can be used as an on-chip instruction RAM. A debugger can download code directly into either instruction cache through JTAG. In addition to downloading code, several cache functions are supported.

Elkhart supports loading the instruction cache during reset and dynamically (without resetting the core). Loading the instruction cache during normal program execution requires a strict handshaking protocol between software running on Elkhart and the debugger.

In the remainder of this section the term 'instruction cache' applies to either main or mini instruction cache.

9.14.1 Mini Instruction Cache Overview

The mini instruction cache is a smaller version of the main instruction cache. The size of the mini instruction cache is proportional to that of the main instruction cache:

A version of the core with a 32KB main instruction cache will have a 2KB mini instruction cache. A version of the core with a 16KB main instruction cache will have a 1KB mini instruction cache.

Refer to the Intel XScale® core implementation option section of the Application Specific Standard Product (ASSP) architecture specification for more details the cache size supported by the ASSP.

The mini instruction cache is a 2-way set associative cache. The 2KB version has 32 sets, the 1KB version has 16 sets. The line size is 8 words. The cache uses the round-robin replacement policy.

The mini instruction cache is virtually addressed and addresses may be remapped by the PID. However, since the debug handler executes in Special Debug State, address translation and PID remapping are turned off. For application code, accesses to the mini instruction cache use the normal address translation and PID mechanisms.

Normal application code is never cached in the mini instruction cache on an instruction fetch. The only way to get code into the mini instruction cache is through the JTAG LDIC function. Code downloaded into the mini instruction cache is essentially locked - it cannot be overwritten by application code running on Elkhart. However, it is not locked against code downloaded through the JTAG LDIC functions.

Application code can invalidate a line in the mini instruction cache using a CP15 Invalidate IC line function to an address that hits in the mini instruction cache. However, a CP15 global invalidate IC function does not affect the mini instruction cache.

The mini instruction cache can be globally invalidated through JTAG by the LDIC Invalidate IC function or by a processor reset when the processor is not in HALT or LDIC mode. A single line in the mini instruction cache can be invalidated through JTAG by the LDIC Invalidate IC-line function.

1. A cache line fill from external memory will never be written into the mini-instruction cache. The only way to load a code into the mini-instruction cache is through JTAG.

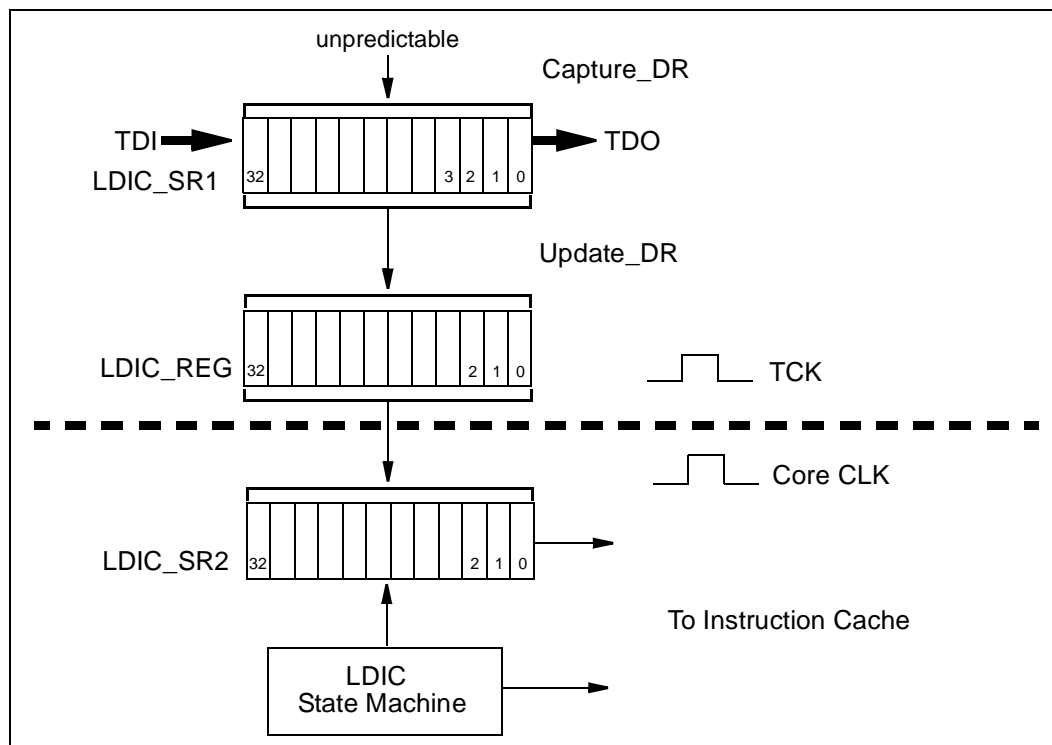
9.14.2 LDIC JTAG Command

The LDIC JTAG instruction selects the JTAG data register for loading code into the instruction cache. The JTAG opcode for this instruction is '00111'. The LDIC instruction must be in the JTAG instruction register in order to load code directly into the instruction cache through JTAG.

9.14.3 LDIC JTAG Data Register

The LDIC JTAG Data Register is selected when the LDIC JTAG instruction is in the JTAG IR. An external host can load and invalidate lines in the instruction cache through this data register.

Figure 9-7. LDIC JTAG Data Register Hardware



The data loaded into LDIC_SR1 during a Capture_DR is unpredictable.

All LDIC functions and data consists of 33 bit packets which are scanned into LDIC_SR1 during the Shift_DR state.

Update_DR parallel loads LDIC_SR1 into LDIC_REG which is then synchronized with the Elkhart clock and loaded into the LDIC_SR2. Once data is loaded into LDIC_SR2, the LDIC State Machine turns on and serially shifts the contents of LDIC_SR2 to the instruction cache.

Note that there is a delay from the time of the Update_DR to the time the entire contents of LDIC_SR2 have been shifted to the instruction cache. Removing the LDIC JTAG instruction from the JTAG IR before the entire contents of LDIC_SR2 are sent to the instruction cache, will result in unpredictable behavior. Therefore, following the Update_DR for the last LDIC packet, the LDIC instruction must remain in the JTAG IR for a minimum of 15 TCKs. This ensures the last packet is correctly sent to the instruction cache.

9.14.4 LDIC Cache Functions

Elkhart supports four cache functions that can be executed through JTAG. Two functions allow an external host to download code into the main instruction cache or the mini instruction cache through JTAG. Two additional functions are supported to allow lines to be invalidated in the instruction cache. The following table shows the cache functions supported through JTAG.

Table 9-19. LDIC Cache Functions

Function	Encoding	Arguments	
		Address	# Data Words
Invalidate IC Line	0b000	VA of line to invalidate	0
Invalidate Mini IC	0b001	-	0
Load Main IC	0b010	VA of line to load	8
Load Mini IC	0b011	VA of line to load	8
RESERVED	0b100-0b111	-	-

Invalidate IC line invalidates the line in the instruction cache containing specified virtual address. If the line is not in the cache, the operation has no effect. It does not take any data arguments.

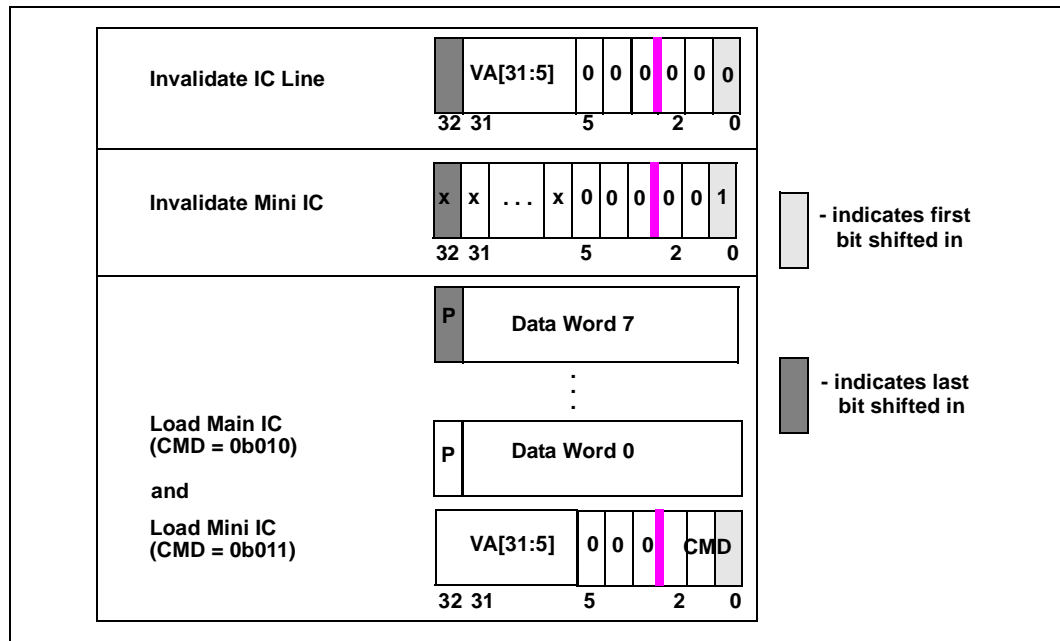
Invalidate Mini IC¹ will invalidate the entire mini instruction cache. It does not effect the main instruction cache. It does not require a virtual address or any data arguments.

Load Main IC and Load Mini IC write one line of data (8 ARM instructions) into the specified instruction cache at the specified virtual address. Load Main IC has been deprecated on the Intel XScale® core. A debugger should only load code into the mini instruction cache.

Each cache function is downloaded through JTAG in 33 bit packets. [Figure 9-8](#) shows the packet formats for each of the JTAG cache functions. Invalidate IC Line and Invalidate Mini IC each require 1 packet. Load Main IC and Load Mini IC each require 9 packets.

1. The LDIC Invalidate Mini IC function does not invalidate the BTB (like the CP15 Invalidate IC function) so software must do this manually where appropriate.

Figure 9-8. Format of LDIC Cache Functions



All packets are 33 bits in length. Bits [2:0] of the first packet specify the function to execute. For functions that require an address, bits[32:6] of the first packet specify an 8-word aligned address (Packet1[32:6] = VA[31:5]). For Load Main IC and Load Mini IC, 8 additional data packets are used to specify 8 ARM instructions to be loaded into the target instruction cache. Bits[31:0] of the data packets contain the data to download. Bit[32] of each data packet is the value of the parity for the data in that packet.

As shown in Figure 9-8, the first bit shifted in TDI is bit 0 of the first packet. After each 33-bit packet, the host must take the JTAG state machine into the Update_DR state. After the host does an Update_DR and returns the JTAG state machine back to the Shift_DR state, the host can immediately begin shifting in the next 33-bit packet.

9.14.5 Loading Instruction Cache During Reset

Code can be downloaded into the instruction cache through JTAG during a processor reset. This feature is used during software debug to download the debug handler prior to starting a debug session. Immediately out of reset, the downloaded handler can intercept the reset vector and turn control of the system to the debugger. The debugger can then initialize the system as necessary and begin the application program.

In general, any code downloaded into the instruction cache through JTAG, must be downloaded to addresses that are not already valid in the instruction cache. Failure to meet this requirement will result in unpredictable behavior by the processor. During a processor reset, the instruction cache is typically invalidated, with the exception of the following cases:

- When LDIC JTAG instruction is loaded in the JTAG IR, neither the mini instruction cache, nor the main instruction cache are invalidated during reset.
- When the Halt Mode bit is set in the DCSR only the mini instruction cache is prevented from being invalidated during reset. The main instruction cache is still be invalidated.

The Figure 9-9 shows the actions necessary to download code into the instruction cache during a cold reset for debug.

Figure 9-9. Code Download During a Cold Reset For Debug

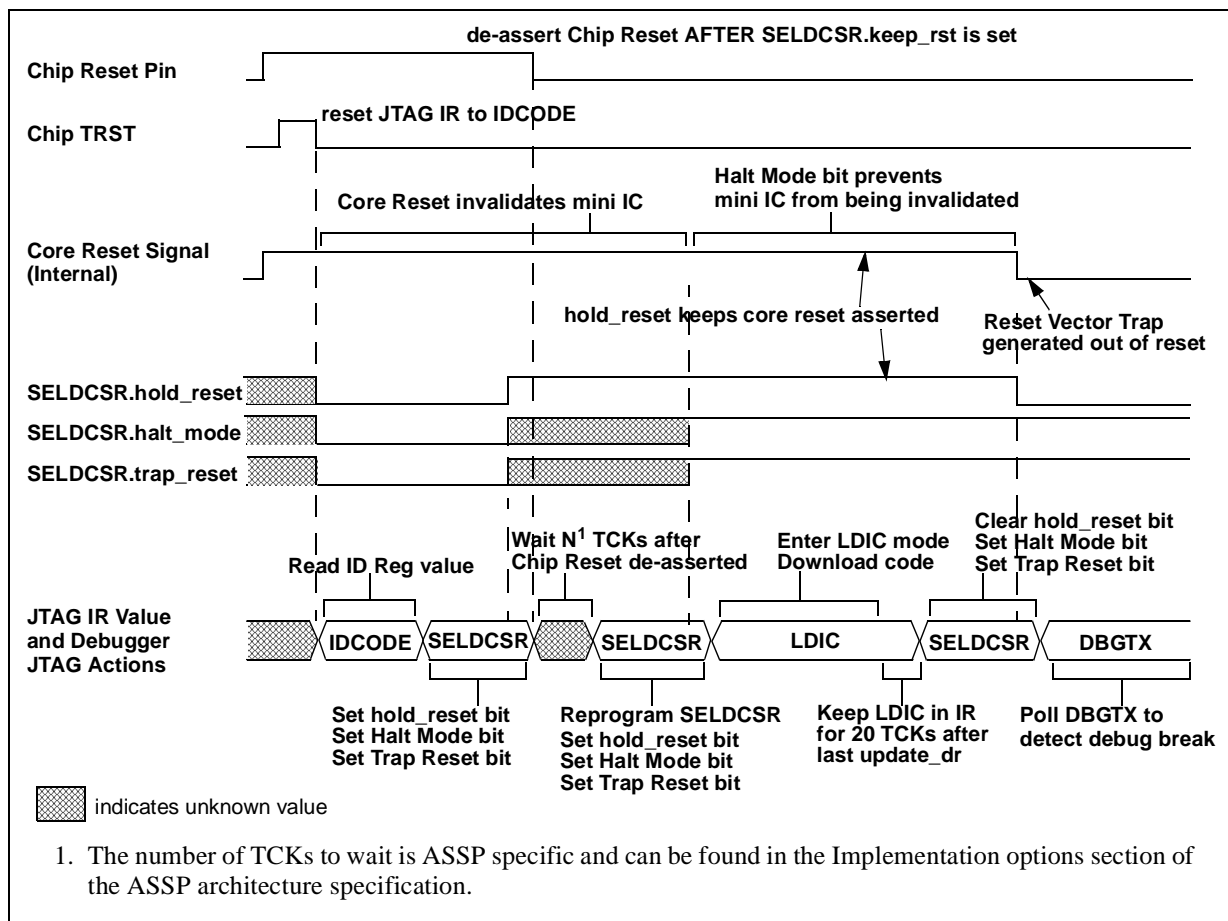


Table 9-20 describes the actions a debugger should take to load code into the mini instruction cache during reset:

Table 9-20. Steps For Loading Mini Instruction Cache During Reset

Step #	Action	Notes
1	Assert Chip Reset and Chip TRST	This resets the JTAG IR to IDCODE and clears the Halt Mode bit in the DCSR, ensuring that the main and mini IC are invalidated.
2	Read ID Register value	
3	Program SELDCSR JTAG register: Halt Mode=1 Trap Reset=1 hold_reset=1	SELDCSR details can be found in Section 9.11.1 . Depending on ASSP implementation, the Halt Mode bit and Trap Reset bit may or may not actually be set to the programmed value. The hold reset bit will be set to the programmed value.
4	Deassert Chip Reset	Internally the core will remain held in reset due to hold_reset being set.
5	Wait N TCKs	N is a ASSP specific number and can be found in the Implementations options section of the ASSP architecture specification. This wait ensures that the core is stable before proceeding.
6	Program SELDCSR JTAG register: Halt Mode=1 Trap Reset=1 hold_reset=1	The SELDCSR instruction must be reloaded into the JTAG IR. Failure to reload the JTAG IR may result in unpredictable behavior. Reprogramming of the SELDCSR JTAG register guarantees that the Halt Mode bit and Trap Reset bit are set before loading the mini instruction cache.
7	Load LDIC JTAG instruction and download the debug handler into mini instruction cache.	Loading into the instruction cache is described in Section 9.14.4, "LDIC Cache Functions"
8	Clock a minimum of 20 TCKs before changing the JTAG IR.	The LDIC JTAG instruction must remain in the JTAG instruction register for at least 20 TCKs following the update_dr for the last cache line, to ensure that line is correctly loaded into the mini instruction cache. Changing the JTAG IR within 20 cycles may result in unpredictable behavior.
9	Program SELDCSR JTAG register: Halt Mode bit = 1 Trap Reset bit = 1 hold_reset = 0	Clearing the hold_reset bit allows the core to come out of reset and begin execution from address 0.
10	poll the DBGTX register	Immediately out of reset, a reset vector trap will occur and the debug handler will begin execution. The debugger must poll DBGTX to identify when this has happened.

9.14.6 Dynamically Loading Instruction Cache After Reset

An debugger can load code into the instruction cache “on the fly” or “dynamically”. This occurs when the debugger downloads code while the core is not held in reset and is useful for expanding the functionality of the debug handler.

Since the debug handler is limited to 2KB (or 1KB w/ 16KB version of core), all possible debug handler functionality cannot fit in the mini instruction cache at one time. The debugger can load key debug handler functionality into the mini instruction cache during reset and download the less often used routines as they are needed, outside of reset.

Loading code dynamically into the mini instruction cache requires strict synchronization between the code running on the core and the debugger. The guidelines for downloading code during program execution must be followed to ensure proper operation of the processor.

To dynamically download code during software debug, there must be a minimal debug handler stub, responsible for doing the handshaking with the debugger, resident in the mini instruction cache. This debug handler stub should be downloaded into the mini instruction cache during processor reset using the method described in [Section 9.14.5](#).

[Figure 9-10](#) shows a high level view of the actions taken by the host and debug handler during dynamic code download and [Table 9-21](#) shows the step-by-step process in more detail.

[Section 9.14.6.1, “Dynamic Download Synchronization Code”](#) provides the details and describes the requirements for implementing the handshaking in the debug handler.

Figure 9-10. Downloading Code in IC During Program Execution

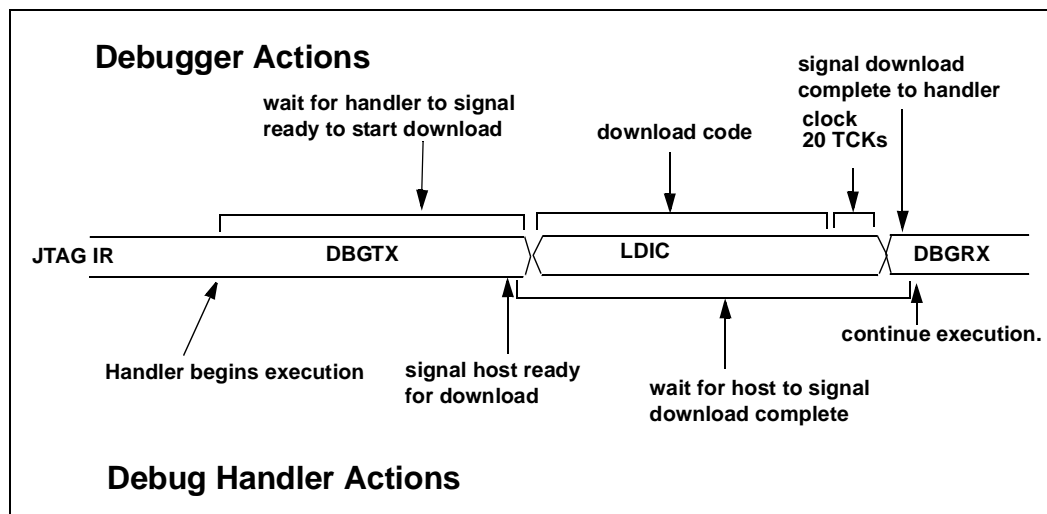


Table 9-21. Steps For Dynamically Loading the Mini Instruction Cache

Step #	Action		Notes
	Debugger	Debug Handler	
1	Poll DBGTX		Debugger must poll DBGTX for an indication from the debug handler that it is safe to begin the download. Refer to Section 9.11.2, "DBGTX JTAG Register" for details on polling DBGTX through JTAG.
2		Write TX,	When the debug handler gets to a safe section of code, it writes TX, allowing the debugger to proceed with the download.
3		Poll RX Read Flag	The handler then begins polling the RX Ready for an indication the debugger has completed the download
34	Detect handler write to TX		When the debugger sees a valid value in TX, it can proceed with the download into the instruction cache.
5	Load LDIC Instr and download code		Debugger loads the LDIC instruction into JTAG IR and downloads code into the instruction cache. For each cache line downloaded, the debugger must invalidate the target line before downloading to that line. Failure to invalidate a line prior to writing it may cause unpredictable operation by the processor. Refer to Section 9.14.4, "LDIC Cache Functions" for details on the invalidate and download functions.
6	clock a minimum of 20 TCKs before changing the JTAG IR.		The LDIC JTAG instruction must remain in the JTAG instruction register for at least 20 TCKs following the update_dr for the last cache line, to ensure that line is correctly loaded into the mini instruction cache. Changing the JTAG IR within 20 cycles may result in unpredictable behavior.
7	Write RX		The completes the handshaking allowing the handler to exit its polling loop. The value written to RX by the debugger is implementation specific.
8		Detect debugger write to RX	The handler exits its polling loop and depending on the implementation of the debug handler, can branch to a fixed address where the code was downloaded or can use the value written to RX by the debugger as the target address to branch to.

Note: The debug handler polling loop must reside in the instruction cache and execute out of the cache while doing the synchronization. The processor should not be doing any code fetches to external memory while code is being downloaded.

9.14.6.1 Dynamic Download Synchronization Code

The following pieces of code are necessary in the debug handler to implement the synchronization used during dynamic code download. The pieces must be ordered in the handler as shown below.

```
# Before the download can start, all outstanding instruction fetches must complete.
# The MCR invalidate IC by line function serves as a barrier instruction in
# the core. All outstanding instruction fetches are guaranteed to complete before
# the next instruction executes.
# NOTE1: the actual address specified to invalidate is implementation defined, but
# must not have any harmful effects.
# NOTE2: The placement of the invalidate code is implementation defined, the only
# requirement is that it must be placed such that by the time the debugger starts
# loading the instruction cache, all outstanding instruction fetches have completed
    mov r5, address
    mcr p15, 0, r5, c7, c5, 1

# The host waits for the debug handler to signal that it is ready for the
# code download. This can be done using the TX register access handshaking
# protocol. The host polls the TR bit through JTAG until it is set, then begins
# the code download. The following MCR does a write to TX, automatically
# setting the TR bit.
# NOTE: The value written to TX is implementation defined.
    mcr p14, 0, r6, c8, c0, 0

# The debug handler waits until the download is complete before continuing. The
# debugger uses the RX handshaking to signal the debug handler when the download
# is complete. The debug handler polls the RR bit until it is set. A debugger write
# to RX automatically sets the RR bit, allowing the handler to proceed.
# NOTE: The value written to RX by the debugger is implementation defined - it can
# be a bogus value signalling the handler to continue or it can be a target address
# for the handler to branch to.
loop:
    mrc    p14, 0, r15, c14, c0, 0      @ handler waits for signal from debugger
    bpl   loop
    mrc    p14, 0, r0, c8, c0, 0       @ debugger writes target address to RX
    bx    r0
```

In a very simple debug handler stub, the above parts may form the complete handler downloaded during reset (with some handler entry and exit code). When a debug exception occurs, routines can be downloaded as necessary. This basically allows the entire handler to be dynamic.

Another possibility is for a more complete debug handler is downloaded during reset. The debug handler may support some operations, such as read memory, write memory, etc. However, other operations, such as reading or writing a group of CP register, can be downloaded dynamically. This method could be used to dynamically download infrequently used debug handler functions, while the more common operations remain static in the mini-instruction cache.

The Intel Debug Handler is a complete debug handler that implements the more commonly used functions, and allows less frequently used functions to be dynamically downloaded.

Performance Considerations

10

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

10.1 Interrupt Latency

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

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

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

Maximum Interrupt Latency can be reduced by:

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

Refer to the Intel XScale® core implementation option section of the ASSP architecture specification for more information on interrupt latency.

10.2 Branch Prediction

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

These instructions -- ARM **B**, ARM **BL** and Thumb **B** -- enter into the branch target buffer when they are “taken” for the first time. (A “taken” branch refers to when they are evaluated to be true.) Once in the branch target buffer, the core dynamically predicts the outcome of these instructions based on previous outcomes. Table 10-1 shows the branch latency penalty when these instructions are correctly predicted and when they are not. A penalty of zero for correct prediction means that the core can execute the next instruction in the program flow in the cycle following the branch.

Table 10-1. Branch Latency Penalty

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

10.3 Addressing Modes

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

10.4 Instruction Latencies

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

The following section explains how to read these tables.

10.4.1 Performance Terms

- Issue Clock (cycle 0)
The first cycle when an instruction is decoded **and** allowed to proceed to further stages in the execution pipeline (i.e., when the instruction is actually issued).
- Cycle Distance from A to B
The cycle distance from cycle **A** to cycle **B** is **(B-A)** -- that is, the number of cycles from the start of cycle **A** to the start of cycle **B**. Example: the cycle distance from cycle 3 to cycle 4 is one cycle.
- Issue Latency
The cycle distance **from** the first issue clock of the current instruction **to** the issue clock of the next instruction. The actual number of cycles can be influenced by cache-misses, resource-dependency stalls, and resource availability conflicts.
- Result Latency
The cycle distance **from** the first issue clock of the current instruction **to** the issue clock of the first instruction that can use the result without incurring a resource dependency stall. The actual number of cycles can be influenced by cache-misses, resource-dependency stalls, and resource availability conflicts.
- Minimum Issue Latency (without Branch Misprediction)
The minimum cycle distance **from** the issue clock of the current instruction **to** the first possible issue clock of the next instruction assuming best case conditions (i.e., that the issuing of the next instruction is not stalled due to a resource dependency stall; the next instruction is immediately available from the cache or memory interface; the current instruction does not incur resource dependency stalls during execution that can not be detected at issue time; and if the instruction uses dynamic branch prediction, correct prediction is assumed).
- Minimum Result Latency
The required minimum cycle distance **from** the issue clock of the current instruction **to** the issue clock of the first instruction that can use the result without incurring a resource dependency stall assuming best case conditions (i.e., that the issuing of the next instruction is not stalled due to a resource dependency stall; the next instruction is immediately available from the cache or memory interface; and the current instruction does not incur resource dependency stalls during execution that can not be detected at issue time).
- Minimum Issue Latency (with Branch Misprediction)
The minimum cycle distance **from** the issue clock of the current branching instruction **to** the first possible issue clock of the next instruction. This definition is identical to *Minimum Issue Latency* except that the branching instruction has been mispredicted. It is calculated by adding *Minimum Issue Latency (without Branch Misprediction)* to the minimum branch latency penalty number from [Table 10-1](#), which is four cycles.

- Minimum Resource Latency

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

For the following code fragment, here is an example of computing latencies:

Example 10-1. Computing Latencies

```

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

Table 10-2 shows how to calculate Issue Latency and Result Latency for each instruction. Looking at the issue column, the **UMLAL** instruction starts to issue on cycle 0 and the next instruction, **ADD**, issues on cycle 2, so the Issue Latency for **UMLAL** is two. From the code fragment, there is a result dependency between the **UMLAL** instruction and the **SUB** instruction. In Table 10-2, **UMLAL** starts to issue at cycle 0 and the **SUB** issues at cycle 5. thus the Result Latency is five.

Table 10-2. Latency Example

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

10.4.2 Branch Instruction Timings

Table 10-3. Branch Instruction Timings (Those predicted by the BTB)

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

Table 10-4. Branch Instruction Timings (Those not predicted by the BTB)

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

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

10.4.3 Data Processing Instruction Timings

Table 10-5. Data Processing Instruction Timings

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

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

10.4.4 Multiply Instruction Timings

Table 10-6. Multiply Instruction Timings (Sheet 1 of 2)

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

Table 10-6. Multiply Instruction Timings (Sheet 2 of 2)

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

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

Table 10-7. Multiply Implicit Accumulate Instruction Timings

Mnemonic	Rs Value (Early Termination)	Minimum Issue Latency	Minimum Result Latency	Minimum Resource Latency (Throughput)
MIA	Rs[31:15] = 0x0000 or Rs[31:15] = 0xFFFF	1	1	1
	Rs[31:27] = 0x0 or Rs[31:27] = 0xF	1	2	2
	all others	1	3	3
MIAXy	N/A	1	1	1
MIAPH	N/A	1	2	2

Table 10-8. Implicit Accumulator Access Instruction Timings

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

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

10.4.5 Saturated Arithmetic Instructions

Table 10-9. Saturated Data Processing Instruction Timings

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

10.4.6 Status Register Access Instructions

Table 10-10. Status Register Access Instruction Timings

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

10.4.7 Load/Store Instructions

Table 10-11. Load and Store Instruction Timings

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

Table 10-12. Load and Store Multiple Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
LDM ^a	2 + numreg ^b	5-18 for load data (4 + numreg for last register in list; 3 + numreg for 2nd to last register in list; 2 + numreg for all other registers in list); 2+ numreg for writeback of base
STM	2 + numreg	2 + numreg for writeback of base

- a. See Table 10-4 for LDM timings when R15 is in the register list
b. numreg is the number of registers in the register list

10.4.8 Semaphore Instructions

Table 10-13. Semaphore Instruction Timings

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

10.4.9 Coprocessor Instructions

Table 10-14. CP15 Register Access Instruction Timings

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

a. MRC to R15 is unpredictable

Table 10-15. CP14 Register Access Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
MRC	8	8
MRC to R15	9	9
MCR	8	N/A
LDC	11	N/A
STC	8	N/A

10.4.10 Miscellaneous Instruction Timing

Table 10-16. Exception-Generating Instruction Timings

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

Table 10-17. Count Leading Zeros Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
CLZ	1	1

10.4.11 Thumb Instructions

In general, the timing of Thumb instructions are the same as their equivalent ARM instructions, except for the cases listed below.

- If the equivalent ARM instruction maps to one in [Table 10-3](#), the “Minimum Issue Latency with Branch Misprediction” goes from 5 to 6 cycles. This is due to the branch latency penalty (see [Table 10-1](#)).
- If the equivalent ARM instruction maps to one in [Table 10-4](#), the “Minimum Issue Latency when the Branch is Taken” increases by 1 cycle. This is due to the branch latency penalty (see [Table 10-1](#)).
- A Thumb BL instruction when $H = 0$ will have the same timing as an ARM data processing instruction.

The mapping of Thumb instructions to ARM instructions can be found in the *ARM Architecture Reference Manual*.

This Page Intentionally Left Blank

Optimization Guide

A

A.1 Introduction

This document contains optimization techniques for achieving the highest performance from the Intel XScale[®] core architecture. It is written for developers who are optimizing compilers or performance analysis tools for Intel XScale[®] core based processors. It can also be used by application developers to obtain the best performance from their assembly language code. The optimizations presented in this chapter are based on the Intel XScale[®] core, and hence can be applied to all products that are based on it.

The Intel XScale[®] core architecture includes a superpipelined RISC architecture with an enhanced memory pipeline. The Intel XScale[®] core instruction set is based on ARM V5TE architecture; however, the core includes new instructions. Code generated for the SA110, SA1100 and SA1110 will execute on Intel XScale[®] core based processors, however to obtain the maximum performance of your application code, it should be optimized for the core using the techniques presented in this document.

A.1.1 About This Guide

This guide considers that you are familiar with the ARM instruction set and the C language. It consists of the following sections:

[Section A.1, "Introduction"](#). Outlines the contents of this guide.

[Section A.2, "The Intel XScale[®] Core Pipeline"](#). This chapter provides an overview of the core pipeline behavior.

[Section A.3, "Basic Optimizations"](#). This chapter outlines basic optimizations that can be applied to the core.

[Section A.4, "Cache and Prefetch Optimizations"](#). This chapter contains optimizations for efficient use of caches. Also included are optimizations that take advantage of the prefetch instruction of the core.

[Section A.5, "Instruction Scheduling"](#). This chapter shows how to optimally schedule code for the core pipeline.

[Section A.6, "Optimizing C Libraries"](#). This chapter contains information relating to optimizations for C library routines.

[Section A.7, "Optimizations for Size"](#). This chapter contains optimizations that reduce the size of the generated code. Thumb optimizations are also included.

A.2 The Intel XScale® Core Pipeline

One of the biggest differences between the Intel XScale® core and StrongARM processors is the pipeline. Many of the differences are summarized in [Figure A-1](#). This section provides a brief description of the structure and behavior of the core pipeline.

A.2.1 General Pipeline Characteristics

While the core pipeline is scalar and single issue, instructions may occupy all three pipelines at once. Out of order completion is possible. The following sections discuss general pipeline characteristics.

A.2.1.1. Number of Pipeline Stages

The Intel XScale® core has a longer pipeline (7 stages versus 5 stages) which operates at a much higher frequency than its predecessors do. This allows for greater overall performance. The longer core pipeline has several negative consequences, however:

- Larger branch misprediction penalty (4 cycles in the core instead of 1 in StrongARM Architecture). This is mitigated by dynamic branch prediction.
- Larger load use delay (LUD) - LUDs arise from load-use dependencies. A load-use dependency gives rise to a LUD if the result of the load instruction cannot be made available by the pipeline in due time for the subsequent instruction. An optimizing compiler should find independent instructions to fill the slot following the load.
- Certain instructions incur a few extra cycles of delay on the core as compared to StrongARM processors (**LDM**, **STM**).
- Decode and register file lookups are spread out over 2 cycles in the core, instead of 1 cycle in predecessors.

A.2.1.2. The Intel XScale® Core Pipeline Organization

The Intel XScale® core single-issue superpipeline consists of a main execution pipeline, MAC pipeline, and a memory access pipeline. These are shown in Figure A-1, with the main execution pipeline shaded.

Figure A-1. The Intel XScale® Core RISC Superpipeline

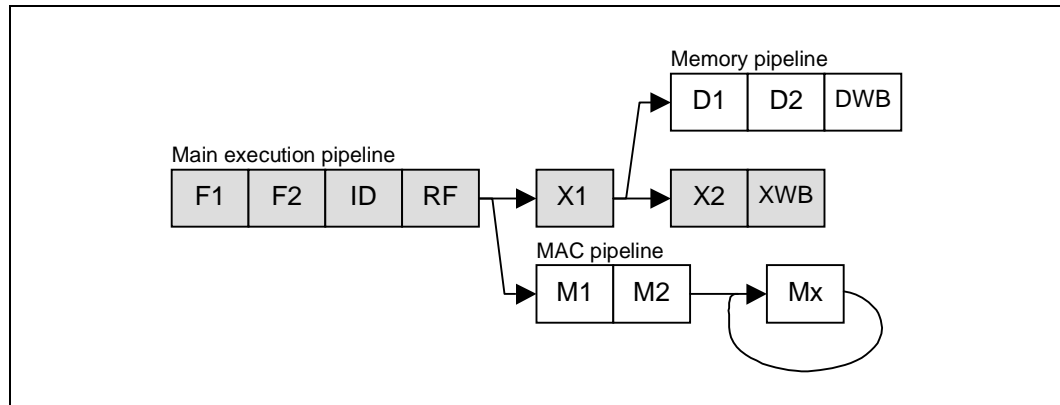


Table A-1 gives a brief description of each pipe-stage.

Table A-1. Pipelines and Pipe stages

Pipe / Pipestage	Description	Covered In
Main Execution Pipeline	Handles data processing instructions	Section A.2.3
IF1/IF2	Instruction Fetch	"
ID	Instruction Decode	"
RF	Register File / Operand Shifter	"
X1	ALU Execute	"
X2	State Execute	"
XWB	Write-back	"
Memory Pipeline	Handles load/store instructions	Section A.2.4
D1/D2	Data Cache Access	"
DWB	Data cache writeback	"
MAC Pipeline	Handles all multiply instructions	Section A.2.5
M1-M5	Multiplier stages	"
MWB (not shown)	MAC write-back - may occur during M2-M5	"

A.2.1.3. Out Of Order Completion

Sequential consistency of instruction execution relates to two aspects: first, to the order in which the instructions are completed; and second, to the order in which memory is accessed due to load and store instructions. The Intel XScale® core preserves a weak processor consistency because instructions may complete out of order, provided that no data dependencies exist.

While instructions are issued in-order, the main execution pipeline, memory, and MAC pipelines are not lock-stepped, and, therefore, have different execution times. This means that instructions may finish out of program order. Short 'younger' instructions may be finished earlier than long 'older' ones. (The term 'to finish' is used here to indicate that the operation has been completed and the result has been written back to the register file.)

A.2.1.4. Register Scoreboarding

In certain situations, the pipeline may need to be stalled because of register dependencies between instructions. A register dependency occurs when a previous MAC or load instruction is about to modify a register value that has not been returned to the register file and the current instruction needs access to the same register. Only the destination of MAC operations and memory loads are scoreboarded. The destinations of ALU instructions are not scoreboarded.

If no register dependencies exist, the pipeline will not be stalled. For example, if a load operation has missed the data cache, subsequent instructions that do not depend on the load may complete independently.

A.2.1.5. Use of Bypassing

The Intel XScale® core pipeline makes extensive use of bypassing to minimize data hazards. Bypassing allows results forwarding from multiple sources, eliminating the need to stall the pipeline.

A.2.2 Instruction Flow Through the Pipeline

The Intel XScale® core pipeline issues a single instruction per clock cycle. Instruction execution begins at the F1 pipestage and completes at the WB pipestage.

Although a single instruction may be issued per clock cycle, all three pipelines (MAC, memory, and main execution) may be processing instructions simultaneously. If there are no data hazards, then each instruction may complete independently of the others.

Each pipestage takes a single clock cycle or machine cycle to perform its subtask with the exception of the MAC unit.

A.2.2.1. ARM* V5TE Instruction Execution

Figure A-1 uses arrows to show the possible flow of instructions in the pipeline. Instruction execution flows from the F1 pipestage to the RF pipestage. The RF pipestage may issue a single instruction to either the X1 pipestage or the MAC unit (multiply instructions go to the MAC, while all others continue to X1). This means that M1 or X1 will be idle.

All load/store instructions are routed to the memory pipeline after the effective addresses have been calculated in X1.

The ARM V5TE bx (branch and exchange) instruction, which is used to branch between ARM and THUMB code, causes the entire pipeline to be flushed (The bx instruction is not dynamically predicted by the BTB). If the processor is in Thumb mode, then the ID pipestage dynamically expands each Thumb instruction into a normal ARM V5TE RISC instruction and execution resumes as usual.

A.2.2.2. Pipeline Stalls

The progress of an instruction can stall anywhere in the pipeline. Several pipestages may stall for various reasons. It is important to understand when and how hazards occur in the core pipeline. Performance degradation can be significant if care is not taken to minimize pipeline stalls.

A.2.3 Main Execution Pipeline

A.2.3.1. F1 / F2 (Instruction Fetch) Pipestages

The job of the instruction fetch stages F1 and F2 is to present the next instruction to be executed to the ID stage. Several important functional units reside within the F1 and F2 stages, including:

- *Branch Target Buffer (BTB)*
- *Instruction Fetch Unit (IFU)*

An understanding of the BTB (See [Chapter 5, “Branch Target Buffer”](#)) and IFU are important for performance considerations. A summary of operation is provided here so that the reader may understand its role in the F1 pipestage.

- Branch Target Buffer (BTB)

The BTB predicts the outcome of branch type instructions. Once a branch type instruction reaches the X1 pipestage, its target address is known. If this address is different from the address that the BTB predicted, the pipeline is flushed, execution starts at the new target address, and the branch's history is updated in the BTB.

- Instruction Fetch Unit (IFU)

The IFU is responsible for delivering instructions to the *instruction decode* (ID) pipestage. One instruction word is delivered each cycle (if possible) to the ID. The instruction could come from one of two sources: instruction cache or fetch buffers.

A.2.3.2. ID (Instruction Decode) Pipestage

The ID pipestage accepts an instruction word from the IFU and sends register decode information to the RF pipestage. The ID is able to accept a new instruction word from the IFU on every clock cycle in which there is no stall. The ID pipestage is responsible for:

- General instruction decoding (extracting the opcode, operand addresses, destination addresses and the offset).
- Detecting undefined instructions and generating an exception.
- Dynamic expansion of complex instructions into sequence of simple instructions. Complex instructions are defined as ones that take more than one clock cycle to issue, such as **LDM**, **STM**, and **SWP**.

A.2.3.3. RF (Register File / Shifter) Pipestage

The main function of the RF pipestage is to read and write to the *register file unit*, or *RFU*. It provides source data to:

- EX for ALU operations
- MAC for multiply operations
- Data Cache for memory writes
- Coprocessor interface

The ID unit decodes the instruction and specifies which registers are accessed in the RFU. Based upon this information, the RFU determines if it needs to stall the pipeline due to a register dependency. A register dependency occurs when a previous instruction is about to modify a register value that has not been returned to the RFU and the current instruction needs to access that same register. If no dependencies exist, the RFU will select the appropriate data from the register file and pass it to the next pipestage. When a register dependency does exist, the RFU will keep track of which register is unavailable and when the result is returned, the RFU will stop stalling the pipe.

The ARM architecture specifies that one of the operands for data processing instructions as the shifter operand, where a 32-bit shift can be performed before it is used as an input to the ALU. This shifter is located in the second half of the RF pipestage.

A.2.3.4. X1 (Execute) Pipestages

The X1 pipestage performs the following functions:

- ALU calculation - the ALU performs arithmetic and logic operations, as required for data processing instructions and load/store index calculations.
- Determine conditional instruction execution - The instruction's condition is compared to the CPSR prior to execution of each instruction. Any instruction with a false condition is cancelled, and will not cause any architectural state changes, including modifications of registers, memory, and PSR.
- Branch target determination - If a branch was mispredicted by the BTB, the X1 pipestage flushes all of the instructions in the previous pipestages and sends the branch target address to the BTB, which will restart the pipeline

A.2.3.5. X2 (Execute 2) Pipestage

The X2 pipestage contains the *program status registers* (PSRs). This pipestage selects what is going to be written to the RFU in the WB cycle: PSRs (MRS instruction), ALU output, or other items.

A.2.3.6. WB (write-back)

When an instruction has reached the write-back stage, it is considered complete. Changes are written to the RFU.

A.2.4 Memory Pipeline

The memory pipeline consists of two stages, D1 and D2. The *data cache unit*, or DCU, consists of the data-cache array, mini-data cache, fill buffers, and writebuffers. The memory pipeline handles load / store instructions.

A.2.4.1. D1 and D2 Pipestage

Operation begins in D1 after the X1 pipestage has calculated the effective address for load/stores. The data cache and mini-data cache returns the destination data in the D2 pipestage. Before data is returned in the D2 pipestage, sign extension and byte alignment occurs for byte and half-word loads.

A.2.5 Multiply/Multiply Accumulate (MAC) Pipeline

The Multiply-Accumulate (MAC) unit executes the multiply and multiply-accumulate instructions supported by the core. The MAC implements the 40-bit accumulator register `acc0` and handles the instructions, which transfer its value to and from general-purpose ARM registers.

The following are important characteristics about the MAC:

- The MAC is not truly pipelined, as the processing of a single instruction may require use of the same datapath resources for several cycles before a new instruction can be accepted. The type of instruction and source arguments determines the number of cycles required.
- No more than two instructions can occupy the MAC pipeline concurrently.
- When the MAC is processing an instruction, another instruction may not enter M1 unless the original instruction completes in the next cycle.
- The MAC unit can operate on 16-bit packed signed data. This reduces register pressure and memory traffic size. Two 16-bit data items can be loaded into a register with one LDR.
- The MAC can achieve throughput of one multiply per cycle when performing a 16 by 32 bit multiply.

A.2.5.1. Behavioral Description

The execution of the MAC unit starts at the beginning of the M1 pipestage, where it receives two 32-bit source operands. Results are completed N cycles later (where N is dependent on the operand size) and returned to the register file. For more information on MAC instruction latencies, refer to [Section 10.4, "Instruction Latencies"](#).

An instruction that occupies the M1 or M2 pipestages will also occupy the X1 and X2 pipestage, respectively. Each cycle, a MAC operation progresses for M1 to M5. A MAC operation may complete anywhere from M2-M5. If a MAC operation enters M3-M5, it is considered committed because it will modify architectural state regardless of subsequent events.

A.3 Basic Optimizations

This chapter outlines optimizations specific to ARM architecture. These optimizations have been modified to suit the core where needed.

A.3.1 Conditional Instructions

The Intel XScale® core architecture provides the ability to execute instructions conditionally. This feature combined with the ability of the core instructions to modify the condition codes makes possible a wide array of optimizations.

A.3.1.1. Optimizing Condition Checks

The Intel XScale® core instructions can selectively modify the state of the condition codes. When generating code for if-else and loop conditions it is often beneficial to make use of this feature to set condition codes, thereby eliminating the need for a subsequent compare instruction. Consider the C code segment:

```
if (a + b)
```

Code generated for the if condition without using an add instruction to set condition codes is:

```
;Assume r0 contains the value a, and r1 contains the value b
    add    r0,r0,r1
    cmp    r0, #0
```

However, code can be optimized as follows making use of add instruction to set condition codes:

```
;Assume r0 contains the value a, and r1 contains the value b
    adds  r0,r0,r1
```

The instructions that increment or decrement the loop counter can also be used to modify the condition codes. This eliminates the need for a subsequent compare instruction. A conditional branch instruction can then be used to exit or continue with the next loop iteration.

Consider the following C code segment:

```
for (i = 10; i != 0; i--)
{
    do something;
}
```

The optimized code generated for the above code segment would look like:

```
L6:
.
.
    subs r3, r3, #1
    bne  .L6
```

It is also beneficial to rewrite loops whenever possible so as to make the loop exit conditions check against the value 0. For example, the code generated for the code segment below will need a compare instruction to check for the loop exit condition.

```
for (i = 0; i < 10; i++)
{
    do something;
}
```

If the loop were rewritten as follows, the code generated avoids using the compare instruction to check for the loop exit condition.

```
for (i = 9; i >= 0; i--)
{
    do something;
}
```

A.3.1.2. Optimizing Branches

Branches decrease application performance by indirectly causing pipeline stalls. Branch prediction improves the performance by lessening the delay inherent in fetching a new instruction stream. The number of branches that can accurately be predicted is limited by the size of the branch target buffer. Since the total number of branches executed in a program is relatively large compared to the size of the branch target buffer; it is often beneficial to minimize the number of branches in a program. Consider the following C code segment.

```
int foo(int a)
{
    if (a > 10)
        return 0;
    else
        return 1;
}
```

The code generated for the if-else portion of this code segment using branches is:

```
    cmp    r0, #10
    ble   L1
    mov   r0, #0
    b     L2
L1:
    mov   r0, #1
L2:
```

The code generated above takes three cycles to execute the else part and four cycles for the if-part assuming best case conditions and no branch misprediction penalties. In the case of the Intel XScale® core, a branch misprediction incurs a penalty of four cycles. If the branch is mispredicted 50% of the time, and if we consider that both the if-part and the else-part are equally likely to be taken, on an average the code above takes 5.5 cycles to execute.

$$\left(\frac{50}{100} \times 4 + \frac{3+4}{2}\right) = 5.5 \quad \text{cycles.}$$

If we were to use the core to execute instructions conditionally, the code generated for the above if-else statement is:

```
    cmp    r0, #10
    movgt r0, #0
    movle r0, #1
```

The above code segment would not incur any branch misprediction penalties and would take three cycles to execute assuming best case conditions. As can be seen, using conditional instructions speeds up execution significantly. However, the use of conditional instructions should be carefully considered to ensure that it does improve performance. To decide when to use conditional instructions over branches consider the following hypothetical code segment:

```
if (cond)
    if_stmt
else
    else_stmt
```


Consider that we have the following data:

- N1_B Number of cycles to execute the if_stmt assuming the use of branch instructions
- N2_B Number of cycles to execute the else_stmt assuming the use of branch instructions
- P1 Percentage of times the if_stmt is likely to be executed
- P2 Percentage of times we are likely to incur a branch misprediction penalty
- N1_C Number of cycles to execute the if-else portion using conditional instructions assuming the if-condition to be true
- N2_C Number of cycles to execute the if-else portion using conditional instructions assuming the if-condition to be false

Once we have the above data, use conditional instructions when:

$$\left(N1_C \times \frac{P1}{100}\right) + \left(N2_C \times \frac{100 - P1}{100}\right) \leq \left(N1_B \times \frac{P1}{100}\right) + \left(N2_B \times \frac{100 - P1}{100}\right) + \left(\frac{P2}{100} \times 4\right)$$

The following example illustrates a situation in which we are better off using branches over conditional instructions. Consider the code sample shown below:

```

cmp    r0, #0
bne   L1
add    r0, r0, #1
add    r1, r1, #1
add    r2, r2, #1
add    r3, r3, #1
add    r4, r4, #1
b     L2
L1:
sub    r0, r0, #1
sub    r1, r1, #1
sub    r2, r2, #1
sub    r3, r3, #1
sub    r4, r4, #1
L2:

```

In the above code sample, the cmp instruction takes 1 cycle to execute, the if-part takes 7 cycles to execute and the else-part takes 6 cycles to execute. If we were to change the code above so as to eliminate the branch instructions by making use of conditional instructions, the if-else part would always take 10 cycles to complete.

If we make the assumptions that both paths are equally likely to be taken and that branches are mis-predicted 50% of the time, the costs of using conditional execution Vs using branches can be computed as follows:

Cost of using conditional instructions:

$$1 + \left(\frac{50}{100} \times 10\right) + \left(\frac{50}{100} \times 10\right) = 11 \quad \text{cycles}$$

Cost of using branches:

$$1 + \left(\frac{50}{100} \times 7\right) + \left(\frac{50}{100} \times 6\right) + \left(\frac{50}{100} \times 4\right) = 9.5 \quad \text{cycles}$$

As can be seen, we get better performance by using branch instructions in the above scenario.

A.3.1.3. Optimizing Complex Expressions

Conditional instructions should also be used to improve the code generated for complex expressions such as the C shortcut evaluation feature. Consider the following C code segment:

```
int foo(int a, int b)
{
    if (a != 0 && b != 0)
        return 0;
    else
        return 1;
}
```

The optimized code for the if condition is:

```
cmp    r0, #0
cmprne r1, #0
```

Similarly, the code generated for the following C segment

```
int foo(int a, int b)
{
    if (a != 0 || b != 0)
        return 0;
    else
        return 1;
}
```

is:

```
cmp    r0, #0
cmpeq  r1, #0
```

The use of conditional instructions in the above fashion improves performance by minimizing the number of branches, thereby minimizing the penalties caused by branch mispredictions. This approach also reduces the utilization of branch prediction resources.

A.3.2 Bit Field Manipulation

The Intel XScale[®] core shift and logical operations provide a useful way of manipulating bit fields. Bit field operations can be optimized as follows:

```
;Set the bit number specified by r1 in register r0
  mov  r2, #1
  orr  r0, r0, r2, asl r1
;Clear the bit number specified by r1 in register r0
  mov  r2, #1
  bic  r0, r0, r2, asl r1
;Extract the bit-value of the bit number specified by r1 of the
;value in r0 storing the value in r0
  mov  r1, r0, asr r1
  and  r0, r1, #1
;Extract the higher order 8 bits of the value in r0 storing
;the result in r1
  mov  r1, r0, lsr #24
```

A.3.3 Optimizing the Use of Immediate Values

The Intel XScale® core **MOV** or **MVN** instruction should be used when loading an immediate (constant) value into a register. Please refer to the *ARM Architecture Reference Manual* for the set of immediate values that can be used in a **MOV** or **MVN** instruction. It is also possible to generate a whole set of constant values using a combination of **MOV**, **MVN**, **ORR**, **BIC**, and **ADD** instructions. The **LDR** instruction has the potential of incurring a cache miss in addition to polluting the data and instruction caches. The code samples below illustrate cases when a combination of the above instructions can be used to set a register to a constant value:

```
;Set the value of r0 to 127
    mov    r0, #127
;Set the value of r0 to 0xfffffeb.
    mvn   r0, #260
;Set the value of r0 to 257
    mov    r0, #1
    orr   r0, r0, #256
;Set the value of r0 to 0x51f
    mov    r0, #0x1f
    orr   r0, r0, #0x500
;Set the value of r0 to 0xf100ffff
    mvn   r0, #0xff, 16
    bic   r0, r0, #0xe, 8
; Set the value of r0 to 0x12341234
    mov    r0, #0x8d, 30
    orr   r0, r0, #0x1, 20
    add   r0, r0, r0, LSL #16 ; shifter delay of 1 cycle
```

Note: It is possible to load any 32-bit value into a register using a sequence of four instructions.

A.3.4 Optimizing Integer Multiply and Divide

Multiplication by an integer constant should be optimized to make use of the shift operation whenever possible.

```
;Multiplication of R0 by 2n
  mov   r0, r0, LSL #n
;Multiplication of R0 by 2n+1
  add   r0, r0, r0, LSL #n
```

Multiplication by an integer constant that can be expressed as $(2^n + 1) \cdot (2^m)$ can similarly be optimized as:

```
;Multiplication of r0 by an integer constant that can be
;expressed as (2n+1)*(2m)
  add   r0, r0, r0, LSL #n
  mov   r0, r0, LSL #m
```

Please note that the above optimization should only be used in cases where the multiply operation cannot be advanced far enough to prevent pipeline stalls.

Dividing an unsigned integer by an integer constant should be optimized to make use of the shift operation whenever possible.

```
;Dividing r0 containing an unsigned value by an integer constant
;that can be represented as 2n
  mov   r0, r0, LSR #n
```

Dividing a signed integer by an integer constant should be optimized to make use of the shift operation whenever possible.

```
;Dividing r0 containing a signed value by an integer constant
;that can be represented as 2n
  mov   r1, r0, ASR #31
  add   r0, r0, r1, LSR #(32 - n)
  mov   r0, r0, ASR #n
```

The add instruction would stall for one cycle. The stall can be prevented by filling in another instruction before add.

A.3.5 Effective Use of Addressing Modes

The Intel XScale® core provides a variety of addressing modes that make indexing an array of objects highly efficient. For a detailed description of these addressing modes please refer to the *ARM Architecture Reference Manual*. The following code samples illustrate how various kinds of array operations can be optimized to make use of these addressing modes:

```
;Set the contents of the word pointed to by r0 to the value
;contained in r1 and make r0 point to the next word
    str    r1,[r0], #4
;Increment the contents of r0 to make it point to the next word
;and set the contents of the word pointed to the value contained
;in r1
    str    r1, [r0, #4]!
;Set the contents of the word pointed to by r0 to the value
;contained in r1 and make r0 point to the previous word
    str    r1,[r0], #-4
;Decrement the contents of r0 to make it point to the previous
;word and set the contents of the word pointed to the value
;contained in r1
    str    r1,[r0, #-4]!
```

A.4 Cache and Prefetch Optimizations

This section considers how to use the various cache memories in all their modes and then examines when and how to use prefetch to improve execution efficiencies.

A.4.1 Instruction Cache

The Intel XScale[®] core has separate instruction and data caches. Only fetched instructions are held in the instruction cache even though both data and instructions may reside within the same memory space with each other. Functionally, the instruction cache is either enabled or disabled. There is no performance benefit in not using the instruction cache. The exception is that code, which locks code into the instruction cache, must itself execute from non-cached memory.

A.4.1.1. Cache Miss Cost

The Intel XScale[®] core performance is highly dependent on reducing the cache miss rate. Refer to the Intel XScale[®] core implementation option section of the ASSP architecture specification for more information on the cycle penalty associated with cache misses. Note that this cycle penalty becomes significant when the core is running much faster than external memory. Executing non-cached instructions severely curtails the processor's performance in this case and it is very important to do everything possible to minimize cache misses.

A.4.1.2. Round-Robin Replacement Cache Policy

Both the data and the instruction caches use a round robin replacement policy to evict a cache line. The simple consequence of this is that at sometime every line will be evicted, assuming a non-trivial program. The less obvious consequence is that predicting when and over which cache lines evictions take place is very difficult to predict. This information must be gained by experimentation using performance profiling.

A.4.1.3. Code Placement to Reduce Cache Misses

Code placement can greatly affect cache misses. One way to view the cache is to think of it as 32 sets of 32 bytes, which span an address range of 1024 bytes. When running, the code maps into 32 blocks modular 1024 of cache space. Any sets, which are overused, will thrash the cache. The ideal situation is for the software tools to distribute the code on a temporal evenness over this space.

This is very difficult if not impossible for a compiler to do. Most of the input needed to best estimate how to distribute the code will come from profiling followed by compiler based two pass optimizations.

A.4.1.4. Locking Code into the Instruction Cache

One very important instruction cache feature is the ability to lock code into the instruction cache. Once locked into the instruction cache, the code is always available for fast execution. Another reason for locking critical code into cache is that with the round robin replacement policy, eventually the code will be evicted, even if it is a very frequently executed function. Key code components to consider for locking are:

- Interrupt handlers
- Real time clock handlers
- OS critical code
- Time critical application code

The disadvantage to locking code into the cache is that it reduces the cache size for the rest of the program. How much code to lock is very application dependent and requires experimentation to optimize.

Code placed into the instruction cache should be aligned on a 1024 byte boundary and placed sequentially together as tightly as possible so as not to waste precious memory space. Making the code sequential also insures even distribution across all cache ways. Though it is possible to choose randomly located functions for cache locking, this approach runs the risk of landing multiple cache ways in one set and few or none in another set. This distribution unevenness can lead to excessive thrashing of the Data and Mini Caches

A.4.2 Data and Mini Cache

The Intel XScale® core allows the user to define memory regions whose cache policies can be set by the user (see [Section 6.2.3, “Cache Policies”](#)). Supported policies and configurations are:

- Non Cacheable with no coalescing of memory writes.
- Non Cacheable with coalescing of memory writes.
- Mini-Data cache with write coalescing, read allocate, and write-back caching.
- Mini-Data cache with write coalescing, read allocate, and write-through caching.
- Mini-Data cache with write coalescing, read-write allocate, and write-back caching.
- Data cache with write coalescing, read allocate, and write-back caching.
- Data cache with write coalescing, read allocate, and write-through caching.
- Data cache with write coalescing, read-write allocate, and write-back caching.

To support allocating variables to these various memory regions, the tool chain (compiler, assembler, linker and debugger), must implement named sections.

The performance of your application code depends on what cache policy you are using for data objects. A description of when to use a particular policy is described below.

The Intel XScale® core allows dynamic modification of the cache policies at run time, however, the operation is requires considerable processing time and therefore should not be used by applications.

If the application is running under an OS, then the OS may restrict you from using certain cache policies.

A.4.2.1. Non Cacheable Regions

It is recommended that non-cache memory (X=0, C=0, and B=0) be used only if necessary as is often necessary for I/O devices. Accessing non-cacheable memory is likely to cause the processor to stall frequently due to the long latency of memory reads.

A.4.2.2. Write-through and Write-back Cached Memory Regions

Write through memory regions generate more data traffic on the bus. Therefore is not recommended that the write-through policy be used. The write back policy must be used whenever possible.

However, in a multiprocessor environment it will be necessary to use a write through policy if data is shared across multiple processors. In such a situation all shared memory regions should use write through policy. Memory regions that are private to a particular processor should use the write back policy.

A.4.2.3. Read Allocate and Read-write Allocate Memory Regions

Most of the regular data and the stack for your application should be allocated to a read-write allocate region. It is expected that you will be writing and reading from them often.

Data that is write only (or data that is written to and subsequently not used for a long time) should be placed in a read allocate region. Under the read-allocate policy if a cache write miss occurs a new cache line will not be allocated, and hence will not evict critical data from the Data cache.

A.4.2.4. Creating On-chip RAM

Part of the Data cache can be converted into fast on chip RAM. Access to objects in the on-chip RAM will not incur cache miss penalties, thereby reducing the number of processor stalls. Application performance can be improved by converting a part of the cache into on chip RAM and allocating frequently allocated variables to it. Due to the core round robin replacement policy, all data will eventually be evicted. Therefore to prevent critical or frequently used data from being evicted it should be allocated to on-chip RAM.

The following variables are good candidates for allocating to the on-chip RAM:

- Frequently used global data used for storing context for context switching.
- Global variables that are accessed in time critical functions such as interrupt service routines.

The on-chip RAM is created by locking a memory region into the Data cache (see [Section 6.4, "Re-configuring the Data Cache as Data RAM"](#) for more details).

When creating the on-chip RAM, care must be taken to ensure that all sets in the on-chip RAM area of the Data cache have approximately the same number of ways locked, otherwise some sets will have more ways locked than the others. This uneven allocation will increase the level of thrashing in some sets and leave other sets under utilized.

For example, consider three arrays arr1, arr2 and arr3 of size 64 bytes each that are being allocated to the on-chip RAM and consider that the address of arr1 is 0, address of arr2 is 1024, and the address of arr3 is 2048. All three arrays will be within the same sets, i.e. set0 and set1, as a result three ways in both sets set0 and set1, will be locked, leaving 29 ways for use by other variables.

This can be overcome by allocating on-chip RAM data in sequential order. In the above example allocating arr2 to address 64 and arr3 to address 128, allows the three arrays to use only 1 way in sets 0 through 5.

A.4.2.5. Mini-data Cache

The mini-data cache is best used for data structures, which have short temporal lives, and/or cover vast amounts of data space. Addressing these types of data spaces from the Data cache would corrupt much if not all of the Data cache by evicting valuable data. Eviction of valuable data will reduce performance. Placing this data instead in Mini-data cache memory region would prevent Data cache corruption while providing the benefits of cached accesses.

A prime example of using the mini-data cache would be for caching the procedure call stack. The stack can be allocated to the mini-data cache so that it's use does not trash the main dcache. This would keep local variables from global data.

Following are examples of data that could be assigned to mini-dcache:

- The stack space of a frequently occurring interrupt, the stack is used only during the duration of the interrupt, which is usually very small.
- Video buffers, these are usual large and can occupy the whole cache.

Over use of the Mini-Data cache will thrash the cache. This is easy to do because the Mini-Data cache only has two ways per set. For example, a loop which uses a simple statement such as:

```
for (i=0; I< IMAX; i++)  
{  
    A[i] = B[i] + C[i];  
}
```

Where A, B, and C reside in a mini-data cache memory region and each is array is aligned on a 1K boundary will quickly thrash the cache.

A.4.2.6. Data Alignment

Cache lines begin on 32-byte address boundaries. To maximize cache line use and minimize cache pollution, data structures should be aligned on 32 byte boundaries and sized to multiple cache line sizes. Aligning data structures on cache address boundaries simplifies later addition of prefetch instructions to optimize performance.

Not aligning data on cache lines has the disadvantage of moving the prefetch address correspondingly to the misalignment. Consider the following example:

```
struct {
    long ia;
    long ib;
    long ic;
    long id;
} tdata[IMAX];

for (i=0, i<IMAX; i++)
{
    PREFETCH(tdata[i+1]);
    tdata[i].ia = tdata[i].ib + tdata[i].ic _tdata[i].id;
    ....
    tdata[i].id = 0;
}
```

In this case if tdata[] is not aligned to a cache line, then the prefetch using the address of tdata[i+1].ia may not include element id. If the array was aligned on a cache line + 12 bytes, then the prefetch would have to be placed on &tdata[i+1].id.

If the structure is not sized to a multiple of the cache line size, then the prefetch address must be advanced appropriately and will require extra prefetch instructions. Consider the following example:

```
struct {
    long ia;
    long ib;
    long ic;
    long id;
    long ie;
} tdata[IMAX];

ADDRESS preadd = tdata

for (i=0, i<IMAX; i++)
{
    PREFETCH(preaddata+=16);
    tdata[I].ia = tdata[I].ib + tdata[I].ic _tdata[I].id +
    tdata[I].ie;
    ....
    tdata[I].ie = 0;
}
```

In this case, the prefetch address was advanced by size of half a cache line and every other prefetch instruction is ignored. Further, an additional register is required to track the next prefetch address.

Generally, not aligning and sizing data will add extra computational overhead.

Additional prefetch considerations are discussed in greater detail in following sections.

A.4.2.7. Literal Pools

The Intel XScale[®] core does not have a single instruction that can move all literals (a constant or address) to a register. One technique to load registers with literals in the core is by loading the literal from a memory location that has been initialized with the constant or address. These blocks of constants are referred to as literal pools. See [Section A.3, “Basic Optimizations”](#) for more information on how to do this. It is advantageous to place all the literals together in a pool of memory known as a literal pool. These data blocks are located in the text or code address space so that they can be loaded using PC relative addressing. However, references to the literal pool area load the data into the data cache instead of the instruction cache. Therefore it is possible that the literal may be present in both the data and instruction caches, resulting in waste of space.

For maximum efficiency, the compiler should align all literal pools on cache boundaries and size each pool to a multiple of 32 bytes (the size of a cache line). One additional optimization would be group highly used literal pool references into the same cache line. The advantage is that once one of the literals has been loaded, the other seven will be available immediately from the data cache.

A.4.3 Cache Considerations

A.4.3.1. Cache Conflicts, Pollution and Pressure

Cache pollution occurs when unused data is loaded in the cache and cache pressure occurs when data that is not temporal to the current process is loaded into the cache. For an example, see [Section A.4.4.2., “Prefetch Loop Scheduling”](#) below.

A.4.3.2. Memory Page Thrashing

Memory page thrashing occurs because of the nature of SDRAM. SDRAMs are typically divided into 4 banks. Each bank can have one selected page where a page address size for current memory components is often defined as 4k. Memory lookup time or latency time for a selected page address is currently 2 to 3 bus clocks. Thrashing occurs when subsequent memory accesses within the same memory bank access different pages. The memory page change adds 3 to 4 bus clock cycles to memory latency. This added delay extends the prefetch distance correspondingly making it more difficult to hide memory access latencies. This type of thrashing can be resolved by placing the conflicting data structures into different memory banks or by paralleling the data structures such that the data resides within the same memory page. It is also extremely important to insure that instruction and data sections are in different memory banks, or they will continually trash the memory page selection.

A.4.4 Prefetch Considerations

The Intel XScale[®] core has a true prefetch load instruction (PLD). The purpose of this instruction is to preload data into the data and mini-data caches. Data prefetching allows hiding of memory transfer latency while the processor continues to execute instructions. The prefetch is important to compiler and assembly code because judicious use of the prefetch instruction can enormously improve throughput performance of the core. Data prefetch can be applied not only to loops but also to any data references within a block of code. Prefetch also applies to data writing when the memory type is enabled as write allocate

The Intel XScale[®] core prefetch load instruction is a true prefetch instruction because the load destination is the data or mini-data cache and not a register. Compilers for processors which have data caches, but do not support prefetch, sometimes use a load instruction to preload the data cache. This technique has the disadvantages of using a register to load data and requiring additional registers for subsequent preloads and thus increasing register pressure. By contrast, the prefetch can be used to reduce register pressure instead of increasing it.

The prefetch load is a hint instruction and does not guarantee that the data will be loaded. Whenever the load would cause a fault or a table walk, then the processor will ignore the prefetch instruction, the fault or table walk, and continue processing the next instruction. This is particularly advantageous in the case where a linked list or recursive data structure is terminated by a NULL pointer. Prefetching the NULL pointer will not fault program flow.

A.4.4.1. Prefetch Distances

Scheduling the prefetch instruction requires understanding the system latency times and system resources which affect when to use the prefetch instruction. Refer to the Intel XScale[®] core implementation option section of the ASSP architecture specification for more information.

A.4.4.2. Prefetch Loop Scheduling

When adding prefetch to a loop which operates on arrays, it may be advantages to prefetch ahead one, two, or more iterations. The data for future iterations is located in memory by a fixed offset from the data for the current iteration. This makes it easy to predict where to fetch the data. The number of iterations to prefetch ahead is referred to as the prefetch scheduling distance. Refer to the Intel XScale[®] core implementation option section of the ASSP architecture specification for more information.

A.4.4.3. Prefetch Loop Limitations

It is not always advantages to add prefetch to a loop. Loop characteristics that limit the use value of prefetch are discussed below.

A.4.4.4. Compute vs. Data Bus Bound

At the extreme, a loop, which is data bus bound, will not benefit from prefetch because all the system resources to transfer data are quickly allocated and there are no instructions that can profitably be executed. On the other end of the scale, compute bound loops allow complete hiding of all data transfer latencies.

A.4.4.5. Low Number of Iterations

Loops with very low iteration counts may have the advantages of prefetch completely mitigated. A loop with a small fixed number of iterations may be faster if the loop is completely unrolled rather than trying to schedule prefetch instructions.

A.4.4.6. Bandwidth Limitations

Overuse of prefetches can usurp resources and degrade performance. This happens because once the bus traffic requests exceed the system resource capacity, the processor stalls. The core data transfer resources are:

- 4 fill buffers
- 4 pending buffers
- 8 half cache line write buffer

SDRAM resources are typically:

- 4 memory banks
- 1 page buffer per bank referencing a 4K address range
- 4 transfer request buffers

Consider how these resources work together. A fill buffer is allocated for each cache read miss. A fill buffer is also allocated each cache write miss if the memory space is write allocate along with a pending buffer. A subsequent read to the same cache line does not require a new fill buffer, but does require a pending buffer and a subsequent write will also require a new pending buffer. A fill buffer is also allocated for each read to a non-cached memory and a write buffer is needed for each memory write to non-cached memory that is non-coalescing. Consequently, a **STM** instruction listing eight registers and referencing non-cached memory will use eight write buffers assuming they don't coalesce and two write buffers if they do coalesce. A cache eviction requires a write buffer for each dirty bit set in the cache line. The prefetch instruction requires a fill buffer for each cache line and 0, 1, or 2 write buffers for an eviction.

When adding prefetch instructions, caution must be asserted to insure that the combination of prefetch and instruction bus requests do not exceed the system resource capacity described above or performance will be degraded instead of improved. The important points are to spread prefetch operations over calculations so as to allow bus traffic to free flow and to minimize the number of necessary prefetches.

A.4.4.7. Cache Memory Considerations

Stride, the way data structures are walked through, can affect the temporal quality of the data and reduce or increase cache conflicts. The data cache and mini-data caches each have 32 sets of 32 bytes. This means that each cache line in a set is on a modular 1K-address boundary. The caution is to choose data structure sizes and stride requirements that do not overwhelm a given set causing conflicts and increased register pressure. Register pressure can be increased because additional registers are required to track prefetch addresses. The effects can be affected by rearranging data structure components to use more parallel access to search and compare elements. Similarly rearranging sections of data structures so that sections often written fit in the same half cache line, 16 bytes for the core, can reduce cache eviction write-backs. On a global scale, techniques such as array merging can enhance the spatial locality of the data.

As an example of array merging, consider the following code:

```
int a_array[NMAX];
int b_array[NMAX];
int ix;

for (i=0; i<NMAX; i++)
{
    ix = b[i];
    if (a[i] != 0)
        ix = a[i];
    do_other calculations;
}
```

In the above code, data is read from both arrays a and b, but a and b are not spatially close. Array merging can place a and b specially close.

```
struct {
    int a;
    int b;
} c_arrays;

int ix;

for (i=0; i<NMAX; i++)
{
    ix = c[i].b;
    if (c[i].a != 0)
        ix = c[i].a;
    do_other_calculations;
}
```

As an example of rearranging often written to sections in a structure, consider the code sample:

```
struct employee {
    struct employee *prev;
    struct employee *next;
    float Year2DatePay;
    float Year2DateTax;
    int ssno;
    int empid;
    float Year2Date401KDed;
    float Year2DateOtherDed;
};
```

In the data structure shown above, the fields Year2DatePay, Year2DateTax, Year2Date401KDed, and Year2DateOtherDed are likely to change with each pay check. The remaining fields however change very rarely. If the fields are laid out as shown above, assuming that the structure is aligned on a 32-byte boundary, modifications to the Year2Date fields is likely to use two write buffers when the data is written out to memory. However, we can restrict the number of write buffers that are commonly used to 1 by rearranging the fields in the above data structure as shown below:

```
struct employee {  
    struct employee *prev;  
    struct employee *next;  
    int ssno;  
    int empid;  
    float Year2DatePay;  
    float Year2DateTax;  
    float Year2Date401KDed;  
    float Year2DateOtherDed;  
};
```

A.4.4.8. Cache Blocking

Cache blocking techniques, such as strip-mining, are used to improve temporal locality of the data. Given a large data set that can be reused across multiple passes of a loop, data blocking divides the data into smaller chunks which can be loaded into the cache during the first loop and then be available for processing on subsequent loops thus minimizing cache misses and reducing bus traffic.

As an example of cache blocking consider the following code:

```
for(i=0; i<10000; i++)
  for(j=0; j<10000; j++)
    for(k=0; k<10000; k++)
      C[j][k] += A[i][k] * B[j][i];
```

The variable A[i][k] is completely reused. However, accessing C[j][k] in the j and k loops can displace A[i][j] from the cache. Using blocking the code becomes:

```
for(i=0; i<10000; i++)
  for(j1=0; j1<100; j1++)
    for(k1=0; k1<100; k1++)
      for(j2=0; j2<100; j2++)
        for(k2=0; k2<100; k2++)
          {
            j = j1 * 100 + j2;
            k = k1 * 100 + k2;
            C[j][k] += A[i][k] * B[j][i];
          }
```

A.4.4.9. Prefetch Unrolling

When iterating through a loop, data transfer latency can be hidden by prefetching ahead one or more iterations. The solution incurs an unwanted side effect that the final interactions of a loop loads useless data into the cache, polluting the cache, increasing bus traffic and possibly evicting valuable temporal data. This problem can be resolved by prefetch unrolling. For example consider:

```
for(i=0; i<NMAX; i++)
{
  prefetch(data[i+2]);
  sum += data[i];
}
```

Interactions i-1 and i, will prefetch superfluous data. The problem can be avoided by unrolling the end of the loop.

```
for(i=0; i<NMAX-2; i++)
{
  prefetch(data[i+2]);
  sum += data[i];
}
sum += data[NMAX-2];
sum += data[NMAX-1];
```

Unfortunately, prefetch loop unrolling does not work on loops with indeterminate iterations.

A.4.4.10. Pointer Prefetch

Not all looping constructs contain induction variables. However, prefetching techniques can still be applied. Consider the following linked list traversal example:

```
while(p) {
    do_something(p->data);
    p = p->next;
}
```

The pointer variable *p* becomes a pseudo induction variable and the data pointed to by *p->next* can be prefetched to reduce data transfer latency for the next iteration of the loop. Linked lists should be converted to arrays as much as possible.

```
while(p) {
    prefetch(p->next);
    do_something(p->data);
    p = p->next;
}
```

Recursive data structure traversal is another construct where prefetching can be applied. This is similar to linked list traversal. Consider the following pre-order traversal of a binary tree:

```
preorder(treeNode *t) {
    if(t) {
        process(t->data);
        preorder(t->left);
        preorder(t->right);
    }
}
```

The pointer variable *t* becomes the pseudo induction variable in a recursive loop. The data structures pointed to by the values *t->left* and *t->right* can be prefetched for the next iteration of the loop.

```
preorder(treeNode *t) {
    if(t) {
        prefetch(t->right);
        prefetch(t->left);
        process(t->data);
        preorder(t->left);
        preorder(t->right);
    }
}
```

Note the order reversal of the prefetches in relationship to the usage. If there is a cache conflict and data is evicted from the cache then only the data from the first prefetch is lost.

A.4.4.11. Loop Interchange

As mentioned earlier, the sequence in which data is accessed affects cache thrashing. Usually, it is best to access data in a contiguous spatially address range. However, arrays of data may have been laid out such that indexed elements are not physically next to each other. Consider the following C code which places array elements in row major order.

```
for(j=0; j<NMAX; j++)
  for(i=0; i<NMAX; i++)
  {
    prefetch(A[i+1][j]);
    sum += A[i][j];
  }
```

In the above example, $A[i][j]$ and $A[i+1][j]$ are not sequentially next to each other. This situation causes an increase in bus traffic when prefetching loop data. In some cases where the loop mathematics are unaffected, the problem can be resolved by induction variable interchange. The above examples becomes:

```
for(i=0; i<NMAX; i++)
  for(j=0; j<NMAX; j++)
  {
    prefetch(A[i][j+1]);
    sum += A[i][j];
  }
```

A.4.4.12. Loop Fusion

Loop fusion is a process of combining multiple loops, which reuse the same data, in to one loop. The advantage of this is that the reused data is immediately accessible from the data cache. Consider the following example:

```
for(i=0; i<NMAX; i++)
{
  prefetch(A[i+1], c[i+1], c[i+1]);
  A[i] = b[i] + c[i];
}
for(i=0; i<NMAX; i++)
{
  prefetch(D[i+1], c[i+1], A[i+1]);
  D[i] = A[i] + c[i];
}
```

The second loop reuses the data elements $A[i]$ and $c[i]$. Fusing the loops together produces:

```
for(i=0; i<NMAX; i++)
{
  prefetch(D[i+1], A[i+1], c[i+1], b[i+1]);
  ai = b[i] + c[i];
  A[i] = ai;
  D[i] = ai + c[i];
}
```

A.4.4.13. Prefetch to Reduce Register Pressure

Prefetch can be used to reduce register pressure. When data is needed for an operation, then the load is scheduled far enough in advance to hide the load latency. However, the load ties up the receiving register until the data can be used. For example:

```
ldr    r2, [r0]
; Process code { not yet cached latency > 60 core clocks }
add   r1, r1, r2
```

In the above case, r2 is unavailable for processing until the add statement. Prefetching the data load frees the register for use. The example code becomes:

```
pld   [r0] ;prefetch the data keeping r2 available for use
; Process code
ldr   r2, [r0]
; Process code {ldr result latency is 3 core clocks}
add   r1, r1, r2
```

With the added prefetch, register r2 can be used for other operations until almost just before it is needed.

A.5 Instruction Scheduling

This chapter discusses instruction scheduling optimizations. Instruction scheduling refers to the rearrangement of a sequence of instructions for the purpose of minimizing pipeline stalls. Reducing the number of pipeline stalls improves application performance. While making this rearrangement, care should be taken to ensure that the rearranged sequence of instructions has the same effect as the original sequence of instructions.

A.5.1 Scheduling Loads

On the Intel XScale® core, an **LDR** instruction has a result latency of 3 cycles assuming the data being loaded is in the data cache. If the instruction after the **LDR** needs to use the result of the load, then it would stall for 2 cycles. If possible, the instructions surrounding the **LDR** instruction should be rearranged

to avoid this stall. Consider the following example:

```
add   r1, r2, r3
ldr   r0, [r5]
add   r6, r0, r1
sub   r8, r2, r3
mul   r9, r2, r3
```

In the code shown above, the **ADD** instruction following the **LDR** would stall for 2 cycles because it uses the result of the load. The code can be rearranged as follows to prevent the stalls:

```
ldr   r0, [r5]
add   r1, r2, r3
sub   r8, r2, r3
add   r6, r0, r1
mul   r9, r2, r3
```

Note that this rearrangement may not be always possible. Consider the following example:

```
cmp   r1, #0
addne r4, r5, #4
subeq r4, r5, #4
ldr   r0, [r4]
cmp   r0, #10
```

In the example above, the **LDR** instruction cannot be moved before the **ADDNE** or the **SUBEQ** instructions because the **LDR** instruction depends on the result of these instructions. Rewrite the above code to make it run faster at the expense of increasing code size:

```
cmp   r1, #0
ldrne r0, [r5, #4]
ldreq r0, [r5, #-4]
addne r4, r5, #4
subeq r4, r5, #4
cmp   r0, #10
```

The optimized code takes six cycles to execute compared to the seven cycles taken by the unoptimized version.

The result latency for an **LDR** instruction is significantly higher if the data being loaded is not in the data cache. To minimize the number of pipeline stalls in such a situation the **LDR** instruction should be moved as far away as possible from the instruction that uses result of the load. Note that this may at times cause certain register values to be spilled to memory due to the increase in register pressure. In such cases, use a preload instruction or a preload hint to ensure that the data access in the **LDR** instruction hits the cache when it executes. A preload hint should be used in cases where we cannot be sure whether the load instruction would be executed. A preload instruction should be used in cases where we can be sure that the load instruction would be executed. Consider the following code sample:

```
; all other registers are in use
    sub    r1, r6, r7
    mul    r3, r6, r2
    mov    r2, r2, LSL #2
    orr    r9, r9, #0xf
    add    r0, r4, r5
    ldr    r6, [r0]
    add    r8, r6, r8
    add    r8, r8, #4
    orr    r8, r8, #0xf
; The value in register r6 is not used after this
```

In the code sample above, the **ADD** and the **LDR** instruction can be moved before the **MOV** instruction. Note that this would prevent pipeline stalls if the load hits the data cache. However, if the load is likely to miss the data cache, move the **LDR** instruction so that it executes as early as possible - before the **SUB** instruction. However, moving the **LDR** instruction before the **SUB** instruction would change the program semantics. It is possible to move the **ADD** and the **LDR** instructions before the **SUB** instruction if we allow the contents of the register r6 to be spilled and restored from the stack as shown below:

```
; all other registers are in use
    str    r6, [sp, #-4]!
    add    r0, r4, r5
    ldr    r6, [r0]
    mov    r2, r2, LSL #2
    orr    r9, r9, #0xf
    add    r8, r6, r8
    ldr    r6, [sp], #4
    add    r8, r8, #4
    orr    r8, r8, #0xf
    sub    r1, r6, r7
    mul    r3, r6, r2
; The value in register r6 is not used after this
```

As can be seen above, the contents of the register r6 have been spilled to the stack and subsequently loaded back to the register r6 to retain the program semantics. Another way to optimize the code above is with the use of the preload instruction as shown below:

```
; all other registers are in use
    add    r0, r4, r5
    pld    [r0]
    sub    r1, r6, r7
    mul    r3, r6, r2
    mov    r2, r2, LSL #2
    orr    r9, r9, #0xf
    ldr    r6, [r0]
    add    r8, r6, r8
    add    r8, r8, #4
    orr    r8, r8, #0xf
; The value in register r6 is not used after this
```


The Intel XScale® core has 4 fill-buffers that are used to fetch data from external memory when a data-cache miss occurs. The core stalls when all fill buffers are in use. This happens when more than 4 loads are outstanding and are being fetched from memory. As a result, the code written should ensure that no more than 4 loads are outstanding at the same time. For example, the number of loads issued sequentially should not exceed 4. Also note that a preload instruction may cause a fill buffer to be used. As a result, the number of preload instructions outstanding should also be considered to arrive at the number of loads that are outstanding.

Similarly, the number of write buffers also limits the number of successive writes that can be issued before the processor stalls. No more than eight stores can be issued. Also note that if the data caches are using the write-allocate with writeback policy, then a load operation may cause stores to the external memory if the read operation evicts a cache line that is dirty (modified). The number of sequential stores may be limited by this fact.

A.5.1.1. Scheduling Load and Store Double (LDRD/STRD)

The Intel XScale® core introduces two new double word instructions: **LDRD** and **STRD**. **LDRD** loads 64-bits of data from an effective address into two consecutive registers, conversely, **STRD** stores 64-bits from two consecutive registers to an effective address. There are two important restrictions on how these instructions may be used:

- the effective address must be aligned on an 8-byte boundary
- the specified register must be even (r0, r2, etc.).

If this situation occurs, using **LDRD/STRD** instead of **LDM/STM** to do the same thing is more efficient because **LDRD/STRD** issues in only one/two clock cycle(s), as opposed to **LDM/STM** which issues in four clock cycles. Avoid **LDRD**s targeting R12; this incurs an extra cycle of issue latency.

The **LDRD** instruction has a result latency of 3 or 4 cycles depending on the destination register being accessed (assuming the data being loaded is in the data cache).

```
add    r6, r7, r8
sub    r5, r6, r9
; The following ldrd instruction would load values
; into registers r0 and r1
ldrd   r0, [r3]
orr    r8, r1, #0xf
mul    r7, r0, r7
```

In the code example above, the **ORR** instruction would stall for 3 cycles because of the 4 cycle result latency for the second destination register of an **LDRD** instruction. The code shown above can be rearranged to remove the pipeline stalls:

```
; The following ldrd instruction would load values
; into registers r0 and r1
ldrd   r0, [r3]
add    r6, r7, r8
sub    r5, r6, r9
mul    r7, r0, r7
orr    r8, r1, #0xf
```

Any memory operation following a **LDRD** instruction (**LDR**, **LDRD**, **STR** and so on) would stall for 1 cycle.

```
; The str instruction below would stall for 1 cycle
ldrd   r0, [r3]
str    r4, [r5]
```

A.5.1.2. Scheduling Load and Store Multiple (LDM/STM)

LDM and **STM** instructions have an issue latency of 2-20 cycles depending on the number of registers being loaded or stored. The issue latency is typically 2 cycles plus an additional cycle for each of the registers being loaded or stored assuming a data cache hit. The instruction following an **ldm** would stall whether or not this instruction depends on the results of the load. A **LDRD** or **STRD** instruction does not suffer from this drawback (except when followed by a memory operation) and should be used where possible. Consider the task of adding two 64-bit integer values. Consider that the addresses of these values are aligned on an 8 byte boundary. This can be achieved using the **LDM** instructions as shown below:

```
; r0 contains the address of the value being copied
; r1 contains the address of the destination location
ldm  r0, {r2, r3}
ldm  r1, {r4, r5}
adds r0, r2, r4
adc  r1, r3, r5
```

If the code were written as shown above, assuming all the accesses hit the cache, the code would take 11 cycles to complete. Rewriting the code as shown below using **LDRD** instruction would take only 7 cycles to complete. The performance would increase further if we can fill in other instructions after **LDRD** to reduce the stalls due to the result latencies of the **LDRD** instructions.

```
; r0 contains the address of the value being copied
; r1 contains the address of the destination location
ldr  r2, [r0]
ldr  r4, [r1]
adds r0, r2, r4
adc  r1, r3, r5
```

Similarly, the code sequence shown below takes 5 cycles to complete.

```
stm  r0, {r2, r3}
add  r1, r1, #1
```

The alternative version which is shown below would only take 3 cycles to complete.

```
strd r2, [r0]
add  r1, r1, #1
```

A.5.2 Scheduling Data Processing Instructions

Most core data processing instructions have a result latency of 1 cycle. This means that the current instruction is able to use the result from the previous data processing instruction. However, the result latency is 2 cycles if the current instruction needs to use the result of the previous data processing instruction for a shift by immediate. As a result, the following code segment would incur a 1 cycle stall for the mov instruction:

```
sub   r6, r7, r8
add   r1, r2, r3
mov   r4, r1, LSL #2
```

The code above can be rearranged as follows to remove the 1 cycle stall:

```
add   r1, r2, r3
sub   r6, r7, r8
mov   r4, r1, LSL #2
```

All data processing instructions incur a 2 cycle issue penalty and a 2 cycle result penalty when the shifter operand is a shift/rotate by a register or shifter operand is RRR. Since the next instruction would always incur a 2 cycle issue penalty, there is no way to avoid such a stall except by re-writing the assembler instruction. Consider the following segment of code:

```
mov   r3, #10
mul   r4, r2, r3
add   r5, r6, r2, LSL r3
sub   r7, r8, r2
```

The subtract instruction would incur a 1 cycle stall due to the issue latency of the add instruction as the shifter operand is shift by a register. The issue latency can be avoided by changing the code as follows:

```
mov   r3, #10
mul   r4, r2, r3
add   r5, r6, r2, LSL #10
sub   r7, r8, r2
```

A.5.3 Scheduling Multiply Instructions

Multiply instructions can cause pipeline stalls due to either resource conflicts or result latencies. The following code segment would incur a stall of 0-3 cycles depending on the values in registers r1, r2, r4 and r5 due to resource conflicts.

```
mul   r0, r1, r2
mul   r3, r4, r5
```

The following code segment would incur a stall of 1-3 cycles depending on the values in registers r1 and r2 due to result latency.

```
mul   r0, r1, r2
mov   r4, r0
```

Note that a multiply instruction that sets the condition codes blocks the whole pipeline. A 4 cycle multiply operation that sets the condition codes behaves the same as a 4 cycle issue operation. Consider the following code segment:

```
muls  r0, r1, r2
add   r3, r3, #1
sub   r4, r4, #1
sub   r5, r5, #1
```

The add operation above would stall for 3 cycles if the multiply takes 4 cycles to complete. It is better to replace the code segment above with the following sequence:

```
mul   r0, r1, r2
add   r3, r3, #1
sub   r4, r4, #1
sub   r5, r5, #1
cmp   r0, #0
```

Please refer to [Section 10.4, "Instruction Latencies"](#) to get the instruction latencies for various multiply instructions. The multiply instructions should be scheduled taking into consideration these instruction latencies.

A.5.4 Scheduling SWP and SWPB Instructions

The **SWP** and **SWPB** instructions have a 5 cycle issue latency. As a result of this latency, the instruction following the **SWP/SWPB** instruction would stall for 4 cycles. **SWP** and **SWPB** instructions should, therefore, be used only where absolutely needed.

For example, the following code may be used to swap the contents of 2 memory locations:

```
; Swap the contents of memory locations pointed to by r0 and r1
ldr  r2, [r0]
swp  r2, [r1]
str  r2, [r1]
```

The code above takes 9 cycles to complete. The rewritten code below, takes 6 cycles to execute:

```
; Swap the contents of memory locations pointed to by r0 and r1
ldr  r2, [r0]
ldr  r3, [r1]
str  r2, [r1]
str  r3, [r0]
```

A.5.5 Scheduling the MRA and MAR Instructions (MRRC/MCRR)

The **MRA (MRRC)** instruction has an issue latency of 1 cycle, a result latency of 2 or 3 cycles depending on the destination register value being accessed and a resource latency of 2 cycles.

Consider the code sample:

```
mra  r6, r7, acc0
mra  r8, r9, acc0
add  r1, r1, #1
```

The code shown above would incur a 1-cycle stall due to the 2-cycle resource latency of an **MRA** instruction. The code can be rearranged as shown below to prevent this stall.

```
mra  r6, r7, acc0
add  r1, r1, #1
mra  r8, r9, acc0
```

Similarly, the code shown below would incur a 2 cycle penalty due to the 3-cycle result latency for the second destination register.

```
mra  r6, r7, acc0
mov  r1, r7
mov  r0, r6
add  r2, r2, #1
```

The stalls incurred by the code shown above can be prevented by rearranging the code:

```
mra  r6, r7, acc0
add  r2, r2, #1
mov  r0, r6
mov  r1, r7
```

The **MAR (MCRR)** instruction has an issue latency, a result latency, and a resource latency of 2 cycles. Due to the 2-cycle issue latency, the pipeline would always stall for 1 cycle following a **MAR** instruction. The use of the **MAR** instruction should, therefore, be used only where absolutely necessary.

A.5.6 Scheduling the MIA and MIAPH Instructions

The **MIA** instruction has an issue latency of 1 cycle. The result and resource latency can vary from 1 to 3 cycles depending on the values in the source register.

Consider the following code sample:

```
mia  acc0, r2, r3
mia  acc0, r4, r5
```

The second **MIA** instruction above can stall from 0 to 2 cycles depending on the values in the registers r2 and r3 due to the 1 to 3 cycle resource latency.

Similarly, consider the following code sample:

```
mia  acc0, r2, r3
mra  r4, r5, acc0
```

The **MRA** instruction above can stall from 0 to 2 cycles depending on the values in the registers r2 and r3 due to the 1 to 3 cycle result latency.

The **MIAPH** instruction has an issue latency of 1 cycle, result latency of 2 cycles and a resource latency of 2 cycles.

Consider the code sample shown below:

```
add  r1, r2, r3
miaph acc0, r3, r4
miaph acc0, r5, r6
mra  r6, r7, acc0
sub  r8, r3, r4
```

The second **MIAPH** instruction would stall for 1-cycle due to a 2-cycle resource latency. The **MRA** instruction would stall for 1-cycle due to a 2-cycle result latency. These stalls can be avoided by rearranging the code as follows:

```
miaph acc0, r3, r4
add  r1, r2, r3
miaph acc0, r5, r6
sub  r8, r3, r4
mra  r6, r7, acc0
```


A.5.7 Scheduling MRS and MSR Instructions

The **MRS** instruction has an issue latency of 1 cycle and a result latency of 2 cycles. The **MSR** instruction has an issue latency of 2 cycles (6 if updating the mode bits) and a result latency of 1 cycle.

Consider the code sample:

```
mrs   r0, cpsr
orr   r0, r0, #1
add   r1, r2, r3
```

The **ORR** instruction above would incur a 1 cycle stall due to the 2-cycle result latency of the **MRS** instruction. In the code example above, the **ADD** instruction can be moved before the **ORR** instruction to prevent this stall.

A.5.8 Scheduling CP15 Coprocessor Instructions

The **MRC** instruction has an issue latency of 1 cycle and a result latency of 3 cycles. The **MCR** instruction has an issue latency of 1 cycle.

Consider the code sample:

```
add   r1, r2, r3
mrc   p15, 0, r7, C1, C0, 0
mov   r0, r7
add   r1, r1, #1
```

The **MOV** instruction above would incur a 2-cycle latency due to the 3-cycle result latency of the **mrc** instruction. The code shown above can be rearranged as follows to avoid these stalls:

```
mrc   p15, 0, r7, C1, C0, 0
add   r1, r2, r3
add   r1, r1, #1
mov   r0, r7
```

A.6 Optimizing C Libraries

Many of the standard C library routines can benefit greatly by being optimized for the core architecture. The following string and memory manipulation routines should be tuned to obtain the best performance from the core architecture (instruction selection, cache usage and data prefetch):

strcat, strchr, strcmp, strcoll, strcpy, strcspn, strlen, strcat, strncmp, strpbrk, strrchr, strspn, strstr, strtok, strxfrm, memchr, memcmp, memcpy, memmove, memset

A.7 Optimizations for Size

For applications such as cell phone software it is necessary to optimize the code for improved performance while minimizing code size. Optimizing for smaller code size will, in general, lower the performance of your application. This chapter contains techniques for optimizing for code size using the core instruction set.

A.7.1 Space/Performance Trade Off

Many optimizations mentioned in the previous chapters improve the performance of ARM code. However, using these instructions will result in increased code size. Use the following optimizations to reduce the space requirements of the application code.

A.7.1.1 Multiple Word Load and Store

The **LDM/STM** instructions are one word long and let you load or store multiple registers at once. Use the **LDM/STM** instructions instead of a sequence of loads/stores to consecutive addresses in memory whenever possible.

A.7.1.2 Use of Conditional Instructions

Using conditional instructions to expand if-then-else statements as described in [Section A.3.1](#), “Conditional Instructions” will result in increasing the size of the generated code. Therefore, do not use conditional instructions if application code space requirements are an issue.

A.7.1.3 Use of PLD Instructions

The preload instruction **PLD** is only a hint, it does not change the architectural state of the processor. Using or not using them will not change the behavior of your code, therefore, you should avoid using these instructions when optimizing for space.

Test Features

B

This chapter gives a brief overview of the Intel XScale[®] core JTAG features. The Intel XScale[®] core provides a baseline set of features from which the ASSP builds upon. A full description of these features can be found in the ASSP architecture specification.

B.1 Overview

The Intel XScale[®] core provides test features compatible with IEEE Standard Test Access Port and Boundary Scan Architecture (IEEE Std. 1149.1). These features include a TAP controller, a 5 or 7 bit instruction register, and test data registers to support software debug. The size of the instruction register depends on which variant of the Intel XScale[®] core is being used. This can be found out by examining the CoreGen field of Coprocessor 15, ID Register (bits 15:13). (See [Table 7-4, “ID Register”](#) on [page 7-81](#) for more details.) A CoreGen value of 0x1 means the JTAG instruction register size is 5 bits and a CoreGen value of 0x2 means the JTAG instruction register size is 7 bits.

The Intel XScale[®] core also provides support for an ASSP defined boundary-scan register, device ID register, and other data test register specific to ASSP implementation.

To avoid confusion and duplication, the description of these features are in the ASSP architecture specification.





This Page Intentionally Left Blank