Intel® 80200 Processor based on Intel® XScale™ Microarchitecture

Developer's Manual

November, 2000

Order Number: 273411-002

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

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

The Intel® 80200 Processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

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

Copyright© Intel Corporation, 2000

*Other brands and names are the property of their respective owners.

ARM and StrongARM are registered trademarks of ARM, Ltd.

Contents

1 Introduction 1-1

1.1 Intel® 80200 Processor based on Intel® XScale™ Microarchitecture High-Level Overview 1-1
1.1.1 ARM* Architecture Compliance 1-1
1.1.2 Features 1-2
1.1.2.1 Multiply/Accumulate (MAC) 1-2
1.1.2.2 Memory Management 1-3
1.1.2.3 Instruction Cache 1-3
1.1.2.4 Branch Target Buffer 1-3
1.1.2.5 Data Cache 1-3
1.1.2.6 Power Management 1-4
1.1.2.7 Interrupt Controller 1-4
1.1.2.8 Bus Controller 1-4
1.1.2.9 Performance Monitoring 1-4
1.1.2.10 Debug 1-4
1.1.2.11 JTAG 1-4
1.2 Terminology and Conventions 1-5
1.2.1 Number Representation 1-5
1.2.2 Terminology and Acronyms 1-5
1.3 Other Relevant Documents 1-6

2 Programming Model 2-1

2.1 ARM* Architecture Compliance 2-1
2.2 ARM* Architecture Implementation Options 2-1
2.2.1 Big Endian versus Little Endian 2-1
2.2.2 26-Bit Code 2-1
2.2.3 Thumb* 2-1
2.2.4 ARM* DSP-Enhanced Instruction Set 2-2
2.2.5 Base Register Update 2-2
2.3 Extensions to ARM* Architecture 2-3
2.3.1 DSP Coprocessor 0 (CP0) 2-3
2.3.1.1 Multiply With Internal Accumulate Format 2-4
2.3.1.2 Internal Accumulator Access Format 2-7
2.3.2 New Page Attributes 2-9
2.3.3 Additions to CP15 Functionality 2-11
2.3.4 Event Architecture 2-12
2.3.4.1 Exception Summary 2-12
2.3.4.2 Event Priority 2-12
2.3.4.3 Prefetch Aborts 2-13
2.3.4.4 Data Aborts 2-14
2.3.4.5 Events from Preload Instructions 2-16
2.3.4.6 Debug Events 2-16

3 Memory Management 3-1

3.1 Overview 3-1
3.2 Architecture Model 3-2
3.2.1 Version 4 vs. Version 5 3-2
3.2.2 Memory Attributes 3-2

3.2.2.1 Page (P) Attribute Bit 3-2
3.2.2.2 Cacheable (C), Bufferable (B), and eXtension (X) Bits 3-2
3.2.2.3 Instruction Cache 3-2
3.2.2.4 Data Cache and Write Buffer 3-3
3.2.2.5 Details on Data Cache and Write Buffer Behavior 3-4
3.2.2.6 Memory Operation Ordering 3-4
3.2.3 Exceptions 3-4
3.3 Interaction of the MMU, Instruction Cache, and Data Cache 3-5
3.4 Control 3-6
3.4.1 Invalidate (Flush) Operation 3-6
3.4.2 Enabling/Disabling 3-6
3.4.3 Locking Entries 3-7
3.4.4 Round-Robin Replacement Algorithm 3-9

4 Instruction Cache 4-1

4.1 Overview 4-1
4.2 Operation 4-2
4.2.1 Operation When Instruction Cache is Enabled 4-2
4.2.2 Operation When The Instruction Cache Is Disabled 4-2
4.2.3 Fetch Policy 4-3
4.2.4 Round-Robin Replacement Algorithm 4-3
4.2.5 Parity Protection 4-4
4.2.6 Instruction Fetch Latency 4-5
4.2.7 Instruction Cache Coherency 4-5
4.3 Instruction Cache Control 4-6
4.3.1 Instruction Cache State at RESET 4-6
4.3.2 Enabling/Disabling 4-6
4.3.3 Invalidating the Instruction Cache 4-7
4.3.4 Locking Instructions in the Instruction Cache 4-8
4.3.5 Unlocking Instructions in the Instruction Cache 4-9

5 Branch Target Buffer 5-1

5.1 Branch Target Buffer (BTB) Operation 5-1
5.1.1 Reset 5-2
5.1.2 Update Policy 5-2
5.2 BTB Control 5-3
5.2.1 Disabling/Enabling 5-3
5.2.2 Invalidation 5-3

6 Data Cache 6-1

6.1 Overviews 6-1
6.1.1 Data Cache Overview 6-1
6.1.2 Mini-Data Cache Overview 6-3
6.1.3 Write Buffer and Fill Buffer Overview 6-4
6.2 Data Cache and Mini-Data Cache Operation 6-5
6.2.1 Operation When Caching is Enabled 6-5
6.2.2 Operation When Data Caching is Disabled 6-5
6.2.3 Cache Policies 6-5
6.2.3.1 Cacheability 6-5
6.2.3.2 Read Miss Policy 6-6

6.2.3.3 Write Miss Policy 6-7
6.2.3.4 Write-Back Versus Write-Through 6-7
6.2.4 Round-Robin Replacement Algorithm 6-8
6.2.5 Parity Protection 6-8
6.2.6 Atomic Accesses 6-8
6.3 Data Cache and Mini-Data Cache Control 6-9
6.3.1 Data Memory State After Reset 6-9
6.3.2 Enabling/Disabling 6-9
6.3.3 Invalidate & Clean Operations 6-9
6.3.3.1 Global Clean and Invalidate Operation 6-10
6.4 Re-configuring the Data Cache as Data RAM 6-12
6.5 Write Buffer/Fill Buffer Operation and Control 6-16

7 Configuration 7-1

7.1 Overview 7-1
7.2 CP15 Registers 7-4
7.2.1 Register 0: ID and Cache Type Registers 7-5
7.2.2 Register 1: Control and Auxiliary Control Registers 7-6
7.2.3 Register 2: Translation Table Base Register 7-8
7.2.4 Register 3: Domain Access Control Register 7-8
7.2.5 Register 4: Reserved 7-8
7.2.6 Register 5: Fault Status Register 7-9
7.2.7 Register 6: Fault Address Register 7-9
7.2.8 Register 7: Cache Functions 7-10
7.2.9 Register 8: TLB Operations 7-12
7.2.10 Register 9: Cache Lock Down 7-13
7.2.11 Register 10: TLB Lock Down 7-14
7.2.12 Register 11-12: Reserved 7-14
7.2.13 Register 13: Process ID 7-15
7.2.13.1 The PID Register Affect On Addresses 7-15
7.2.14 Register 14: Breakpoint Registers 7-16
7.2.15 Register 15: Coprocessor Access Register 7-17
7.3 CP14 Registers 7-18
7.3.1 Registers 0-3: Performance Monitoring 7-18
7.3.2 Register 4-5: Reserved 7-18
7.3.3 Registers 6-7: Clock and Power Management 7-19
7.3.4 Registers 8-15: Software Debug 7-20

8 System Management 8-1

8.1 Clocking 8-1
8.2 Processor Reset 8-3
8.2.1 Reset Sequence 8-3
8.2.2 Reset Effect on Outputs 8-4
8.3 Power Management 8-5
8.3.1 Invocation 8-5
8.3.2 Signals Associated with Power Management 8-5

9 Interrupts 9-1

9.1 Introduction 9-1
9.2 External Interrupts 9-1

9.3 Programmer Model 9-2
9.3.1 INTCTL 9-3
9.3.2 INTSRC 9-4
9.3.3 INTSTR 9-5

10 External Bus 10-1

10.1 General Description 10-1
10.2 Signal Description 10-3
10.2.1 Request Bus 10-4
10.2.1.1 Intel® 80200 Processor Use of the Request Bus 10-4
10.2.2 Data Bus 10-6
10.2.3 Critical Word First 10-7
10.2.4 Configuration Pins 10-8
10.2.5 Multimaster Support 10-9
10.2.6 Abort 10-11
10.2.7 ECC 10-12
10.2.8 Big Endian System Configuration 10-13
10.3 Examples 10-14
10.3.1 Simple Read Word 10-14
10.3.2 Read Burst, No Critical Word First 10-15
10.3.3 Read Burst, Critical Word First Data Return 10-16
10.3.4 Word Write 10-17
10.3.5 Two Word Coalesced Write 10-18
10.3.5.1 Write Burst 10-19
10.3.6 Write Burst, Coalesced 10-20
10.3.7 Pipelined Accesses 10-21
10.3.8 Locked Access 10-22
10.3.9 Aborted Access 10-23
10.3.10 Hold 10-24

11 Bus Controller 11-1

11.1 Introduction 11-1
11.2 ECC 11-1
11.3 Error Handling 11-2
11.3.1 Bus Aborts 11-2
11.3.2 ECC Errors 11-3
11.4 Programmer Model 11-5
11.4.1 BCU Control Registers 11-5
11.4.2 ECC Error Registers 11-9

12 Performance Monitoring 12-1

12.1 Overview 12-1
12.2 Clock Counter (CCNT; CP14 - Register 1) 12-2
12.3 Performance Count Registers (PMN0 - PMN1; CP14 - Register 2 and 3, Respectively) 12-3
12.3.1 Extending Count Duration Beyond 32 Bits 12-3
12.4 Performance Monitor Control Register (PMNC) 12-4
12.4.1 Managing PMNC 12-5
12.5 Performance Monitoring Events 12-6
12.5.1 Instruction Cache Efficiency Mode 12-7
12.5.2 Data Cache Efficiency Mode 12-8

12.5.3 Instruction Fetch Latency Mode 12-8
12.5.4 Data/Bus Request Buffer Full Mode 12-9
12.5.5 Stall/Writeback Statistics 12-9
12.5.6 Instruction TLB Efficiency Mode 12-10
12.5.7 Data TLB Efficiency Mode 12-10
12.6 Multiple Performance Monitoring Run Statistics 12-11
12.7 Examples 12-12

13 Software Debug 13-1

13.1 Definitions 13-1
13.2 Debug Registers 13-1
13.3 Introduction 13-2
13.3.1 Halt Mode 13-2
13.3.2 Monitor Mode 13-2
13.4 Debug Control and Status Register (DCSR) 13-3
13.4.1 Global Enable Bit (GE) 13-4
13.4.2 Halt Mode Bit (H) 13-4
13.4.3 Vector Trap Bits (TF,TI,TD,TA,TS,TU,TR) 13-5
13.4.4 Sticky Abort Bit (SA) 13-5
13.4.5 Method of Entry Bits (MOE) 13-5
13.4.6 Trace Buffer Mode Bit (M) 13-5
13.4.7 Trace Buffer Enable Bit (E) 13-5
13.5 Debug Exceptions 13-6
13.5.1 Halt Mode 13-6
13.5.2 Monitor Mode 13-8
13.6 HW Breakpoint Resources 13-9
13.6.1 Instruction Breakpoints 13-9
13.6.2 Data Breakpoints 13-10
13.7 Software Breakpoints 13-11
13.8 Transmit/Receive Control Register (TXRXCTRL) 13-12
13.8.1 RX Register Ready Bit (RR) 13-13
13.8.2 Overflow Flag (OV) 13-14
13.8.3 Download Flag (D) 13-14
13.8.4 TX Register Ready Bit (TR) 13-15
13.8.5 Conditional Execution Using TXRXCTRL 13-15
13.9 Transmit Register (TX) 13-16
13.10 Receive Register (RX) 13-16
13.11 Debug JTAG Access 13-17
13.11.1 SELDCSR JTAG Command 13-17
13.11.2 SELDCSR JTAG Register 13-18
13.11.2.1 DBG.HLD_RST 13-19
13.11.2.2 DBG.BRK 13-20
13.11.2.3 DBG.DCSR 13-20
13.11.3 DBGTX JTAG Command 13-20
13.11.4 DBGTX JTAG Register 13-21
13.11.5 DBGRX JTAG Command 13-21
13.11.6 DBGRX JTAG Register 13-22
13.11.6.1 RX Write Logic 13-23
13.11.6.2 DBGRX Data Register 13-24
13.11.6.3 DBG.RR 13-24

13.11.6.4 DBG.V 13-25
13.11.6.5 DBG.RX 13-25
13.11.6.6 DBG.D 13-25
13.11.6.7 DBG.FLUSH 13-25
13.11.7 Debug JTAG Data Register Reset Values 13-25
13.12 Trace Buffer 13-26
13.12.1 Trace Buffer CP Registers 13-26
13.12.1.1 Checkpoint Registers 13-26
13.12.1.2 Trace Buffer Register (TBREG) 13-27
13.13 Trace Buffer Entries 13-28
13.13.1 Message Byte 13-28
13.13.1.1 Exception Message Byte 13-29
13.13.1.2 Non-exception Message Byte 13-30
13.13.1.3 Address Bytes 13-31
13.13.2 Trace Buffer Usage 13-32
13.14 Downloading Code in the ICache 13-34
13.14.1 LDIC JTAG Command 13-34
13.14.2 LDIC JTAG Data Register 13-35
13.14.3 LDIC Cache Functions 13-36
13.14.4 Loading IC During Reset 13-38
13.14.4.1 Loading IC During Cold Reset for Debug 13-39
13.14.4.2 Loading IC During a Warm Reset for Debug 13-41
13.14.5 Dynamically Loading IC After Reset 13-43
13.14.5.1 Dynamic Code Download Synchronization 13-45
13.14.6 Mini Instruction Cache Overview 13-46
13.15 Halt Mode Software Protocol 13-47
13.15.1 Starting a Debug Session 13-47
13.15.1.1 Setting up Override Vector Tables 13-47
13.15.1.2 Placing the Handler in Memory 13-48
13.15.2 Implementing a Debug Handler 13-49
13.15.2.1 Debug Handler Entry 13-49
13.15.2.2 Debug Handler Restrictions 13-49
13.15.2.3 Dynamic Debug Handler 13-50
13.15.2.4 High-Speed Download 13-52
13.15.3 Ending a Debug Session 13-53
13.16 Software Debug Notes/Errata 13-54

14 Performance Considerations 14-1

14.1 Interrupt Latency 14-1
14.2 Branch Prediction 14-2
14.3 Addressing Modes 14-2
14.4 Instruction Latencies 14-3
14.4.1 Performance Terms 14-3
14.4.2 Branch Instruction Timings 14-4
14.4.3 Data Processing Instruction Timings 14-5
14.4.4 Multiply Instruction Timings 14-6
14.4.5 Saturated Arithmetic Instructions 14-8
14.4.6 Status Register Access Instructions 14-8
14.4.7 Load/Store Instructions 14-8
14.4.8 Semaphore Instructions 14-9
14.4.9 Coprocessor Instructions 14-9

14.4.10 Miscellaneous Instruction Timing 14-9
14.4.11 Thumb* Instructions 14-9

A Compatibility: Intel® 80200 Processor vs. SA-110 ..... A-1

A.1 Introduction A-1
A.2 Summary A-1
A.3 Architecture Deviations A-3
A.3.1 Read Buffer A-3
A.3.2 26-bit Mode A-3
A.3.3 Cacheable (C) and Bufferable (B) Encoding A-3
A.3.4 Write Buffer Behavior A-4
A.3.5 External Aborts A-4
A.3.6 Performance Differences A-5
A.3.7 System Control Coprocessor A-5
A.3.8 New Instructions and Instruction Formats A-5
A.3.9 Augmented Page Table Descriptors A-5

B Optimization Guide B-1

B.1 Introduction B-1
B.1.1 About This Guide B-1
B.2 Intel® 80200 Processor Pipeline B-2
B.2.1 General Pipeline Characteristics B-2
B.2.1.1. Number of Pipeline Stages B-2
B.2.1.2. Intel® 80200 Processor Pipeline Organization B-3
B.2.1.3. Out Of Order Completion B-4
B.2.1.4. Register Scoreboarding B-4
B.2.1.5. Use of Bypassing B-4
B.2.2 Instruction Flow Through the Pipeline B-5
B.2.2.1. ARM* V5 Instruction Execution B-5
B.2.2.2. Pipeline Stalls B-5
B.2.3 Main Execution Pipeline B-6
B.2.3.1. F1 / F2 (Instruction Fetch) Pipestages B-6
B.2.3.2. ID (Instruction Decode) Pipestage B-6
B.2.3.3. RF (Register File / Shifter) Pipestage B-7
B.2.3.4. X1 (Execute) Pipestages B-7
B.2.3.5. X2 (Execute 2) Pipestage B-7
B.2.3.6. WB (write-back) B-7
B.2.4 Memory Pipeline B-8
B.2.4.1. D1 and D2 Pipestage B-8
B.2.5 Multiply/Multiply Accumulate (MAC) Pipeline B-8
B.2.5.1. Behavioral Description B-8
B.3 Basic Optimizations B-9
B.3.1 Conditional Instructions B-9
B.3.1.1. Optimizing Condition Checks B-9
B.3.1.2. Optimizing Branches B-10
B.3.1.3. Optimizing Complex Expressions B-12
B.3.2 Bit Field Manipulation B-13
B.3.3 Optimizing the Use of Immediate Values B-14
B.3.4 Optimizing Integer Multiply and Divide B-15
B.3.5 Effective Use of Addressing Modes B-16
B.4 Cache and Prefetch Optimizations B-17

B.4.1 Instruction Cache B-17
B.4.1.1. Cache Miss Cost B-17
B.4.1.2. Round Robin Replacement Cache Policy B-17
B.4.1.3. Code Placement to Reduce Cache Misses B-17
B.4.1.4. Locking Code into the Instruction Cache B-18
B.4.2 Data and Mini Cache B-19
B.4.2.1. Non Cacheable Regions B-19
B.4.2.2. Write-through and Write-back Cached Memory Regions B-19
B.4.2.3. Read Allocate and Read-write Allocate Memory Regions B-20
B.4.2.4. Creating On-chip RAM B-20
B.4.2.5. Mini-data Cache B-21
B.4.2.6. Data Alignment B-22
B.4.2.7. Literal Pools B-23
B.4.3 Cache Considerations B-24
B.4.3.1. Cache Conflicts, Pollution and Pressure B-24
B.4.3.2. Memory Page Thrashing B-24
B.4.4 Prefetch Considerations B-25
B.4.4.1. Prefetch Distances in the Intel® 80200 Processor B-25
B.4.4.2. Prefetch Loop Scheduling B-27
B.4.4.3. Prefetch Loop Limitations B-27
B.4.4.4. Compute vs. Data Bus Bound B-27
B.4.4.5. Low Number of Iterations B-27
B.4.4.6. Bandwidth Limitations B-28
B.4.4.7. Cache Memory Considerations B-29
B.4.4.8. Cache Blocking B-31
B.4.4.9. Prefetch Unrolling B-31
B.4.4.10. Pointer Prefetch B-32
B.4.4.11. Loop Interchange B-33
B.4.4.12. Loop Fusion B-33
B.4.4.13. Prefetch to Reduce Register Pressure B-34
B.5 Instruction Scheduling B-35
B.5.1 Scheduling Loads B-35
B.5.1.1. Scheduling Load and Store Double (LDRD/STRD) B-37
B.5.1.2. Scheduling Load and Store Multiple (LDM/STM) B-38
B.5.2 Scheduling Data Processing Instructions B-39
B.5.3 Scheduling Multiply Instructions B-40
B.5.4 Scheduling SWP and SWPB Instructions B-41
B.5.5 Scheduling the MRA and MAR Instructions (MRRC/MCRR) B-42
B.5.6 Scheduling the MIA and MIAPH Instructions B-43
B.5.7 Scheduling MRS and MSR Instructions B-44
B.5.8 Scheduling CP15 Coprocessor Instructions B-44
B.6 Optimizing C Libraries B-45
B.7 Optimizations for Size B-45
B.7.1 Space/Performance Trade Off B-45
B.7.1.1. Multiple Word Load and Store B-45
B.7.1.2. Use of Conditional Instructions B-45
B.7.1.3. Use of PLD Instructions B-45

C Test Features C-1

C.1 Introduction C-1
C.2 JTAG - IEEE1149.1 ..... C-1
C.2.1 Boundary Scan Architecture C-2

C.2.2 TAP Pins C-3
C.2.3 Instruction Register (IR) C-4
C.2.3.1. Boundary-Scan Instruction Set C-4
C.2.4 TAP Test Data Registers C-6
C.2.4.1. Device Identification Register C-6
C.2.4.2. Bypass Register C-6
C.2.4.3. Boundary-Scan Register C-6
C.2.5 TAP Controller C-7
C.2.5.1. Test Logic Reset State C-8
C.2.5.2. Run-Test/Idle State C-8
C.2.5.3. Select-DR-Scan State C-8
C.2.5.4. Capture-DR State C-8
C.2.5.5. Shift-DR State C-9
C.2.5.6. Exit1-DR State C-9
C.2.5.7. Pause-DR State C-9
C.2.5.8. Exit2-DR State C-9
C.2.5.9. Update-DR State C-10
C.2.5.10. Select-IR Scan State C-10
C.2.5.11. Capture-IR State C-10
C.2.5.12. Shift-IR State C-10
C.2.5.13. Exit1-IR State C-11
C.2.5.14. Pause-IR State C-11
C.2.5.15. Exit2-IR State C-11
C.2.5.16. Update-IR State C-11
C.2.5.17. Boundary-Scan Example C-12

Figures

1-1 Intel® 80200 Processor based on Intel® XScale™ Microarchitecture Features 1-2
3-1 Example of Locked Entries in TLB 3-9
4-1 Instruction Cache Organization 4-1
4-2 Locked Line Effect on Round Robin Replacement 4-8
5-1 BTB Entry 5-1
5-2 Branch History 5-2
6-1 Data Cache Organization 6-2
6-2 Mini-Data Cache Organization 6-3
6-3 Locked Line Effect on Round Robin Replacement 6-15
8-1 Reset Sequence 8-3
8-2 Pin State at Reset 8-4
8-3 Waiting for the Wake-up Interrupt 8-6
9-1 Interrupt Controller Block Diagram 9-2
10-1 Typical System 10-1
10-2 Alternate Configuration 10-2
10-3 Big Endian Lane Swapping on a 64-bit Bus 10-13
10-4 Basic Read Timing 10-14
10-5 Read Burst, No CWF 10-15
10-6 Read Burst, CWF 10-16
10-7 Basic Word Write 10-17
10-8 Two Word Coalesced Write 10-18
10-9 Four Word Eviction Write 10-19
10-10 Four Word Coalesced Write Burst 10-20
10-11 Pipeline Example 10-21
10-12 Locked Access 10-22
10-13 Aborted Access 10-23
10-14 Hold Assertion 10-24
13-1 SELDCSR Hardware 13-18
13-2 SELDCSR Data Register 13-19
13-3 DBGTX Hardware 13-21
13-4 DBGRX Hardware 13-22
13-5 RX Write Logic 13-23
13-6 DBGRX Data Register 13-24
13-7 Message Byte Formats 13-28
13-8 Indirect Branch Entry Address Byte Organization 13-31
13-9 High Level View of Trace Buffer 13-32
13-10 LDIC JTAG Data Register Hardware 13-35
13-11 Format of LDIC Cache Functions 13-37
13-12 Code Download During a Cold Reset For Debug 13-39
13-13 Code Download During a Warm Reset For Debug 13-41
13-14 Downloading Code in IC During Program Execution 13-43
B-1 Intel® 80200 Processor RISC Superpipeline B-3
C-1 Test Access Port Block Diagram C-2
C-2 TAP Controller State Diagram C-7
C-3 JTAG Example C-13
C-4 Timing Diagram Illustrating the Loading of Instruction Register C-14
C-5 Timing Diagram Illustrating the Loading of Data Register C-15

Tables

2-1 Multiply with Internal Accumulate Format 2-4
2-2 MIA{<cond>} acc0, Rm, Rs 2-4
2-3 MIAPH{<cond>} acc0, Rm, Rs 2-5
2-4 MIAxy{<cond>} acc0, Rm, Rs 2-6
2-5 Internal Accumulator Access Format 2-7
2-6 MAR{<cond>} acc0, RdLo, RdHi 2-8
2-7 MRA{<cond>} RdLo, RdHi, acc0 2-8
2-8 First-level Descriptors 2-10
2-9 Second-level Descriptors for Coarse Page Table 2-10
2-10 Second-level Descriptors for Fine Page Table 2-10
2-11 Exception Summary 2-12
2-12 Event Priority 2-12
2-13 Intel® 80200 Processor Encoding of Fault Status for Prefetch Aborts 2-13
2-14 Intel® 80200 Processor Encoding of Fault Status for Data Aborts 2-14
3-1 Data Cache and Buffer Behavior when X = 0 3-3
3-2 Data Cache and Buffer Behavior when X = 1 3-3
3-3 Memory Operations that Impose a Fence 3-4
3-4 Valid MMU & Data/mini-data Cache Combinations 3-5
7-1 MRC/MCR Format 7-2
7-2 LDC/STC Format 7-3
7-3 CP15 Registers 7-4
7-4 ID Register 7-5
7-5 Cache Type Register 7-5
7-6 ARM* Control Register 7-6
7-7 Auxiliary Control Register 7-7
7-8 Translation Table Base Register 7-8
7-9 Domain Access Control Register 7-8
7-10 Fault Status Register 7-9
7-11 Fault Address Register 7-9
7-12 Cache Functions 7-10
7-13 TLB Functions 7-12
7-14 Cache Lockdown Functions 7-13
7-15 Data Cache Lock Register 7-13
7-16 TLB Lockdown Functions 7-14
7-17 Accessing Process ID 7-15
7-18 Process ID Register 7-15
7-19 Accessing the Debug Registers 7-16
7-20 Coprocessor Access Register 7-17
7-21 CP14 Registers 7-18
7-22 Accessing the Performance Monitoring Registers 7-18
7-23 PWRMODE Register 7-19
7-24 Clock and Power Management 7-19
7-25 CCLKCFG Register 7-19
7-26 Accessing the Debug Registers 7-20
8-1 Reset CCLK Configuration 8-1
8-2 Software CCLK Configuration 8-2
8-3 Low Power Modes 8-5
8-4 PWRSTATUS[1:0] Encoding 8-5

9-1 Interrupt Control Register (CP13 register 0) 9-3
9-2 Interrupt Source Register (CP13, register 4) 9-4
9-3 Interrupt Steer Register (CP13, register 8) 9-5
10-1 Intel® 80200 Processor based on Intel® XScale™ Microarchitecture Bus Signals 10-3
10-2 Requests on a 64-bit Bus 10-4
10-3 Requests on a 32-bit Bus 10-5
10-4 Return Order for 8-Word Burst, 64-bit Data Bus 10-7
10-5 Return Order for 8-Word Burst, 32-bit Data Bus 10-7
11-1 BCU Response to ECC Errors 11-3
11-2 BCUCTL (Register 0) 11-5
11-3 BCUMOD (Register 1) 11-7
11-4 ELOG0, ELOG1(Registers 4, 5) 11-9
11-5 ECAR0, ECAR1(Registers 6, 7) 11-9
11-6 ECTST (Register 8) 11-10
12-1 Clock Count Register (CCNT) 12-2
12-2 Performance Monitor Count Register (PMN0 and PMN1) 12-3
12-3 Performance Monitor Control Register (CP14, register 0) 12-4
12-4 Performance Monitoring Events 12-6
12-5 Some Common Uses of the PMU 12-7
13-1 Debug Control and Status Register (DCSR) 13-3
13-2 Event Priority 13-6
13-3 Instruction Breakpoint Address and Control Register (IBCRx) 13-9
13-4 Data Breakpoint Register (DBRx) 13-10
13-5 Data Breakpoint Controls Register (DBCON) 13-10
13-6 TX RX Control Register (TXRXCTRL) 13-12
13-7 Normal RX Handshaking 13-13
13-8 High-Speed Download Handshaking States 13-13
13-9 TX Handshaking 13-15
13-10 TXRXCTRL Mnemonic Extensions 13-15
13-11 TX Register 13-16
13-12 RX Register 13-16
13-13 DEBUG Data Register Reset Values 13-25
13-14 CP 14 Trace Buffer Register Summary 13-26
13-15 Checkpoint Register (CHKPTx) 13-26
13-16 TBREG Format 13-27
13-17 Message Byte Formats 13-28
13-18 LDIC Cache Functions 13-36
14-1 Minimum Interrupt Latency 14-1
14-2 Branch Latency Penalty 14-2
14-3 Latency Example 14-4
14-4 Branch Instruction Timings (Those predicted by the BTB) 14-4
14-5 Branch Instruction Timings (Those not predicted by the BTB) 14-5
14-6 Data Processing Instruction Timings 14-5
14-7 Multiply Instruction Timings 14-6
14-8 Multiply Implicit Accumulate Instruction Timings 14-7
14-9 Implicit Accumulator Access Instruction Timings 14-7
14-10 Saturated Data Processing Instruction Timings 14-8
14-11 Status Register Access Instruction Timings 14-8
14-12 Load and Store Instruction Timings 14-8
14-13 Load and Store Multiple Instruction Timings 14-8

14-14 Semaphore Instruction Timings 14-9
14-15 CP15 Register Access Instruction Timings 14-9
14-16 CP14 Register Access Instruction Timings 14-9
14-17 SWI Instruction Timings 14-9
14-18 Count Leading Zeros Instruction Timings 14-9
A-1 C and B encoding A-3
B-1 Pipelines and Pipe stages B-3
C-1 TAP Controller Pin Definitions C-3
C-2 JTAG Instruction Set C-4
C-3 IEEE Instructions C-5
C-4 JTAG ID Register Value C-6

Introduction 1

1.1 Intel® 80200 Processor based on Intel® XScale™ Microarchitecture High-Level Overview

The Intel® 80200 processor based on Intel® XScale™ microarchitecture, is the next generation in the Intel® StrongARM* processor family (compliant with ARM* Architecture V5TE). It is designed for high performance and low-power; leading the industry in mW/MIPs. The Intel® 80200 processor integrates a bus controller and an interrupt controller around a core processor, with intended embedded markets such as: handheld devices, networking, remote access servers, etc. This technology is ideal for internet infrastructure products such as network and I/O processors, where ultimate performance is critical for moving and processing large amounts of data quickly.

The Intel® 80200 processor 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 Intel® 80200 processor 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.
- cache locking.
- and a pipelined external bus.

All these features improve the efficiency of the external bus.

The Intel® 80200 processor 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.

1.1.1 ARM* Architecture Compliance

ARM* Version 5 (V5) Architecture added floating point instructions to ARM* Version 4. The Intel® 80200 processor implements the integer instruction set architecture of ARM V5, but does not provide hardware support of the floating point instructions.

The Intel® 80200 processor provides the Thumb* instruction set (ARM* V5T) and the ARM* V5E DSP extensions.

Backward compatibility with the first generation of Intel® StrongARM* products is maintained for user-mode applications. Operating systems may require modifications to match the specific hardware features of the Intel® 80200 processor and to take advantage of the performance enhancements added to the Intel® 80200 processor.

1.1.2 Features

Figure 1-1 shows the major functional blocks of the Intel® 80200 processor. The following sections give a brief, high-level overview of these blocks.

Figure 1-1. Intel® 80200 Processor based on Intel® XScale™ Microarchitecture Features

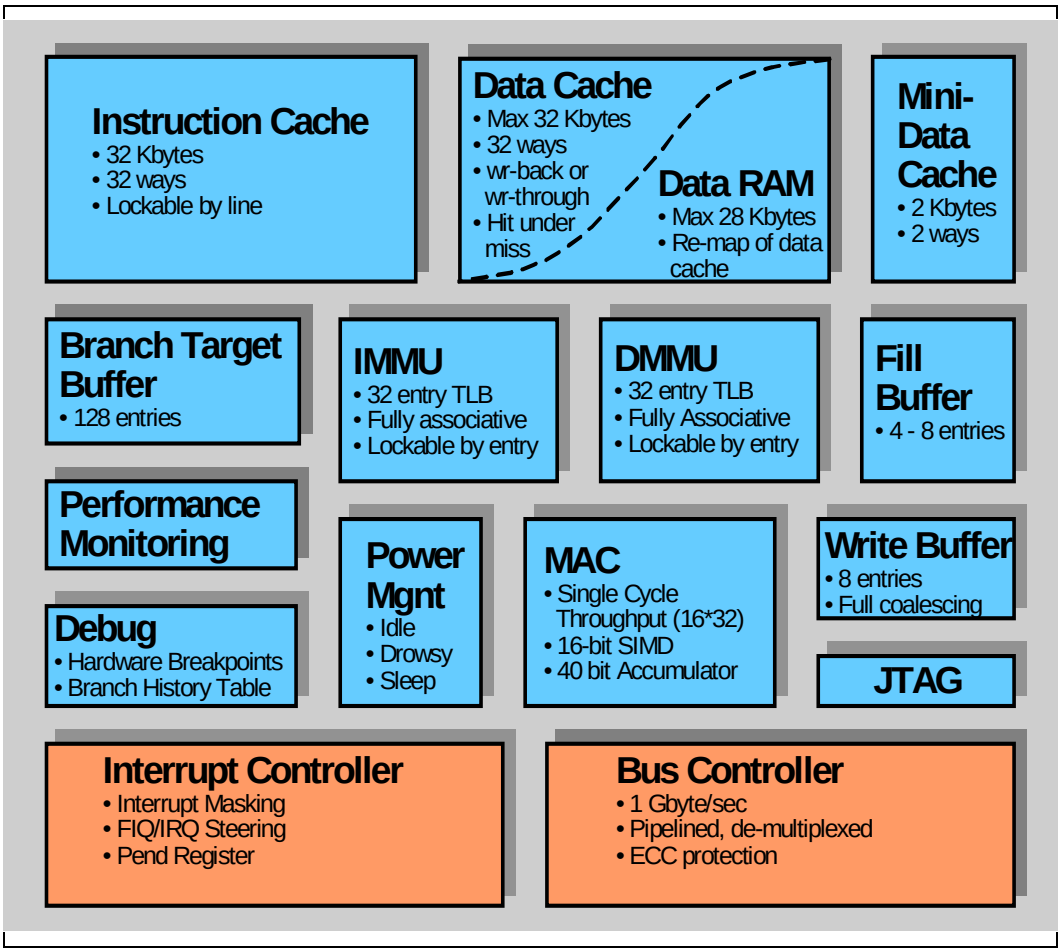

1.1.2.1 Multiply/Accumulate (MAC)

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

See Section 2.3, “Extensions to ARM* Architecture” on page 2-3 for more details.

1.1.2.2 Memory Management

The Intel® 80200 processor 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.

1.1.2.3 Instruction Cache

The Intel® 80200 processor implements a 32-Kbyte, 32-way set associative instruction cache with a line size of 32 bytes. All requests that "miss" the instruction cache generate a 32-byte read request to external memory. A mechanism to lock critical code within the cache is also provided.

Chapter 4, "Instruction Cache"discusses this in more detail.

1.1.2.4 Branch Target Buffer

The Intel® 80200 processor 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.

1.1.2.5 Data Cache

The Intel® 80200 processor implements a 32-Kbyte, a 32-way set associative data cache and a 2-Kbyte, 2-way set associative mini-data cache. Each cache has a line size of 32 bytes, supports write-through or write-back caching.

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

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

The Intel® 80200 processor 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 6-12 for more information on this.

1.1.2.6 Power Management

The Intel® 80200 processor supports three low power modes: idle, drowsy, and sleep. These modes are discussed in Section 8.3, "Power Management" on page 8-5.

1.1.2.7 Interrupt Controller

An interrupt controller is implemented on the Intel® 80200 processor that provides masking of interrupts and the ability to steer interrupts to FIQ or IRQ. It is accessed through Coprocessor 13 registers. See Chapter 9, "Interrupts"for more detail.

1.1.2.8 Bus Controller

The Intel® 80200 processor supports a pipelined external bus that runs at 100 MHz. The data bus is 32/64 bits with ECC protection. The bus controller can be configured to provide critical word first on load operations, enhancing overall system performance. The bus controller has four request queues, where all four requests can be active on the pipelined external bus.

Chapter 10, "External Bus" describes the external bus protocol and Chapter 11, "Bus Controller" covers the aspects of ECC protection. The bus controller registers are accessed via coprocessor 13.

1.1.2.9 Performance Monitoring

Two performance monitoring counters have been added to the Intel® 80200 processor that can be configured to monitor various events in the Intel® 80200 processor. These events allow a software developer to measure cache efficiency, detect system bottlenecks and reduce the overall latency of programs.

Chapter 12, "Performance Monitoring"discusses this in more detail.

1.1.2.10 Debug

The Intel® 80200 processor 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 13, "Software Debug"discusses this in more detail.

1.1.2.11 JTAG

Testability is supported on the Intel® 80200 processor through the Test Access Port (TAP) Controller implementation, which is based on IEEE 1149.1 (JTAG) Standard Test Access Port and Boundary-Scan Architecture. The purpose of the TAP controller is to support test logic internal and external to the Intel® 80200 processor such as built-in self-test, boundary-scan, and scan.

Appendix C.2 discusses this in more detail.

1.2 Terminology and Conventions

1.2.1 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.

1.2.2 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	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 gets written back to external memory. After this operation, the line is still valid and both dirty bits are deasserted.
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.
Reserved	*A reserved* 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.

1.3 Other Relevant Documents

- *Intel® 80200 Processor based on Intel® XScale™ Microarchitecture Datasheet*, Intel Order # *273414*
- *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.
- *ARM Architecture Reference Manual* Document Number: ARM DDI 0100B
  This document describes Version 4 of the ARM Architecture.
- *Intel® XScale™ Microarchitecture Programming Reference Manual*, Intel Order # 273436
- *Intel® 80312 I/O Companion Chip Developer's Manual*, Intel Order # 273410
- *StrongARM SA-1100 Microprocessor Developer's Manual*, Intel Order # 278088
- *StrongARM SA-110 Microprocessor Technical Reference Manual*, Intel Order #278058

Programming Model 2

This chapter describes the programming model of the Intel® 80200 processor based on Intel® XScale™ microarchitecture, namely the implementation options and extensions to the ARM* Version 5 architecture.

The ARM* Architecture Version 5TE Specification (ARM DDI 0100E) describes Version 5TE of the ARM Architecture, including the Thumb* ISA and ARM DSP-Enhanced ISA.

2.1 ARM* Architecture Compliance

The Intel® 80200 processor implements the integer instruction set architecture specified in ARM* Version 5TE. T refers to the Thumb instruction set and E refers to the DSP-Enhanced instruction set.

ARM* Version 5 introduces a few more architecture features over Version 4, 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® 80200 processor 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 occur. Note that this bit takes effect even if the MMU is disabled.

2.2.2 26-Bit Code

The Intel® 80200 processor does not support 26-bit code.

2.2.3 Thumb*

The Intel® 80200 processor supports the Thumb instruction set.

2.2.4 ARM* DSP-Enhanced Instruction Set

The Intel® 80200 processor implements 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<=16x16+32
- SMLAWy 32<=32x16+32
- SMLALxy 64<=16x16+64
- SMULxy 32<=16x16
- SMULWy 32<=32x16
- 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® 80200 processor 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 never generates 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 generation an alignment exception when the address bits [2:0] = 0b100.

MCRR and MRRC are only supported on the Intel® 80200 processor when directed to coprocessor 0 and are used to access the internal accumulator. See Section 2.3.1.2 for more information. Access to any other coprocessor besides 0x0 are undefined.

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 is not updated. This holds for all load and store instructions. This behavior matches that of the first generation Intel® StrongARM* processor and is referred to in the ARM V5 architecture as the *Base Restored Abort Model*.

2.3 Extensions to ARM* Architecture

The Intel® 80200 processor made a few extensions to the ARM Version 5 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 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 attribute specifying ECC for 1Meg regions was also added.
- Additional functionality has been added to coprocessor 15. Coprocessor 14 was also created.
- Enhancements were made to the Event Architecture, instruction cache and data cache parity error exceptions, breakpoint events, and imprecise external data aborts.

2.3.1 DSP Coprocessor 0 (CP0)

The Intel® 80200 processor 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 new instructions.

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 causes an undefined exception. (See Section 7.2.15, “Register 15: Coprocessor Access Register” on page 7-17 for more details). Note that only privileged software can set this bit in the Coprocessor Access Register.

The 40-bit accumulator needs 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 2-4 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	Intel® 80200 processor 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	Intel® 80200 processor only implements acc0; access to any other acc has 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® 80200 processor defines a single 40-bit accumulator referred to as *acc0;* future implementations may define multiple internal accumulators.The Intel® 80200 processor 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	0	0	0	0	Rs				0	0	0	0	0	0	0	1	Rm			

```
Operation:  if ConditionPassed(<cond>) then
                acc0 = (Rm[31:0] * Rs[31:0])[39:0] + acc0[39:0]
Exceptions: none
Qualifiers  Condition Code
            No condition code flags are updated
Notes:      Early termination is supported. Instruction timings can be found
            in Section 14.4.4, "Multiply Instruction Timings" on page 14-6.
            Specifying R15 for register Rs or Rm has unpredictable results.
            acc0 is defined to be 0b000 on 80200.
```

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 are 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	1 0 0 0	Rs	0 0 0 0	0 0 0	1	Rm

```
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]
Exceptions: none
Qualifiers  Condition Code
            S bit is always cleared; no condition code flags are updated
Notes:      Instruction timings can be found
            in Section 14.4.4, "Multiply Instruction Timings" on page 14-6.
            Specifying R15 for register Rs or Rm has unpredictable results.
            acc0 is defined to be 0b000 on 80200
```

The **MIAPH** instruction performs two16-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

```
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]

Exceptions: none
Qualifiers  Condition Code
            S bit is always cleared; no condition code flags are updated
Notes:      Instruction timings can be found
            in Section 14.4.4, "Multiply Instruction Timings" on page 14-6.
            Specifying R15 for register Rs or Rm has unpredictable results.
            acc0 is defined to be 0b000 on 80200.
```

The **MIAxy** instruction performs one16-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 selects bits [31:16] of the register which is specified in the mnemonic as T (for top). A value of 0x0 selects 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 are 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® 80200 processor defines a new instruction format for accessing internal accumulators in CP0. Table 2-5, "Internal Accumulator Access Format" on page 2-7 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 Intel® StrongARM* registers and an internal accumulator. The *acc* field specifies 1 of 8 internal accumulators to transfer data to/from. The Intel® 80200 processor 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 7-17 for more details).

The Intel® 80200 processor implements two instructions **MAR** and **MRA** that move two Intel® StrongARM* registers to acc0 and move acc0 to two Intel® StrongARM* 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	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 is sign extended. On a write to the acc, the lower 8 bits of this register is 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	Intel® 80200 processor only implements acc0; access to any other acc 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** produces 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

```
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 14.4.4, "Multiply Instruction Timings" on page 14-6
            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

```
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 14.4.4, "Multiply Instruction Timings" on page 14-6
            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® 80200 processor 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 3-2.

The Intel® 80200 processor retains ARM definitions of the C and B encoding when X = 0, which is different than the first generation Intel® 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) generates 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® 80200 processor also added a P bit in the first level descriptors to identify which pages of memory are protected with ECC.

A descriptor with the P bit set indicates the corresponding page in memory is ECC protected. If the BCUs ECC mode is enabled (see Chapter 11, “Bus Controller”) then writes to such a page are accompanied with an ECC and reads are validated by an ECC.

Bit 1 in the Control Register (coprocessor 15, register 1, opcode=1) enables ECC protection 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 2-10, Table 2-9, “Second-level Descriptors for Coarse Page Table” on page 2-10 and Table 2-10, “Second-level Descriptors for Fine Page Table” on page 2-10. Two second-level descriptor formats have been defined for Intel® 80200 processor, 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 P bit controls ECC.

The TEX (Type Extension) field is present in several of the descriptor types. In the Intel® 80200 processor, only the LSB of this field is used; this is called the X bit.

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® 80200 processor, 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 each Intel® StrongARM* processor. For the Intel® 80200 processor, the sequence -- called CPWAIT -- is shown in Example 2-1 on page 2-11.

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 emitting CPWAIT only after the sequence of MCR instructions.

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 takes 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 Intel® 80200 processor 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 13, "Software Debug" for more details

2.3.4.2 Event Priority

The Intel® 80200 processor 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 13, "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® 80200 processor 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. Intel® 80200 Processor 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® 80200 processor: 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 has advanced beyond the instruction that caused the data abort.

On the Intel® 80200 processor 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, "Intel® 80200 Processor Encoding of Fault Status for Data Aborts" on page 2-14.
  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. Intel® 80200 Processor 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 Second level	0b01100 0b01110	invalid valid	valid valid
	Translation	Section Page	0b00101 0b00111	invalid valid	valid valid
	Domain	Section Page	0b01001 0b01011	valid valid	valid valid
	Permission	Section Page	0b01101 0b01111	valid valid	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.

The Intel® 80200 processor generates external data aborts on multi-bit ECC errors and when the **Abort** pin is asserted on memory transactions. (See Chapter 11, "Bus Controller" for more details.) An external data abort can occur on non-cacheable loads, reads into the cache, cache evictions, or stores to external memory.

Although the Intel® 80200 processor 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 that are difficult for an abort handler to recover. Both external data aborts and data cache parity errors may result in corrupted data in the targeted registers. Because these faults are imprecise, it is possible that the corrupted data has been used before the Data Abort fault handler is invoked. Because of this, software should treat imprecise data aborts as unrecoverable.

Note that 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 is 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 happens during the NOPs; the data abort handler sees 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 2-15.

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.
```

Of course, 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 is 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 is 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 never causes 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 **PLD**s speculatively. For example, Example 2-3 on page 2-16 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 does not cause a fault. Rather, it is ignored and the loop terminates 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 13.5, “Debug Exceptions” on page 13-6.

Memory Management 3

This chapter describes the memory management unit implemented in the Intel® 80200 processor based on Intel® XScale™ microarchitecture, and is compliant with the ARM* Architecture V5TE.

3.1 Overview

The Intel® 80200 processor implements the Memory Management Unit (MMU) Architecture specified in the *ARM Architecture Reference Manual*. To accelerate virtual to physical address translation, the Intel® 80200 processor 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® 80200 processor allows system software to associate various attributes with regions of memory:

- cacheable
- bufferable
- line allocate policy
- write policy
- I/O
- mini Data Cache
- Coalescing
- ECC-Protected

See Section 3.2.2, “Memory Attributes” on page 3-2 for a description of page attributes and Section 2.3.2, “New Page Attributes” on page 2-9 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 7-15 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 2-9.

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 Intel® 80200 processor attributes take advantage of encoding space in the descriptors that was formerly reserved.

3.2.2.1 Page (P) Attribute Bit

The P bit specifies that the associated memory should be protected with ECC. The P bit is only present in the first level descriptors. Thus, ECC memory is specified with a 1 megabyte granularity.

If the MMU is disabled, ECC is disabled for all memory accesses. If the MMU is enabled, ECC is enabled for a region of memory if:

- its P bit in the first level descriptor for that virtual memory is set and
- the BCU has ECC enabled (see Chapter 11, "Bus Controller")

Accesses to memory for page walks do not use the MMU. For these accesses, ECC is enabled if:

- the CP15 Auxiliary Control Register enables it (see Section 7.2.2, "Register 1: Control and Auxiliary Control Registers" on page 7-6) and
- the BCU has ECC enabled (see Chapter 11, "Bus Controller")

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 does not fill a cache entry. If the C bit is set, then fetches from the associated memory region are 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 continues executing after a data access if no dependency on that access is encountered. With this setting, the processor stalls 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 are 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 do 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 are 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 a 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 do not cause a line to be allocated.

If read/write-allocate is in effect, load or store operations that miss the cache requests 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 6-5.

3.2.2.6 Memory Operation Ordering

A *fence* memory operation (memop) is one that guarantees all memops issued prior to the fence executes 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 3-4 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 2-12.

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 causes 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 7-12 for a listing of commands supported by the Intel® 80200 processor.

Globally invalidating a TLB does not affect locked TLB entries. However, the invalidate-entry operations can invalidate individual locked entries. In this case, the locked entry remains in the TLB, but never "hits" 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 3-6.

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 2-11
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 7-13 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 ensures proper operation. Software can also accomplish this by invalidating all entries, as shown in Example 3-2 on page 3-7.

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 3-7.

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 2-14), 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 is 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 3-8.

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 wraps 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 31can never be locked. If the lock pointer is at entry 31, a lock operation updates the TLB entry with the translation and ignore the lock. In this case, the round-robin pointer stays at entry 31.

Figure 3-1. Example of Locked Entries in TLB

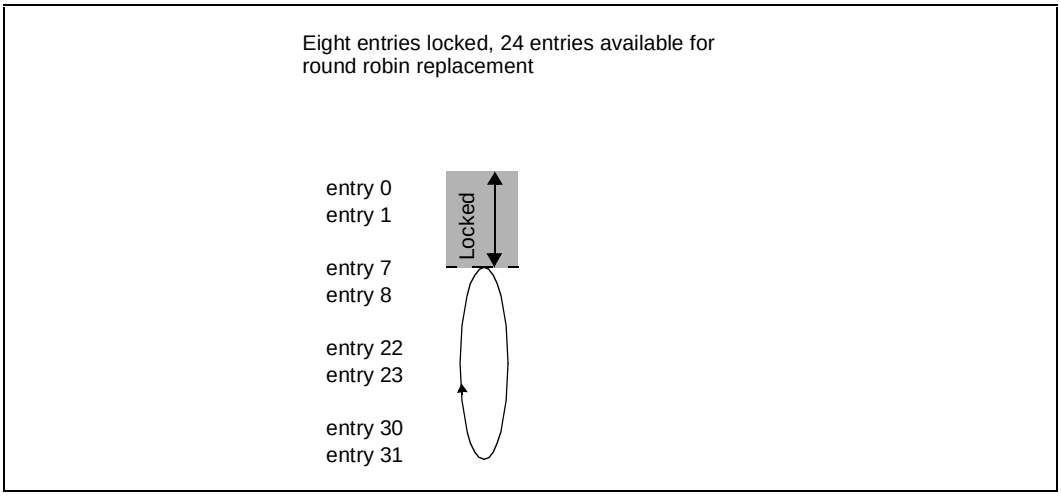

Instruction Cache 4

The Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) 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 a ***32-Kbyte, 32-way set associative cache***; this means there are 32 sets with each set containing 32 ways. Each way of a set contains eight 32-bit words and one valid bit, which is referred to as a line. The replacement policy is a round-robin algorithm and the cache also supports the ability to lock code in at a line granularity.

Figure 4-1. Instruction Cache Organization

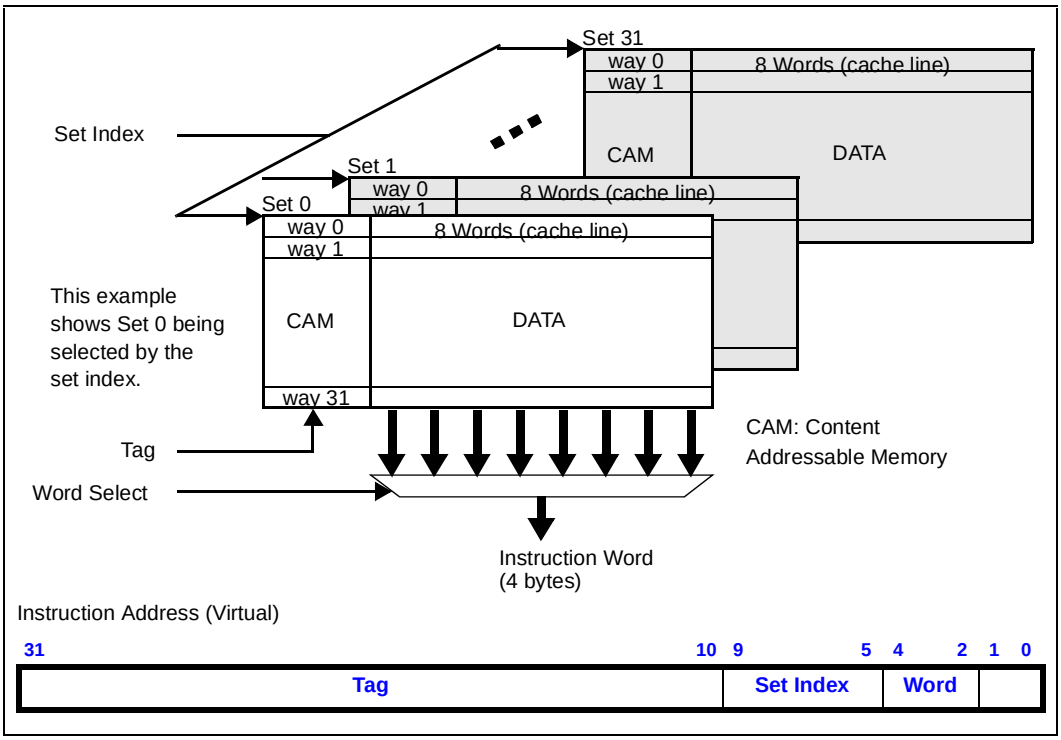

The instruction cache is virtually addressed and virtually tagged.

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 7-15 for a description of the PID register.

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 is 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 is 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 are always 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 is written into the instruction cache if cacheable and if the instruction cache is enabled. The line chosen for update in the cache is controlled by the round-robin replacement algorithm. This update may evict a valid line at that location.
6. Once the cache is updated, the eight valid bits of the fetch buffer are invalidated.

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 4-8 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 Intel® 80200 processor attempts to execute the instruction. Before servicing the exception, hardware place a notification of the error in the Fault Status Register (Coprocessor 15, register 5).

A software exception handler can recover from an instruction cache parity error. This can be accomplished by invalidating the instruction cache and the branch target buffer and then returning to the instruction that caused the prefetch abort exception. A simplified code example is shown in Example 4-1 on page 4-4. 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

Because the Intel® 80200 processor core is clocked at a multiple of the external bus clock, and the two clocks are truly asynchronous, an exact fetch latency is difficult to derive. In general, if a fetch can be directly issued (no other memory accesses are intervening), then the delay to the first instruction is approximately (8 + W) bus clocks, where W is number of memory wait states.

As an example: in a system with 2-wait-state memory (W = 2), an unoccluded fetch would require about 10 bus clocks to get the first instruction. If this system were running with a core/bus clock ratio of 6, then the core would perceive this as a latency of about 60 cycles.

These numbers are best case and assume that no other active memory transactions exist. Refer to Chapter 10, "External Bus" for more information on External Bus signal definitions and request timings.

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 4-7 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 7-10 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 7-13.

There is an inherent delay from the execution of the instruction cache invalidate command to where the next instruction sees 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® 80200 processor also supports invalidating an individual line from the instruction cache. See Table 7-12, “Cache Functions” on page 7-10 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 ignores the lock command if software is trying to lock all the lines in a particular set (i.e., ways 28-31can never be locked). When this happens, the line is still allocated into the cache, but the lock is ignored. The round-robin pointer stays 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 7-13 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:

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

Failure to follow these requirements produces 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 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

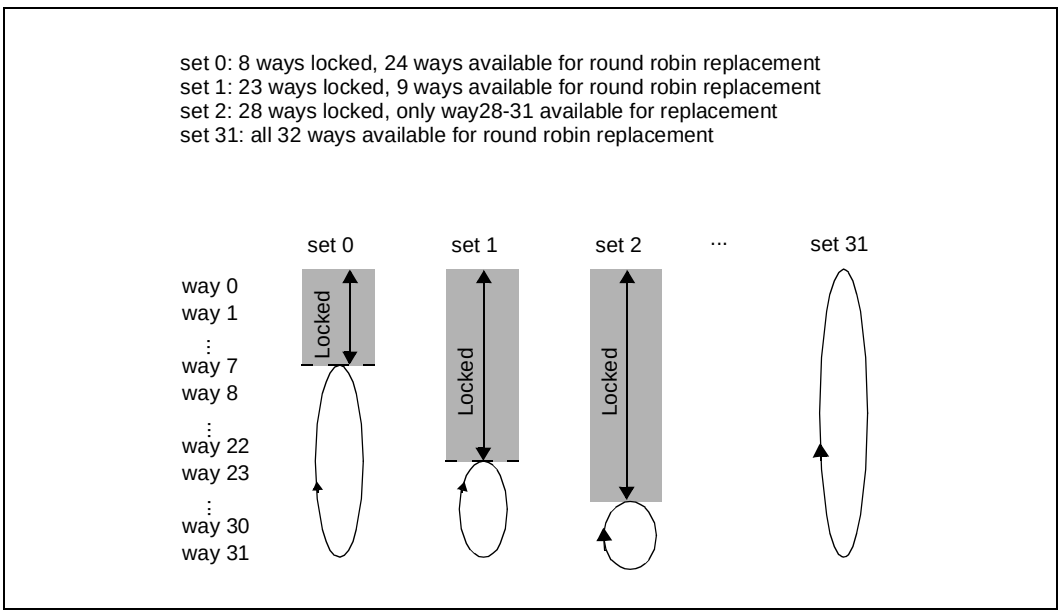

Software can lock down several different routines located at different memory locations. This may cause some sets to have more locked lines than others as shown in Figure 4-2.

Example 4-4 on page 4-9 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 2-12).

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® 80200 processor 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 7-13 for the exact command.)

Branch Target Buffer 5

Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) uses dynamic branch prediction to reduce the penalties associated with changing the flow of program execution. The Intel® 80200 processor 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

TAG	DATA	
Branch Address[31:9,1]	Target Address[31:1]	History Bits[1:0]

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, contends 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 5-2 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 14, “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

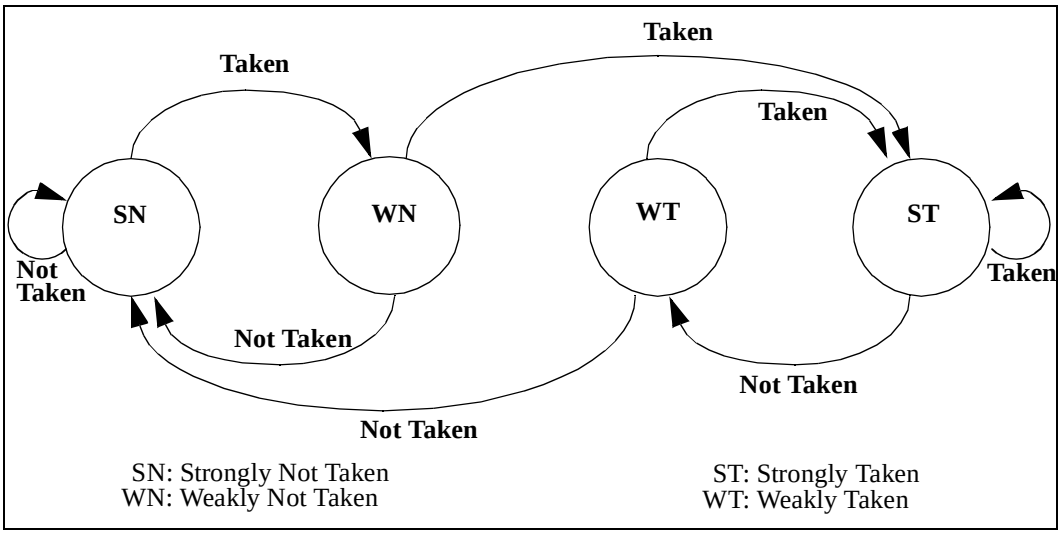

5.1.1 Reset

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 is 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 out of 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 ensures 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 7-10.
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.

Data Cache 6

The Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) data cache enhances performance by reducing the number of data accesses to and from external memory. There are two data cache structures in the Intel® 80200 processor, a 32 Kbyte data cache and a 2 Kbyte mini-data cache. An eight entry write buffer and a four entry fill buffer are also implemented to decouple the Intel® 80200 processor instruction execution from external memory accesses, which increases overall system performance.

6.1 Overviews

6.1.1 Data Cache Overview

The data cache is a ***32-Kbyte, 32-way set associative cache***; this means there are 32 sets with each set containing 32 ways. Each way of a set contains 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 6-2 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 allocates 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

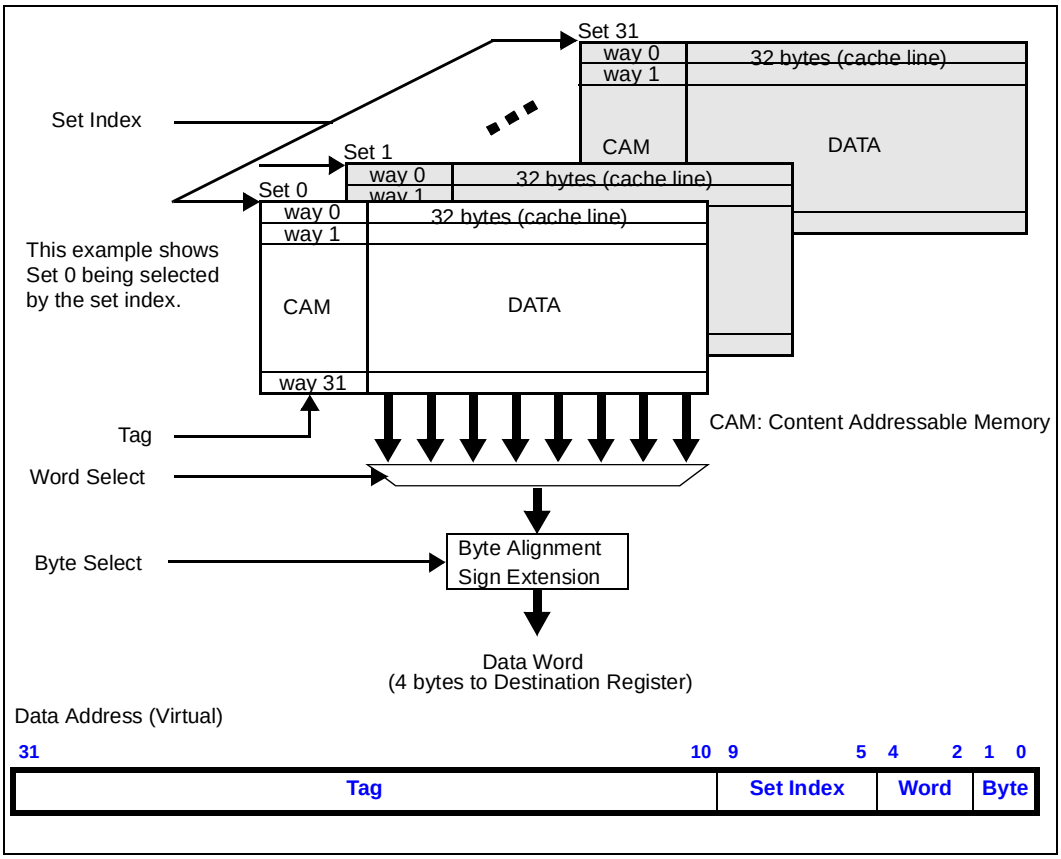

6.1.2 Mini-Data Cache Overview

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

Figure 6-2, "Mini-Data Cache Organization" on page 6-3 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

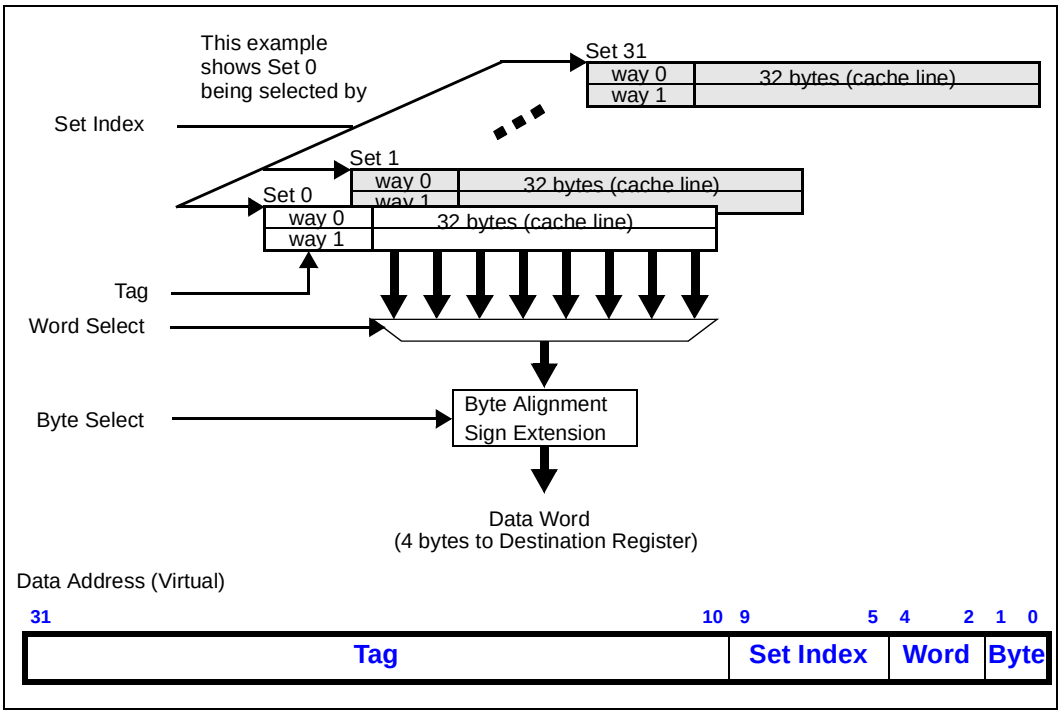

6.1.3 Write Buffer and Fill Buffer Overview

The Intel® 80200 processor employs an eight entry write buffer, each entry containing 16 bytes. Stores to external memory are first placed in the write buffer and subsequently taken out when the bus is available.

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

The fill buffer holds the external memory request information for a data cache or mini-data cache fill or non-cacheable read request. Up to four 32-byte read request operations can be outstanding in the fill buffer before the Intel® 80200 processor 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 6-6 and Section 6.2.3.3, “Write Miss Policy” on page 6-7 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 returns 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 does 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 6-5) load operation misses the cache:

1. The fill buffer is checked to see if an outstanding fill request already exists for that line.

   If so, the current request is placed in the pending buffer and waits until the previously requested fill completes, after which it accesses the cache again, to obtain the request data and returns it to the destination register.

   If there is no outstanding fill request for that line, the current load request is placed in the fill buffer and a 32-byte external memory read request is made. If the pending buffer or fill buffer is full, the Intel® 80200 processor stalls 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 6-8). 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 are 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, obtains 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 requests 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 Intel® 80200 processor stalls until an entry is available.

2. The 32-bytes of data can be returned back to the Intel® 80200 processor in any word order, i.e, the eight words in the line can be returned in any order. Note that it does not matter, for performance reasons, which order the data is returned to the Intel® 80200 processor 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 6-8). 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 are written back to external memory as a 4 word burst operation. This write operation is 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 causes 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® 80200 processor 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 does 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 6-12 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 sets bit 10 of the Fault Status Register register.

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

A data abort due to a data/mini-data cache parity error may not be recoverable if the data address that caused the abort occurred on a line in the cache that has a write-back caching policy. Prior updates to this line may be lost; in this case the software exception handler should perform a "clean and clear" operation on the data cache, ignoring subsequent parity errors, and restart the offending process. This operation is shown in Section 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. If the swap operation is directed to external memory the BCU performs a locked set of memory operations (see Chapter 11, "Bus Controller").

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.

Example 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 Chapter 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 & 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 7-10 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 7-13.

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 will evict any dirty data in the cache back to external memory. Example 6-2 shows how the 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 produces unpredictable results.

The line-allocate command will not operate on the mini Data Cache, so system software must clean this cache by reading 2KByte 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 the global clean operation depends on the number of dirty lines in the 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 always hits the cache unless it is invalidated. Once a line is locked into the data cache it is no longer available for cache allocation on a line fill. Up to 28 lines in each set can be reconfigured as data RAM, such that the maximum data RAM size is 28 Kbytes.

Hardware does not support locking lines into the mini-data cache; any attempt to do this produces 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 6-13 and Example 6-4 on page 6-14. The difference between these two routines is that Example 6-3 on page 6-13 actually requests the entire line of data from external memory and Example 6-4 on page 6-14 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 6-14 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 produces 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 drain ensures predictable results if a programmer tries to lock more than 28 lines in a set; the tag gets 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
PLD [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.
DRAIN
SUBS R0, R0, #1; Decrement loop count
BNE LOOP1

        ; Turn off data cache locking
DRAIN
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],#4   ;
STRD R4, [R1],#4   ;
STRD R4, [R1],#4   ;
STRD R4, [R1],#4   ;

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 7-13 for the exact command.) Once enabled, any new lines allocated into the data cache are locked down.

Note that the **PLD** instruction does 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 6-15 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

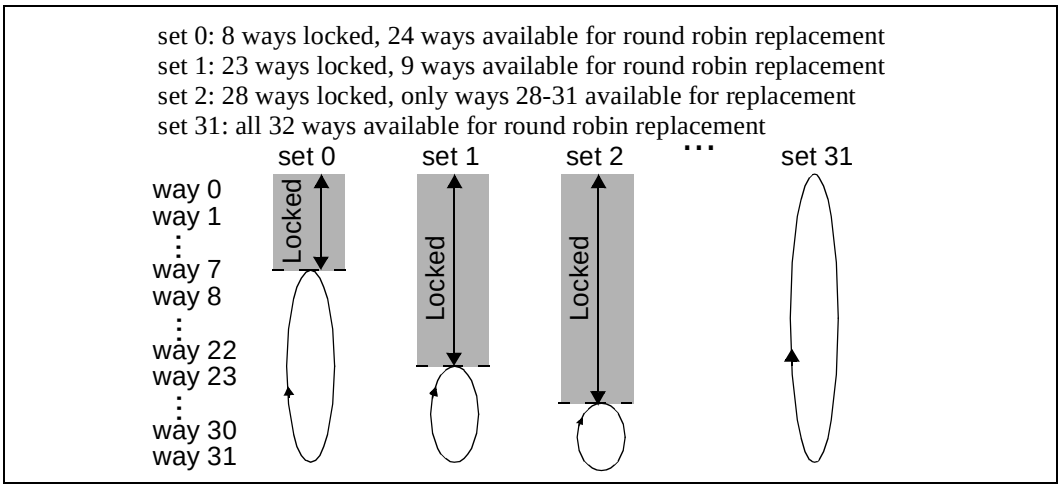

Software can lock down data located at different memory locations. This may cause some sets to have more locked lines than others as shown in Figure 6-3.

Lines are unlocked in the data cache by performing an unlock operation. See Section 7.2.10, "Register 9: Cache Lock Down" on page 7-13 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 Intel® 80200 processor does not refetch such data, which results 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 1.2.2, “Terminology and Acronyms” on page 1-5 for a definition of coalescing.

The write buffer is always enabled, which means, stores to external memory are 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 occurs 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 Intel® 80200 processor data requests to external memory including the write operations in the bus controller have completed. See Table 7-12, “Cache Functions” on page 7-10 for the exact command.

Writes to a region marked non-cacheable/non-bufferable (page attributes C, B, and X all 0) causes 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 7-10.

Configuration 7

This chapter describes the System Control Coprocessor (CP15) and coprocessor 14 (CP14). CP15 configures the MMU, caches, buffers and other system attributes. Where possible, the definition of CP15 follows the definition in the first generation Intel® StrongARM* products. CP14 contains the performance monitor registers and the trace buffer 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 causes an undefined instruction exception.

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

Coprocessors on the Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) do not support access via **CDP**, **MRRC**, or **MCRR** instructions. An attempt to execute these instructions results 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 2-11.

Like certain other ARM* architecture products, the Intel® 80200 processor 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 7-15. Privileged code needs to be aware of this facility because, when interacting with CP15, some addresses are modified by the PID and others are not.

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

The format of **MRC** and **MCR** is shown in Table 7-1.

cp_num is defined for CP15, CP14, CP13 and CP0. CP13 contains the interrupt controller and bus controller registers and is described in Chapter 9, "Interrupts"and Chapter 11, "Bus Controller," respectively. CP0 supports instructions specific for DSP and is described in Chapter 2, "Programming Model." Access to all other coprocessors on the Intel® 80200 processor causes an undefined exception.

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	0b1111 = CP15 0b1110 = CP14 0x1101 = CP13 0x0000 = CP0
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** is shown in Table 7-2. **LDC** and **STC** follow the programming notes in the *ARM Architecture Reference Manual*.

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.

Table 7-2. LDC/STC 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	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 *ARM Architecture Reference Manual*.	-
22	**N** - should be 0 for Intel® 80200 processors. Setting this bit to 1 has 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	0b1111 = Undefined Exception 0b1110 = CP14 0b1101 = CP13
7:0	8-bit word offset	-

7.2 CP15 Registers

Table 7-3 lists the CP15 registers implemented in the Intel® 80200 processor.

Table 7-3. CP15 Registers

Register (CRn)	Opcode_2	Access	Description
0	0	Read / Write-Ignored	ID
0	1	Read / Write-Ignored	Cache Type
1	0	Read / Write	Control
1	1	Read / Write	Auxiliary Control
2	0	Read / Write	Translation Table Base
3	0	Read / Write	Domain Access Control
4	-	Unpredictable	Reserved
5	0	Read / Write	Fault Status
6	0	Read / Write	Fault Address
7	0	Read-unpredictable / Write	Cache Operations
8	0	Read-unpredictable / Write	TLB Operations
9	0	Read / Write	Cache Lock Down
10	0	Read / Write	TLB Lock Down
11 - 12	-	Unpredictable	Reserved
13	0	Read / Write	Process ID (PID)
14	0	Read / Write	Breakpoint Registers
15	0	Read / Write	(CRm = 1) CP Access

7.2.1 Register 0: ID and Cache Type Registers

Register 0 houses two read-only registers that are used for part identification: an ID register and a cache type register.

The ID Register is selected when *opcode_2*=0. This register returns the code for the Intel® 80200 processor: 0x69052000 for A0 stepping/revision. The low order four bits of the register are the chip revision number and will be incremented for future steppings.

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	0	0	1	0	0	0	0	0	0	0	0	0	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 5
15:4	Read / Write Ignored	Part Number (Implementation Specified) Intel® 80200 processor: 0x200 Bits[15:12] refer to the processor generation. Bits[11:8] refer to the implementation Bits[7:4] used for implementation derivatives
3:0	Read / Write Ignored	Revision number for the processor (Implementation Specified) A0 stepping = 0b0000 B0 stepping = 0b0010

The Cache Type Register is selected when *opcode_2*=1 and describes the present Intel® 80200 processor cache.

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	1	1	0	1	0	1	0	1	0	0	0	0	1	1	0	1	0	1	0	1	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 = 0b110 = 32 kB
17:15	Read / Write Ignored	Data cache associativity = 0b101 = 32
14	Read-as-Zero / Write Ignored	Reserved
13:12	Read / Write Ignored	Data cache line length = 0b10 = 8 words/line
11:9	Read-as-Zero / Write Ignored	Reserved
8:6	Read / Write Ignored	Instruction cache size = 0b110 = 32 kB
5:3	Read / Write Ignored	Instruction cache associativity = 0b101 = 32 kB
2	Read-as-Zero / Write Ignored	Reserved
1:0	Read / Write Ignored	Instruction cache line length = 0b10 = 8 words/line

7.2.2 Register 1: Control and Auxiliary Control Registers

Register 1 is made up of two registers, one that is compliant with ARM Version 5 and is referenced by *opcode_2* = 0x0, and the other which is specific to Intel® StrongARM* and is referenced by *opcode_2* = 0x1.

The Exception Vector Relocation bit (bit 13 of the ARM control register) allows the vectors to be mapped into high memory rather than their default location at address 0. This bit is readable and writable by software. If the MMU is enabled, the exception vectors are accessed via the usual translation method involving the PID register (see Section 7.2.13, "Register 13: Process ID" on page 7-15) 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: writeable 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 *ARM Architecture Reference Manual* for more information.
8	Read / Write	**System Protection (S)** This selects the access checks performed by the memory management unit. See the *ARM Architecture Reference Manual* 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 Intel® 80200 processor Control Register, are used to control the allocation policy for the mini-data cache and whether it uses write-back caching or write-through caching.

The configuration of the mini-data cache 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: writeable 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 are coalesced in the write buffer as long as coalescing is globally enabled (bit 0 of this register). 0b00 = Write back, Read allocate 0b01 = Write back, Read/Write allocate 0b10 = Write through, Read allocate 0b11 = Unpredictable
3:2	Read-Unpredictable / Write-as-Zero	Reserved
1	Read / Write	**Page Table Memory Attribute (P)** If set, page table accesses are protected by ECC. See Chapter 11, “Bus Controller” 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 *ARM Architecture Reference Manual*.

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 2-12. Bit 9 indicates that a debug event occurred and the exact source of the event is found in the debug control and status register (CP14, register 10). When bit 9 is set, the domain and extended status field are undefined.

Upon entry into the prefetch abort or data abort handler, hardware updates this register with the source of the exception. Software is not required to clear these fields.

Table 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 2-12
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

All the functions defined in the first generation of Intel® StrongARM* appear here. The Intel® 80200 processor adds other functions as well. 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 Intel® 80200 processor does not check permissions on addresses supplied for cache or TLB functions. Because only privileged software may execute these functions, full accessibility is assumed. Cache functions do not generate any of the following:

- translation faults
- domain faults
- permission faults

The invalidate instruction cache line command does not invalidate the BTB. 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 D Cache Line function reads from the data cache, it is capable of generating a parity fault. The other operations do not generate parity faults.

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 generates 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 never gets evicted. However, if a valid store is made to that line it is marked as “dirty” and gets written back to external memory if another line is allocated to the same cache location. This eviction produces unpredictable results if the line-allocate command used a virtual address that mapped to non-existent memory.

  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 is generated during normal program execution. This is the case when line-allocate is used to clean/invalidate the entire cache.

  — The line-allocate operation is used only on a cache region destined to be locked. When the region is unlocked, it must be invalidated before making another data access.

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 cache, instruction TLB, and data TLB. The entry to lock is specified by the virtual address in Rd. The data cache locking mechanism follows a different procedure than the others. The data cache is placed in lock down mode such that all subsequent fills to the data cache result in that line being locked in, as controlled by Table 7-15.

Lock/unlock operations on a disabled cache 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 7-14. 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: writeable 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® 80200 processor supports the remapping of virtual addresses through a Process ID (PID) register. This remapping occurs before the instruction cache, instruction TLB, data cache and data TLB are accessed. The PID register controls when virtual addresses are remapped and to what value.

The PID register is a 7-bit value that is ORed with bits 31:25 of the virtual address when they are zero. This effectively remaps the address to one of 128 “slots” in the 4 Gbytes of address space. 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 Intel® 80200 processor 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 7-16) 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® 80200 processor contains two instruction breakpoint address registers (IBCR0 and IBCR1), one data breakpoint address register (DBR0), one configurable data mask/address register (DBR1), and one data breakpoint control register (DBCON). The Intel® 80200 processor also supports a 256 entry, trace buffer that records program execution information. The registers to control the trace buffer are located in CP14.

Refer to Chapter 13, "Software Debug" for more information on these features of the Intel® 80200 processor.

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 CP13 for the Intel® 80200 processor. 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 results 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** - Program to zero for future compatibility
15:14	Read-as-Zero/Write-as-Zero	**Reserved** - Program to zero for future compatibility
13:0	Read / Write	**Coprocessor Access Rights**- Each bit in this field corresponds to the access rights for each coprocessor. For each bit: 0 = Access denied. Attempts to access corresponding coprocessor generates an undefined exception. 1 = Access allowed. Includes read and write accesses. Setting any of bits 12:1 has an undefined effect.

7.3 CP14 Registers

Table 7-21 lists the CP14 registers implemented in the Intel® 80200 processor.

Table 7-21. CP14 Registers

Register (CRn)	Access	Description
0-3	Read / Write	Performance Monitoring Registers
4-5	Unpredictable	Reserved
6-7	Read / Write	Clock and Power Management
8-15	Read / Write	Software Debug

7.3.1 Registers 0-3: Performance Monitoring

The performance monitoring unit 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 12, "Performance Monitoring", along with a description on how to use the performance monitoring facility.

Opcode_2 and CRm should be zero.

Table 7-22. Accessing the Performance Monitoring Registers

Function	CRn (Register #)	Instruction
Read PMNC	0b0000	MRC p14, 0, Rd, c0, c0, 0
Write PMNC	0b0000	MCR p14, 0, Rd, c0, c0, 0
Read CCNT	0b0001	MRC p14, 0, Rd, c1, c0, 0
Write CCNT	0b0001	MCR p14, 0, Rd, c1, c0, 0
Read PMN0	0b0010	MRC p14, 0, Rd, c2, c0, 0
Write PMN0	0b0010	MCR p14, 0, Rd, c2, c0, 0
Read PMN1	0b0011	MRC p14, 0, Rd, c3, c0, 0
Write PMN1	0b0011	MCR p14, 0, Rd, c3, c0, 0

7.3.2 Register 4-5: Reserved

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

7.3.3 Registers 6-7: Clock and Power Management

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

Three low power modes are supported that are entered upon executing the functions listed in Table 7-24. To enter any of these modes, write the appropriate data to CP14, register 7 (PWRMODE). Software may read this register, but since software only runs during ACTIVE mode, it always reads 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: writeable bits set to 0

Bits	Access	Description
31:2	Read-unpredictable / Write-as-Zero	Reserved
1:0	Read / Write	**Mode (M)** 0 = ACTIVE 1 = IDLE 2 = DROWSY 3 = SLEEP

Software can change the core clock frequency by writing to the CP 14 register 6, CCLKCFG. This function waits for all the Intel® 80200 processor initiated memory requests to complete and informs the PLL to change the core clock frequency. This function completes when the PLL is re-locked. Software can read CCLKCFG to determine current operating frequency.

Table 7-24. Clock and Power Management

Function	Data	Instruction
Go to IDLE	1	MCR p14, 0, Rd, c7, c0, 0
Go to DROWSY	2	MCR p14, 0, Rd, c7, c0, 0
Go to SLEEP	3	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. The value in this field is multiplied by REFCLK to obtain core clock. See Table 8-2.

7.3.4 Registers 8-15: Software Debug

Software debug is supported by address breakpoint registers (Coprocessor 15, register 14), serial communication over the JTAG interface and a trace buffer. Registers 8 and 9 are used for the serial interface and registers 10 through 13 support a 256 entry trace buffer. Register 14 and 15 are the debug link register and debug SPSR (saved program status register). These registers are explained in more detail in Chapter 13, "Software Debug".

Opcode_2 and CRm should be zero.

Table 7-26. Accessing the Debug Registers

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

System Management 8

This chapter describes the clocking and power management features of the Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) along with reset details. Main features include a software controlled internal clock frequency and three low power modes:

- idle
- drowsy
- sleep

8.1 Clocking

CLK is the input reference clock for the Intel® 80200 processor. CLK accepts an input clock frequency of 33 to 66 MHz. The Intel® 80200 processor uses an internal PLL to lock to the input clock and multiplies the frequency by a variable multiplier to produce a high-speed core clock (CCLK). This multiplier is initially configured by the PLL configuration pin (PLLCFG) and can be changed anytime later by software.Table 8-1 shows the possible clock multipliers immediately after the reset sequence. PLLCFG can select between a multiplier of three and six. PLLCFG is sampled when RESET# transitions from low to high.

Table 8-1. Reset CCLK Configuration

PLLCFG Value (Sampled at Deassertion of RESET#)	Initial CLK Multiplier
0	3
1	6[a]

a. Lower speed grade parts may not provide this; instead they substitute the highest supported multiplier.

MCLK is the input memory clock for the Intel® 80200 processor. MCLK is asynchronous with respect to CCLK and supports frequencies up to 100 MHz. The ratio of MCLK to CCLK must be 1:3 or less. For example, if CCLK is 200 MHz, MCLK is restricted to 66 MHz or less.

Normally, CLK and MCLK support 40%/60% duty cycle inputs. At higher input frequencies, MCLK may impose a stricter requirement, such as 50%/50%. See the latest *Intel® 80200 Processor based on Intel® XScale™ Datasheet* for details.

Software has the ability to change the frequency of CCLK without having to reset the core. This feature allows software to conserve power by matching the core frequency to the current workload. Register CCLKCFG (see Section 7.3.3, “Registers 6-7: Clock and Power Management” on page 7-19) controls the clock multiplier.

Table 8-2. Software CCLK Configuration

CCLKCFG[3:0] (Coprocessor 14, register 6)	Multiplier for CLK	Example: CCLK Frequency (MHZ), assuming CLK Frequency of 66MHz
0 (reserved)	Unpredictable	
1	3	200
2	4	266
3	5	333
4	6	400
5	7	466
6	8	533
7	9	600
8	10	666
9	11	733
10-15 (reserved)	Unpredictable	

The Intel® 80200 processor supports low voltage operation with a supply as low as 0.95 V. At lower voltages, not all CCLK configurations are available. See the Intel® 80200 processor Datasheet for voltage/frequency information.

Changing CCLK frequency is similar to entering a low power mode. First, the core is stalled waiting for all processing to complete, second the new configuration is programmed into CCLKCFG, and then finally the core waits for the PLL to re-lock. The exact code sequence is shown in Equation 8-1. If there are no external bus transactions, this procedure takes approximately one thousand CLK cycles, the same time it takes to transition out of reset.

After the Intel® 80200 processor resets, the value in CCLKCFG reflects the effect of the PLLCFG pin: CCLKCFG contains either one or four.

Example 8-1. CCLK Modification Procedure

```
MOV R1, #7; New CCLKCFG value
MCR P14,0,R1,C6,C0,0; Change core clock frequency and wait for PLL
        ; to re-lock
```

8.2 Processor Reset

The RESET# pin must be asserted when CLK and power are applied to the processor. CLK, MCLK, and power must be present and stable before RESET# can be deasserted.

To ensure reset, RESET# must be asserted for at least 32 MCLK cycles once both clocks and the power are stable. Reset pulses shorter than this have an undefined effect.

To simplify external reset circuitry, the Intel® 80200 processor has a Schmitt trigger and an internal pull-up resistor on RESET#. This allows a board to implement power-on reset simply by connecting a 0.1 µF capacitor between RESET# and VSS_P.

TRST#, the JTAG reset pin, must be asserted simultaneously with RESET#. It is permissible to deassert it when RESET# is deasserted, or TRST# may remain asserted (tied to its active state). Like RESET#, TRST# has a Schmitt trigger and internal pull-up, so it is permissible to tie these together to a single external capacitor.

Hardware debug solutions may impose special requirements on RESET# and TRST#. Before designing a board, consult the guidelines provided by the hardware debug solutions you may be utilizing.

8.2.1 Reset Sequence

The output pin RESETOUT# is asserted when RESET# is asserted. The deassertion of RESET# triggers an internal reset timer that keeps the Intel® 80200 processor in a reset state until the PLL has locked. CCLK is not running during this reset state and RESETOUT# remains asserted. When the PLL is locked and stable, RESETOUT# is deasserted informing the system that the Intel® 80200 processor has completed reset and CCLK is running. The time from the deasserting of RESET# to the deassertion of RESETOUT# is approximately two thousand CLK cycles.

MCLK (input memory clock) must be stable, at a minimum, 10 CLK cycles before the deassertion of RESET#. After RESETOUT# is deasserted, MCLK need only be present when the Intel® 80200 processor external bus is active.

Figure 8-1. Reset Sequence

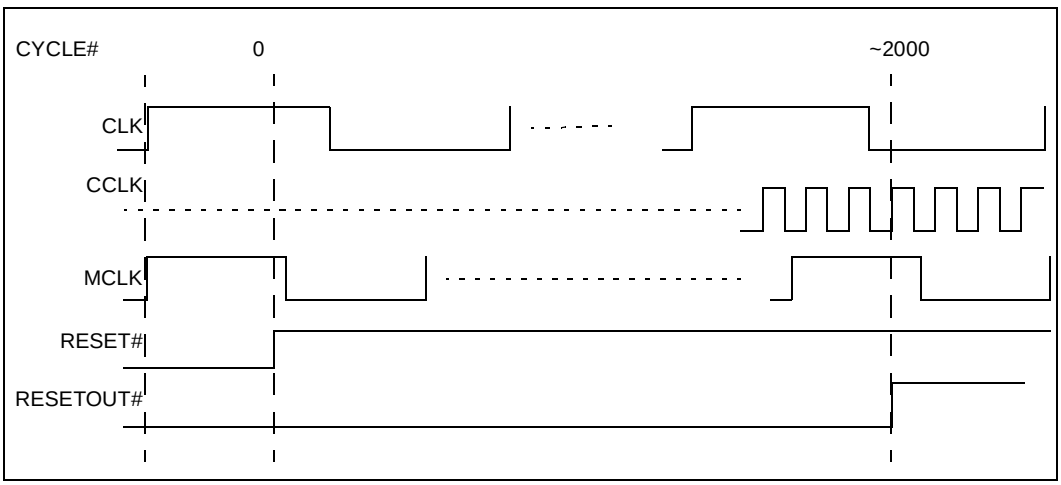

8.2.2 Reset Effect on Outputs

After RESETOUT# is asserted, the processor's output pins are driven to a well-defined state. Critical bus signals receive a '0' or '1' value, as shown in Figure 8-2. This figure also illustrates that HOLD is acknowledged during the reset sequence. Output pins only transition if a valid MCLK is present.

Figure 8-2. Pin State at Reset

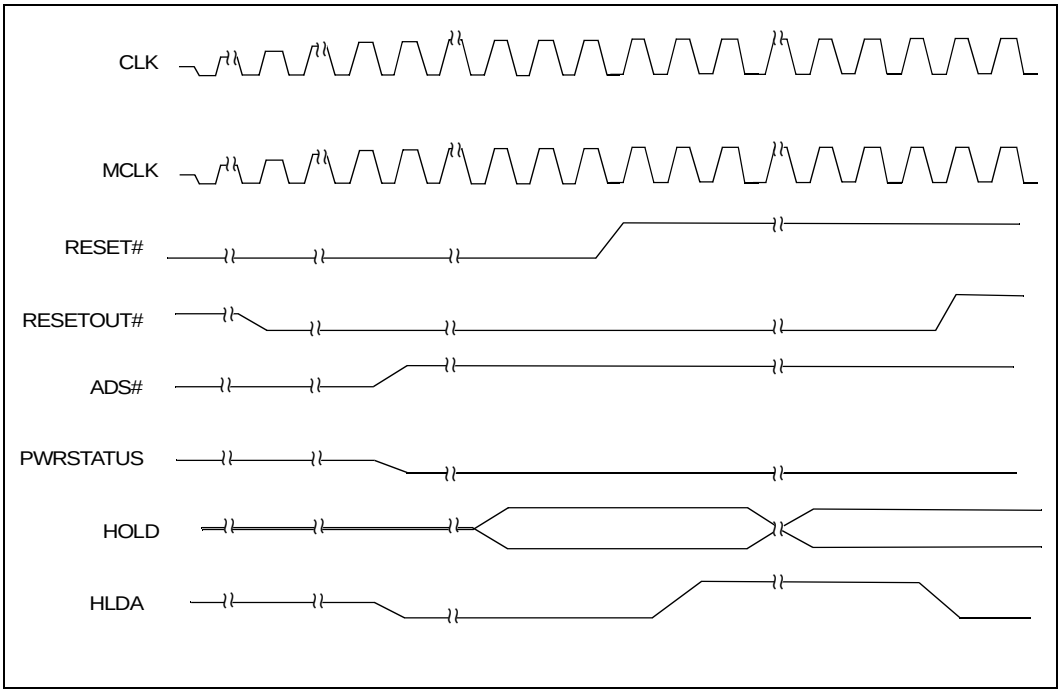

8.3 Power Management

The Intel® 80200 processor provides low power modes: idle, drowsy and sleep, which are listed in increasing power saving order. Table 8-3 describes the attributes of each low power mode.

Table 8-3. Low Power Modes

Low Power Mode	PLL	Architectural State	Wakeup Method
Idle	On	Retained	FIQ#/IRQ#
Drowsy	Off	Retained	FIQ#/IRQ#
Sleep	Off	Must be saved prior to entering	RESET#

8.3.1 Invocation

The Intel® 80200 processor provides a simple instruction to enter into low power mode. See Table 7-24, “Clock and Power Management” on page 7-19 for exact commands. This instruction waits for all processing to complete, asserts the **PWRSTATUS** output pins and waits for a wake up event to transition back into the normal running mode. For idle and drowsy, the wake-up event is the assertion of **FIQ#** or **IRQ#** pins. If the interrupt is masked, the Intel® 80200 processor still wakes-up but won’t service the interrupt.

The only way to exit Sleep mode is to go through the reset sequence. Sleep mode provides the greatest power savings since the Intel® 80200 processor supply voltage can be reduced to zero.

Example 8-2. Entering Drowsy Mode

```
LDR R0, #2
MCR P14,0,R0,C7,C0,0; transition to Drowsy mode
```

8.3.2 Signals Associated with Power Management

PWRSTATUS[1:0] is a 2-bit output from the Intel® 80200 processor. It carries information about the current power mode on the part. Table 8-4 shows the encoding of the **PWRSTATUS[1:0]** signals.

Table 8-4. PWRSTATUS[1:0] Encoding

PWRSTATUS[1:0]	Power Mode
00	Normal
01	Idle
10	Drowsy
11	Sleep

The external interrupt pins may be used to exit Idle or Drowsy mode. If MCLK is toggling, asserting FIQ# or IRQ# wakes the Intel® 80200 processor up, even if the interrupt is disabled.

If interrupts are enabled, it takes some time after the processor is woken up. The exact timing depends on the CCLK:MCLK ratio and implementation details, and can not be reliably predicted. If software needs to guarantee it does not proceed until an interrupt is taken, it should poll for an asserted interrupt after it is woken up. Typical code to accomplish this task is shown in Example 8-3 on page 8-6.

Figure 8-3. Waiting for the Wake-up Interrupt

```
MOV     R0, #2
MOV     R1, #0xDF
MSR     CPSR_c, R1                        ; disable interrupts
MCR     P14,0,R0,C7,C0,0                  ; transition to Drowsy mode

;; Get here when we leave drowsy mode. Now, wait
;; until the interrupt that woke us has been registered.

loop:
MRC     P13, 0, R15, C4, C0, 0     ; get INTSRC into flags
BMI     sawInterrupt               ; if bit 31 set, then we got an FIQ
BNE     loop                       ; if bit 30 set, then we got an IRQ

sawInterrupt:
;; Get here because FIQ or IRQ asserted and registered

MOV     R1, #0x1F
MSR     CPSR_c, R1                        ; enable interrupts

;; Before we get to the next instruction, the interrupt
;; will be taken.
```

As with all external interrupts, the interrupt source must keep the wake-up interrupt asserted until told otherwise by software running on the Intel® 80200 processor. See Chapter 9, “Interrupts”, for more information. After an interrupt is asserted, the Intel® 80200 processor takes approximately two thousand CLK cycles to exit Drowsy mode, and approximately 10 CLK cycles to exit Idle mode.

The JTAG clock must be stopped during drowsy and sleep mode. Drive a ‘0’ into the JTAG clock when not toggling it.

Interrupts 9

9.1 Introduction

The Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) supports a variety of external and internal interrupt sources. The Interrupt Control Unit (ICU) controls how the Intel® 80200 processor reacts to these interrupts. Ultimately, all interrupt sources are combined into one of two internal interrupts: IRQ and FIQ. These interrupts correspond to the IRQ and FIQ described in the *ARM Architecture Reference Manual*.

The two interrupt signals that enter the chip are FIQ# (fast interrupt) and IRQ# (normal interrupt). These signals must be asserted and held low to interrupt the processor.

The internal interrupt sources originate in the Bus Controller Unit (see Chapter 11, "Bus Controller") and the Performance Monitoring Unit (see Chapter 12, "Performance Monitoring"). To allow flexible system design, these interrupts may be steered under software control to act equivalently to either FIQ or IRQ.

All interrupts are *level sensitive*: interrupt sources must keep asserting the interrupt signal until software causes the source to deassert it.

All interrupt sources are individually maskable with the ICUs Interrupt Control register (INTCTL). Additionally, all interrupts may be quickly disabled by altering the F and I bits in the CPSR as specified in the *ARM Architecture Reference Manual*.

When software running on the Intel® 80200 processor is vectored to an Interrupt Service Routine (ISR), it may query the ICUs Interrupt Source register (INTSRC) to quickly determine the source of the interrupt.

9.2 External Interrupts

The two external interrupts, FIQ# and IRQ#, go through synchronization logic before being sampled by the ICU.

External interrupts must be held asserted until cleared at the interrupting source by software. The Intel® 80200 processor does not latch the external interrupt signals. Enabled interrupts that are deasserted before software enters the interrupt service routine causes *UNPREDICTABLE* behavior.

9.3 Programmer Model

Software has access to three registers in the ICU. INTCTL is used to enable or disable (mask) individual interrupts. As mentioned, masking of all interrupts may still be accomplished via the CPSR register in the core. INTSRC is a read-only register that records all currently active interrupt sources. Even if an interrupt is masked, software may use INTSRC to test for its source. INTSTR is used to direct internal interrupts to either FIQ or IRQ.

Figure 9-1. Interrupt Controller Block Diagram

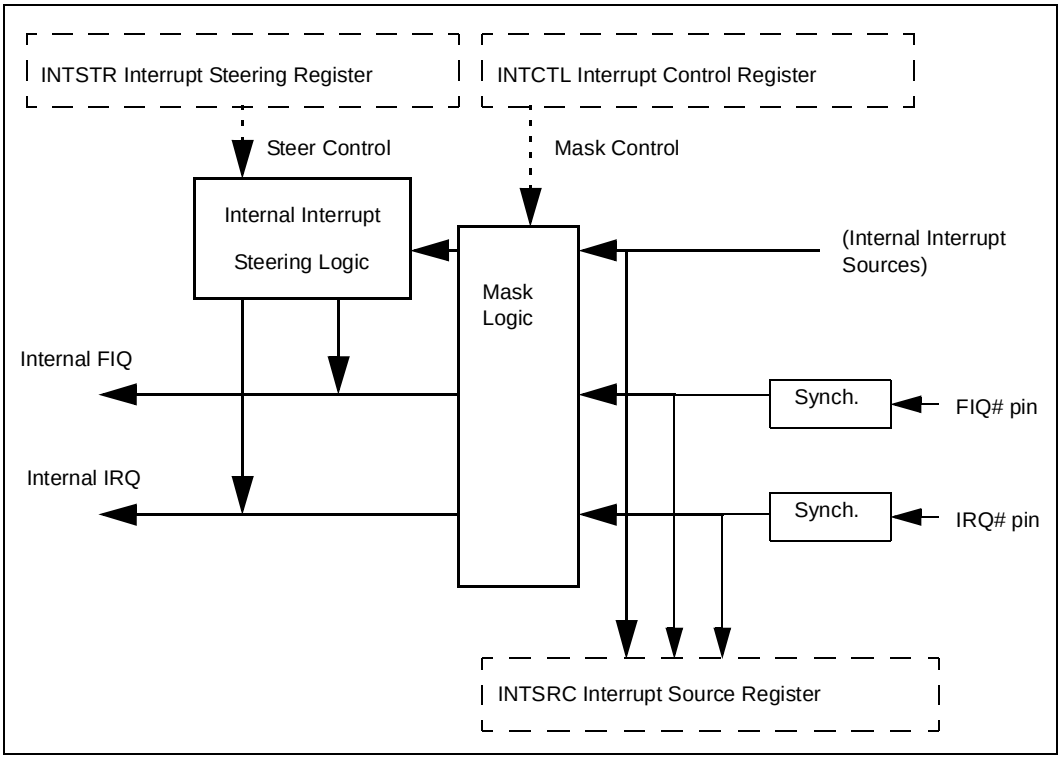

The ICU registers reside in Coprocessor 13 (CP13). They may be accessed/manipulated with the MCR, MRC, STC, and LDC instructions. The *CRn* field of the instruction denotes the register number to be accessed. The *opcode_1*, *opcode_2*, and *CRm* fields of the instruction should be zero. Most systems restricts access to CP13 to privileged processes. To control access to CP13, use the Coprocessor Access Register (see Section 7.2.15).

An instruction that modifies an ICU register is guaranteed to take effect before the next instruction executes. For example, if an instruction masks an interrupt source, subsequent instructions execute in an environment in which the masked interrupt does not occur.

The details of the ICU registers are discussed in the following sections.

9.3.1 INTCTL

INTCTL is used to specify what interrupts are disabled (masked).

Table 9-1. Interrupt Control Register (CP13 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
																												BM	PM	IM	FM

reset value: writeable bits set to 0

Bits	Access	Description
31:4	Read-unpredictable / Write-as-Zero	Reserved
3	Read / Write	**BM** -BCU Mask Controls whether BCU interrupts are enabled 0 = disable interrupt 1 = enable interrupt
2	Read / Write	**PM** -PMU Mask Controls whether PMU interrupts are enabled 0 = disable interrupt 1 = enable interrupt
1	Read / Write	**IM** - IRQ# mask Enables external interrupts from the IRQ# pin 0 = disable interrupt 1 = enable interrupt
0	Read / Write	**FM** -FIQ# mask Enables external interrupts from the FIQ# pin 0 = disable interrupt 1 = enable interrupt

9.3.2 INTSRC

The Interrupt Source register (INTSRC) indicates which interrupts are active.

This register may be used by an ISR to determine quickly the source of an interrupt. Even if an interrupt is masked with INTCTL, software may still detect whether it is asserted by reading its bit from INTSRC.

Table 9-2. Interrupt Source Register (CP13, register 4)

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
FI	II	BI	PI																												

reset value: undefined

Bits	Access	Description
31	Read / Write-ignored	**FI** - FIQ# active Holds the state of the synchronized FIQ# signal. 0 = not interrupting 1 = interrupting
30	Read / Write-ignored	**II** - IRQ# active Holds the state of the synchronized IRQ# signal. 0 = not interrupting 1 = interrupting
29	Read / Write-ignored	**BI** - BCU Interrupt Active Holds the state of the BCU interrupt request. 0 = BCU not interrupting 1 = BCU interrupting
28	Read / Write-ignored	**PI** -PMU Interrupt Active Holds the state of the PMU interrupt request. 0 = not interrupting 1 = interrupting
27:0	Read-unpredictable / Write-ignored	Reserved

Note that memory buffering and external logic on FIQ# and IRQ# could cause INTSRC.II and INTSRC.FI to remain asserted for several cycles after the interrupt source has been commanded to stop asserting. An ISR should ensure that the interrupt source is quelled, or a spurious recursive entry to the ISR may result when interrupts are enabled.

Example 9-1 on page 9-4, shows how software might wait for the FIQ# signal to be deasserted before proceeding.

Example 9-1. Waiting for FIQ# Deassertion

```
waitForNoFIQ:
MRC P13, 0, R15, C4, C0, 0; get high bits of INTSRC into flags
BMI waitForNoFIQ; if FI bit set, try again
```

9.3.3 INTSTR

Systems may have differing priorities for the various interrupt cases; the ICU allows system designers to associate each internal interrupt source with one of the two internal interrupts: FIQ and IRQ. This association is called *steering*.

INTSTR is used to specify how internal interrupt sources should be steered

Table 9-3. Interrupt Steer Register (CP13, register 8)

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
																														BS	PS

reset value: writeable bits set to 0

Bits	Access	Description
31:2	Read-unpredictable / Write-as-Zero	Reserved
1	Read / Write	**BS** - BCU Steering If BCU interrupts are enabled, this bit steers them. 0 = BCU interrupts directed to internal IRQ 1 = BCU interrupts directed to internal FIQ
0	Read / Write	**PS** - PMU Steering If PMU interrupts are enabled, this bit steers them. 0 = PMU interrupts directed to internal IRQ 1 = PMU interrupts directed to internal FIQ

External Bus 10

10.1 General Description

The Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) bus is a split bus, with separate request and data buses. It is designed primarily as the memory and I/O bus for the Intel® 80200 processor, not as a general purpose multi-master bus, although it is possible to have several masters on it efficiently. It is deeply pipelined and works well with deeply pipelined memory technologies, with the memory sharing the data bus with the Intel® 80200 processor, or with the memory on a separate chipset bus.

Figure 10-1 shows a typical system; this configuration allows the chipset to avoid having a second 64-bit memory bus in addition to the 64-bit Intel® 80200 processor data bus, with significant savings in chipset cost.

Figure 10-1. Typical System

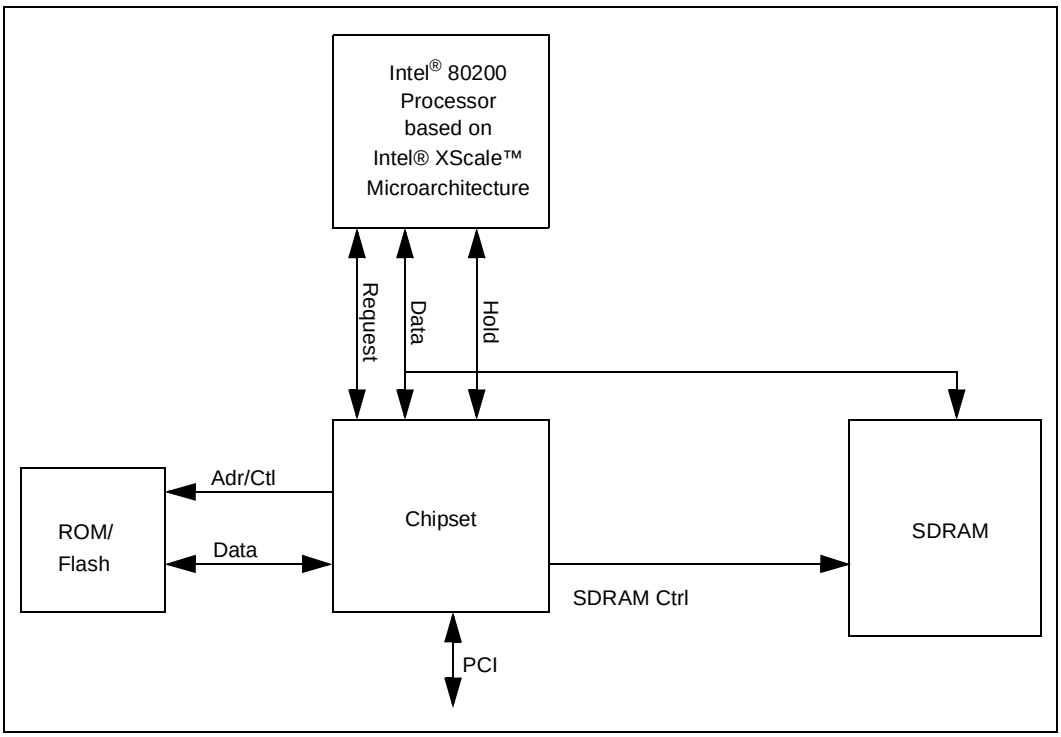

An alternate configuration with a separate memory bus is also possible, shown in Figure 10-2. All signals on this bus, data and request, are sampled on the rising edge of **MCLK**. **MCLK** is created by the system and is an input to the Intel® 80200 processor. **MCLK** is asynchronous with respect to the Intel® 80200 processor core frequency and any other Intel® 80200 processor input clocks. **MCLK** in the configuration shown in the figure would also need to be the SDRAM clock. **MCLK** frequencies of up to 100 MHz are supported. A 50% duty cycle is required. **MCLK** must be one third or less of the internal clock frequency of the core, however. An Intel® 80200 processor system running the core at 200 MHz would be allowed a maximum **MCLK** of 66 MHz, for example. This constraint comes from the design of the low latency synchronization logic in the Intel® 80200 processor bus controller.

Figure 10-2. Alternate Configuration

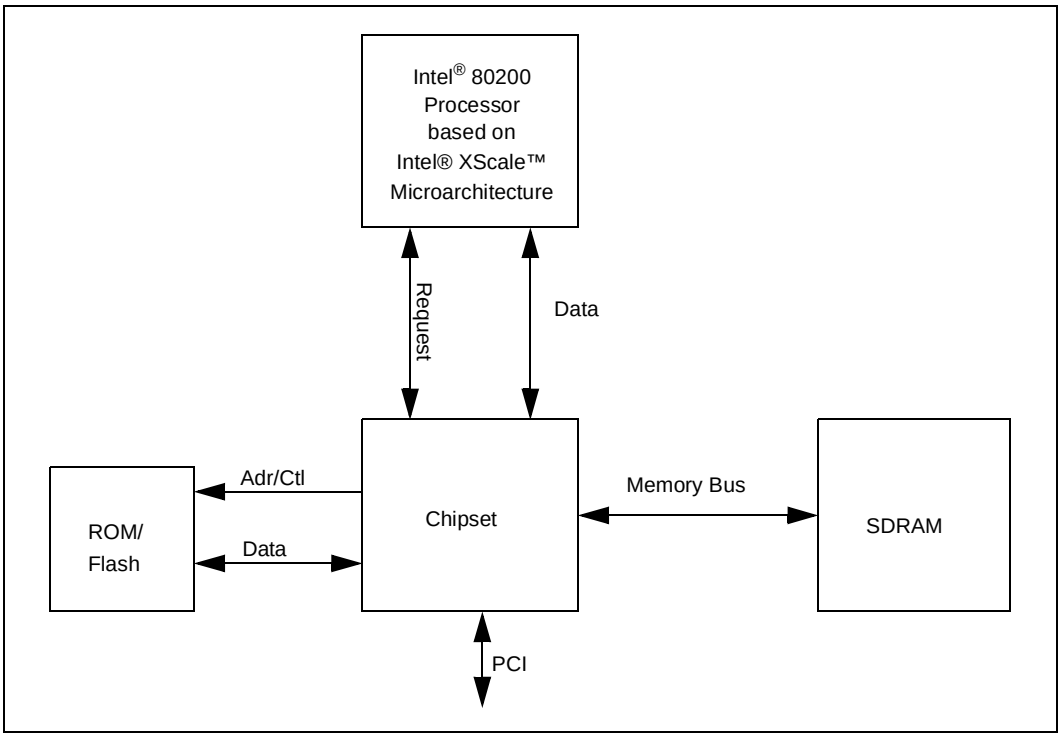

10.2 Signal Description

Table 10-1. Intel® 80200 Processor based on Intel® XScale™ Microarchitecture Bus Signals

Signal	Width	I/O	Function
MCLK	1	I	bus clock (note: all bus activity is triggered by the rising edge of this clock).
Request Bus			
ADS#/LEN[2]	1	O	During the first cycle of the issue phase: this signal indicates the start of a bus request. During the second cycle of the issue phase: this signal is the MSB of a value which indicates the length of the transaction.
Lock/LEN[1]	1	O	During the first cycle of the issue phase: this signal indicates whether the current transaction is part of an atomic read-write pair. During the second cycle of the issue phase: this signal is the middle bit of a value which indicates the length of the transaction.
W/R#/LEN[0]	1	O	During the first cycle of the issue phase: this signal indicates whether the current transaction is a write (**W/R#** = 1) or a read (**W/R#** = 0). During the second cycle of the issue phase: this signal is the LSB of a value which indicates the length of the transaction.
A	16	O	During the first cycle of the issue phase: this signal carries the upper 16 bits of the address for the access. During the second cycle of the issue phase: this signal carries the lower 16 bits of the address.
Data Bus			
D	64	I/O	Data Bus
BE#	8	O	Byte Enables for writes (timing same as Data)
DCB	8	I/O	Data Check Bits for ECC (timing same as Data)
CWF	1	I	Critical Word First: Indicates the order in which the current 32-byte read burst is returning. See Section 10.2.3, "Critical Word First" on page 10-7 for more information on this signal. This pin must be asserted at the same time as the DValid of the first cycle of returning data for a given read transaction. It is ignored at all other times.
DValid	1	I	Indicates that two cycles later there is an Intel® 80200 Processor data cycle on **D, DCB,** and possibly **BE#** (either read sampled by the Intel® 80200 processor or write driven by the Intel® 80200 processor)
Abort	1	I	Asserted with **DValid**, indicates that the Intel® 80200 processor transaction next in order on the Data bus has been aborted (timing same as **DValid**)
Multimaster Support			
Hold	1	I	Input that tells the Intel® 80200 processor to float the following pins: **ADS#**, **A**, **W/R#**, **Lock**
HldA	1	O	Asserted when the Intel® 80200 processor has completed the transition to hold mode in response to **Hold**.
Configuration			
CWF/ DBusWidth	1 (shared with CWF)	I	When the Intel® 80200 processor is in reset, this pin functions as **DBusWidth**: the Intel® 80200 processor samples this pin to find the bus width in use. If it is 1, then the system is operating with a 32-bit bus, otherwise the Intel® 80200 processor uses a 64-bit bus. This pin is sampled by the Intel® 80200 processor while RESETOUT# is asserted.

10.2.1 Request Bus

The request bus issues read or write requests from the Intel® 80200 processor or other bus master to the chipset or memory controller. Each request takes two **MCLK** cycles. All signals should be sampled on the rising edge of **MCLK**. No data is ever transferred on the request bus.

On the first cycle, **ADS#/LEN[2]**, **Lock/LEN[1]**, and **W/R#/LEN[0]** are used to carry the **ADS#**, **Lock**, and **W/R#** signals. A valid request is indicated by the **ADS#** signal being asserted low. On that same clock edge, the sampled value of **A** is the most-significant 16 bits of the 32-bit address of the request, **W/R#** indicates whether the request is a read or write from the Intel® 80200 processor, and the **Lock** pin indicates whether there is an atomic pair of operations outstanding.

On the second cycle of a request, **ADS#/LEN[2]**, **Lock/LEN[1]**, and **W/R#/LEN[0]** are used to carry **LEN[2:0]**. **LEN** is used to indicate the number of data bytes associated with the request. See Table 10-2 and Table 10-3 for information on how this signal is encoded. **A** has the least significant 16 bits of the address for the request.

Lock is asserted for the read request that begins the atomic pair, and is asserted for each new request up to (but not including) the write operation that ends the atomic pair. It is possible that instruction fetches and instruction MMU page table walk requests occurs between the read and the write of an atomic operation, but no other requests are made. Details of the lock mechanism and its interaction with the multimaster logic is described below.

A bus master may have at most four requests outstanding at any time. That is: a bus master may issue up to four requests before receiving back any data. For more information, see Section 10.3.7, “Pipelined Accesses” on page 10-21.

10.2.1.1 Intel® 80200 Processor Use of the Request Bus

The possible sizes and alignments of read and write requests that can be issued are shown in Table 10-2 and Table 10-3.

Table 10-2. Requests on a 64-bit Bus

LEN[1]	Number of Data Bytes	Number of Data Bus Cycles	Used for Reads?	Used for Writes?	Address Alignment
000	1	1	Y	Y	Any Address
001	2	1	Y	Y	A[0] = “0”
010	4	1	Y	Y	A[1:0] = “00”
011	8	1	Y[2]	Y	A[2:0] = “000”
100	Not Used in 64-bit Bus Mode				
101	16	2	N	Y	A[3:0] = “0000”
110	32	4	Y	N	A[1:0] = “00”[3]
111	Not Used in 64-bit Bus Mode				

1. LEN of 000, 001, or 010 should be treated as a LEN of 011 if directed to a slave that implements ECC.
2. An 8-byte read will only occur if ECC is enabled; it will occur as part of read-modify-write transactions (see Section 11.2, “ECC” on page 11-1). If ECC is not enabled, the Intel® 80200 processor will never perform an 8-byte read.
3. On a 32-byte load, A[4:2] carries information for Critical Word First logic. See Section 10.2.3, “Critical Word First” on page 10-7 for more information.

Table 10-3. Requests on a 32-bit Bus

LEN	# Data Bytes	# Data Bus Cycles	Used for Reads?	Used for Writes?	Address Alignment
000	1	1	Y	Y	Any Address
001	2	1	Y	Y	A[0] = “0”
010	4	1	Y	Y	A[1:0] = “00”
011	8	2	N	Y	A[1:0] = “00”
100	12	3	N	Y	A[3:0] = “0X00”
101	16	4	N	Y	A[3:0] = “0000”
110	32	8	Y	N	A[1:0] = “00”[1]
111	Not Used in 32-bit Bus Mode				

1. On a 32-byte load, A[4:2] carries Critical Word First logic information see Section 10.2.3, “Critical Word First.

In addition to the alignment constraints listed above, read transactions never cross a 32-byte boundary, and write transactions never cross a 16-byte boundary.

Some write case explanations. Byte and short writes are caused by non-cacheable non-bufferable store commands in the Intel® 80200 processor. The four word write can be caused by eviction of dirty data in a cache line (the Intel® 80200 processor has two dirty bits for each eight-word cache line and evicts halves separately as necessary), or from the write buffer. For all writes not from write buffer (non-cacheable writes and cache line evictions), the writes are simple and the byte enables on the data bus are asserted for the contiguous bytes specified by the address and the length specified in the request.

Writes coming from write buffers can look somewhat different. Due to coalescing in the write buffers, it is possible to get a single write request on the bus writing out a non-contiguous byte pattern. The write buffers temporarily hold outgoing store data in 4-word aligned blocks, and later stores (byte, short, or word) to same block can be merged into (or overwrite) previous data.

Once a given write buffer is next in line for access to the bus, merging to it stops, and whatever pattern of bytes is valid in that write buffer determines the type of write transaction sent out. The byte enables on the bus indicates which bytes within that word are valid and need to be written.

Even if only one byte is valid in the coalesce buffer, a word store goes out. The byte enables are only asserted for the one byte, however. This means that a single byte write request to address 0x2402 can be requested in two valid ways: a non-cacheable **strb** instruction causes a write of **Len** 0x0 (byte) to **A** 0x2402, with only one byte enable asserted when the data is driven; whereas a coalesce buffer drain could cause a write of **Len** 0x2 (word) to **A** 0x2400, with only one byte enable asserted. The same applies for two byte stores.

Notice that it is possible on a 32-bit bus for a 3-4 word write transaction to go out which requires 3-4 data cycles on the bus, but during one or more of the middle data cycles no byte enables are asserted. The first and last data cycle always has at least one byte enable valid. Even though it seems inefficient to waste data cycles with no byte enables, on average the merging of writes in the write buffers can be a big performance gain.

Some examples: a pair of non-cacheable bufferable byte stores to addresses 0x2401 and 0x2403 might be merged in the write buffer. On the bus, they show up as a write of length “010” (word) to address 0x2400. When data is driven out on the data bus, however, only byte enables for bytes 1 and 3 would be asserted.

A pair of non-cacheable bufferable byte stores to addresses 0x2401 and 0x2409 could be merged in the write buffers and require a three word write. On a 64-bit bus, there would be two data cycles. Both data cycles have only byte enable 1 asserted. On a 32-bit bus, there would be three data cycles. The first and third would have 1-byte enable asserted.

10.2.2 Data Bus

Some time after a request is made on the request bus, data must be transferred for that request on the data bus. Each request has a corresponding transaction (one or more cycles) on the data bus. Data bus transactions must occur in the same order as the requests were made. The delay between a request going out and the data coming back to or being driven from the bus master is arbitrary. No explicit wait-state insertion is needed.

All data on the 64-bit data bus is read or written in its natural location within an aligned 64-bit memory block. In little endian mode, 64-bit bus (big endian and 32-bit busses are covered later) bits 7-0 of **D** always correspond to a byte with low address bits (2:0) of "000", and bits 63-56 always correspond to a byte with low address bits of "111". As an example, a word (32-bit) read to address 0x24004 would need to be returned to the core with the most significant byte on bits [63:56] and the least significant byte on bits [39:32]. For a byte write to location 0x3703, the valid data byte would be driven out on bits [31:24] of **D** (and only bit 3 of the BE# would be asserted).

Each data transaction consists of one or more data cycles. Each data cycle begins with the assertion of **DValid** to indicate to the Intel® 80200 processor that the next cycle of data is going to be transferred, followed two clock cycles later by the data transfer (read or write data) on the **D**, **BE#** (for writes only), and **DCB** buses.

All lines in **D** must be driven during a data transaction: on writes the Intel® 80200 processor will drive them, on reads they must be driven by the addressed slave. If the read request was for less than a full bus-width of data (example: a byte read), the other lines should be driven with some binary value.

If a read is directed to a slave that implements ECC, a full bus-width of valid data (64 bits) must be returned, without regard for the requested size. For example, even if just a byte is requested from ECC memory, the memory should still return eight bytes of data. This restriction guarantees that the Intel® 80200 processor will be able to compute ECC on the data; see Section 10.2.7 for more information about ECC.

DValid for a transaction may be asserted no earlier than the clock after the transaction request. In other words, if the Intel® 80200 processor asserts **ADS#** at cycle N, it must not receive a matching **DValid** until at least cycle N+2. See Figure 10-11 for an example of this minimum-wait timing.

For a single-word read returning to the core, the transaction would consist of the **DValid** signal being asserted high and sampled on clock edge *n*, and the data being sampled by the Intel® 80200 processor from the **D** bus on clock edge *n+2* (see Figure 10-4).

For a read burst request multiple data cycles are needed. On a 64-bit bus an eight word burst read would be four data cycles. The data cycles are independent and can occur back-to-back, or spread out with any delay between the cycles. Each cycle consists of **DValid** being asserted followed two cycles later by the corresponding data. This can be overlapped, such that **DValid** for one data cycle is being asserted while the **D** for another data cycle is being driven on the bus. A typical burst would be **DValid** asserted for four contiguous cycles (n, n+1, n+2, n+3) with the data for that burst being driven in four contiguous cycles delayed two cycles from the **DValid** (n+2, n+3, n+4, n+5), as shown in Figure 10-5.

If the part has been configured to use a 32-bit data bus, then **D**[63:32] and **DCB**[7:0] are floated on all writes. Also, bits in **BE#**[7:4] drives a binary ('1' or '0') value. It is suggested that **DCB** and the upper bits of **D** be pulled-down on 32-bit data bus systems.

10.2.3 Critical Word First

The **CWF** signal is only used during read bursts of eight words (**Len** = 6). **CWF** needs to be driven at the same time as **DValid** of the first data cycle of the transaction. This bit indicates to the requesting master what order the data is returning in. The Intel® 80200 processor uses this sort of transaction to fill a cache line.

There are two acceptable return orders that may be signalled with **CWF**:

- (**CWF** = 0) starting with the words 0 and 1 (64-bit bus) or word 0 (32-bit bus) and sequentially returning the data
- (**CWF** = 1) returning the word or pair of words that contain the byte indicated by the request address on the first cycle, sequentially returning the next highest memory location word or pair of words until the end of the eight word aligned block, then returning the lowest word or pair of words and incrementing back up the word or pair of words just below the word or pair of words pointed to by the request address.

The allowable orders are spelled out below. In this description, A,B,C, and D refer to pairs of 32-bit words in the aligned 8-word block (A=0,1 B=2,3 C=4,5 D=6,7). In the 32-bit bus section, a-h are referring to the eight words in the aligned 8-word block. Note that the Intel® 80200 processor will align all 8-word burst transactions on a 64-bit address. Specifically, if the instruction or data address requested by the program is odd word aligned, the Intel® 80200 processor will emit a 64-bit aligned address by setting A[2] = 0x0.

Table 10-4. Return Order for 8-Word Burst, 64-bit Data Bus

CWF	A[4:3]	Return Order
0	XX	ABCD
1	00	ABCD
1	01	BCDA
1	10	CDAB
1	11	DABC

Table 10-5. Return Order for 8-Word Burst, 32-bit Data Bus

CWF	A[4:2][1]	Return Order
0	XXX	abcdefgh
1	000	abcdefgh
1	010	cdefghab
1	100	efghabcd
1	110	ghabcdef

1. The Intel® 80200 processor will never generate an 8-word burst with A[2] = 0x1. The addresses for these requests will be 64-bit aligned.

The timing of write transactions on the data bus is similar to the timing of reads. After a write request has gone out, two cycles before the chipset or memory is ready to receive the data from the Intel® 80200 processor the **DValid** is asserted. Two cycles later the Intel® 80200 processor drives the data onto the data bus. For a write whose data all fits within one aligned 64-bit memory block (for a 64-bit bus) or a 32-bit aligned memory block (for a 32-bit bus) there is a single data cycle. For a burst write, there can be up to two (64-bit bus) or four (32-bit bus) data cycles, for a maximum write burst of four 32-bit words.

There are eight byte enables (**BE#**) associated with the **D** bus. Each byte enable corresponds to one byte of the bus. During a write cycle, the byte enables for each byte that is being written is asserted low. More detail on write transactions are given below.

Eight check bits, **DCB,** are also provided as part of the data bus. These bits are used for ECC. Section 10.2.7, "ECC" on page 10-12, has more information on how the Intel® 80200 processor uses ECC.

The data bus pins **D**, **BE#**, and **DCB** are not driven by the Intel® 80200 processor except when explicitly requested by a **DValid** assertion for a write two clock edges earlier. This means that when the chipset is not getting write data from the Intel® 80200 processor it can use that bus for other purposes or allow its use by other masters.

Between a read and a write data cycle on the data bus there should be one turnaround cycle to avoid contention on the bus. The Intel® 80200 processor expects this cycle and system operation is not guaranteed without it. It is up to the chipset or SDRAM controller to control the data bus cycles.

Because the Intel® 80200 processor transactions on the data bus must occur in the order they were requested, **DValid** can be used for both read and write data cycles. Both the Intel® 80200 processor and the chipset have enough information to know if the Intel® 80200 processoris driving or sampling the data bus for any given transaction.

10.2.4 Configuration Pins

DBusWidth, which is on the **CWF** pin at reset, indicates the data bus is either 32 bits wide or 64 bits wide. If the pin is sampled as '0' during reset, the Intel® 80200 processor assumes a 64-bit bus. If the pin is '1' at reset, a 32-bit bus is assumed.

10.2.5 Multimaster Support

Simple multimaster support is supplied with the **Hold** pin. The **Hold** pin causes the Intel® 80200 processor to stop issuing new requests as soon as possible (see below for timing) and to float the following pins: **A**, **ADS#/LEN[2]**, **W/R#/LEN[0]**, and **Lock/LEN[1]**. Before floating **ADS#**, the Intel® 80200 processor drives it to an inactive state (high).

Simultaneously with floating the affected signals, the Intel® 80200 processor asserts **HldA**.

When **HldA** is asserted, it is up to the chipset to make sure that the floating signals are driven or sustained.

Once asserted, the level on the **Hold** pin must be maintained until **HldA** has been asserted by the Intel® 80200 processor. If **Hold** is deasserted before **HldA** asserts, the results are unpredictable.

The delay to asserted **HldA** after **Hold** is sampled may vary depending on the state of the bus. If **Hold** is asserted on clock edge *n*, the Intel® 80200 processor guarantees it asserts **HldA** by clock edge *n+6*.

The Intel® 80200 processor drives the floated signals two cycles after **Hold** is deasserted. That is, if **Hold** is deasserted at clock edge *n*, then the Intel® 80200 processor drives the floating signals to a valid level on clock edge *n+2*. The implication of this is that the signals are not carrying valid data that may be sampled until clock edge *n+3*. **HldA** is deasserted at the same time the signals are taken out of float.

Once **Hlda** is deasserted, external hardware must wait at least one cycle before asserting **Hold** again. That is, **Hold** may only be asserted at cycle *n* if **Hlda** is not asserted at cycle *n-1*.

If the Intel® 80200 processor is in a low power mode (see Section 8.3, “Power Management” on page 8-5) then the timing for entering and exiting hold mode may be faster than described above. If Hold functionality is desired during one of these modes, the MCLK clock must be toggling.

During reset, the Intel® 80200 processor honors requests for Hold.

Note that the data bus is not affected at all by the **Hold** pin or the floating of the request bus. While the Intel® 80200 processor is held off the issue bus, data can continue to be returned to the Intel® 80200 processor on the data bus, and write data can be requested from the Intel® 80200 processor by the memory controller. When write data is not being requested by the chipset from the Intel® 80200 processor, the Intel® 80200 processor automatically floats the data bus and associated signals: **D**, **BE#**, **DCB**. This means that write data from another master or read data to another master can be driven onto the data buses without informing or requesting permission from the Intel® 80200 processor. The chipset owns the data bus and completely controls who gets to drive it.

Care must be taken, however, to not assert the Intel® 80200 processor **DValid** pin any time other than two cycles before the next valid Intel® 80200 processor data cycle.

This system allows memory accesses from the Intel® 80200 processor to be pipelined with memory accesses from another master. An Intel® 80200 processor memory access can be issued, followed by a memory access from another master, followed by another Intel® 80200 processor memory access, all before the data cycles of the first access begin. The data cycles for those transactions can then occur sequentially (except for any required turnaround cycles) on the data busses. This would of course require that the other master and the chipset supported this pipelining.

A simpler but lower performance method would be to assert **Hold** to the Intel® 80200 processor, wait for all outstanding transactions to complete, grant the issue bus to the alternate master (using the issue bus pins with the Intel® 80200 processor bus protocols, or whatever protocol the alternate master required) and give the bus back to the Intel® 80200 processor only once the alternate bus master is completely finished.

The **Lock** signal is active in transactions which require atomicity on a read/write pair. The minimum level of memory granularity over which a lock can be held should be at least 32 bytes. That is: if a master on the bus asserts **Lock** on a particular address, the naturally aligned 32-byte block that contains that address should be considered protected from access by other bus masters. It is permissible, of course, to consider all of memory locked and to hold off all accesses by other bus masters.

The interaction of **Hold** with the **Lock** signal is interesting. If the **Lock** pin is asserted by the Intel® 80200 processor, the Intel® 80200 processor is executing an ARM* *Swap* instruction, which does a read from a memory location followed atomically by a write to the same location. This is used for updating semaphores in shared memory. Until a request appears that does not have the **Lock** signal asserted, the chipset should not allow accesses of any kind to the memory location of the read that asserted the lock.

It is possible that the chipset may assert **Hold** to allow another master on the bus on clock edge *n* and the Intel® 80200 processor issues a read request and asserts **Lock** on the same clock edge *n*. In this case, the chipset should not let another master access memory at this time -- the Intel® 80200 processor is stalled waiting for access to the bus. However, the Intel® 80200 processor continues to respect the **Hold** pin and floats the request bus as it normally would. This allows the chipset to have a guaranteed delay between **Hold** assertion and the Intel® 80200 processor floating the pins.

In the general case, the **Hold** pin should be deasserted a cycle or two later (speed here is not critical, as long as no other master is allowed to initiate a memory request) and the Intel® 80200 processor continues on with the atomic pair of requests. Once the write request that deasserts **Lock** is issued, the chipset can reassert **Hold** and give the bus to another master.

If the system designer knows that the other requesting master is not accessing the same 32-byte memory region as the locked read, the chipset may choose to not deassert **Hold**, and can continue on with the multimaster request.

Another possibility is for the chipset to accept the read with **Lock** request and store it into the chipset queues, but to delay execution of the read with **Lock** until after the transactions from the other bus master to avoid a semaphore conflict.

Any of these strategies work as long as there are no accesses to the 32-byte memory region of the locked read after the read has executed and before the next write request is executed (which deasserts **Lock**).

10.2.6 Abort

If for any reason a request made by the Intel® 80200 processor can not be completed, it must be aborted. At the same time as the assertion **DValid** for any data cycle of any transaction, **Abort** can be asserted. This has the effect of ending that transaction at that data cycle. The Intel® 80200 processor saves the address of the aborted transaction and take an exception.

If **Abort** is asserted at the same time as **DValid** of the first data cycle of a given bus transaction, no data is sampled or driven by the Intel® 80200 processor off of the **D** bus two cycles after the **Abort** is asserted. That transaction is finished as soon as the **Abort** signal is sampled and no further data is transferred.

On a transaction with multiple data cycles, **Abort** can be asserted along with **DValid** at any time during the transaction. Data is read from or driven to the data bus two cycles after each of the pre-abort **DValid**s.

For a write transaction, no data is driven at the clock edge two cycles after **Abort** is sampled. For a read transaction, external logic must drive the data buses (**D** and **DCB**) to a valid level two cycles after **Abort** is sampled. At that point the transaction is cancelled and no further data or **DValid**s are expected for that transaction.

For burst transactions, care must be taken to end the transaction when the **Abort** is asserted. If a cache line fill (four data cycles on a 64-bit bus) has **Abort** asserted along with the first **DValid**, the Intel® 80200 processor ends the transaction and expect that none of the other data cycles for the burst occurs. If the chipset or memory were to assert **DValid** at that point and continue to send data, the Intel® 80200 processor would assume that data was associated with the next read request after the cache line fill that was aborted. If there had been no read request after the cache line fill, those extra data cycles would cause unpredictable results.

Abort must not be asserted for two consecutive cycles. In other words, back-to-back aborts are not permitted, and causes incorrect operation if they occur. The Data Bus portion of the bus must have a dead cycle after any abort cycle. That is: **DValid** must not be asserted the cycle after **Abort** has been asserted.

10.2.7 ECC

Software running on the Intel® 80200 processor may configure pages in memory as being ECC protected. For such pages, the Intel® 80200 processor checks the ECC code associated with read data, and generates an ECC code to associate with write data. The ECC code for a data element is transported over the **DCB** bus. For 64-bit wide memories, an additional eight bits of width are required to hold the ECC code. ECC is not supported in 32-bit memory systems.

If system software has indicated a memory region is ECC-protected, then the Intel® 80200 processor will always process that memory at bus-width granularity (64-bits). Read data should only be returned in multiples of 64-bits (8-, 16-, and 32-bit read requests should be answered with a full 64 bits of data). Narrow writes to ECC-protected memory will be seen as read-modify-write cycles of bus-width granularity, with **Lock** asserted during the transaction.

If the Intel® 80200 processor is accessing a region of memory for which ECC is not enabled, it drives zeroes on **DCB** during writes; receivers should ignore such values. On reads to such a memory region, the Intel® 80200 processor ignores the value on **DCB**, but this bus should be driven to a valid level (all zeroes for example).

Note for 32-bit data bus systems: the system can satisfy this requirement on **DCB** by tieing its bits low through pull-down resistors.

To ensure that a memory location has a valid ECC byte, software must write to an address before it ever reads from it. This would not seem a problem, but the Intel® 80200 processor aggressive memory architecture can make it a challenge. When the data cache receives a write request, it may respond by reading the associated line from memory, and then updating the specific address (data cache operation is described in Chapter 6, “Data Cache”). This line-fill operation could be the first access to that memory. So, a write to memory could trigger an ECC fault!

The solution is to write zeroes to memory before ECC is enabled. The Intel® 80200 processor writes zero into the **DCB** bus for this kind of transaction, and the ECC code for 64-bits of zero is zero.

It is possible for the Intel® 80200 processor to perform a write in which all byte-enables are not asserted (i.e., **BE#**[7:0] = “11111111”). Systems that contains data memory for ECC should interpret this to mean no data bytes or ECC bytes should be updated in response to the write. Because the Intel® 80200 processor only ever performs bus-width sized writes in ECC systems, a simple solution to this requirement is to wire **BE#**[0] to the byte-enable of both the LSB and the ECC byte.

To summarize the behavior of **BE#** on ECC-protected memory:

- Normal writes are 64-bits wide, so all **BE#** lines are asserted
- Reads do not activate the **BE#** lines, and should be answered with at least 64 bits of data
- If an ECC RMW operation detects an error on the read phase, it deasserts all **BE#** lines during the write phase.

10.2.8 Big Endian System Configuration

The Intel® 80200 processor supports execution in a big endian system. A system is said to be big endian if multi-byte values are accessed with the MSB at lower addresses. The endian orientation of a system is only evident when software performs sub-word sized accesses.

To operate an Intel® 80200 processor in a big endian system, software running on the Intel® 80200 processor must configure the device appropriately, and the board hosting the Intel® 80200 processor must swap the byte lanes in each word of the **D** bus. The requisite arrangement for a 64-bit bus is shown in Figure 10-3.

Figure 10-3. Big Endian Lane Swapping on a 64-bit Bus

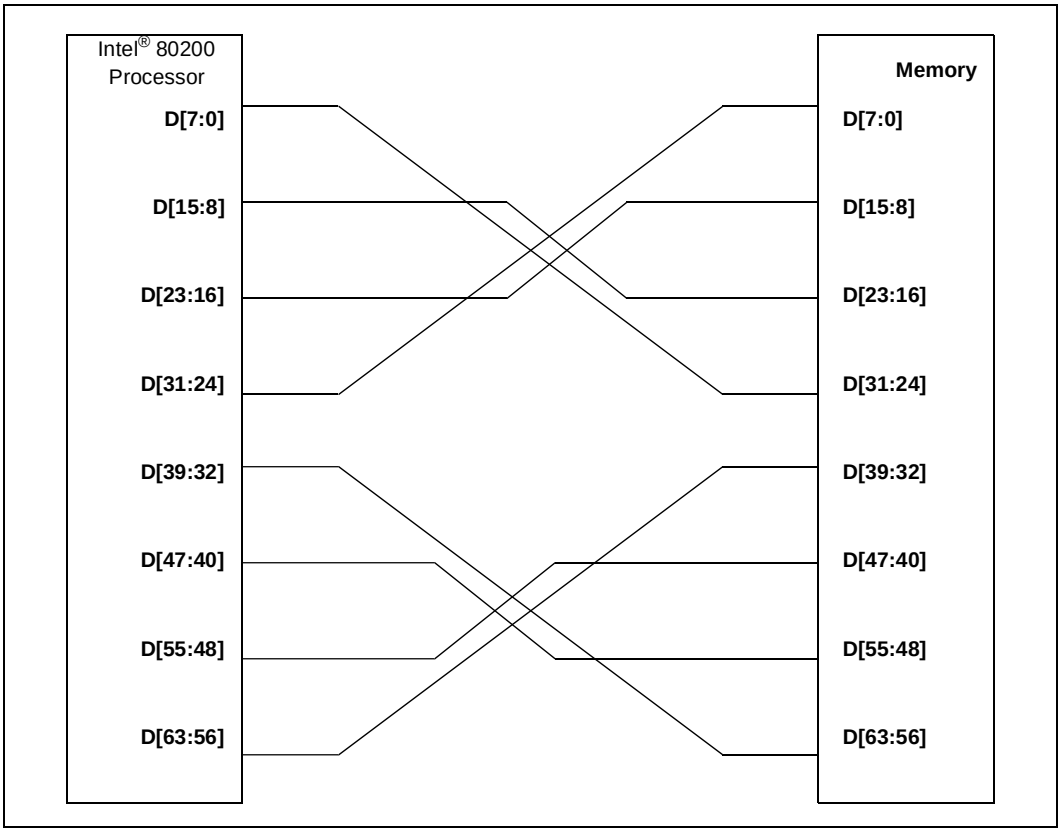

The lines in **DCB** should be wired the same for either endian configuration. Entities that create and check ECC should always interpret the contents of memory in the system format (either big- or little- endian).

The processor uses the Byte Enable lines (**BE#**) to indicate valid bytes during writes. In big-endian systems, these lines should be swapped also. That is, bits in **BE#**[3:0] should be swapped, and bits in **BE#**[7:4] should be swapped. This has the effect of always associating **BE#**[0] with **D**[7:0], **BE#**[1] with **D**[15:8], and so on.

10.3 Examples

All examples assume a 64-bit bus, in a little endian system.

10.3.1 Simple Read Word

In Figure 10-4, a read request for one word at address 0x240 is issued at time 10 ns. ADS# is asserted low at that clock edge, 0x240 is driven on A, W/R# is driven low to indicate a read request, and 0x2 is driven onto the Len bus to indicate that the access if for four bytes. Some time later (four clocks in this case), DValid is asserted to indicate the next sequential data cycle is occurring. Two clock edges later the data word from 0x240 is driven onto D[31:0]. The other half of D can be any value. The ECC or Parity bits associated with the data are driven onto DCB at the same time as the data.

Figure 10-4. Basic Read Timing

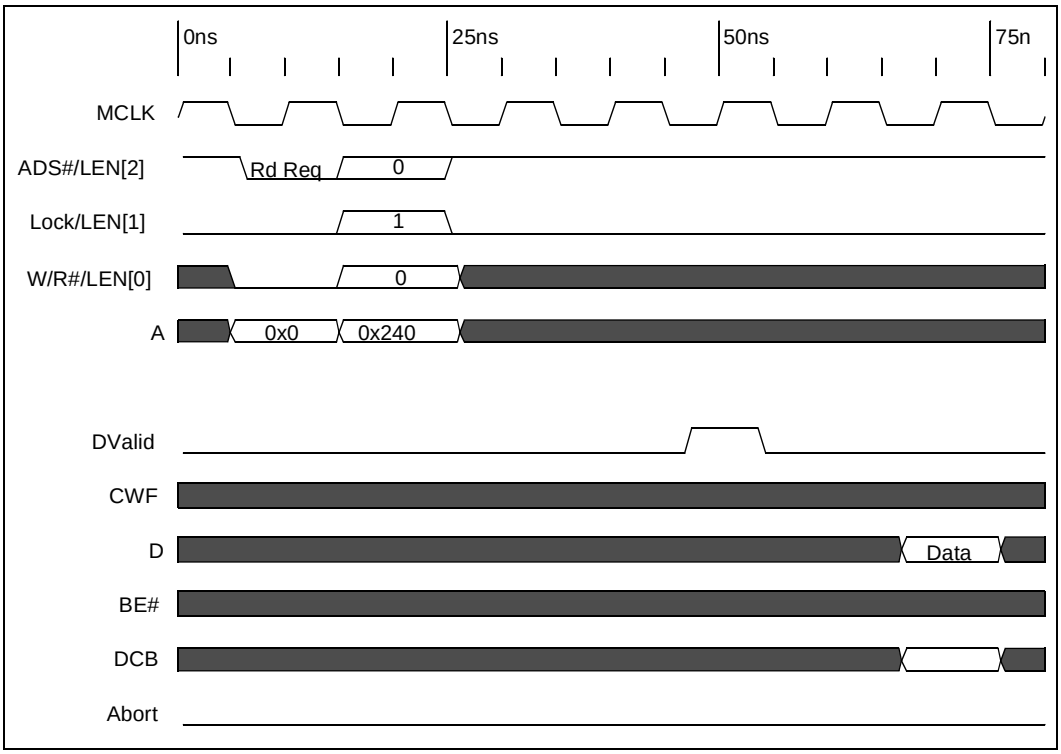

10.3.2 Read Burst, No Critical Word First

In Figure 10-5 the request goes out the same as the last example, with the address 0x248 this time and the length 0x6, indicating an eight word cache line fill. The first data cycle begins at 50 ns with **DValid** being asserted with **CWF** low to indicate that this burst starts at the lowest word pair and return sequentially. Notice that the data returning is the eight word block beginning at 0x240, not 0x249. The low five bits of the address are not used for determining what data to return in a 32-byte read request (cache line fill). **DValid** stays high for three cycles, drops for a cycle, and then is asserted for one more cycle.

Two clocks after each clock where **DValid** is sampled high, the next sequential pair of data words is driven on the bus. This data can come back to back in sequential cycles (as it likely would from a burst SDRAM, for example), or can be spaced further apart, as shown by the last data cycle here. Each data cycle in the transaction is independent of the others in timing, as long as the order of cycles in the transaction is maintained consistent with the **CWF** value asserted on the first data cycle.

Figure 10-5. Read Burst, No CWF

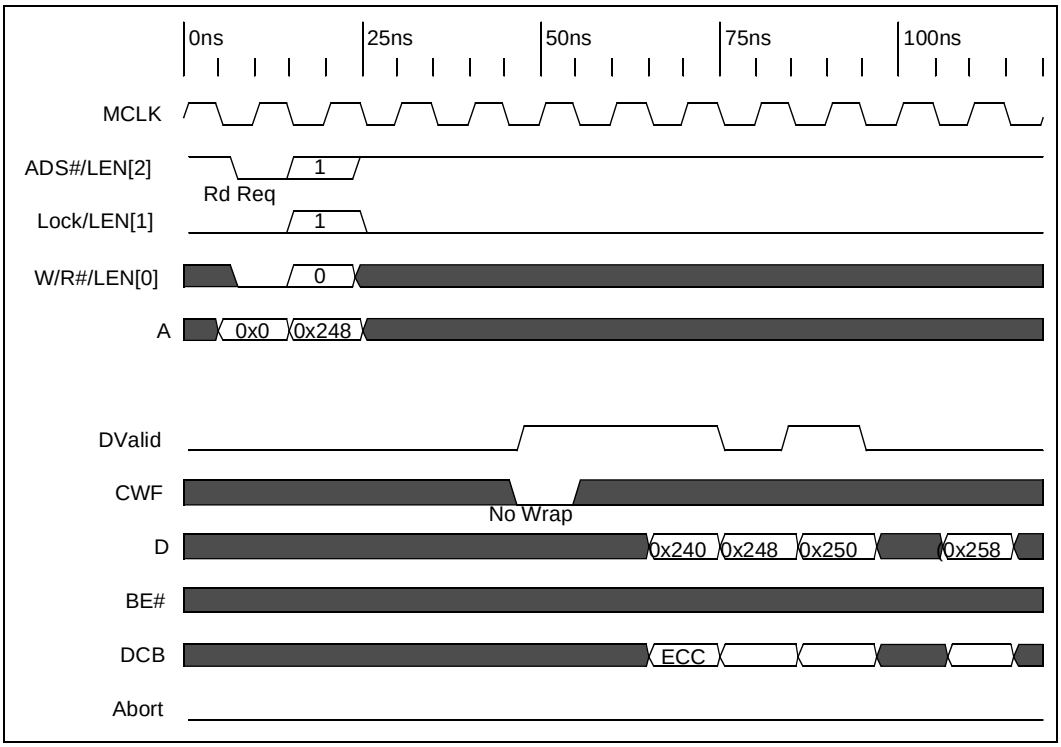

10.3.3 Read Burst, Critical Word First Data Return

Figure 10-6 is the same as the last with one difference: **CWF** is asserted high on the first data cycle of the return data. This indicates that the data is returning critical word first. In this case, since the address requested was 0x248, the word pair containing that byte starting at 0x248 is returned first. The data is then returned sequentially through the end of the cache line and starting over at the beginning.

Figure 10-6. Read Burst, CWF

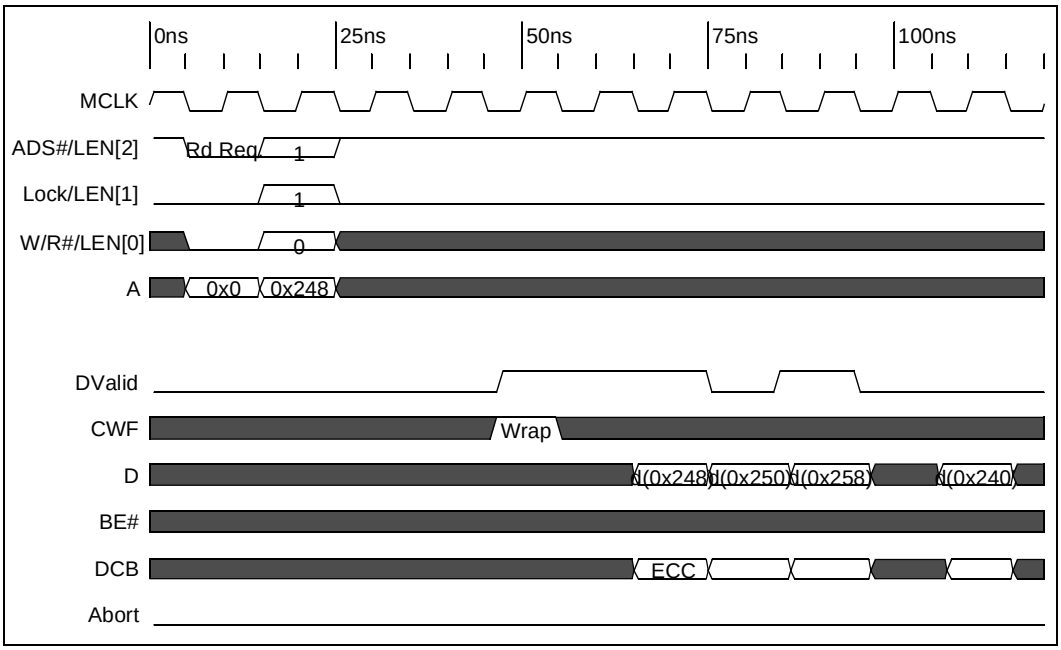

10.3.4 Word Write

Figure 10-7 shows a 32-bit write request to address 0x240. **W/R#** is high when **ADS#** is asserted low. Two cycles before the write data needs to be on the bus for the SDRAM, **DValid** is asserted by the chipset to the Intel® 80200 processor to tell the Intel® 80200 processor the data is needed. Two cycles later the Intel® 80200 processor drives the data onto the **D** bus (the lower 32 bits in this case) along with the appropriate check bits and byte enables. In this case the low four byte enables are asserted (low) and the upper four are deasserted (high) because the write data is on the low four bytes of the bus.

Figure 10-7. Basic Word Write

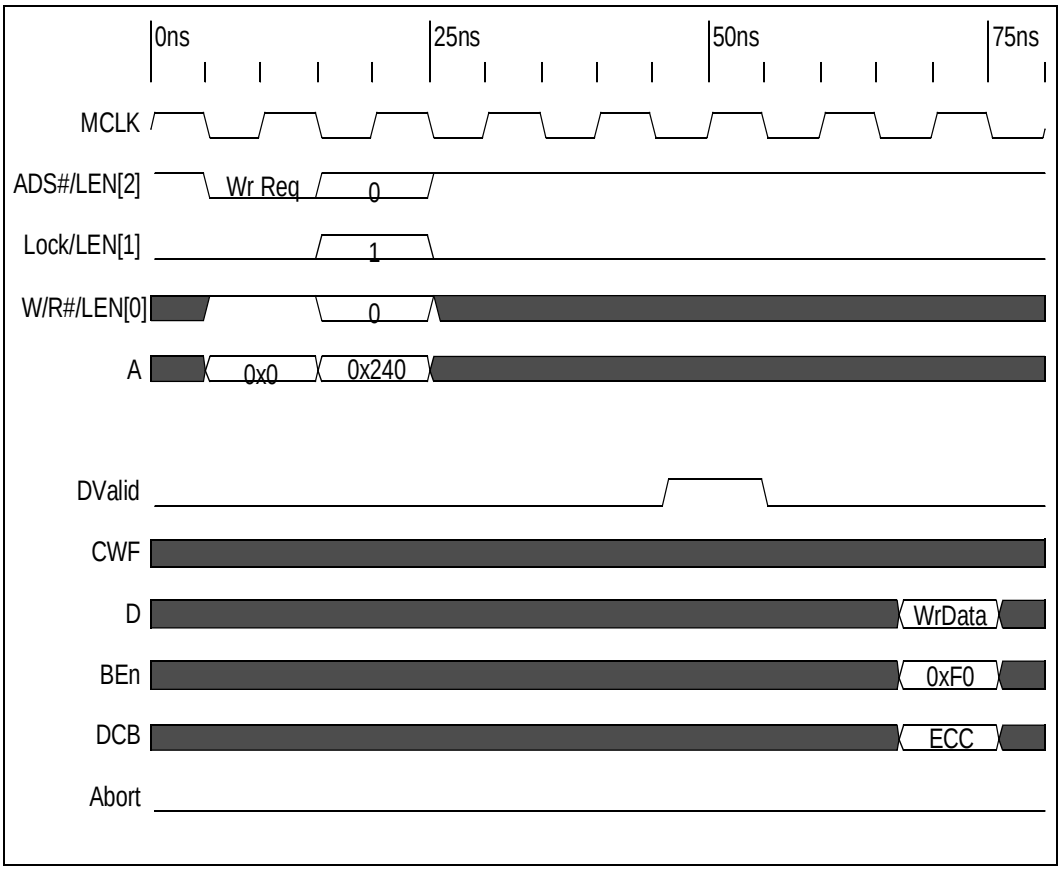

10.3.5 Two Word Coalesced Write

In Figure 10-8, two store byte instructions from the instruction stream have been coalesced into a single write command in the write buffer. The bytes were stored to addresses 0x240 and 0x247. The request is the same as the basic write word case except now the length is 0x3, indicating a two word write. When the chipset or memory needs the data, **DValid** is asserted and two cycles later the data is driven. In this case, however, only **BE#** bits 0 and 7 are asserted, indicating that the first and last byte of the bus have valid data to be stored and the rest must not be written.

Figure 10-8. Two Word Coalesced Write

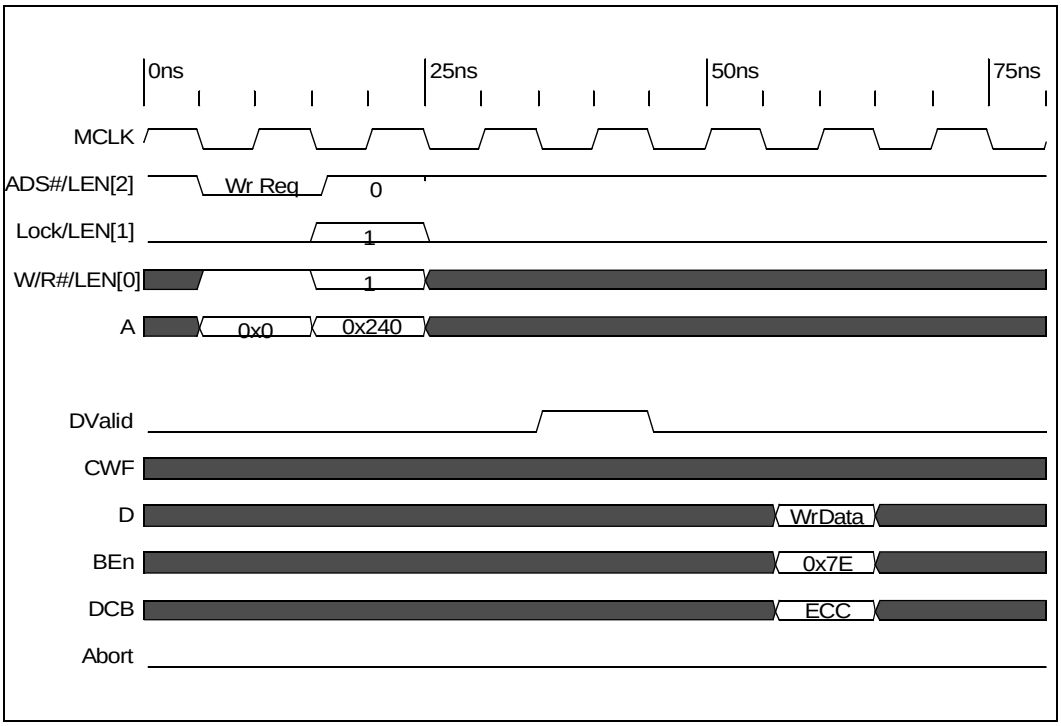

10.3.5.1 Write Burst

Figure 10-9 shows a four word write caused by the eviction of a half cache line. In this case, the **Len** is 0x5 indicating four words. **DValid** is asserted for two consecutive cycles here, but the two cycles could be spread out. In this case the Intel® 80200 processor drives the data as requested, along with **BE#** of 0x00 each cycle, indicating that all the bytes are being written.

Figure 10-9. Four Word Eviction Write

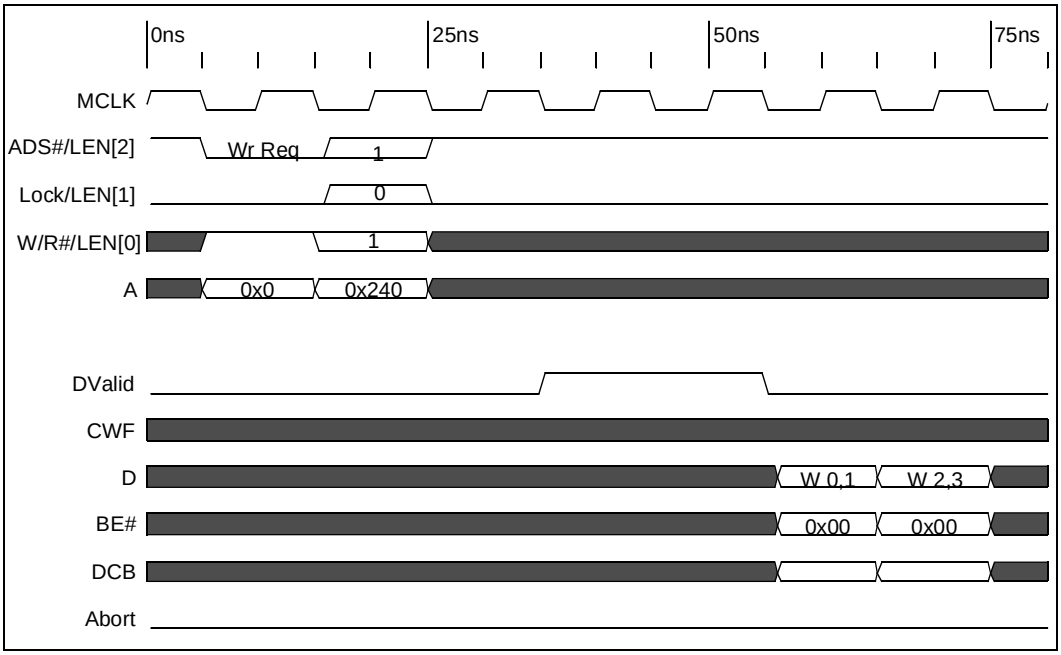

10.3.6 Write Burst, Coalesced

Figure 10-10 shows a four word cache write caused by store requests coalesced in a write buffer. The **Len** is 0x5 indicating four words. **DValid** is asserted for two consecutive cycles. The Intel® 80200 processor drives the data as requested, but this time the byte enables are not all zeroes. The byte enables here are asserted low only for those bytes that were stored by the instruction stream. Any possible combination of byte enables can occur, with only the requirement that the first and last data cycles have at least one byte enable asserted. (If the first or last data cycle had no bytes to store, the transaction would have been issued as a shorter transaction).

Figure 10-10. Four Word Coalesced Write Burst

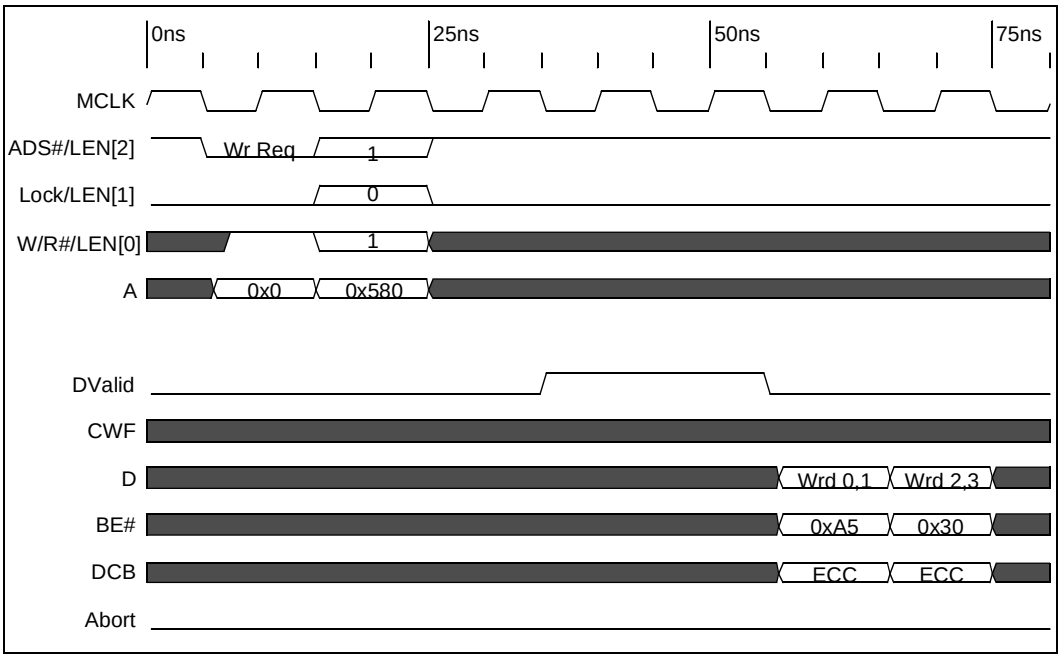

10.3.7 Pipelined Accesses

The example in Figure 10-11 demonstrates the four deep pipelined nature of this bus. In this example, the Intel® 80200 processor is bus limited and is issuing requests as quickly as it can. Before time 0ns, there are no outstanding transactions. Two reads (A and B) followed by a write (C) and another read (D) are all requested before 85 ns in this timing diagram.

Because the Intel® 80200 processor may have up to four outstanding transactions, and in this example only three are outstanding at time 85 ns, it can send another request (for E) at time 90 ns. If none of the transactions had completed, the E transaction would have been delayed.

Figure 10-11. Pipeline Example

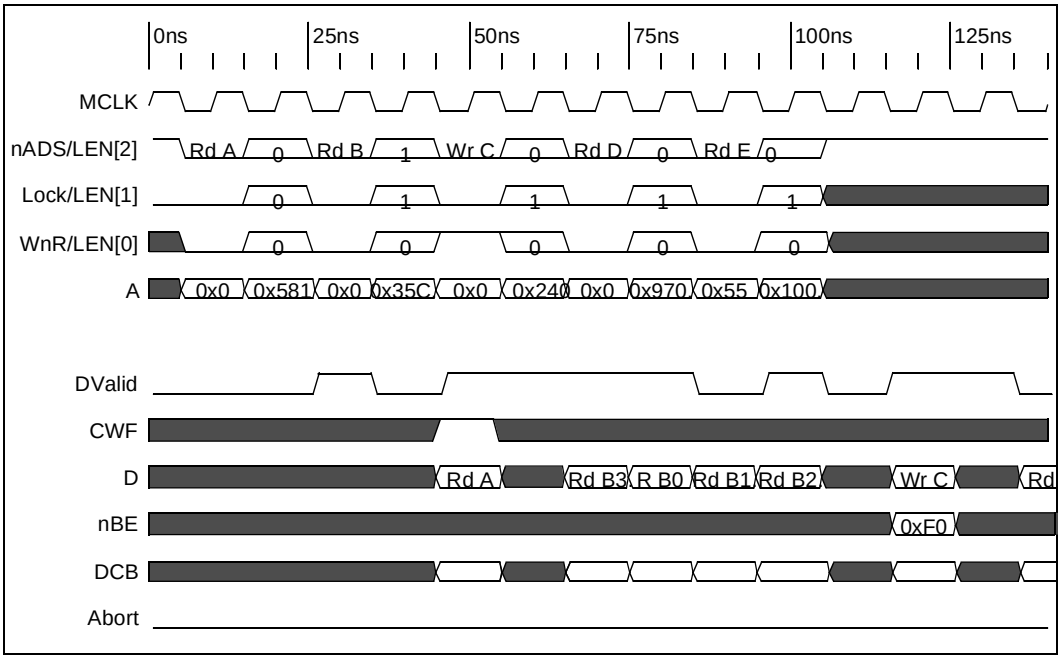

10.3.8 Locked Access

An example of a locked access is shown in Figure 10-12. Here the processor is doing an atomic read/write to address 0x240, denoted as *A* in the figure. The **Lock** signal, which is valid at the positive edge of **MCLK** when **ADS#** is asserted, are asserted for each request from the read of *A* just prior to the matching write of *A*.

Note that an intervening read cycle to location 0x12349988 also occurs, perhaps to read a page table element. This is legal behavior and must be accommodated by all entities on the bus. The bus makes no guarantee as to how many cycles may elapse between a locked read and its corresponding locked write. It is guaranteed that no writes intervenes during that period, although an arbitrary number of reads may occur.

Figure 10-12. Locked Access

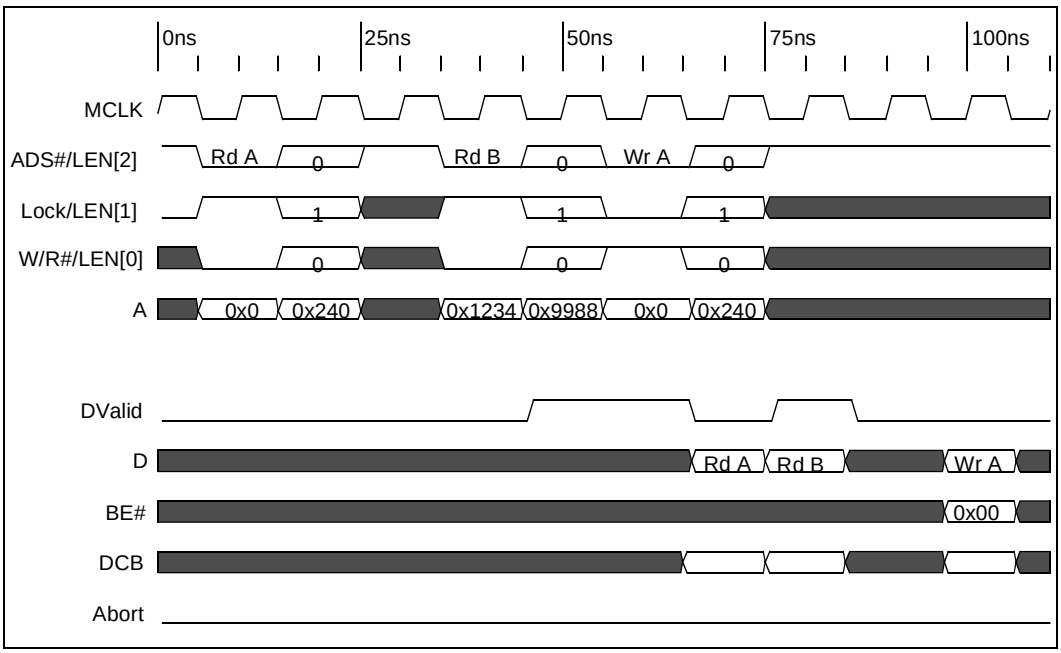

10.3.9 Aborted Access

As discussed in Section 10.2.6, “Abort” on page 10-11, any request from the Intel® 80200 processor can be aborted by the chipset or memory. This might occur if there was a PCI error, or if a request was issued to unimplemented memory. Figure 10-13 shows an aborted read. Read A is issued at time 10 ns for 32 bytes of data. At 50 ns **DValid** goes high to indicate the beginning of the first data cycle. **Abort** is not asserted along with this **DValid** assertion. At 60 ns, **DValid** for the second data cycle is asserted, and this time **Abort** is asserted. This indicates that this transaction is aborted and no more data cycles begin for this transaction. Notice, however, that the first data cycle begun with the **DValid** assertion at 50 ns is still ongoing, and the Intel® 80200 processor latches that data and return it to the core as valid data.

If transaction A was a write request rather than a read, the Intel® 80200 processor would drive data onto the bus at time 70 ns, as requested by the **DValid** with **Abort** NOT asserted at 50 ns. No data would be driven at time 80 ns because the **DValid** at time 60 ns had **Abort** asserted high.

When the **Abort** data cycle occurs, the Intel® 80200 processor ends that transaction and expect no further data on it. When **DValid** next goes high at time 80 ns, the Intel® 80200 processor expects a data cycle associated with the next transaction, in this case the first data cycle of read B.

If an aborted data burst cannot be stopped by the memory system, it is sufficient to allow it to complete with **DValid** deasserted. This is a slight bandwidth hit, but aborts should be rare.

If, during a locked access, any read access encounters an **Abort**, the Intel® 80200 processor still emits the final unlocking write. If the final unlocking write is an **Abort**, the bus is still unlocked with the write. The Intel® 80200 processor never fails to deassert **Lock**.

External logic must leave a gap of at least one **MCLK** cycle between aborts to the Intel® 80200 processor.

Figure 10-13. Aborted Access

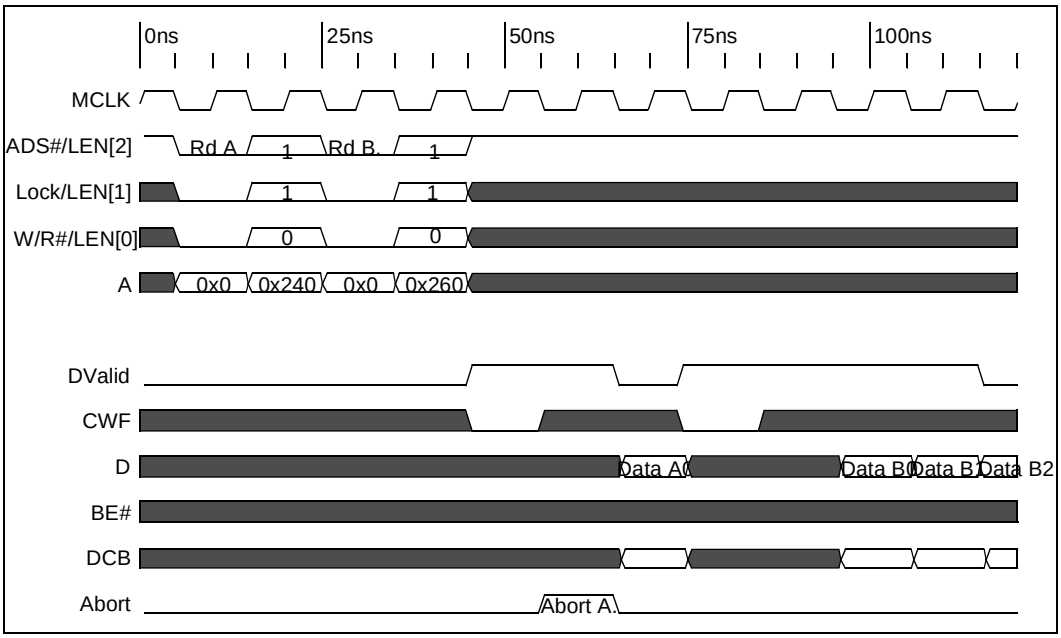

10.3.10 Hold

Figure 10-14 shows an example of hold being asserted to stop new transactions being issued. The Intel® 80200 processor floats the issue bus pins and issues no transactions until **HldA** is deasserted.

The **Hold** signal assertion does not affect the data bus, which continues to operate normally. Read data for requests A and B continue to return.

"Rd C" shows the Intel® 80200 processor requesting another access after **Hold** has been deasserted. If **Hold** had continued to be asserted, another bus master could take control of the request bus.

Figure 10-14. Hold Assertion

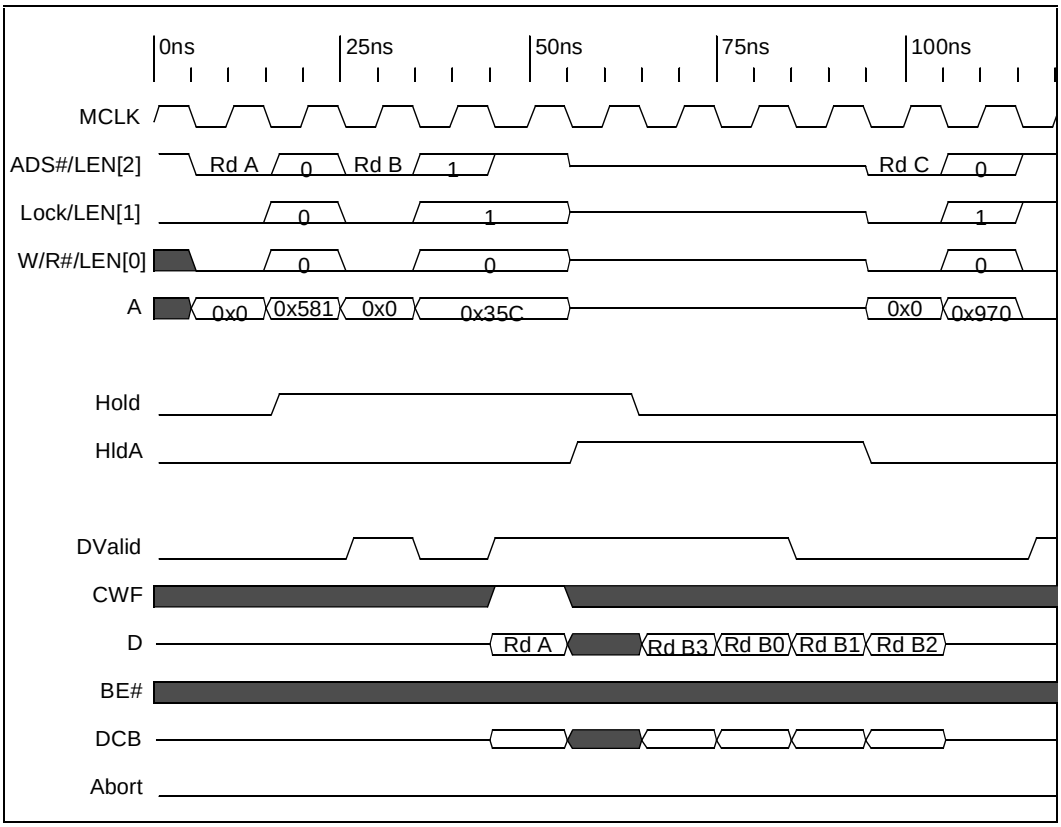

Bus Controller 11

11.1 Introduction

The Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) Bus Controller Unit (BCU) is responsible for accessing off-chip memory. It initiates bus cycles as documented in Chapter 10, “External Bus”.

The BCU is capable of queuing four outstanding transactions. This improves the performance of the processor, because it does not need to wait for the result of a memory transaction before initiating another.

If enabled by software, the BCU can protect data with an Error Correcting Code (ECC).

The BCU has software-accessible state in the form of coprocessor registers. All BCU registers reside in Coprocessor 13 (CP13).

11.2 ECC

System software has the ability to request ECC checking on memory accesses. If the MMU determines that a region of memory is protected by ECC, the BCU is responsible for checking and generating ECC on accesses to that region. See Chapter 3, “Memory Management”, for a description of the MMU.

The Intel® 80200 processor data bus width is configured at reset-time to either 32 or 64 bits. Bus configuration is detailed in Chapter 10, “External Bus”. The Intel® 80200 processor only supports ECC in the 64-bit mode.

When ECC is enabled for a memory region, the BCU never performs sub bus-width (64 bits) writes to that region. If directed by the core to perform a sub bus-width write, the BCU performs a bus-width read, merge in the appropriate bytes, and then perform a bus-width write with all byte-enables asserted. This read-modify-write (RMW) is performed as an atomic transaction on the external bus. This RMW behavior affects performance and can be avoided by specifying the region as read/write-allocate and write-back cacheable in the MMU.

When writing to ECC-protected memory, the BCU always calculates the correct ECC bits and writes them to memory at the same time as the data.

The ECC algorithm supported by the Intel® 80200 processor uses eight bits to protect the data bus. It can detect and correct any one-bit error, and it can detect any two-bit error.

Whenever the BCU checks ECC, it generates a *syndrome* -- a bitwise exclusive-OR of the expected ECC versus the actual code received on the bus. Should an error be detected, this syndrome is available to software to aid diagnosis.

11.3 Error Handling

The BCU is able to detect and respond to two classes of errors: bus aborts and ECC errors.

Information about errors is captured in a set of programmer-accessible registers: ELOG0, ELOG1, and ECAR0, ECAR1. The ELOGx registers log general information about an error, while the ECARx registers capture the address associated with an error.

11.3.1 Bus Aborts

A bus abort occurs when the **Abort** pin is asserted during an external bus transaction. This is described in detail in Chapter 10, “External Bus”. In response to a bus abort, the BCU causes an exception in the currently running software. Note that if the exception raised is an External Data Abort, then it is an *imprecise* exception -- it is not necessarily related to the instruction that just executed.

The BCU attempts the additional response of logging the error into a register. If register ELOG0 is not already being used to record an error, the BCU logs error information into registers ELOG0 and ECAR0. If ELOG0 already has error information, but ELOG1 does not, the BCU use ELOG1 and ECAR1 to log the error. If both ELOG0 and ELOG1 are in use, the BCU sets the error overflow bit (register BCUCTL, bit EV). A description of these registers is in Section 11.4.1 and Section 11.4.2.

See Chapter 2, “Programming Model”, for more discussion of Data Aborts.

11.3.2 ECC Errors

An ECC error occurs when the BCU reads data and notices that the associated ECC bits do not match the data. This could also happen as a result of the RMW that the BCU performs on sub bus-width writes. A single transaction on the bus could result in multiple ECC errors, as the BCU checks each bus-width entity as it is received. Table 11-1 summarizes the BCU error response for ECC errors. The response is tailorable; some of the actions may be disabled, such as requesting a core interrupt. See Section 11.4.1, "BCU Control Registers" on page 11-5, for information on the **EE**, and **SR** bits, and how the BCU error response may be altered.

Table 11-1. BCU Response to ECC Errors

Event	EE	SR	Response
Read with 1-bit error	0	-	No correction, no notification
	1	0	• Correction if BCUCTL.SC = 1
	1	1	• Correction if BCUCTL.SC = 1 Request core interrupt
Read with multi-bit error	0	-	No notification
	1	-	• If transaction is a data read: imprecise data abort • If transaction is an instruction fetch: prefetch abort • If transaction is an MMU operation: precise data abort
RMW with 1-bit error from the read cycle	0	-	Does not occur -- RMW only if BCUCTL.EE = 1
	1	0	• Correction if BCUCTL.SC = 1 • Merge store data with data from read • Write updated data to memory
	1	1	• Correction if BCUCTL.SC = 1 • Merge store data with data from read • Write updated data to memory • Request core interrupt
RMW with multi-bit error from the read cycle	0	-	Does not occur -- RMW only if BCUCTL.EE = 1
	1	-	• Merge store data with data from read • Write updated data to memory with deasserted byte enables • If transaction is a data read: imprecise data abort • If transaction is an instruction fetch: prefetch abort • If transaction is an MMU operation: precise data abort

In all cases, the BCU attempts the additional response of logging the error into a register. If register ELOG0 is not already being used to record an error, the BCU logs error information into registers ELOG0 and ECAR0. If ELOG0 already has error information, but ELOG1 does not, the BCU uses ELOG1 and ECAR1 to log the error. If both ELOG0 and ELOG1 are in use, the BCU sets the error overflow bit (register BCUCTL, bit EV). A description of these registers is in Section 11.4.1 and Section 11.4.2.

When an error is detected during a RMW, the BCU writes the corrupted data back to memory because it needs to finish the atomic transaction. If it did not do this, an external bus agent might consider the bus to be permanently locked. The BCU does not assert any of the byte enables (nBE) when writing the corrupted data back. A memory system should use this as an indication to not update the DCB value or any data bytes.

Error reporting may be enabled with the BCUCTL register, described in Section 11.4.1. If enabled, single bit errors cause the BCU to assert an interrupt to the Interrupt Controller Unit (ICU). If the interrupt is not enabled in the ICU, it is not propagated to the core. This interrupt may be cleared by software by writing to the BCU Control Register (see Section 11.4.1) If a masked interrupt is not cleared by software, it interrupts the core if software ever unmasks it.

Although single-bit errors can be corrected by the BCU, the software may choose to take the additional step of *scrubbing* the offending memory location. To accomplish this, the software needs to write the correct data back to the location. The BCU logs sufficient information in its registers (ELOGx and ECARx) to enable software to re-issue the load that uncovered the problem.

Single-bit errors detected during a RMW are corrected and written back if single-bit-correction is enabled. The BCU still reports the error if enabled, but software does not need to scrub the location.

Single bit errors may also be detected in the received ECC itself. In this case, the BCU doesn't need to modify the received data, but it still reports the error if enabled. Software can scrub the memory's ECC by simply writing to it. The BCU automatically generates and writes the correct ECC value.

If a multi-bit error is detected, the BCU requests an exception of the core. The exact exception triggered depends on the data's destination. For example, a code fetch that received a multi-bit error would result in a Prefetch Abort if the processor attempted to execute the code.

If the BCU receives a bus abort (see Section 10.2.6, "Abort" on page 10-11 and Section 11.3.1, "Bus Aborts" on page 11-2) and an ECC error on the same cycle, it ignores the ECC information for the cycle and process the bus abort normally.

Any ECC error that elicits a Prefetch or Data abort from the BCU is the final error associated with a data burst. The BCU accepts the remaining data from the bus, but ignores it and does not perform ECC checks.

11.4 Programmer Model

The BCU registers reside in Coprocessor 13 (CP13). They may be accessed/manipulated with the MCR, MRC, STC, and LDC instructions. The *CRn* field of the instruction denotes the register number to be accessed. Field *CRm* must be set to 1. The *opcode_1, and opcode_2* fields of the instruction should be zero. Access to CP13 may be controlled using the Coprocessor Access Register (see Section 7.2.13, "Register 13: Process ID" on page 7-15).

An instruction that modifies a BCU register is guaranteed to take effect before the next instruction executes.

11.4.1 BCU Control Registers

The BCU Control Register (BCUCTL) allows software to view and control the behavior of the BCU.

Table 11-2. BCUCTL (Register 0) (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
TP	EV	E1	E0																									EE	SC		SR

reset value: all implemented bits are 0

Bits	Access	Description
31	Read / Write-ignored	**TP** - Transactions Pending Indicates whether the BCU is idle 0 = no memory transactions pending 1 = one or more transactions pending
30	Read / Write	**EV** - Error Overflow Read Values: 0 = no unlogged errors have occurred 1 = errors have occurred beyond those logged in ELOG0 and ELOG1 Write Values: 0 = no change 1 = clear this bit
29	Read / Write	**E1** - ELOG1 is valid Read Values: 0 = contents of ELOG1 should be disregarded 1 = Error occurred and is logged in ELOG1 Write Values: 0 = no change 1 = clear this bit
28	Read / Write	**E0** - ELOG0 is valid Read Values: 0 = contents of ELOG0 should be disregarded 1 = Error occurred and is logged in ELOG0 Write Values: 0 = no change 1 = clear this bit
27:4	Read-unpredictable / Write-as-0	reserved

Table 11-2. BCUCTL (Register 0) (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
TP	EV	E1	E0																									EE	SC		SR

reset value: all implemented bits are 0

Bits	Access	Description
3	Read / Write	**EE** - ECC Enable 0 = disable ECC generation and checking 1 = enable ECC
2	Read / Write	**SC** - Single bit Correct Enable 0 = disable single bit error correction 1 = enable single bit error correction
1	Read-unpredictable/ Write-as-1	**Reserved.**
0	Read / Write	**SR** - Single bit Error Reporting Enable 0 = disable single bit error reporting 1 = enable single bit error reporting

BCUCTL.TP allows software to determine if the BCU has any pending memory transactions. This may be used to ensure that all memory operations have completed before attempting to modify system state. For example, the code in Example 11-1 simply waits until the BCU is idle.

Example 11-1. Loop to Wait on BCU

```
; Wait for BCU to finish all outstanding operations
waitLoop:
   MRC P13, 0, R15, C0, C1, 0 ; read BCUCTL, update condition code flags
   BMI waitLoop               ; try again if BCU is busy (bit 31 set)
;
; Get here when BCU is no longer busy
```

Of course, this sort of code should be in a cacheable region, or it may keep the BCU eternally busy fetching the code!

If BCUCTL.EE is set, then the BCU performs ECC generation/checking as described in Table 11-1. The other control bits (SC, SR) should only be modified while the EE bit is cleared.

To ensure correct operation, hardware waits until all pending operations are completed before allowing the EE bit to take affect. Code similar to that shown in Example 11-2 is recommended for enabling ECC.

Example 11-2. Enabling ECC

```
MACRO FLUSHALL
    MCR P15, 0, R0, C7, C10, 4     ; Drain buffers
ENDM

enableECC:
    FLUSHALL                       ; Finish pending memory operations
    MOV R0, #0xA                   ; Set bits 3 and 1
    MCR P13, 0, R0, C0, C1, 0      ; Set BCUCTL -- enable ECC
```

When ECC is enabled, the BCU only generates an interrupt on a single-bit error if BCUCTL.SR is set. When ECC is enabled, the BCU always generates an abort on a multi-bit error.

The BCU repairs single bit errors if BCUCTL.SC is set. It is recommended that this bit always be set; running with this bit cleared could cause software to operate on corrupted data before the ECC-error detect interrupt is received.

If BCUCTL.EE is zero, then the BCU ignores the ECC bits on reads and drives all zeroes to the ECC bits on writes. This is the fastest mode for the BCU, as ECC error detection/correction incurs an additional two MCLK cycles.

If error reporting is enabled in BCUCTL, and any of the bits BCUCTL.EV, BCUCTL.E1, BCUCTL.E0 are set, then the BCU asserts its interrupt to the ICU on a single-bit error. These bits are set by the BCU when it detects an error and saves the error information to ELOGx and ECARx. If BCUCTL.En is set (n = 0,1), the BCU does not alter the contents of registers ELOGn and ECARn. This ensures that software can use these registers until it clears BCUCTL.En.

These bits may be used by an ISR to quickly determine how many errors have occurred, and to locate the error information in other registers (ELOGx and ECARx). To clear the BCU's interrupt, an ISR must ensure that all of these bits are cleared. Pseudocode for a typical ISR that handles BCU Errors is shown in Example 11-3.

Example 11-3. Handling BCU Errors

```
if(BCUCTL.EV)
    ...         more errors that couldn't be logged -- react to them here
    BCUCTL.EV = 1          clear the EV bit -- error is handled

if(BCUCTL.E1)
    ...         react to error here
    BCUCTL.E1 = 1          clear the E1 bit -- error is handled

if(BCUCTL.E0)
    ...         react to error here
    BCUCTL.E0 = 1          clear the E0 bit -- error is handled
```

It is possible for more errors to occur while clearing the existing errors. In this case, the ISR is reinvoked as soon as BCU interrupts are unmasked.

The BCU Control Register (BCUCTL) allows software to view and control the behavior of the BCU.

Table 11-3. BCUMOD (Register 1)

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
																															AF

reset value: all implemented bits are 0

Bits	Access	Description
0	Read-unpredictable / Write	**AF** - Aligned Fetch 0 = On a 32-byte read, BCU requests an aligned block 1 = BCU can request 32-byte reads on any 4-byte aligned address.

BCUMOD.AF affects the behavior of the BCU when it is reading a 32-byte block (a cache line-fill). If this bit is ‘0’, then the BCU always emits the 32-byte aligned address of the cache line when requesting it. If this bit is ‘1’, then the BCU emits the address of the “critical word” in the cache line when requesting it. This latter setting allows external logic to implement CWF logic (as detailed in Section 10.2.3) which will usually yield higher performance.

11.4.2 ECC Error Registers

Table 11-4. ELOG0, ELOG1(Registers 4, 5)

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
RW	ET																							syn							

reset value: undefined

Bits	Access	Description
31	Read / Write-ignored	**RW** - indicates the direction of the errant transfer 0 = read error 1 = write error
30:29	Read / Write-ignored	**ET** - Error Type 00 = single bit ECC error 01 = multi-bit ECC error 10 = bus abort 11 = reserved
28:8	Read-unpredictable / Write-as-0	reserved
7:0	Read / Write-ignored	**syn** - the syndrome value that indicated the error. This field has an undefined value if ET is greater than 1.

Table 11-5. ECAR0, ECAR1(Registers 6, 7)

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
addr																															

reset value: undefined

Bits	Access	Description
31:0	Read / Write-ignored	**addr** - the physical address that yielded an error

The contents of these registers should only be considered valid if the corresponding bit in register BCUCTL is set. When an error is detected, the BCU selects a free ELOGx/ECARx register pair and updates it with information relevant to the error. It then sets BCUCTL.Ex.

ECARx holds the *physical* address associated with the transaction that caused the error. If the transaction was a multiple-cycle burst, then only the initial address is captured, the actual error occurred at some point during the burst. Because a burst can cover up to 32 bytes, software may only know that the error was in a range: ECARx .. ECARx + 31.

System software bears the burden of translating this address to a logical one, if needed. Since changes to the page tables may make this a non-trivial exercise, systems that respond to errors may wish to delay page table updates until the BCU is quiescent (as determined by the BCUCTL.TP bit).

For ECC errors, the lower bits of this register are always zero, because ECC always operates on 64-bit items. This means, ECC errors have an ECARx register in which bits 2..0 is zero.

For bus aborts, which could occur on byte-sized transactions, all address bits are recorded.

The BCU does not write to these ELOGx/ECARx registers unless the corresponding BCUCTL.Ex bit is cleared, either by reset or by software.

Table 11-6. ECTST (Register 8)

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
																								mask							

reset value: all implemented bits are 0

Bits	Access	Description
31:8	Read-unpredictable / Write-as-0	reserved
7:0	Read / Write	**mask** - When writing, the BCU exclusive-ORs this mask with the ECC code it generates

Software can generate data with incorrect ECC values for Validation purposes. By setting a bit in ECTST.mask to ‘1’, the corresponding bit in all generated ECCs is inverted during a write to memory. Subsequent reads of that memory generate an ECC error.

Performance Monitoring 12

This chapter describes the performance monitoring facility of the Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE). The events that are monitored can provide performance information for compiler writers, system application developers and software programmers.

12.1 Overview

The Intel® 80200 processor hardware provides two 32-bit performance counters that allow two unique events to be monitored simultaneously. In addition, the Intel® 80200 processor implements a 32-bit clock counter that can be used in conjunction with the performance counters; its sole purpose is to count the number of core clock cycles which is useful in measuring total execution time.

The Intel® 80200 processor can monitor either occurrence events or duration events. When counting occurrence events, a counter is incremented each time a specified event takes place and when measuring duration, a counter counts the number of processor clocks that occur while a specified condition is true. If any of the three counters overflow, an IRQ or FIQ is generated if it is enabled. (IRQ or FIQ selection is programmed in the interrupt controller.) Each counter has its own interrupt 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® 80200 processor clock counter can be used to measure the executing time of an application. This information combined with a duration event can feedback a percentage of time the event occurred with respect to overall execution time.

Each of the three counters and the performance monitoring control register are accessible through Coprocessor 14 (CP14), registers 0-3. Refer to Section 7.3.1, “Registers 0-3: Performance Monitoring” on page 7-18 for more details on accessing these registers with **MRC**, **MCR**, **LDC**, and **STC** coprocessor instructions. Access is allowed in privileged mode only.

12.2 Clock Counter (CCNT; CP14 - Register 1)

The format of CCNT is shown in Table 12-1. 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 causes it to roll over to zero and set the overflow flag (bit 6) in PMNC. An IRQ or FIQ is reported if it is enabled via bit 6 in the PMNC register.

Table 12-1. 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 causes it to roll over to zero and generate an IRQ or FIQ if enabled.

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

There are two 32-bit event counters; their format is shown in Table 12-2. 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 12-2. 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 causes it to roll over to zero and generate an IRQ or FIQ interrupt if enabled.

12.3.1 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 enables 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}).

12.4 Performance Monitor Control Register (PMNC)

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

- controls which events PMN0 and PMN1 monitors
- 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 12-3 shows the format of the PMNC register.

Table 12-3. Performance Monitor Control Register (CP14, register 0) (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
				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 12-4 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 12-4 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

Table 12-3. Performance Monitor Control Register (CP14, register 0) (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
				evtCount1								evtCount0									flag				inten			D	C	P	E

reset value: E **and** inten **are 0, others unpredictable**

Bits	Access	Description
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 64th 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

12.4.1 Managing PMNC

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

- An interrupt is reported when a counter overflow flag is set and its associated interrupt enable bit is set in the PMNC register. The interrupt remains asserted until software clears the overflow flag by writing a one to the flag that is set. Note that the interrupt unit (Chapter 9, “Interrupts”) 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.

12.5 Performance Monitoring Events

Table 12-4 lists events that may be monitored by the PMU. Each of the Performance Monitor Count Registers (PMN0 and PMN1) can count any listed event. Software selects which event is counted by each PMNx register by programming the evtCountx fields of the PMNC register.

Table 12-4. Performance Monitoring Events

Event Number (evtCount0 or evtCount1)	Event Definition
0x0	Instruction cache miss requires fetch from external memory.
0x1	Instruction cache cannot deliver an instruction. This could indicate an ICache miss or an ITLB miss. This event occurs every cycle in which the condition is present.
0x2	Stall due to a data dependency. This event occurs every cycle in which the condition is present.
0x3	Instruction TLB miss.
0x4	Data TLB miss.
0x5	Branch instruction executed, branch may or may not have changed program flow.
0x6	Branch mispredicted. (B and BL instructions only.)
0x7	Instruction executed.
0x8	Stall because the data cache buffers are full. This event occurs every cycle in which the condition is present.
0x9	Stall because the data cache buffers are full. This event occurs 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. This event occurs any time the PC is changed by software and there is not a mode change. For example, a **mov** instruction with PC as the destination triggers this event. Executing a **swi** from User mode does not trigger this event, because it incurs a mode change.
0x10	The BCU received a new memory request from the core.
0x11	The BCUs request queue is full. This event takes place each clock cycle in which the condition is met. A high incidence of this event indicates the BCU is often waiting for transactions to complete on the external bus.
0x12	The number of times the BCU queues were drained due to a Drain Write Buffer command or an I/O transaction as identified by C = 0 and B = 0 (cacheable and bufferable page attribute bits).
0x13	Reserved, unpredictable results.
0x14	The BCU detected an ECC error, but no ELOG register was available in which to log the error. (See Section 11.4.2, “ECC Error Registers” on page 11-9 for a description of the ELOG registers).
0x15	BCU detected a 1-bit error while reading data from the bus. This event may be counted even if reporting of 1-bit errors is disabled. See Section 11.3, “Error Handling” on page 11-2 for a description of 1-bit errors.
0x16	RMW cycle occurred due to narrow write on ECC-protected memory (see Section 11.2, “ECC” on page 11-1 for a description of ECC and RMW cycles).
all others	Reserved, unpredictable results

Some typical combination of counted events are listed in this section and summarized in Table 12-5. In this section, we call such an event combination a *mode*.

Table 12-5. Some Common Uses of the PMU

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

12.5.1 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.

12.5.2 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** each count as several accesses to the data cache depending on the number of registers specified in the register list. **LDRD** registers 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.

12.5.3 Instruction Fetch Latency Mode

PMN0 accumulates the number of cycles when the instruction-cache is not able to deliver an instruction to the Intel® 80200 processor due to an instruction-cache miss or instruction-TLB miss. This event means that the processor core is stalled.

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

Statistics derived from these two events:

- *The average number of cycles the processor stalled waiting for an instruction fetch from external memory to return.* This is calculated by dividing PMN0 by PMN1. If the average is high then the Intel® 80200 processor may be starved of the bus external to the Intel® 80200 processor.
- *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.

12.5.4 Data/Bus Request Buffer Full Mode

The Data Cache has buffers available to service cache misses or uncacheable accesses. For every memory request that the Data Cache receives from the processor core, a buffer is speculatively allocated in case an external memory request is required or temporary storage is needed for an unaligned access. If no buffers are available, the Data Cache will stall the processor core. How often the Data Cache stalls depends on the performance of the bus external to the Intel® 80200 processor and what the memory access latency is for Data Cache miss requests to external memory. If the Intel® 80200 processor memory access latency is high, possibly due to starvation, these Data Cache buffers becomes full. This performance monitoring mode is provided to see if the Intel® 80200 processor is being starved of the bus external to the Intel® 80200 processor, which effects the performance of the application running on the Intel® 80200 processor.

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® 80200 processor may be starved of the bus external to the Intel® 80200 processor.
- *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.

12.5.5 Stall/Writeback Statistics

When an instruction requires the result of a previous instruction and that result is not yet available, the Intel® 80200 processor 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 14, "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.

12.5.6 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 does 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.

12.5.7 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** each count as several accesses to the data TLB depending on the number of registers specified in the register list. **LDRD** registers 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 does not increment.

The statistic derived from these two events is:

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

12.6 Multiple Performance Monitoring Run Statistics

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

12.7 Examples

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 generates an IRQ interrupt.

Example 12-1. Configuring the Performance Monitor

```
; Configure PMNC with the following values:
;   evtCount0 = 7, evtCount1 = 0instruction 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 12-2. 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>
SUBSPC,R14,#4          ; return from interrupt
```

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

Example 12-3. 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.

Software Debug 13

This chapter describes software debug and related features in the Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with ARM* Architecture V5TE), namely:

- debug modes, registers and exceptions
- a serial debug communication link via the JTAG interface
- a trace buffer
- a mechanism to load the instruction cache through JTAG
- Debug Handler SW requirements and suggestions

13.1 Definitions

debug handler -is an event handler that runs on the Intel® 80200 processor, when a debug event occurs.

debugger - is software that runs on a host system outside of the Intel® 80200 processor.

13.2 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 cause an undefined instruction trap). The CRn field specifies the number of the register to access. The CRm, opcode_1, and opcode_2 fields are not used and should be set to 0.

Software access to all debug registers must be done in privileged mode. User mode access generates 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.

13.3 Introduction

The Intel® 80200 processor debug unit, when used with a debugger application, allows software running on a the Intel® 80200 processor 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 the Intel® 80200 processor, one of two debug modes can be entered:

- Halt mode
- Monitor mode

13.3.1 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.

13.3.2 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.

13.4 Debug Control and Status Register (DCSR)

The DCSR register is the main control register for the debug unit. Table 13-1 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 13.11.2, SELDCSR JTAG Register for details about accessing DCSR through JTAG.

Table 13-1. 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							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
29:24	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 13-1. 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							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 Asserted 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

13.4.1 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.

13.4.2 Halt Mode Bit (H)

The Halt Mode bit configures the debug unit for either halt mode or monitor mode.

13.4.3 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 the Intel® 80200 processor 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.

13.4.4 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 13.5.1, 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.

13.4.5 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 13-2) in the MOE field.

13.4.6 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.

13.4.7 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 13.12, Trace Buffer.

13.5 Debug Exceptions

A debug exception causes the processor to re-direct execution to a debug event handling routine. The Intel® 80200 processor debug architecture defines the following debug exceptions:

- instruction breakpoint
- data breakpoint
- software breakpoint
- external debug break
- exception vector trap
- trace-buffer full 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 13-2 shows the priority of debug exceptions relative to other processor exceptions.

Table 13-2. Event Priority

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

13.5.1 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 top Section 13.14, "Downloading Code in the ICache" on page 13-34 for details about downloading code into the instruction cache.

During Halt mode, software running on the Intel® 80200 processor 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, the processor takes the following actions:

- disables the trace buffer
- sets DCSR.moe encoding
- processor enters a Special Debug State (SDS)
- for data breakpoints, trace buffer full break, and external debug break:
  R14_dbg = PC of the next instruction to execute + 4
  for instruction breakpoints and software breakpoints and vector traps:
  R14_dbg = PC of the aborted instruction + 4
- SPSR_dbg = CPSR
- CPSR[4:0] = 0b10101 (DEBUG mode)
- CPSR[5] = 0
- CPSR[6] = 1
- CPSR[7] = 1
- PC = 0x0[1]

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 13.6, HW Breakpoint Resources) and the DCSR (see Table 13-1, “Debug Control and Status Register (DCSR)” on page 13-3).
- 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.

1. When the vector table is relocated (CP15 Control Register[13] = 1), the debug vector is relocated to 0xffff0000

13.5.2 Monitor Mode

In monitor mode, the processor handles debug exceptions like normal ARM exceptions. If debug functionality is enabled (DCSR[31] = 1) 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_abt = PC of the next instruction to execute + 4 (for Data Aborts)
  R14_abt = PC of the faulting instruction + 4 (for Prefetch Aborts)
- SPSR_abt = CPSR
- CPSR[4:0] = 0b10111 (ABORT mode)
- CPSR[5] = 0
- CPSR[6] = unchanged
- CPSR[7] = 1
- PC = 0xc (for Prefetch Aborts),
  PC = 0x10 (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 immediately generate a debug exception. Any pending debug breaks are cleared out when any type of debug exception occurs.

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.

13.6 HW Breakpoint Resources

The Intel® 80200 processor 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 XXX 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.

13.6.1 Instruction Breakpoints

The Debug architecture defines two instruction breakpoint registers (IBCR0 and IBCR1). The format of these registers is shown in Table 13-3., 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 13-3. 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																															E
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 generates 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).

13.6.2 Data Breakpoints

The Intel® 80200 processor debug architecture defines two data breakpoint registers (DBR0, DBR1). The format of the registers is shown in Table 13-4.

Table 13-4. 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 13-5. 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 do 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 is not 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.

13.7 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 13.5, “Debug Exceptions” on page 13-6.

13.8 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 13.11.4, DBGTX JTAG Register) and DBGRX (Section 13.11.6, 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 13-6. 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
RR	OV	D	TR																												

reset value: 0x00000000

Bits	Access	Description
31	SW Read-only / Write-ignored JTAG Write-only	**RR** RX Register Ready
30	SW Read / Write	**OV** RX overflow sticky flag
29	SW Read-only/ Write-ignored JTAG Write-only	**D** High-speed download flag
28	SW Read-only/ Write-ignored JTAG Write-only	**TR** TX Register Ready
27:0	Read-as-Zero / Write-ignored	Reserved

13.8.1 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 13-7 shows the normal handshaking used to access the RX register.

Table 13-7. Normal RX Handshaking

Debugger Actions
Debugger wants to send data to debug handler. Before writing new data to the RX register, the debugger polls RR through JTAG until the bit is cleared. 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.
Debug Handler Actions
Debug handler is expecting data from the debugger. The debug handler polls the RR bit until it is set, indicating data in the RX register is valid. 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.

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 13-8 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 13-8. High-Speed Download Handshaking States

Debugger Actions
Debugger wants to transfer code into the Intel® 80200 processor system memory. Prior to starting download, the debugger must polls RR bit until it is clear. Once the RR bit is clear, indicating the debug handler is ready, the debugger starts the download. 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. 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. 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 13.8.2, Overflow Flag (OV) for more details on the overflow condition). After completing the download, the debugger clears the D bit allowing the debug handler to exit the download loop.
Debug Handler Actions
Debug is handler in a routine waiting to write data out to memory. The routine loops based on the D bit in TXRXCTRL. 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.

13.8.2 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 is 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 remains set, until cleared by a write to TXRXCTRL 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.

13.8.3 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.

13.8.4 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 13-9 shows the handshaking used to access the TX register.

Table 13-9. 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.

13.8.5 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 directly updates the condition codes in the CPSR. The debug handler can then conditionally execute based on each CC bit. Table 13-10 shows the mnemonic extension to conditionally execute based on whether the TXRXCTRL bit is set or clear.

Table 13-10. 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:   mcr     p14, 0, r15, c14, c0, 0# read the handshaking bit in TXRXCTRL
        mcrmi   p14, 0, r0, c9, c0, 0  # if RX is valid, read it
        bpl     loop                   # if RX is not valid, loop
```

13.9 Transmit Register (TX)

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

Table 13-11. 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																															

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 13-9, “TX Handshaking” on page 13-15.

13.10 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 13-12. 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																															

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 13.8.1, RX Register Ready Bit (RR).

13.11 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 13.14, Downloading Code in the ICache. The other three JTAG instructions are described in this section.

SELDCSR, DBGTX and DBGRX use a common 36-bit shift register (DBG_SR). New data is shifted in and captured data out through the DBG_SR. In the UPDATE_DR state, the new data shifted into the appropriate data register.

13.11.1 SELDCSR JTAG Command

The ‘SELDCSR’ JTAG instruction selects the DCSR JTAG data register. The JTAG opcode is ‘01001’. 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.

13.11.2 SELDCSR JTAG Register

Placing the "SELDCSR" JTAG instruction in the JTAG IR, selects the DCSR JTAG Data register (Figure 13-1), allowing the debugger to access the DCSR, generate an external debug break, set the hold_rst signal, which is used when loading code into the instruction cache during reset.

Figure 13-1. SELDCSR Hardware

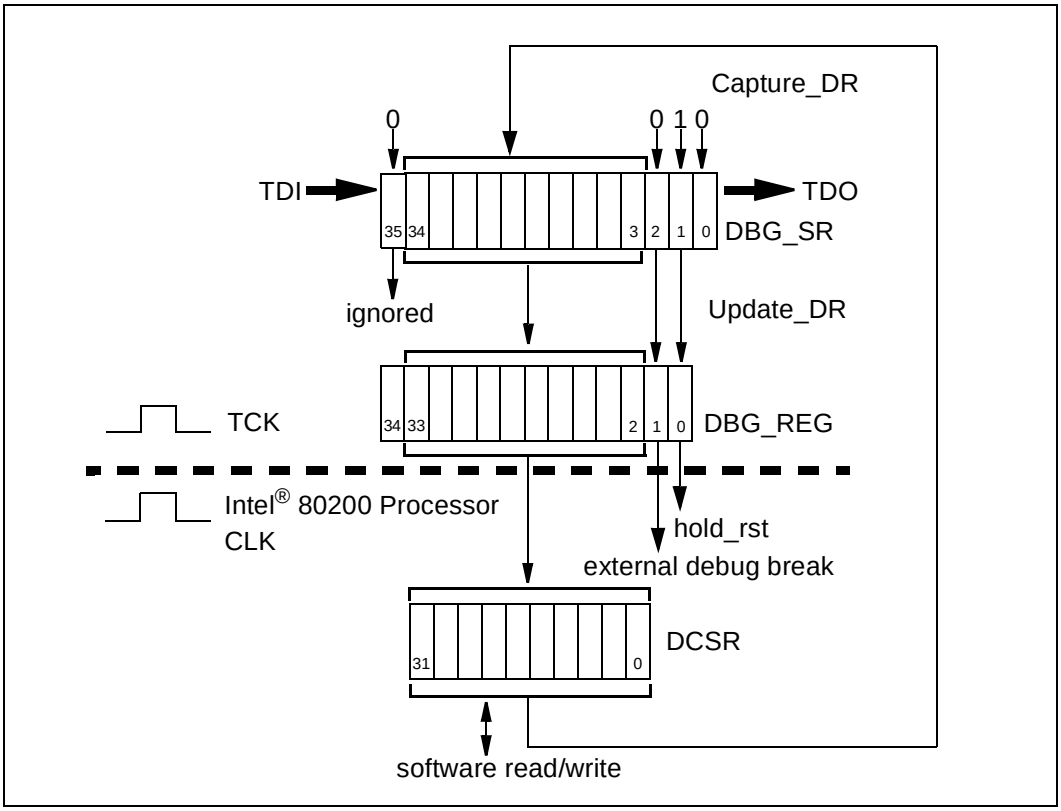

A Capture_DR loads the current DCSR value into DBG_SR[34:3]. The other bits in DBG_SR are loaded as shown in Figure 13-1.

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 DBG_REG[33:2]. This value is then loaded into the actual DCSR register. All bits defined as JTAG writable in Table 13-1, "Debug Control and Status Register (DCSR)" on page 13-3 are updated.

An external host and the debug handler running on the Intel® 80200 processor 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.

Figure 13-2. SELDCSR Data Register

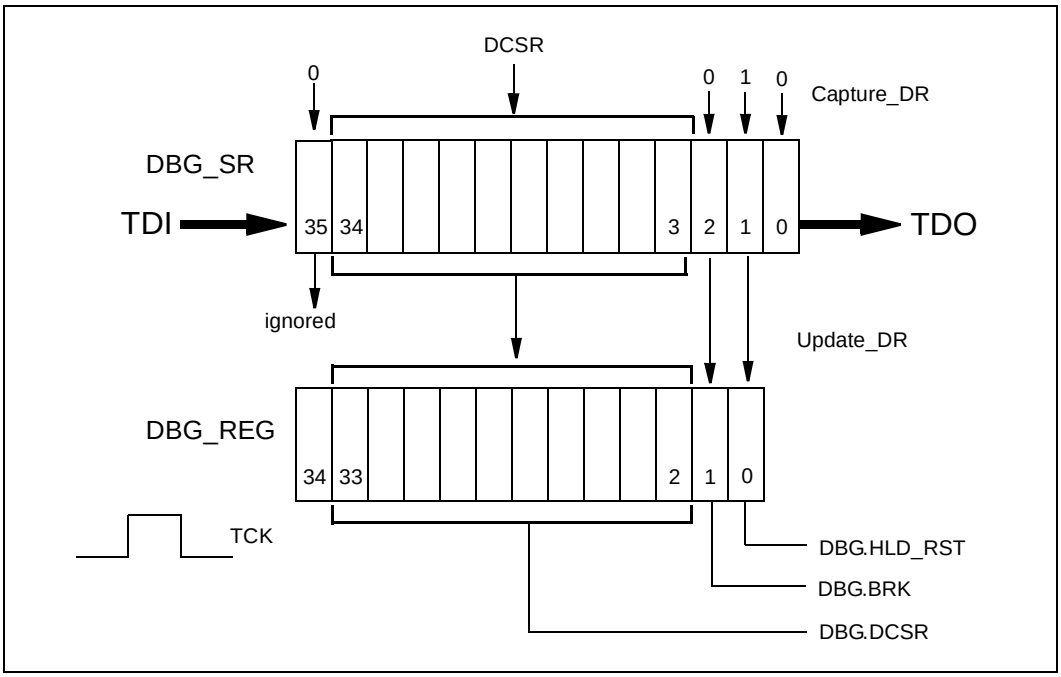

13.11.2.1 DBG.HLD_RST

The debugger uses DBG.HLD_RST when loading code into the instruction cache during a processor reset. Details about loading code into the instruction cache are in Section 13.14, Downloading Code in the ICache.

The debugger must set DBG.HLD_RST before or during assertion of the reset pin. Once DBG.HLD_RST is set, the reset pin can be de-asserted, and the processor internally remains 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 DBG.HLD_RST. This takes the processor out of reset, and execution begins at the reset vector.

A debugger sets DBG.HLD_RST 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 sets the DBG.HLD_RST, clear DBG.BRK, 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 DBG.BRK.

DBG.HLD_RST can only be cleared by scanning in a ‘0’ to DBG_SR[1] and scanning in the appropriate values for the DCSR and DBG.BRK.

13.11.2.2 DBG.BRK

DBG.BRK allows the debugger to generate an external debug break and asynchronously re-direct execution to a debug handling routine.

A debugger sets an external debug break by scanning data into the DBG_SR with DBG_SR[2] set and the desired value to set the DCSR JTAG writable bits in DBG_SR[34:3].

Once an external debug 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, 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. This minimizes breakpoint skid, by reducing the number of instructions that can execute after the external debug break is requested. However, the processor continues to process any instructions which may have already begun execution. Debug mode is not entered until all processor activity has ceased in an orderly fashion.

13.11.2.3 DBG.DCSR

The DCSR is updated with the value loaded into DBG.DCSR following an Update_DR. Only bits specified as writable by JTAG in Table 13-1 are updated.

13.11.3 DBGTX JTAG Command

The ‘DBGTX’ JTAG instruction selects the DBGTX JTAG data register. The JTAG opcode for this instruction is ‘0b10000’. Once the DBGTX data register is selected, the debugger can receive data from the debug handler.

13.11.4 DBGTX JTAG Register

The DBGTX JTAG instruction selects the Debug JTAG Data register (Figure 13-3). The debugger uses the DBGTX data register to poll for breaks (internal and external) to debug mode and once in debug mode, to read data from the debug handler.

Figure 13-3. DBGTX Hardware

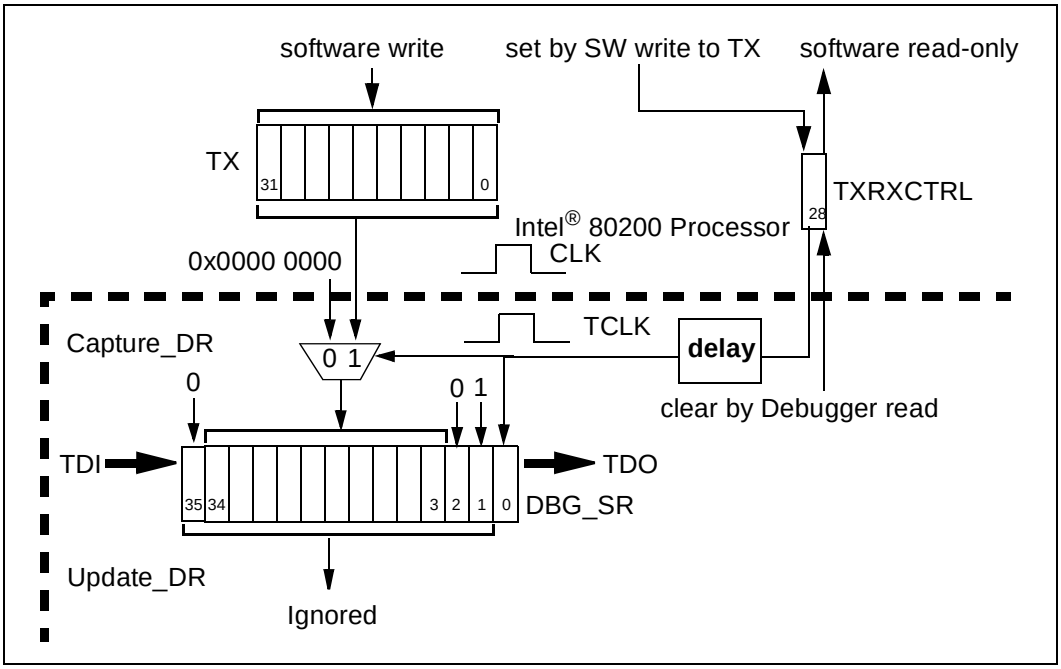

A Capture_DR loads the TX register value into DBG_SR[34:3] and TXRXCTRL[28] into DBG_SR[0]. The other bits in DBG_SR are loaded as shown in Figure 13-1.

The captured TX value is scanned out during the Shift_DR state.

Data scanned in is ignored on an Update_DR.

A ‘1’ captured in DBG_SR[0] indicates the captured TX data is valid. After doing a Capture_DR, the debugger must place the JTAG state machine in the Shift_DR state to guarantee that a debugger read clears TXRXCTRL[28].

13.11.5 DBGRX JTAG Command

The ‘DBGRX’ JTAG instruction selects the DBGRX JTAG data register. The JTAG opcode for this instruction is ‘0b00010’. Once the DBGRX data register is selected, the debugger can send data to the debug handler through the RX register.

13.11.6 DBGRX JTAG Register

The DBGRX JTAG instruction selects the DBGRX JTAG Data register. The debugger uses the DBGRX data register to send data or commands to the debug handler.

Figure 13-4. DBGRX Hardware

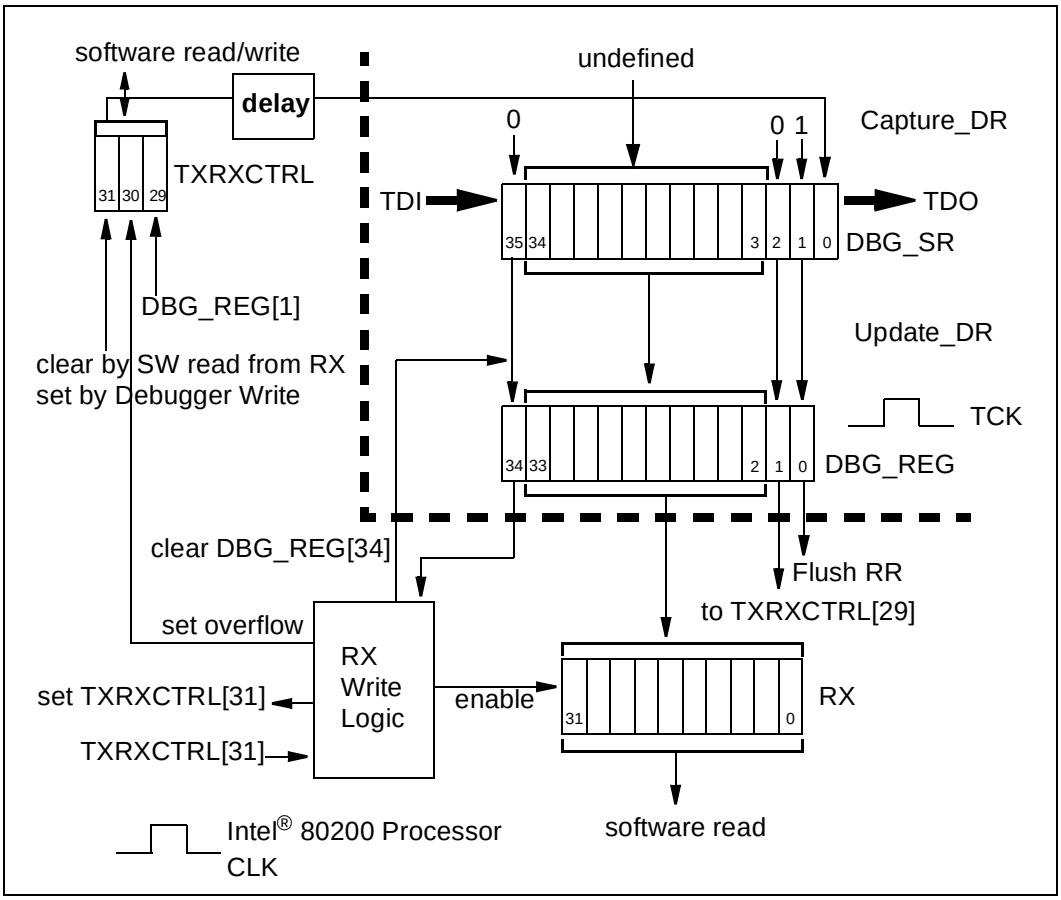

A Capture_DR loads TXRXCTRL[31] into DBG_SR[0]. The other bits in DBG_SR are loaded as shown in Figure 13-4.

The captured data is scanned out during the Shift_DR state.

Care must be taken while scanning in data. While polling TXRXCTRL[31], incorrectly setting DBG_SR[35] or DBG_SR[1] may cause unpredictable behavior following an Update_DR.

Update_DR parallel loads DBG_SR[35:1] into DBG_REG[34:0]. Whether the new data gets written to the RX register or an overflow condition is detected depends on the inputs to the RX write logic.

13.11.6.1 RX Write Logic

The RX write logic (Figure 13-6) serves 4 functions:

1) Enable the debugger write to RX - the logic ensures only new, valid data from the debugger is written to RX. In particular, when the debugger polls TXRXCTRL[31] to see whether the debug handler has read the previous data from RX. The JTAG state machine must go through Update_DR, which should not modify RX.

2) Clear DBG_REG[34] - mainly to support high-speed download. During high-speed download, the debugger continuously scan in a data to send to the debug handler and sets DBG_REG[34] to signal the data is valid. Since DBG_REG[34] is never cleared by the debugger in this case, the ‘0’ to ‘1’ transition used to enable the debugger write to RX would not occur.

3) Set TXRXCTRL[31] - When the debugger writes new data to RX, the logic automatically sets TXRXCTRL[31], signalling to the debug handler that the data is valid.

4) Set the overflow flag (TXRXCTRL[30] - 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 overwrite the previous data before the handler can read it. The logic sets the overflow flag when the previous data has not been read yet, and the debugger has just written new data to RX.

Figure 13-5. RX Write Logic

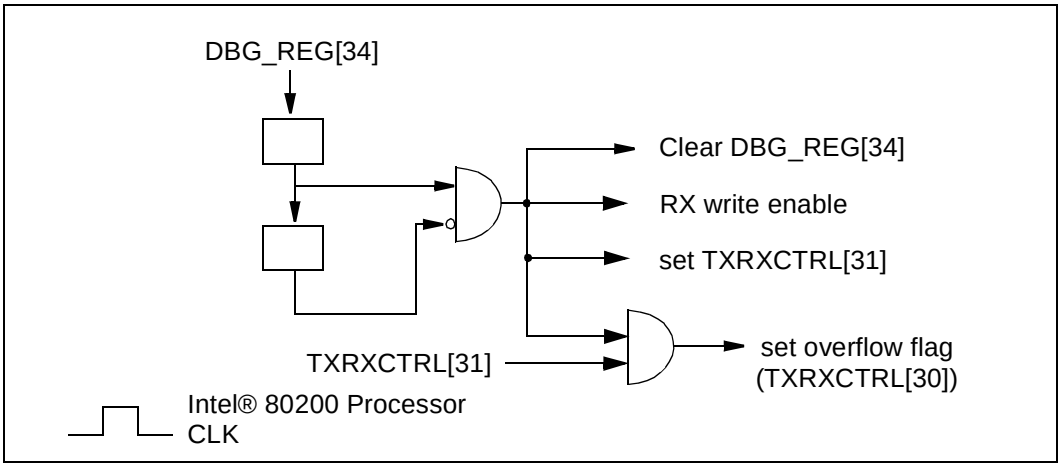

13.11.6.2 DBGRX Data Register

The bits in the DBGRX data register (Figure 13-6) are used by the debugger to send data to the processor. The data register also contains a bit to flush previously written data and a high-speed download flag.

Figure 13-6. DBGRX Data Register

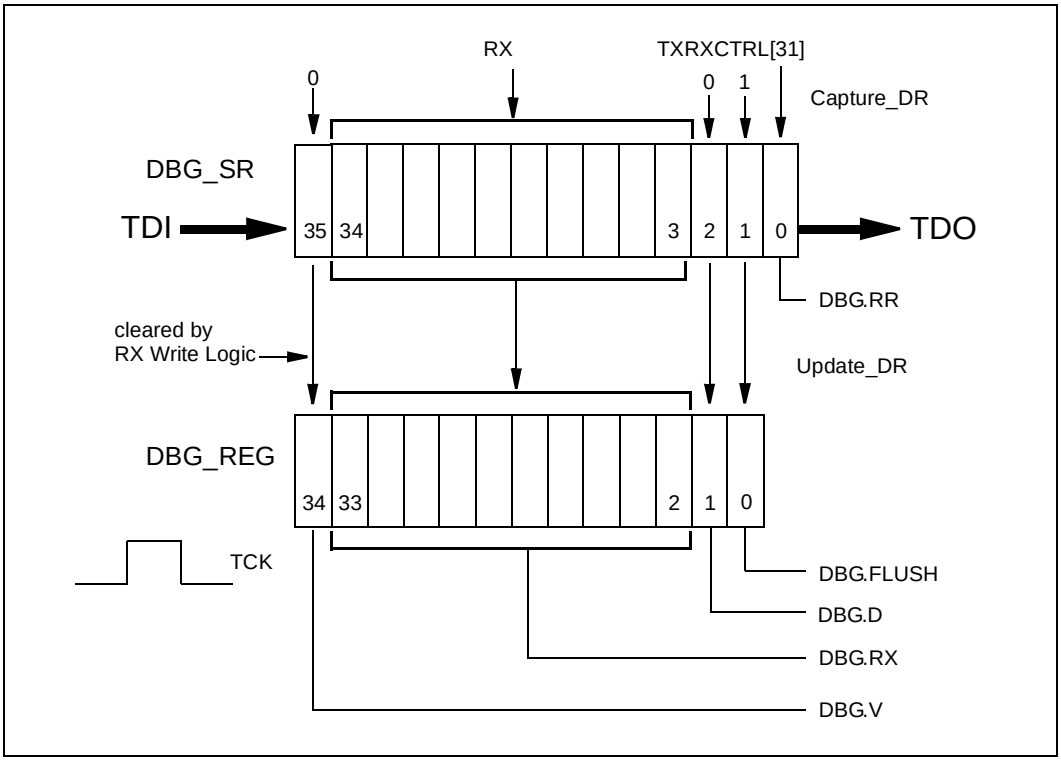

13.11.6.3 DBG.RR

The debugger uses DBG.RR as part of the synchronization that occurs between the debugger and debug handler for accessing RX. This bit contains the value of TXRXCTRL[31] after a Capture_DR. The debug handler automatically sets TXRXCTRL[31] by doing a write to RX.

The debugger polls DBG.RR to determine when the handler has read the previous data from RX.

The debugger sets TXRXCTRL[31] by setting the DBG.V bit.

13.11.6.4 DBG.V

The debugger sets this bit to indicate the data scanned into DBG_SR[34:3] is valid data to write to RX. DBG.V is an input to the RX Write Logic and is also cleared by the RX Write Logic.

When this bit is set, the data scanned into the DBG_SR is written to RX following an Update_DR. If DBG.V is not set and the debugger does an Update_DR, RX is unchanged.

This bit does not affect the actions of DBG.FLUSH or DBG.D.

13.11.6.5 DBG.RX

DBG.RX is written into the RX register based on the output of the RX Write Logic. Any data that needs to be sent from the debugger to the processor must be loaded into DBG.RX with DBG.V set to 1. DBG.RX is loaded from DBG_SR[34:3] when the JTAG enters the Update_DR state.

DBG.RX is written to RX following an Update_DR when the RX Write Logic enables the RX register.

13.11.6.6 DBG.D

DBG.D is provided for use during high speed download. This bit is written directly to TXRXCTRL[29]. The debugger sets DBG.D when downloading a block of code or data to the Intel® 80200 processor system memory. The debug handler then uses TXRXCTRL[29] as a branch flag to determine the end of the loop.

Using DBG.D 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.

13.11.6.7 DBG.FLUSH

DBG.FLUSH allows the debugger to flush any previous data written to RX. Setting DBG.FLUSH clears TXRXCTRL[31].

13.11.7 Debug JTAG Data Register Reset Values

Upon asserting TRST, the DEBUG data register is reset. Assertion of the reset pin does not affect the DEBUG data register. Table 13-13 shows the reset and TRST values for the data register. Note: these values apply for DBG_REG for SELDCSR, DBGTX and DBGRX.

Table 13-13. DEBUG Data Register Reset Values

Bit	TRST	RESET
DBG_REG[0]	0	unchanged
DBG_REG[1]	0	unchanged
DBG_REG[33:2]	unpredictable	unpredictable
DBG_REG[34]	0	unchanged

13.12 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.

13.12.1 Trace Buffer CP Registers

CP14 defines three registers (see Table 13-14) for use with the trace buffer. These CP14 registers are accessible using MRC, MCR, LDC and STC (CDP to any CP14 registers causes an undefined instruction trap). The CRn field specifies the number of the register to access. The CRm, opcode_1, and opcode_2 fields are not used and should be set to 0.

Table 13-14. CP 14 Trace Buffer Register Summary

CP14 Register Number	Register Name
11	Trace Buffer Register (TBREG)
12	Checkpoint 0 Register (CHKPT0)
13	Checkpoint 1 Register (CHKPT1)

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

13.12.1.1 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 <= 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 13-15. 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 the Intel® 80200 processor 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 13-17., 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 are typically 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.

13.12.1.2 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 13-16 shows the format of the trace buffer register.

Table 13-16. 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

13.13 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.

13.13.1 Message Byte

There are two message formats, (exception and non-exception) as shown in Figure 13-7.

Figure 13-7. Message Byte Formats

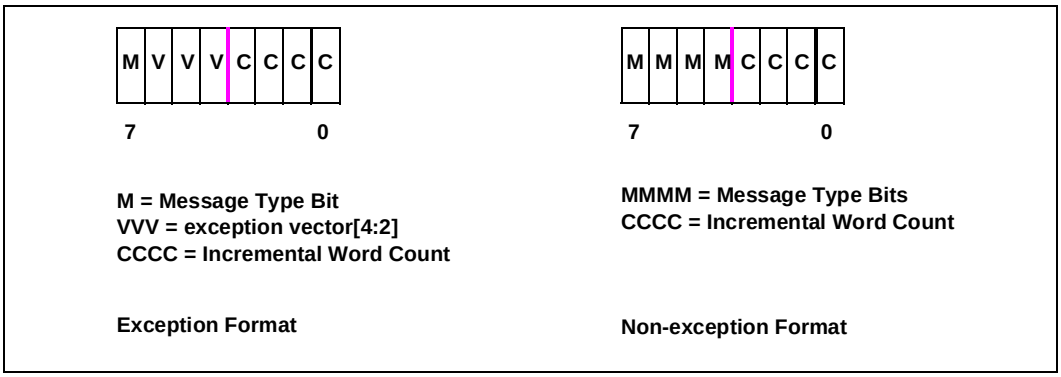

Table 13-17 shows all of the possible trace messages.

Table 13-17. 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.

13.13.1.1 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 host SW 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 is 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).

13.13.1.2 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 13-17).

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 is the 16th instruction to execute). This is shown in Example 13-1. 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 13-1. Rollover Messages Examples

```
count = 5
BL label1
count = 0          ←  branch message placed in trace buffer after branch executes
MOV                   count = 0b0101
count = 1
MOV
count = 2
MOV
...

count = 14
MOV
count = 15
MOV                ←  rollover message placed in trace buffer after 16th instruction executes
count = 0             count = 0b1111
```

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.

13.13.1.3 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 13-8.

Figure 13-8. Indirect Branch Entry Address Byte Organization

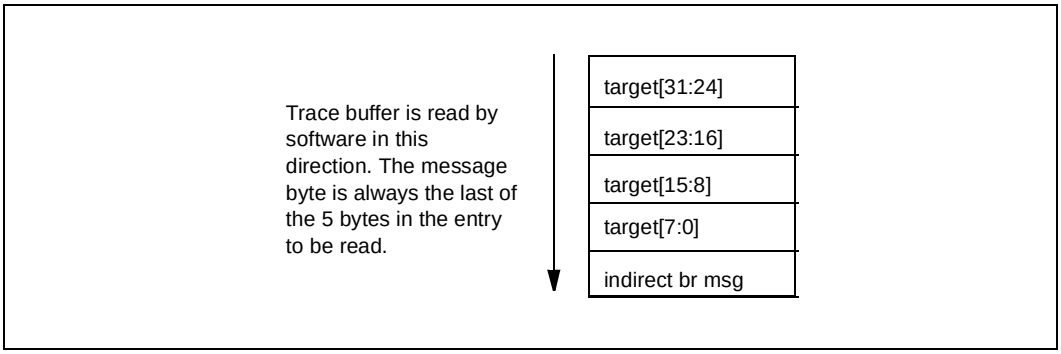

13.13.2 Trace Buffer Usage

The Intel® 80200 processor 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 is always 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 the Intel® 80200 processor) before the buffer can be parsed. Figure 13-9 is a high level view of the trace buffer.

Figure 13-9. High Level View of Trace Buffer

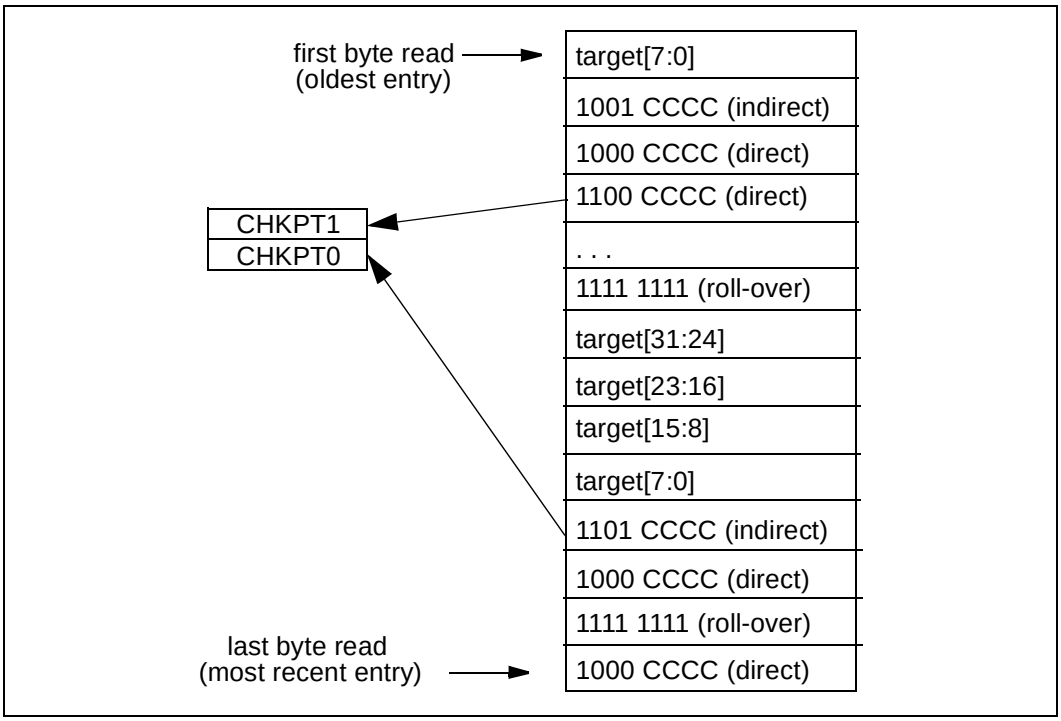

The trace buffer must be initialized prior to its initial usage, then again prior to each subsequent usage. Initialization is done be reading the entire trace buffer. The process of reading the trace buffer also clears it out (all entries are set to 0b0000 0000), so when the trace buffer has been used 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 is the first non-zero entry read out.

Reading 4 or less consecutive “0b0000 0000” entries requires a bit more intelligence in the host SW. The host SW 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 13.13.1.3, 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 host SW 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 is 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 host SW 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 host SW would have to start at its known address (the first checkpoint) which is half way through the buffer and work forward from there.

13.14 Downloading Code in the ICache

On the Intel® 80200 processor, a 2K mini instruction cache, physically separate[1] from the 32K main instruction cache can be used as an on-chip instruction RAM. An external host can download code directly into either instruction cache through JTAG. In addition to downloading code, several cache functions are supported.

The Intel® 80200 processor supports loading the instruction cache during reset and during program execution. Loading the instruction cache during normal program execution requires a strict handshaking protocol between software running on the Intel® 80200 processor and the external host.

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

13.14.1 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.

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

13.14.2 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 13-10. LDIC JTAG Data Register Hardware

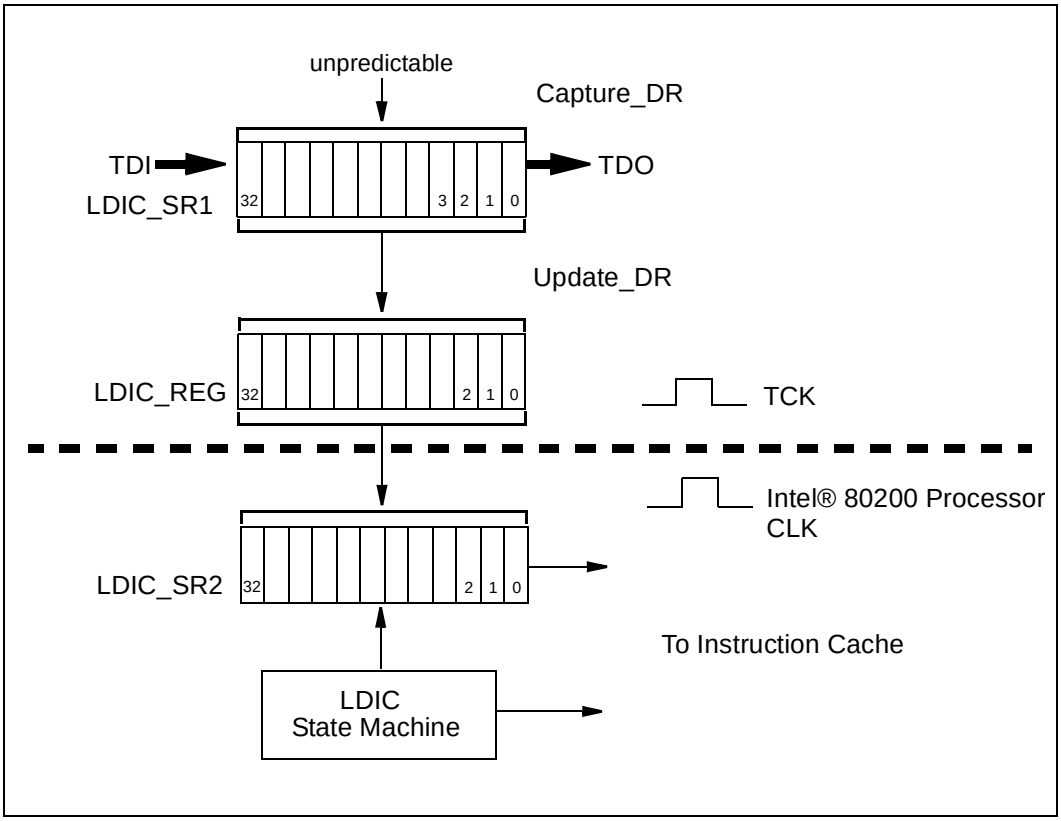

The data loaded into LDIC_SR1 during a Capture_DR is unpredictable.

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 Intel® 80200 processor 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 if 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, results 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.

13.14.3 LDIC Cache Functions

The Intel® 80200 processor 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 13-18. 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 IC[1] invalidates 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.

Each cache function is downloaded through JTAG in 33 bit packets. Figure 13-11 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 13-11. Format of LDIC Cache Functions

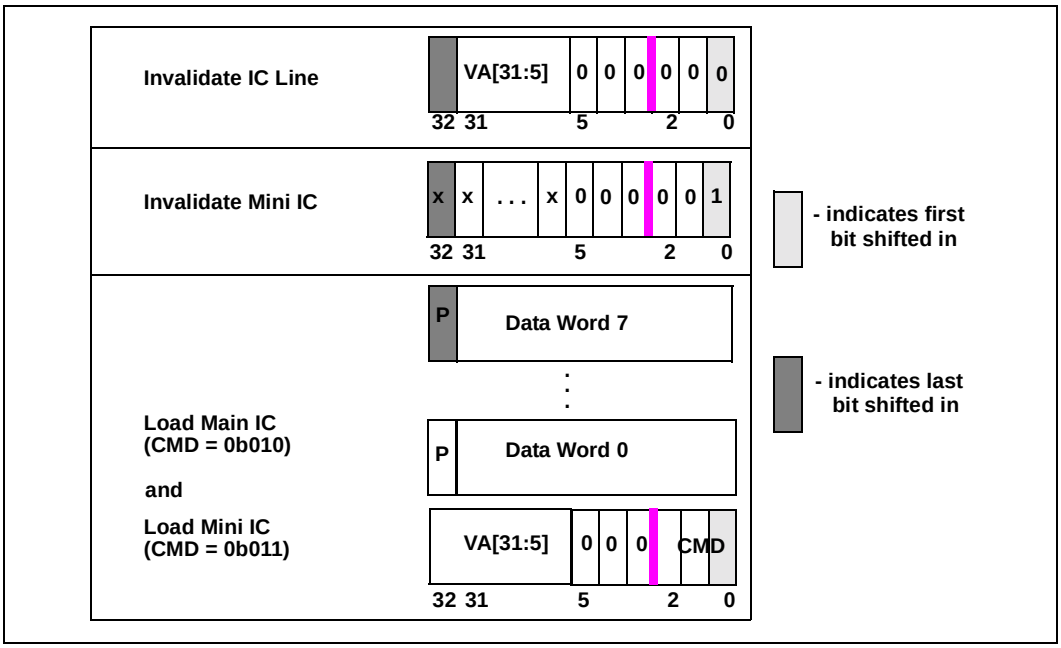

All packets are 33 bits in length. Bits [2:0] of the first packet specify the function to execute. For functions that require an address, bits[32:6] of the first packet specify an 8-word aligned address (Packet1[32:6] = VA[31:5]). For Load Main IC and Load Mini IC, 8 additional data packets are used 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 13-11, 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.

13.14.4 Loading IC 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 an application program. The downloaded handler can then intercept the reset vector and do any necessary setup before the application code executes

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 results in unpredictable behavior by the processor. During a processor reset, the instruction cache is typically invalidated, with the exception of the following modes:

- LDIC mode: active when LDIC JTAG instruction is loaded in the JTAG IR; prevents the mini instruction cache and the main instruction cache from being invalidated during reset.
- HALT mode: active when the Halt Mode bit is set in the DCSR; prevents only the mini instruction cache from being invalidated; main instruction cache is invalidated by reset.

During a cold reset (in which both a processor reset and a JTAG reset occurs) it can be guaranteed that the instruction cache is invalidated since the JTAG reset takes the processor out of any of the modes listed above.

During a warm reset, if a JTAG reset does not occur, the instruction cache is not invalidated by reset when any of the above modes are active. This situation requires special attention if code needs be downloaded during the warm reset.

Note that while Halt Mode is active, reset can invalidate the main instruction cache. Thus debug handler code downloaded during reset can only be loaded into the mini instruction cache. However, code can be dynamically downloaded into the main instruction cache. (refer to Section 13.14.5, Dynamically Loading IC After Reset).

The following sections describe the steps necessary to ensure code is correctly downloaded into the instruction cache.

13.14.4.1 Loading IC During Cold Reset for Debug

The Figure 13-12 shows the actions necessary to download code into the instruction cache during a cold reset for debug.

NOTE: In the Figure 13-12 hold_rst is a signal that gets set and cleared through JTAG When the JTAG IR contains the SELDCSR instruction, the hold_rst signal is set to the value scanned into DBG_SR[1].

Figure 13-12. Code Download During a Cold Reset For Debug

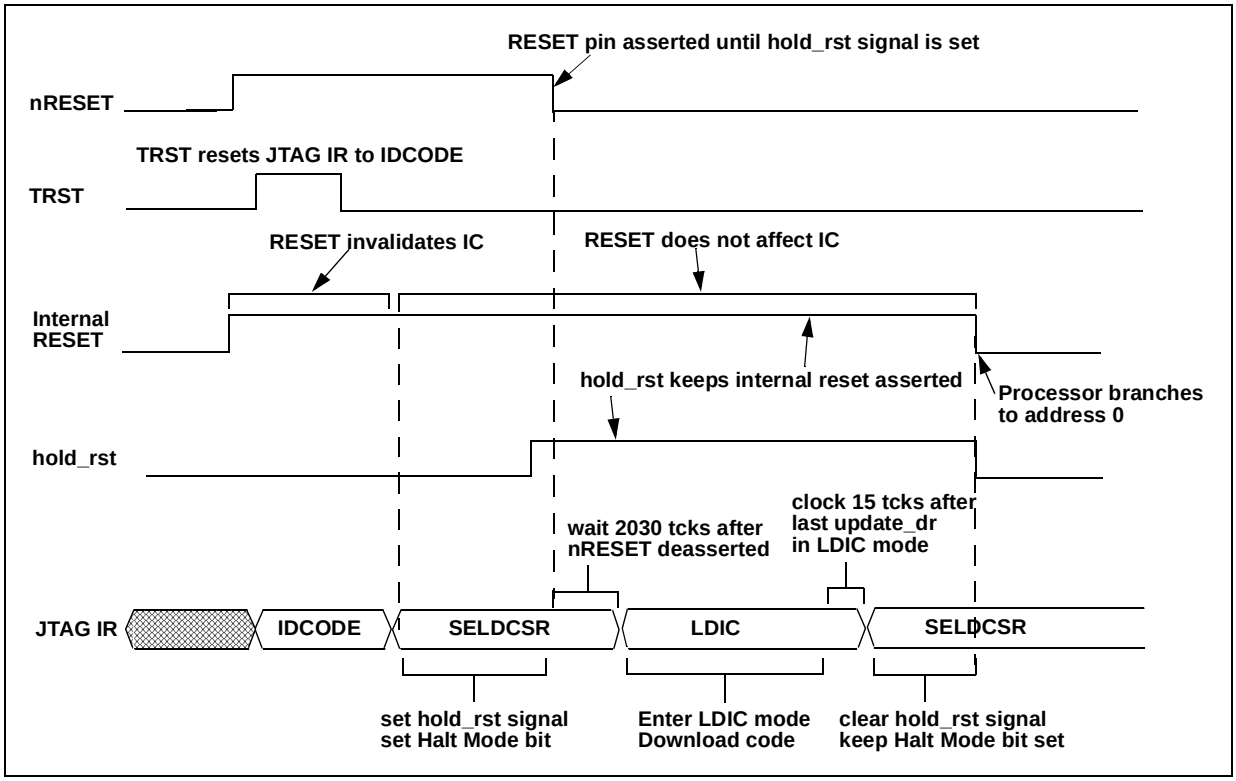

An external host should take the following steps to load code into the instruction cache following a cold reset:

- Assert the nRESET and TRST pins: This resets the JTAG IR to IDCODE and invalidates the instruction cache (main and mini).
- Load the SELDCSR JTAG instruction into JTAG IR and scan in a value to set the Halt Mode bit in DCSR and to set the hold_rst signal. For details of the SELDCSR, refer to Section 13.11.2.
- After hold_rst is set, de-assert the nRESET pin. Internally the processor remains held in reset.
- After nRESET is de-asserted, wait 2030 TCKs.
- Load the LDIC JTAG instruction into JTAG IR.
- Download code into instruction cache in 33-bit packets as described in Section 13.14.3, LDIC Cache Functions.
- After code download is complete, clock a minimum of 15 TCKs following the last update_dr in LDIC mode.
- Place the SELDCSR JTAG instruction into the JTAG IR and scan in a value to clear the hold_rst signal. The Halt Mode bit must remain set to prevent the instruction cache from being invalidated.
- When hold_rst is cleared, internal reset is de-asserted, and the processor executes the reset vector at address 0.

An additional issue for debug is setting up the reset vector trap. This must be done before the internal reset signal is de-asserted. As described in Section 13.4.3, the Halt Mode and the Trap Reset bits in the DCSR must be set prior to de-asserting reset in order to trap the reset vector. There are two possibilities for setting up the reset vector trap:

- The reset vector trap can be set up before the instruction cache is loaded by scanning in a DCSR value that sets the Trap Reset bit in addition to the Halt Mode bit and the hold_rst signal; OR
- The reset vector trap can be set up after the instruction cache is loaded. In this case, the DCSR should be set up to do a reset vector trap, with the Halt Mode bit and the hold_rst signal remaining set.

In either case, when the debugger clears the hold_rst bit to de-assert internal reset, the debugger must set the Halt Mode and Trap Reset bits in the DCSR.

13.14.4.2 Loading IC During a Warm Reset for Debug

Loading the instruction cache during a warm reset may be a slightly different situation than during a cold reset. For a warm reset, the main issue is whether the instruction cache gets invalidated by the processor reset or not. There are several possible scenarios:

- While reset is asserted, TRST is also asserted.

  In this case the instruction cache is invalidated, so the actions taken to download code are identical to those described in Section 13.14.4.1

- When reset is asserted, TRST is not asserted, but the processor is not in Halt Mode.

  In this case, the instruction cache is also invalidated, so the actions are the same as described in Section 13.14.4.1, after the LDIC instruction is loaded into the JTAG IR.

- When reset is asserted, TRST is not asserted, and the processor is in Halt Mode.

  In this last scenario, the mini instruction cache does not get invalidated by reset, since the processor is in Halt Mode. This scenario is described in more detail in this section.

In the last scenario described above is shown in Figure 13-14.

Figure 13-13. Code Download During a Warm Reset For Debug

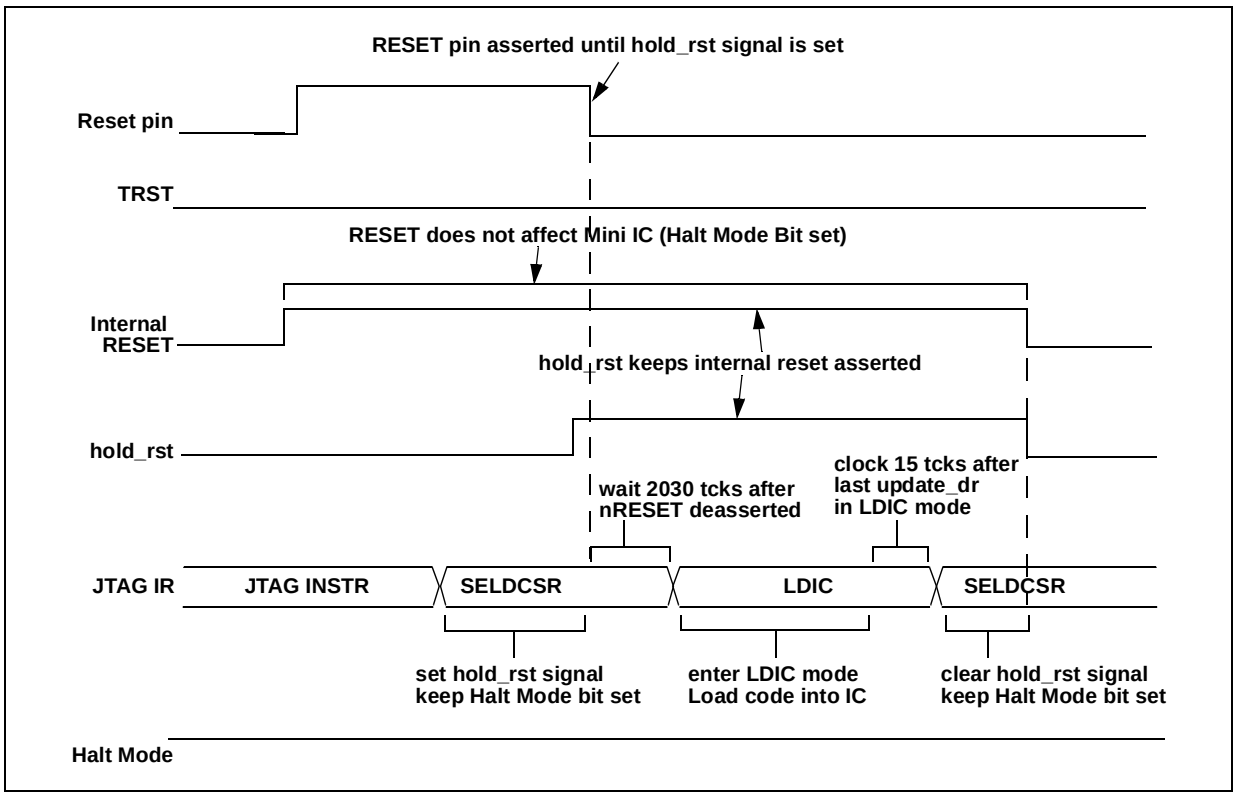

As shown in the figure, reset does not invalidate the instruction cache because of the processor is in Halt Mode. Since the instruction cache was not invalidated, it may contain valid lines. The host must avoid downloading code to virtual addresses that are already valid in the instruction cache (mini IC or main IC), otherwise the processor may behave unpredictably.

There are several possible solutions that ensure code is not downloaded to a VA that already exists in the instruction cache.

1) Since the mini instruction cache was not invalidated, any code previously downloaded into the mini IC is valid in the mini IC, so it is not necessary to download the same code again.

If it is necessary to download code into the instruction cache then:

2) Assert TRST. This clears the Halt Mode bit allowing the instruction cache to be invalidated.

3) Clear the Halt Mode bit through JTAG. This allows the instruction cache to be invalidated by reset.

4) Place the LDIC JTAG instruction in the JTAG IR, then proceed with the normal code download, using the Invalidate IC Line function before loading each line. This requires 10 packets to be downloaded per cache line instead of the 9 packets described in Section 13.14.3

13.14.5 Dynamically Loading IC After Reset

An external host can load code into the instruction cache “on the fly” or “dynamically”. This occurs when the host downloads code while the processor is not being reset. However, this requires strict synchronization between the code running on the Intel® 80200 processor and the external host. The guidelines for downloading code during program execution must be followed to ensure proper operation of the processor. The description in this section focuses on using a debug handler running on the Intel® 80200 processor to synchronize with the external host, but the details apply for any application that is running while code is dynamically downloaded.

To dynamically download code during software debug, there must be a minimal debug handler stub, responsible for doing the handshaking with the host, resident in the instruction cache. This debug handler stub should be downloaded into the instruction cache during processor reset using the method described in Section 13.14.4. Section 13.14.5.1, Dynamic Code Download Synchronization describes the details for implementing the handshaking in the debug handler.

Figure 13-14 shows a high level view of the actions taken by the host and debug handler during dynamic code download.

Figure 13-14. Downloading Code in IC During Program Execution

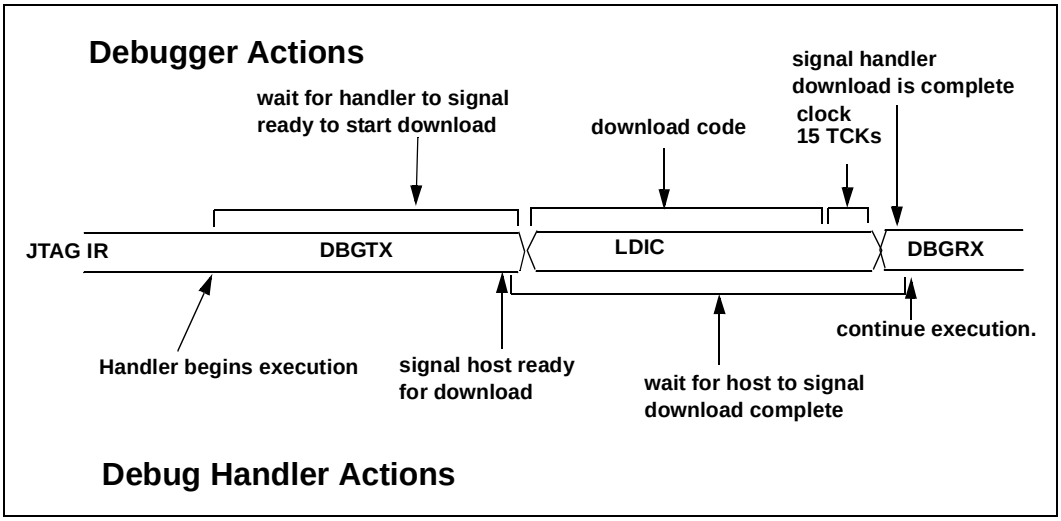

The following steps describe the details for downloading code:

- Since the debug handler is responsible for synchronization during the code download, the handler must be executing before the host can begin the download. The debug handler execution starts when the application running on the Intel® 80200 processor generates a debug exception or when the host generates an external debug break.
- While the DBGTX JTAG instruction is in the JTAG IR (see Section 13.11.3, DBGTX JTAG Command), the host polls DBG_SR[0], waiting for the debug handler to set it.
- When the debug handler gets to the point where it is OK to begin the code download, it writes to TX, which automatically sets DBG_SR[0]. This signals the host it is OK to begin the download. The debug handler then begins polling TXRXCTRL[31] waiting for the host to clear it through the DBGRX JTAG register (to indicate the download is complete).
- The host writes LDIC to the JTAG IR, and downloads the code. For each line downloaded, the host must invalidate the target line before downloading code to that line. Failure to invalidate a line prior to writing it may cause unpredictable operation by the processor.
- When the host completes its download, the host must wait a minimum of 15 TCKs, then switch the JTAG IR to DBGRX, and complete the handshaking (by scanning in a value that sets DBG_SR[35]). This clears TXRXCTL[31] and allows the debug handler code to exit the polling loop. The data scanned into DBG_SR[34:3] is implementation specific.
- After the handler exits the polling loop, it branches to the downloaded code.

Note that this debug handler stub 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.

13.14.5.1 Dynamic Code Download Synchronization

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 80200. 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.

13.14.6 Mini Instruction Cache Overview

The mini instruction cache is a smaller version of the main instruction cache (Refer to Chapter 4 for more details on the main instruction cache). It is a 2KB, 2-way set associative cache. There are 32 sets, each containing two ways; each way contains 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 the Intel® 80200 processor. 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.

13.15 Halt Mode Software Protocol

This section describes the overall debug process in Halt Mode. It describes how to start and end a debug session and details for implementing a debug handler. Intel provides a standard Debug Handler that implements some of the techniques in this chapter. The Intel Debug Handler itself is a a document describing additional handler implementation techniques and requirements.

13.15.1 Starting a Debug Session

Prior to starting a debug session in Halt Mode, the debugger must download code into the instruction cache during reset, via JTAG. (Section 13.14, Downloading Code in the ICache). This downloaded code should consist of:

- a debug handler;
- an override default vector table;
- an override relocated vector table (if necessary).

While the processor is still in reset, the debugger should set up the DCSR to trap the reset vector. This causes a debug exception to occur immediately when the processor comes out of reset. Execution is redirected to the debug handler allowing the debugger to perform any necessary initialization. The reset vector trap is the only debug exception that can occur with debug globally disabled (DCSR[31]=0). Therefore, the debugger must also enable debug prior to existing the handler to ensure all subsequent debug exceptions correctly break to the debug handler.

13.15.1.1 Setting up Override Vector Tables

The override default vector table intercepts the reset vector and branches to the debug handler when a debug exception occurs. If the vector table is relocated, the debug vector is relocated to address 0xffff0000. Thus, an override relocated vector table is required to intercept vector 0xffff0000 and branch to the debug handler.

Both override vector tables also intercept the other debug exceptions, so they must be set up to either branch to a debugger specific handler or go to the application’s handlers.

It is possible that the application modifies its vector table in memory, so the debugger may not be able to set up the override vector table to branch to the application’s handlers. The Debug Handler may be used to work around this problem by reading memory and branching to the appropriate address. Vector traps can be used to get to the debug handler, or the override vector tables can redirect execution to a debug handler routine that examines memory and branches to the application’s handler.

13.15.1.2 Placing the Handler in Memory

The debug handler is not required to be placed at a specific pre-defined address. However, there are some limitations on where the handler can be placed due to the override vector tables and the 2-way set associative mini instruction cache.

In the override vector table, the reset vector must branch to the debug handler using:

- a direct branch, which limits the start of the handler code to within 32 MB of the reset vector, or
- an indirect branch with a data processing instruction. The data processing instruction creates an address using immediate operands and then branches to the target. An LDR to the PC does not work because the debugger cannot set up data in memory before starting the debug handler.

The 2-way set associative limitation is due to the fact that when the override default and relocated vector tables are downloaded, they take up both ways of Set 0 (w/ addresses 0x0 and 0xffff0000). Therefore, debug handler code can not be downloaded to an address that maps into Set 0, otherwise it overwrites one of the vector tables (avoid addresses w/ lower 12 bits=0).

The instruction cache 2-way set limitation is not a problem when the reset vector uses a direct branch, since the branch offset can be adjusted accordingly. However, it makes using indirect branches more complicated. Now, the reset vector actually needs multiple data processing instructions to create the target address and branch to it.

One possibility is to set up vector traps on the non-reset exception vectors. These vector locations can then be used to extend the reset vector.

Another solution is to have the reset vector do a direct branch to some intermediate code. This intermediate code can then uses several instructions to create the debug handler start address and branch to it. This would require another line in the mini instruction cache, since the intermediate code must also be downloaded. This method also requires that the layout of the debug handler be well thought out to avoid the intermediate code overwriting a line of debug handler code, or vice versa.

For the indirect branch cases, a temporary scratch register may be necessary to hold intermediate values while computing the final target address. DBG_r13 can be used for this purpose (see Section 13.15.2.2, Debug Handler Restrictions for restrictions on DBG_r13 usage).

13.15.2 Implementing a Debug Handler

The debugger uses the debug handler to examine or modify processor state by sending commands and reading data through JTAG. The API between the debugger and debug handler is specific to a debugger implementation. Intel provides a standard debug handler and API which can be used by third-party vendors. Issues and details for writing a debug handler are discussed in this section and in the Intel Debug Handler.

13.15.2.1 Debug Handler Entry

When the debugger requests an external debug break or is waiting for an internal break, it should poll the TR bit through JTAG to determine when the processor has entered Debug Mode. The debug handler entry code must do a write to TX to signal the debugger that the processor has entered Debug Mode. The write to TX sets the TR bit, signalling the host that a debug exception has occurred and the processor has entered Debug Mode. The value of the data written to TX is implementation defined (debug break message, contents of register to save on host, etc.).

13.15.2.2 Debug Handler Restrictions

The Debug Handler executes in Debug Mode which is similar to other privileged processor modes, however, there are some differences. Following are restrictions on Debug Handler code and differences between Debug Mode and other privileged modes.

- The processor is in Special Debug State following a debug exception, and thus has special functionality as described in Section 13.5.1, Halt Mode.
- Although address translation and PID remapping are disabled for instruction accesses (as defined in Special Debug State), data accesses use the normal address translation and PID remapping mechanisms.
- Debug Mode does not have a dedicated stack pointer, DBG_r13. Although DBG_r13 exists, it is not a general purpose register. Its contents are unpredictable and should not be relied upon across any instructions or exceptions. However, DBG_r13 can be used, by data processing (non RRX) and MCR/MRC instructions, as a temporary scratch register.
- The following instructions should not be executed in Debug Mode, they may result in unpredictable behavior:

  LDM

  LDR w/ Rd=PC

  LDR w/ RRX addressing mode

  SWP

  LDC

  STC
- The handler executes in Debug Mode and can be switched to other modes to access banked registers. The handler must not enter User Mode; any User Mode registers that need to be accessed can be accessed in System Mode. Entering User Mode may cause unpredictable behavior.

13.15.2.3 Dynamic Debug Handler

On the Intel® 80200 processor, the debug handler and override vector tables reside in the 2 KB mini instruction cache, separate from the main instruction cache. A “static” Debug Handler is downloaded during reset. This is the base handler code, necessary to do common operations such as handler entry/exit, parse commands from the debugger, read/write ARM registers, read/write memory, etc.

Some functions may require large amounts of code or may not be used very often. As long as there is space in the mini-instruction cache, these functions can be downloaded as part of the static Debug Handler. However, if space is limited, the debug handler also has a dynamic capability that allows a function to be downloaded when it is needed. There are three methods for implementing a dynamic debug handler (using the mini instruction cache, main instruction cache, or external memory). Each method has their limitations and advantages. Section 13.14.5, Dynamically Loading IC After Reset describes how do dynamically load the mini or main instruction cache.

1. using the Mini IC

   The static debug handler can support a command which can have functionality dynamically mapped to it. This dynamic command does not have any specific functionality associated with it until the debugger downloads a function into the mini instruction cache. When the debugger sends the dynamic command to the handler, new functionality can be downloaded, or the previously downloaded functionality can be used.

   There are also variations in which the debug handler supports multiple dynamic commands, each mapped to a different dynamic function; or a single dynamic command that can branch to one of several downloaded dynamic functions based on a parameter passed by the debugger.

   Debug Handlers that allow code to be dynamically downloaded into the mini instruction cache must be carefully written to avoid inadvertently overwriting a critical piece of debug handler code. Dynamic code is downloaded to the way pointed to by the round-robin pointer. Thus, it is possible for critical debug handler code to be overwritten, if the pointer does not select the expected way.

   To avoid this problem, the debug handler should be written to avoid placing critical code in either way of a set that is intended for dynamic code download. This allows code to be downloaded into either way, and the only code that is overwritten is the previously downloaded dynamic function. This method requires that space within the mini instruction cache be allocated for dynamic download, limiting the space available for the static Debug Handler. Also, the space available may not be suitable for a larger dynamic function.

   Once downloaded, a dynamic function essentially becomes part of the Debug Handler. Since it is in the mini instruction cache, it does not get overwritten by application code. It remains in the cache until it is replaced by another dynamic function or the lines where it is downloaded are invalidated.

2. Using the Main IC

   The steps for downloading dynamic functions into the main instruction cache is similar to downloading into the mini instruction cache. However, using the main instruction cache has its advantages.

   Using the main instruction cache eliminates the problem of inadvertently overwriting static Debug Handler code by writing to the wrong way of a set, since the main and mini instruction caches are separate. The debug handler code does not need to be specially mapped out to avoid this problem. Also, space for dynamic functions does not need to be allocated in the mini instruction cache and dynamic functions are not limited to the size allocated.

   The dynamic function can actually be downloaded anywhere in the address space. The debugger specifies the location of the dynamic function by writing the address to RX when it signals to the handler to continue. The debug handler then does a branch-and-link to that address.

   If the dynamic function is already downloaded in the main instruction cache, the debugger immediately downloads the address, signalling the handler to continue.

   The static Debug Handler only needs to support one dynamic function command. Multiple dynamic functions can be downloaded to different addresses and the debugger uses the function's address to specify which dynamic function to execute.

   Since the dynamic function is being downloaded into the main instruction cache, the downloaded code may overwrite valid application code, and conversely, application code may overwrite the dynamic function. The dynamic function is only guaranteed to be in the cache from the time it is downloaded to the time the debug handler returns to the application (or the debugger overwrites it).

- External memory

  Dynamic functions can also we downloaded to external memory (or they may already exist there). The debugger can download to external memory using the write-memory commands. Then the debugger executes the dynamic command using the address of the function to identify which function to execute. This method has the many of the same advantages as downloading into the main instruction cache.

  Depending on the memory system, this method could be much slower than downloading directly into the instruction cache. Another problem is the application may write to the memory where the function is downloaded. If it can be guaranteed that the application does not modify the downloaded dynamic function, the debug handler can save the time it takes to re-download the code. Otherwise, the ensure the application does not corrupt the dynamic functions, the debugger should re-download any dynamic functions it uses.

For all three methods, the downloaded code executes in the context of the debug handler. The processor is in Special Debug State, so all of the special functionality applies.

The downloaded functions may also require some common routines from the static debug handler, such as the polling routines for reading RX or writing TX. To simplify the dynamic functions, the debug handler should define a set of registers to contain the addresses of the most commonly used routines. The dynamic functions can then access these routines using indirect branches (BLX). This helps reduce the amount of code in the dynamic function since common routines do not need to be replicated within each dynamic function.

13.15.2.4 High-Speed Download

Special debug hardware has been added to support a high-speed download mode to increase the performance of downloads to system memory (vs. writing a block of memory using the standard handshaking).

The basic assumption is that the debug handler can read any data sent by the debugger and write it to memory, before the debugger can send the next data. Thus, in the time it takes for the debugger to scan in the next data word and do an Update_DR, the handler is already in its polling loop, waiting for it. Using this assumption, the debugger does not have to poll RR to see whether the handler has read the previous data - it assumes the previous data has been consumed and immediately starts scanning in the next data word.

The pitfall is when the write to memory stalls long enough that the assumption fails. In this case the download with normal handshaking can be used (or high-speed download can still be used, but a few extra TCKs in the Pause_DR state may be necessary to allow a little more time for the store to complete).

The hardware support for high-speed download includes the Download bit (DCSR[29]) and the Overflow Flag (DCSR[30]).

The download bit acts as a branch flag, signalling to the handler to continue with the download. This removes the need for a counter in the debug handler.

The overflow flag indicates that the debugger attempted to download the next word before the debugger read the previous word.

More details on the Download bit, Overflow flag and high-speed download, in general, can be found in Section 13.8, Transmit/Receive Control Register (TXRXCTRL).

Following is example code showing how the Download bit and Overflow flag are used in the debug handler:

```
hs_write_word_loop:
hs_write_overflow:
    bl      read_RX                @ read data word from host

    @@ read TXRXCTRL into the CCs
    mrc     p14, 0, r15, c14, c0, 0
    bcc     hs_write_done      @ if D bit clear, download complete, exit loop.
    beq     hs_write_overflow  @ if overflow detected, loop until host clears D bit

    str     r0, [r6], #4       @ store only if there is no overflow.

    b       hs_write_word_loop     @ get next data word

hs_write_done:
    @@ after the loop, if the overflow flag was set, return error message to host
    moveq   r0, #OVERFLOW_RESPONSE
    beq     send_response
    b       write_common_exit
```

13.15.3 Ending a Debug Session

Prior to ending a debug session, the debugger should take the following actions:

- Clear the DCSR (disable debug, exit Halt Mode, clear all vector traps, disable the trace buffer)
- turn off all breakpoints;
- invalidate the mini instruction cache;
- invalidate the main instruction cache;
- invalidate the btb;

These actions ensure that the application program executes correctly after the debugger has been disconnected.

13.16 Software Debug Notes/Errata

1. Trace buffer message count value on data aborts:

   LDR to non-PC that aborts gets counted in the exception message. But an LDR to the PC that aborts does not get counted on exception message.

2. SW Note on data abort generation in Special Debug State.

   1) Avoid code that could generate precise data aborts.

   2) If this cannot be done, then handler needs to be written such that a memory access is followed by 1 nops. In this case, certain memory operations must be avoided - LDM, STM, STRD, LDC, SWP.

3. Data abort on Special Debug State:

   When write-back is on for a memory access that causes a data abort, the base register is updated with the write-back value. This is inconsistent with normal (non-SDS) behavior where the base remains unchanged if write-back is on and a data abort occurs.

4. Trace Buffer wraps around and loses data in Halt Mode when configured for fill-once mode:

   It is possible to overflow (and lose) data from the trace buffer in fill-once mode, in Halt Mode. When the trace buffer fills up, it has space for 1 indirect branch message (5 bytes) and 1 exception message (1 byte).

   If the trace buffer fills up with an indirect branch message and generates a trace buffer full break at the same time as a data abort occurs, the data abort has higher priority, so the processor first goes to the data abort handler. This data abort is placed into the trace buffer without losing any data.

   However, if another imprecise data abort is detected at the start of the data abort handler, it has higher priority than the trace buffer full break, so the processor goes back to the data abort handler. This 2nd data abort also gets written into the trace buffer. This causes the trace buffer to wrap-around and one trace buffer entry is lost (oldest entry is lost). Additional trace buffer entries can be lost if imprecise data aborts continue to be detected before the processor can handle the trace buffer full break (which turns off the trace buffer).

   This trace buffer overflow problem can be avoided by enabling vector traps on data aborts.

5. TXRXCTRL.RR prevents TX register from being updated (even if TXRXCTRL.TR is clear). This is to be fixed on B-step.

   The problem is that there is incorrect (and unnecessary) interaction between the RX ready (RR) flag and writing the TX register. The debug handler looks at the TX ready bit before writing to the TX register. If this bit is clear, then the handler should be able to write to the TX register. However, in the current implementation even if the TR bit is clear, if the RR bit is set, TX is unchanged when the handler writes to it. It is OK to prevent a write to TX when the TR bit is set (since the host has not read the previous data in the TX, and we don’t want a write to TX to overwrite previous data).

6. The TXRXCTRL.OV bit (overflow flag) does not get set during high-speed download when the handler reads the RX register at the same time the debugger writes to it.

   If the debugger writes to RX at the same time the handler reads from RX, the handler read returns the newly written data and the previous data is lost. However, in this specific case, the overflow flag does not get set, so the debugger is unaware that the download was not successful.

Performance Considerations 14

This chapter describes relevant performance considerations that compiler writers, application programmers and system designers need to be aware of to efficiently use Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE). Performance numbers discussed here include interrupt latency, branch prediction, and instruction latencies.

14.1 Interrupt Latency

Table 14-1 shows the *Minimum Interrupt Latency* for the Intel® 80200 processor, which is 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.

Table 14-1. Minimum Interrupt Latency

# MCLK Clock Cycles	Description
3	***Minimum Interrupt Latency.*** This is measured from the assertion of IRQ or FIQ interrupt pin to the execution of the first instruction of the interrupt event handler.

Note: This number assumes that the interrupt vector is resident in the instruction cache. The Intel® 80200 processor does provide the capability to lock the vector and the interrupt service routine into the instruction cache.

Many parameters can affect this best-case performance:

- instruction currently executing: could be as bad as a 16-register LDM
- fault status: processor could fault just when the interrupt arrives
- stalls: processor could be waiting for data from a load, doing a page table walk, etc.
- bus ratio: the best case assumes a 3:1 core:bus ratio. Higher ratios would slightly improve performance

14.2 Branch Prediction

The Intel® 80200 processor 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 is predicted not taken and incur a branch latency penalty.

These instructions -- ARM **B,** ARM **BL** and Thumb **B --** enter into the branch target buffer when they are "taken" for the first time. (A "taken" branch refers to when they are evaluated to be true.) Once in the branch target buffer, the Intel® 80200 processor dynamically predicts the outcome of these instructions based on previous outcomes. Table 14-2 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 Intel® 80200 processor can execute the next instruction in the program flow in the cycle following the branch.

Table 14-2. 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. 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

14.3 Addressing Modes

All load and store addressing modes implemented in the Intel® 80200 processor do not add to the instruction latencies numbers.

14.4 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.

14.4.1 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 14-2, 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 14-1. Computing Latencies

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

Table 14-3 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 14-3, **UMLAL** starts to issue at cycle 0 and the **SUB** issues at cycle 5. thus the Result Latency is five.

Table 14-3. 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

14.4.2 Branch Instruction Timings

Table 14-4. 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 14-5. Branch Instruction Timings (Those not predicted by the BTB)

Mnemonic	Minimum Issue Latency when Branch Not Taken	Minimum Issue Latency when Branch Taken
BLX(1)	N/A	5
BLX(2)	1	5
BX	1	5
Data Processing Instruction with PC as the destination	Same as Table 14-6	4 + numbers in Table 14-6
LDR PC, <>	2	8
LDM with PC in register list	3 + numreg[1]	10 + max (0, numreg-3)

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

14.4.3 Data Processing Instruction Timings

Table 14-6. 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[1]	Minimum Issue Latency	Minimum Result Latency[1]
ADC	1	1	2	2
ADD	1	1	2	2
AND	1	1	2	2
BIC	1	1	2	2
CMN	1	1	2	2
CMP	1	1	2	2
EOR	1	1	2	2
MOV	1	1	2	2
MVN	1	1	2	2
ORR	1	1	2	2
RSB	1	1	2	2
RSC	1	1	2	2
SBC	1	1	2	2
SUB	1	1	2	2
TEQ	1	1	2	2
TST	1	1	2	2

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

14.4.4 Multiply Instruction Timings

Table 14-7. Multiply Instruction Timings (Sheet 1 of 2)

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

Table 14-7. Multiply Instruction Timings (Sheet 2 of 2)

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

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

Table 14-8. Multiply Implicit Accumulate Instruction Timings

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

Table 14-9. Implicit Accumulator Access Instruction Timings

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

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

14.4.5 Saturated Arithmetic Instructions

Table 14-10. Saturated Data Processing Instruction Timings

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

14.4.6 Status Register Access Instructions

Table 14-11. Status Register Access Instruction Timings

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

14.4.7 Load/Store Instructions

Table 14-12. Load and Store Instruction Timings

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

Table 14-13. Load and Store Multiple Instruction Timings

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

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

14.4.8 Semaphore Instructions

Table 14-14. Semaphore Instruction Timings

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

14.4.9 Coprocessor Instructions

Table 14-15. CP15 Register Access Instruction Timings

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

Table 14-16. CP14 Register Access Instruction Timings

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

14.4.10 Miscellaneous Instruction Timing

Table 14-17. SWI Instruction Timings

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

Table 14-18. Count Leading Zeros Instruction Timings

Mnemonic	Minimum Issue Latency	Minimum Result Latency
CLZ	1	1

14.4.11 Thumb* Instructions

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

Compatibility: Intel® 80200 Processor vs. SA-110 A

This appendix highlights the differences between the first generation Intel® StrongARM* technology (SA-110) and the Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE).

A.1 Introduction

The Intel® 80200 processor architecture has been defined to be compatible with SA-110 where possible, however, there are some features not supported on the Intel® 80200 processor or the definition of them has been modified. The following sections discuss these deviations.

A programmer who is developing an application for SA-110 and wishes to migrate to the Intel® 80200 processor must be aware of these architecture differences. Use of these architecture features should be avoided or isolated in developing the application so that migrating to the Intel® 80200 processor can occur with minimal effort.

A.2 Summary

Various features of the SA-110 and the Intel® 80200 processor are outlined in this section. Subsequent sections give more details.

Feature	SA-110	Intel® 80200 Processor
Intel Extensions		
40-bit accumulator access instructions (MRA, MAR)		•
Cache preload (PLD)		•
New multiply instructions for packed data (MIA, MIAPH, MIAxy)		•
New Load/Store Consecutive (LDRD/STRD)		•
ARM v5		
Sticky overflow flag in SPSR for saturated math		•
DSP extensions (SMLAxy, SMLAWy, SMLALxy, SMULxy, SMULWy, QADD, QDADD, QSUB, QDSUB, CLZ)		•
ARM* / Thumb* transfer instructions		•
Thumb		•
Floating Point Instructions		
Tiny pages		•
26-bit code	•	
Big Endian	•	•
Little Endian	•	•

Feature / Parameter	Brief Description or Note	SA-110	Intel® 80200 Processor
Main Execution Pipeline	Scalar, in-order execution, single issue	•	•
• RISC Superpipeline	Pipeline with more than usual number of pipe stages. Allows greater operating frequency.	(5 stage)	• (7 stage)
• Out-of order completion	Instructions may finish out of program order.		•
• Concurrent execution in 3 pipes	Instructions may occupy the MAC, ALU, and data-cache pipes concurrently.		•
• ALU	Number of cycles ALU takes to complete instruction.	1	1[a]
• Register Scoreboarding	Allows instructions in different pipelines to execute as long as there are no data hazards.		•
• Dynamic branch prediction	128 entry branch table address cache holds history of branches taken and not taken.		•
• Branch misprediction penalty	(in cycles)	1	4
MAC Pipeline	Handles all multiply operations		•
• Dedicated 40-bit accumulator	Allows 256 accumulates before scaling of data is necessary.		•
• 1 cycle 16x32 MAC sustained			•
• Early terminate	MAC may finish instruction and return results in any of the MAC pipestages.	n/a	•
Instruction Cache		•	•
• Geometry / Replacement policy	Both the SA-110 and the Intel® 80200 processor are 32-way set associative with round-robin replacement. They differ in size.	16K	32K
• Lockable by line	Instructions may be locked into the instruction cache with line granularity, preventing future eviction.		•
• Fill buffers	Buffer incoming external memory operations.		•
Data Cache / Mini Data Cache			
• Replacement Policy		round-robin	round-robin
• Primary Data Cache Geometry	Both SA-110 and the Intel® 80200 processor are 32-way set associative with round-robin replacement. They differ in size.	16K, 32 ways[b]	32K, 32 way
• Mini Data Cache Geometry	Set associative	512 bytes, 2 ways	2K, 2 ways
• Data RAM	Software can re-map portions of the data cache into data-RAM on a line granularity.		•
• Hit under miss	Allow accesses to the data cache while there are outstanding miss requests to external memory		•
• Write-back	Store operation that “hits” cache is not written to external memory.	•	•
• Write-through	Store operations are written to external memory even if cache is “hit”.		•
• Fill Buffer	Buffer incoming external memory operations.	•	•
• Pending Buffer	Collects memory requests that “hit” an outstanding load		•
• Write Buffer	8 entry write buffer	•	•
• Write Buffer coalescing	Number of entries that a new store request can coalesce to in the write buffer	last entry	All entries

a. A 32-bit shift and ALU operation takes 1.5 cycles.
b. SA1100

A.3 Architecture Deviations

A.3.1 Read Buffer

A Read Buffer is not supported on the Intel® 80200 processor and the definition of CP15 register 9 has changed from controlling the read buffer (on SA-110) to one that controls cache/TLB lock down (on the Intel® 80200 processor).

The functionality of the Read Buffer on the Intel® 80200 processor can be realized with the existing architecture features of the Intel® 80200 processor. The Read Buffer allowed applications to prefetch data into the Read Buffer for future use and did not stall SA-110 during the data prefetch. The Intel® 80200 processor provides a PLD instruction that preloads 32 bytes of data into the data cache. This instruction combined with the support for hit-under-miss provides similar functionality to the Read Buffer. Note that this is for cacheable data only, so software needs to re-map non-cacheable read-only data to cacheable before issuing the PLD instruction.

If the targeted memory region is being shared by another hardware entity, be sure to issue a cache-invalidate (see Section 7.2.8, "Register 7: Cache Functions" on page 7-10s) before the PLD. This ensures an incoherent copy doesn't already exist in the Intel® 80200 processor cache.

A.3.2 26-bit Mode

SA-110 supported 26-bit mode, the Intel® 80200 processor does not.

A.3.3 Cacheable (C) and Bufferable (B) Encoding

Table A-1 describes the differences in the encoding of the C and B bits for data accesses. The Intel® 80200 processor now follows the ARM definition of the C and B bits (when X=0). The main difference occurs when cacheable and non-bufferable data is specified (C=1, B=0); SA-110 uses this encoding for the mini-data cache and the Intel® 80200 processor uses this encoding to specify write-through caching. Another subtle difference is for C=0, B=1, where the Intel® 80200 processor coalesces and stores in the write buffer and SA-110 does not.

Table A-1. C and B encoding

Encoding	SA-110 Function	Intel® 80200 Processor Function
C=1,B=1	Cacheable in data cache; store misses can coalesce in write buffer	Cacheable in data cache, store misses can coalesce in write buffer
C=1,B=0	Cacheable in mini-data cache; store misses can coalesce in write buffer	Cacheable in data cache, with a write-through policy. Store misses can coalesce in write buffer
C=0,B=1	Non-cacheable; no coalescing in write buffer, but can wait in write buffer	Non-cacheable; stores can coalesce in the write buffer
C=0,B=0	Non-cacheable; no coalescing in the write buffer, SA-110 stalls until this transaction is done	Non-cacheable, no coalescing in the write buffer, Intel® 80200 processor does stall until the operation is complete.

A.3.4 Write Buffer Behavior

Definition of 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.

There is a difference in how stores are coalesced to existing entries in the write buffer. When coalescing is enabled, SA-110 only coalesces to the last entry placed in the write buffer. The Intel® 80200 processor can coalesce to any entry in the write buffer. The Intel® 80200 processor also added a global coalesce disable bit located in the Control Register (CP15, register 1, opcode_2=1).

Another difference between SA-110 and the Intel® 80200 processor is that the write buffer is always enabled on the Intel® 80200 processor. Bit 3 of the Control Register (CP15, register 1, opcode_2=0) was used in SA-110 to enable/disable the write buffer. For the Intel® 80200 processor, this bit is always set to 1.

Memory references are rearranged if that would cause incorrect program behavior (see Section 6.5, “Write Buffer/Fill Buffer Operation and Control” on page 6-16).

A.3.5 External Aborts

External aborts are imprecise exceptions on the Intel® 80200 processor. External aborts may be generated by external memory when, for example, there is a parity error detected during a memory access. Since the Intel® 80200 processor continues instruction execution during external memory requests, the PC that is saved in R14 when the exception is reported may not be the PC of the offending instruction. Many instructions may have executed after the offending instruction.

SA-110 always stalls the processor when there was an external load request or when an external write request occurs with C=0 and B=0. External aborts detected on these requests would be precise, meaning the PC that is saved in R14_ABORT when the exception is reported is the address of the offending instruction. The Intel® 80200 processor also stalls the processor for these requests but an external abort on these requests would not be precise. The value in R14_ABORT when the exception is reported would not be that of the offending instruction.

Software relying on this feature of SA-110 may not be compatible with the Intel® 80200 processor.

A.3.6 Performance Differences

There exists significant performance differences in program execution between SA-110 and the Intel® 80200 processor. If an SA-110 application had operations that had specific timing relationships, these relationships would not hold for the Intel® 80200 processor. In all typical applications, the Intel® 80200 processor performance greatly exceeds that of SA-110.

The following is a list of all the new features introduced in the Intel® 80200 processor that significantly improve the instruction per cycle (IPC) number seen in SA-110. Any application written for SA-110 will encounter these performance enhancing features.

- Data cache is non-blocking, which means execution does not stall for every data cache miss.
- Multiply instructions execute in parallel with other non-multiply instructions.
- The instruction and data cache size doubled.
- A branch target buffer was added to reduce branch latency.

Overall, the instruction timings specified in Section 14.4, “Instruction Latencies” on page 14-3 are similar to the SA-110.

There are however, a few cases that may negatively impact the IPC number.

- Worst case branch latency increased from two to five cycles.
- The minimum result latency of loads that zero extend the result data increased from two to three cycles.
- The minimum result latency of an ALU operation followed by a shift operation increased from one to two cycles.

Keep in mind that the previous discussion was focusing on IPC differences. There is also the frequency difference, which is targeted to give the Intel® 80200 processor a 3X performance boost over SA-110 alone.

A.3.7 System Control Coprocessor

Additional bits and registers were added to CP15 to support the added functionality of the Intel® 80200 processor. Also, some system resources are controlled from CP14 registers. See Chapter 7, “Configuration”, for more information.

A.3.8 New Instructions and Instruction Formats

Chapter 2, “Programming Model”, discusses new instructions and instruction formats. These instructions would have generated Undefined faults on the SA-110.

A.3.9 Augmented Page Table Descriptors

Chapter 2, “Programming Model”, discusses how the Intel® 80200 processor augments the SA-110 descriptors.

Optimization Guide B

B.1 Introduction

This appendix contains optimization techniques for achieving the highest performance from the Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE). It is written for developers who are optimizing compilers or performance analysis tools for the Intel® 80200 processor 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® 80200 processor core, and hence can be applied to all products that are based on the Intel® 80200 processor core.

The Intel® 80200 processor architecture includes a superpipelined RISC architecture with an enhanced memory pipeline. The Intel® 80200 processor instruction set is based on ARM* V5TE architecture; however, the Intel® 80200 processor includes additional instructions. Code generated for the SA-110, SA-1100 and SA-1110 execute on the Intel® 80200 processors, however to obtain the maximum performance of your application code, it should be optimized for the Intel® 80200 processor architecture using the techniques presented in this document.

B.1.1 About This Guide

This guide assumes that you are familiar with the Intel® StrongARM* instruction set and the C language. It consists of the following sections:

Section B.1, “Introduction”. Outlines the contents of this guide.

Section B.2, “Intel® 80200 Processor Pipeline”. This chapter provides an overview of the Intel® 80200 processor pipeline behavior.

Section B.3, “Basic Optimizations”. This chapter outlines basic Intel® StrongARM* optimizations that can be applied to the Intel® 80200 processors.

Section B.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 Intel® 80200 processor.

Section B.5, “Instruction Scheduling”. This chapter shows how to optimally schedule code for the Intel® 80200 processor pipeline.

Section B.6, “Optimizing C Libraries”. This chapter contains information relating to optimizations for C library routines.

Section B.7, “Optimizations for Size”. This chapter contains optimizations that reduce the size of the generated code. Thumb* optimizations are also included.

B.2 Intel® 80200 Processor Pipeline

One of the biggest differences between the Intel® 80200 processor and first-generation Intel® StrongARM* processors is the pipeline. Many of the differences are summarized in Figure B-1. This section provides a brief description of the structure and behavior of the Intel® 80200 processor pipeline.

B.2.1 General Pipeline Characteristics

While the Intel® 80200 processor 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.

B.2.1.1. Number of Pipeline Stages

The Intel® 80200 processor 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 the Intel® 80200 processor pipeline has several negative consequences, however:

- Larger branch misprediction penalty (4 cycles in the Intel® 80200 processor instead of 1 in Intel® StrongARM*). 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 Intel® 80200 processor as compared to first generation Intel® StrongARM* processors (**LDM**, **STM**).
- Decode and register file lookups are spread out over 2 cycles in the Intel® 80200 processor, instead of 1 cycle in predecessors.

B.2.1.2. Intel® 80200 Processor Pipeline Organization

The Intel® 80200 processor single-issue superpipeline consists of a main execution pipeline, MAC pipeline, and a memory access pipeline. These are shown in Figure B-1, with the main execution pipeline shaded.

Figure B-1. Intel® 80200 Processor RISC Superpipeline

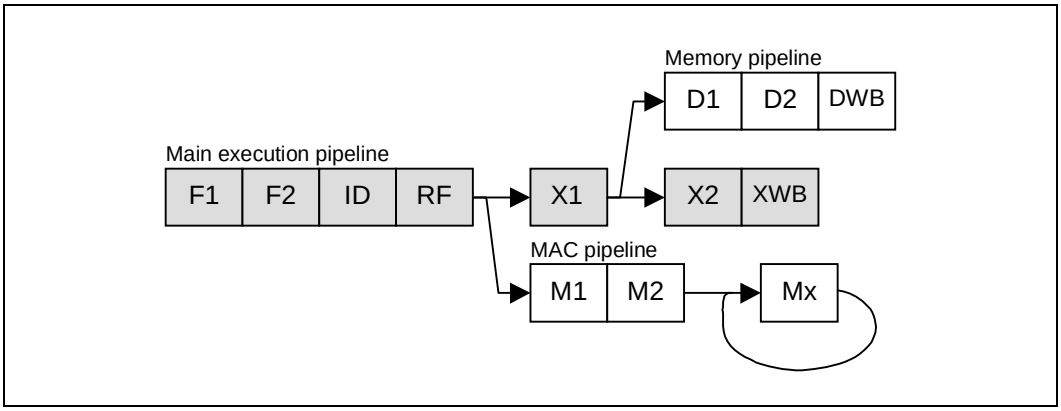

Table B-1 gives a brief description of each pipe-stage.

Table B-1. Pipelines and Pipe stages

Pipe / Pipestage	Description	Covered In
Main Execution Pipeline	Handles data processing instructions	Section B.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 B.2.4
D1/D2	Data Cache Access	“
DWB	Data cache writeback	“
MAC Pipeline	Handles all multiply instructions	Section B.2.5
M1-M5	Multiplier stages	“
MWB (not shown)	MAC write-back - may occur during M2-M5	“

B.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® 80200 processor 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.)

B.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 is not stalled. For example, if a load operation has missed the data cache, subsequent instructions that do not depend on the load may complete independently.

B.2.1.5. Use of Bypassing

The Intel® 80200 processor pipeline makes extensive use of bypassing to minimize data hazards. Bypassing allows results forwarding from multiple sources, eliminating the need to stall the pipeline.

B.2.2 Instruction Flow Through the Pipeline

The Intel® 80200 processor 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.

B.2.2.1. ARM* V5 Instruction Execution

Figure B-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 is idle.

All load/store instructions are routed to the memory pipeline after the effective addresses have been calculated in X1.

The ARM v5 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 v5 RISC instruction and execution resumes as usual.

B.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 Intel® 80200 processor pipeline. Performance degradation can be significant if care is not taken to minimize pipeline stalls.

B.2.3 Main Execution Pipeline

B.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 fill buffers.

B.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**.

B.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 selects the appropriate data from the register file and pass it to the next pipestage. When a register dependency does exist, the RFU keeps track of which register is unavailable and when the result is returned, the RFU stops 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.

B.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 does 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 restarts the pipeline

B.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.

B.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.

B.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.

B.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.

B.2.5 Multiply/Multiply Accumulate (MAC) Pipeline

The Multiply-Accumulate (MAC) unit executes the multiply and multiply-accumulate instructions supported by the Intel® 80200 processor core. The MAC implements the 40-bit Intel® 80200 processor 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.

B.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 14.4, “Instruction Latencies”.

An instruction that occupies the M1 or M2 pipestages 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 modifies architectural state regardless of subsequent events.

B.3 Basic Optimizations

This chapter outlines optimizations specific to ARM architecture. These optimizations have been modified to suit the Intel® 80200 processor architecture where needed.

B.3.1 Conditional Instructions

The Intel® 80200 processor architecture provides the ability to execute instructions conditionally. This feature combined with the ability of the Intel® 80200 processor instructions to modify the condition codes makes possible a wide array of optimizations.

B.3.1.1. Optimizing Condition Checks

Intel® 80200 processor instructions can selectively modify condition codes state. 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 need for a subsequent compare instruction. Consider 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, the code can be optimized as follows making use of the add instruction to set the 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 needs 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;
}
```

B.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 Intel® 80200 processors, a branch misprediction incurs a penalty of four cycles. If the branch is mispredicted 50% of the time, and if we assume 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.

$$\left(\frac{50}{100} \times 4 + \frac{3+4}{2}\right) = 5.5 \qquad cycles.$$

If we were to use Intel® 80200 processors 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
```

Assume 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:

$$\left(N1_C \times \frac{P1}{100}\right) + \left(N2_C \times \frac{100 - P1}{100}\right) \le \left(N1_B \times \frac{P1}{100}\right) + \left(N2_B \times \frac{100 - P1}{100}\right) + \left(\frac{P2}{100} \times 4\right)$$

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:

$$1 + \left(\frac{50}{100} \times 10\right) + \left(\frac{50}{100} \times 10\right) = 11 \qquad cycles$$

Cost of using branches:

$$1 + \left(\frac{50}{100} \times 7\right) + \left(\frac{50}{100} \times 6\right) + \left(\frac{50}{100} \times 4\right) = 9.5 \qquad cycles$$

As can be seen, we get better performance by using branch instructions in the above scenario.

B.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
    cmpne 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.

B.3.2 Bit Field Manipulation

The Intel® 80200 processor 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
```

B.3.3 Optimizing the Use of Immediate Values

The Intel® 80200 processor **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 0xffffefb.
    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 that it is possible to load any 32-bit value into a register using a sequence of four instructions.

B.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 2^n
    mov   r0, r0, LSL #n
;Multiplication of R0 by 2^n+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 (2^n+1)*(2^m)
    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 2^n
    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 2^n
    mov   r1, r0, ASR #31
    add   r0, r0, r1, LSR #(32 - n)
    mov   r0, r0, ASR #n
```

The add instruction would stall for 1 cycle. The stall can be prevented by filling in another instruction before add.

B.3.5 Effective Use of Addressing Modes

The Intel® 80200 processor 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]!
```

B.4 Cache and Prefetch Optimizations

This chapter 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.

B.4.1 Instruction Cache

The Intel® 80200 processor 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.

B.4.1.1. Cache Miss Cost

The Intel® 80200 processor performance is highly dependent on reducing the cache miss rate. When an instruction cache miss occurs, the timing to retrieve the next instruction is the same as that for retrieving data for the data cache. Section B.4.4.1., “Prefetch Distances in the Intel® 80200 Processor” provides a more detailed explanation of the required time. Using the same assumptions as those used for the data caches, the result is it takes about 60 to 90 core cycles to retrieve the first instruction. Once the first 8-byte word is read, it takes another six core cycles to read in the next two instructions or a total of 78 to 108 clocks to fill a cache line. If the new instructions each execute in one core cycle, then the processor is stalled for 4 cycles waiting for the next pair of instructions. Further, if the next pair of instructions each execute in one cycle each, the processor is again stalled for 4 more cycles. From this it is clear that executing non-cached instructions severely curtails the processor's performance. It is very important to do everything possible to minimize cache misses.

B.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 is 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.

B.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, thrashes 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 comes from profiling followed by compiler based two pass optimizations.

B.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 is 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

B.4.2 Data and Mini Cache

The Intel® 80200 processor 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® 80200 processor processor 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.

B.4.2.1. Non Cacheable Regions

It is recommended that non-cacheable memory (X=0, C=0, and B=0) be used only if necessary or 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.

B.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 is 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.

B.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 write and read 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 is not allocated, and hence does not evict critical data from the Data cache.

B.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 does 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 Intel® 80200 processor round robin replacement policy, all data is eventually 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). If the data in the on-chip RAM is to be initialized to zero, then the locking process can be speed up by using the CP15 prefetch zero function. This function does not generate external memory references. See the Intel® 80200 processor reference manual for more information on how to do this.

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 have more ways locked than the others. This uneven allocation increases 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 assume that the address of arr1 is 0, address of arr2 is 1024, and the address of arr3 is 2048. All three arrays are within the same sets, i.e. set0 and set2, as a result three ways in both sets set0 and set1, are locked, leaving 28 ways for use by other variables.

This can 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 8.

B.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 reduces 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 thrashes 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 quickly thrashes the cache.

B.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 halve 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 requires 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(predata+=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 adds extra computational overhead.

Additional prefetch considerations are discussed in greater detail in following sections.

B.4.2.7. Literal Pools

The Intel® 80200 processor 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 Intel® 80200 processor 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 B.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 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 are available immediately from the data cache.

B.4.3 Cache Considerations

B.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 B.4.4.2., "Prefetch Loop Scheduling" below.

B.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 continually trash the memory page selection.

B.4.4 Prefetch Considerations

The Intel® 80200 processor 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 Intel® 80200 processor. 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® 80200 processor 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 Intel® 80200 processor prefetch can be used to reduce register pressure instead of increasing it.

The Intel® 80200 processor prefetch load is a hint instruction and does not guarantee that the data is loaded. Whenever the load would cause a fault or a table walk, then the processor ignores 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 does not fault program flow.

B.4.4.1. Prefetch Distances in the Intel® 80200 Processor

Scheduling the prefetch instruction requires understanding the system latency times and system resources which affect when to use the prefetch instruction. This section considers three timing elements:

N_{cwf} critical word first
N_{clxfer} full cache line transfer time
$N_{subissue}$ subsequent prefetch issue time to insure uninterrupted transfers

The memory latency times presented here assume typical SDRAM that is currently available and working with the Intel® 80200 processor. It is assumed that the SDRAM supports “Critical Word First” transfers. That is when a cache line is being transferred, the first word transferred corresponds to the one needed by the processor immediately as opposed to transferring the data from lowest address first.

The cycle times assume that the core is running six times a fast as the memory transfer bus. Further, the example values presented here apply to the current processor implantation, (TBD processor name) and are different for future implementations.

N_{cwf} is the number of core cycles required to transfer the first critical word of a prefetch or load operation:

$$N_{cwf} = N_{lookup} + N_{cwfxfer}$$

Where:

N_{lookup} This is the number of core clocks required for the processor to issue a memory transfer request to the SDRAM plus the time the SDRAM requires to locate the data.

$$N_{lookup} = N_{processor} + N_{memwait} + N_{mempagewait}$$

The Intel® 80200 processor needs seven bus clocks to process a memory request to the SDRAM ($N_{processor}$). Typical SDRAM needs 2 to 3 bus clocks to select the memory locations provided that the current SDRAM memory page is selected ($N_{memwait}$). If the current SDRAM memory page is not selected, then an additional 3 to 4 bus cycles are required to lookup the memory data locations ($N_{mempagewait}$). Thus the lookup time can range from 9 to 14 bus clock cycles. Translating this to core cycles at a ratio of six to one means between 54 and 84 core clocks.

$N_{cwfxfer}$ This is the number of core clocks required to transfer the first critical word of a cache line fill operation. It takes one bus clock to transfer the first word if the data is in the lower word address of the transfer and one additional core clock if the word is in the upper word address range of the transfer. Thus for the examples presented here this would be 6 or 7 core clock cycles.

N_{cwf} for the Intel® 80200 processor works out to be 60 instructions assuming 2 wait state SDRAM and that the current SDRAM memory page is selected. The second 64 bits of data are available at the next bus cycle or 6 core clocks

N_{clxfer} is the minimal number of cycles to prefetch ahead for an entire cache line:

$$N_{clxfer} = N_{lookup} + N_{linexfer}$$

Where:

$N_{linexfer}$ This is the number of core clocks required to transfer one complete cache line. The Intel® 80200 processor requires 4 bus cycles to transfer four 64 bit words of a full cache line. Given the six to one core to bus clock ratio this would be 24 core clock cycles.

N_{clxfer} works out to be about 78 cycles for the Intel® 80200 processor when using 2 bus cycle wait state

$N_{subissue}$ This is the maximum number of core clocks that a subsequent bus transfer request must be made to guarantee that transfer takes place immediately after the previous request has completed its transfer. If a transfer is not made in this time, then idle bus cycles occur reducing efficiently. This time transfer time of the previous request. If the previous transfer was for a full cache line read or write, then this would take 24 core cycles at a six to one ratio between core and bus clocks. If the previous operation was for a half cache line, then this would be done in 12 core clocks.

Consider the following code sample:

```
    add   r1, r1, #1
; Sequence of instructions that use r2. These instructions leave r3 unchanged.
    ldr   r2, [r3]
    add   r3, r3, #4
    mov   r4, r3
    sub   r2, r2, #1
```

The sub instruction above would stall if the data being loaded misses the cache. These stalls can be avoided by using a pld instruction well ahead as shown below. The number of instructions required to insure a stall does not occur is proportional to N_{cwf} for a given system.

```
    pld   [r3]
    add   r1, r1, #1
; Sequence of instructions that use r2. These instructions leave r3 unchanged.
    ldr   r2, [r3]
    add   r3, r3, #4
    mov   r4, r3
    sub   r2, r2, #1
```

B.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 refereed to as the prefetch scheduling distance or psd. For the Intel® 80200 processor this can be calculated as:

$$psd = floor\left(\frac{(N_{lookup} + N_{linexfer} \times N_{pref} + N_{hwlinexfer} \times N_{evict})}{(CPI \times N_{inst})}\right)$$

Where:

N_{pref} Is the number of cache lines to be prefetched for both reading and writing.

N_{evict} Is the number of cache half line evictions caused by the loop.

N_{inst} Is the number of instructions executed in one iteration of the loop

$N_{hwlinexfer}$ This is the number of core clocks required to write half a cache line as would happen if only one of the cache line dirty bits were set when a line eviction occurred. For the Intel® 80200 processor this takes 2 bus clocks or 12 core clocks.

CPI This is the average number of core clocks per instruction.

The *psd* number provided by the above equation is a good starting point, but may not be the most ideal consideration. Estimating N_{evict} is very difficult from static code. However, if the operational data uses the mini-data cache and if the loop operations should overflow the mini-data cache, then a first order estimate of N_{evict} would be the number of bytes written pre loop iteration divided by a half cache line size of 16 bytes. Cache overflow can be estimated by the number of cache lines transferred each iteration and the number of expected loop iterations. N_{evict} and CPI can be estimated by profiling the code using the performance monitor “cache write-back” event count.

B.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.

B.4.4.4. Compute vs. Data Bus Bound

At the extreme, a loop, which is data bus bound, does 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.

B.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.

B.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 Intel® 80200 processor 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 are 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 also requires 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 uses 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 is 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.

B.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 Intel® 80200 processor 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 Intel® 80200 processor, 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;
};
```

B.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 subsequence 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; j<100; j++)
        for(k1=0; k<100; k++)
            for(j2=0; j<100; j++)
                for(k2=0; k<100; k++)
                {
                    j = j1 * 100 + j2;
                    k = k1 * 100 + k2;
                    C[j][k] += A[i][k] * B[j][i];
                }
```

B.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 affect 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, prefetches superfluous data. The problem can be avoid 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.

B.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.

B.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];
    }
```

B.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];
}
```

B.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.

B.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.

B.5.1 Scheduling Loads

On the Intel® 80200 processor, 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 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 code sample above, **ADD** and **LDR** instruction can be moved before the **MOV** instruction. Note this would prevent pipeline stalls if the load hits the data cache. However, if load is likely to miss 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
```

Intel® 80200 processor has four fill-buffers used to fetch data from external memory when a data-cache miss occurs. Intel® 80200 processor stalls when all fill buffers are in use. This happens when more than four loads are outstanding and are being fetched from memory. As a result, code written should ensure no more than four loads are outstanding at same time. For example, number of loads issued sequentially should not exceed four. Also note, a preload instruction may cause fill buffer to be used. As a result, number of preload instructions outstanding should also be considered to arrive at number of loads that are outstanding.

Similarly, number of write buffers also limits number of successive writes issued before the processor stalls. No more than eight stores can be issued. Also note, if data caches are using write-allocate with writeback policy, then a load operation may cause stores to external memory if read operation evicts a cache line that is dirty (modified). The number of sequential stores may be limited by this fact.

B.5.1.1. Scheduling Load and Store Double (LDRD/STRD)

The Intel® 80200 processor 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]
```

B.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. Assume 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
    ldrd  r2, [r0]
    ldrd  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
```

B.5.2 Scheduling Data Processing Instructions

Most Intel® 80200 processor 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 RRX. 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
```

B.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 14.4, “Instruction Latencies” to get the instruction latencies for various multiply instructions. The multiply instructions should be scheduled taking into consideration these instruction latencies.

B.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]
```

B.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.

B.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
```

B.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.

B.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
```

B.6 Optimizing C Libraries

Many of the standard C library routines can benefit greatly by being optimized for the Intel® 80200 processor architecture. The following string and memory manipulation routines should be tuned to obtain the best performance from the Intel® 80200 processor architecture (instruction selection, cache usage and data prefetch):

strcat, strchr, strcmp, strcoll, strcpy, strcspn, strlen, strncat, strncmp, strpbrk, strrchr, strspn, strstr, strtok, strxfrm, memchr, memcmp, memcpy, memmove, memset

B.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, in general, lowers the performance of your application. This chapter contains techniques for optimizing for code size using the Intel® 80200 processor instruction set.

B.7.1 Space/Performance Trade Off

Many optimizations mentioned in the previous chapters improve the performance of ARM code. However, using these instructions results in increased code size. Use the following optimizations to reduce the space requirements of the application code.

B.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.

B.7.1.2. Use of Conditional Instructions

Using conditional instructions to expand if-then-else statements as described in Section B.3.1, “Conditional Instructions” results in increasing the size of the generated code. Therefore, do not use conditional instructions if application code space requirements are an issue.

B.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 C

The Intel® 80200 processor based on Intel® XScale™ microarchitecture (compliant with the ARM* Architecture V5TE) implements Design For Test (DFT) techniques to ensure quality and reliability. This appendix describes those techniques.

C.1 Introduction

Testing VLSI circuits is critical for achieving high outgoing quality levels. Unfortunately the cost of testing is already one of the largest portions of the final product cost and is getting more expensive as the complexity of VLSI chips grows. Reducing test cost is thus one of the critical issues that must be resolved in order to make a cost competitive VLSI chip, and area and/or delay minimization at the cost of increased test time may not be a good design choice.

Testability is the ability or ease with which tests can be generated for and applied to CUT (Circuit Under Test). Design For Testability (DFT) stands for having design decisions early in the design cycle that ease the testing, test generation, and fault grading process for a given circuit.

C.2 JTAG - IEEE1149.1

The goal of JTAG also known as IEEE1149.1 is to ensure that chips containing a common denominator of DFT circuitry makes testing of boards containing these chips significantly less costly and more effective. JTAG consists of TAP controller, Boundary-Scan register, instruction and data registers, and dedicated signals including **TDI**, **TDO**, **TCK**, **TMS**, and **TRST#**.

The Intel® 80200 processor provides test features compatible with IEEE Standard Test Access Port and Boundary Scan Architecture (IEEE Std. 1149.1). JTAG ensures that components function correctly, connections between components are correct, and components interact correctly on the printed circuit board.

C.2.1 Boundary Scan Architecture

Boundary scan test logic consists of a Boundary-Scan register and support logic. These are accessed through a Test Access Port (TAP). The TAP provides a simple serial interface that allows all processor signal pins to be driven and/or sampled, thereby providing the direct control and monitoring of processor pins at the system level.

This mode of operation is valuable for design debugging and fault diagnosis since it permits examination of connections not normally accessible to the test system. The following subsections describe the boundary scan test logic elements: TAP pins, instruction register, test data registers, and TAP controller. Figure C-1 illustrates how these pieces fit together to form the JTAG unit.

To ensure that the processor does not enter an invalid state during boundary scan, it must have received a valid reset (on **RESET#**) including a valid clock input (on **clk**).

Figure C-1. Test Access Port Block Diagram

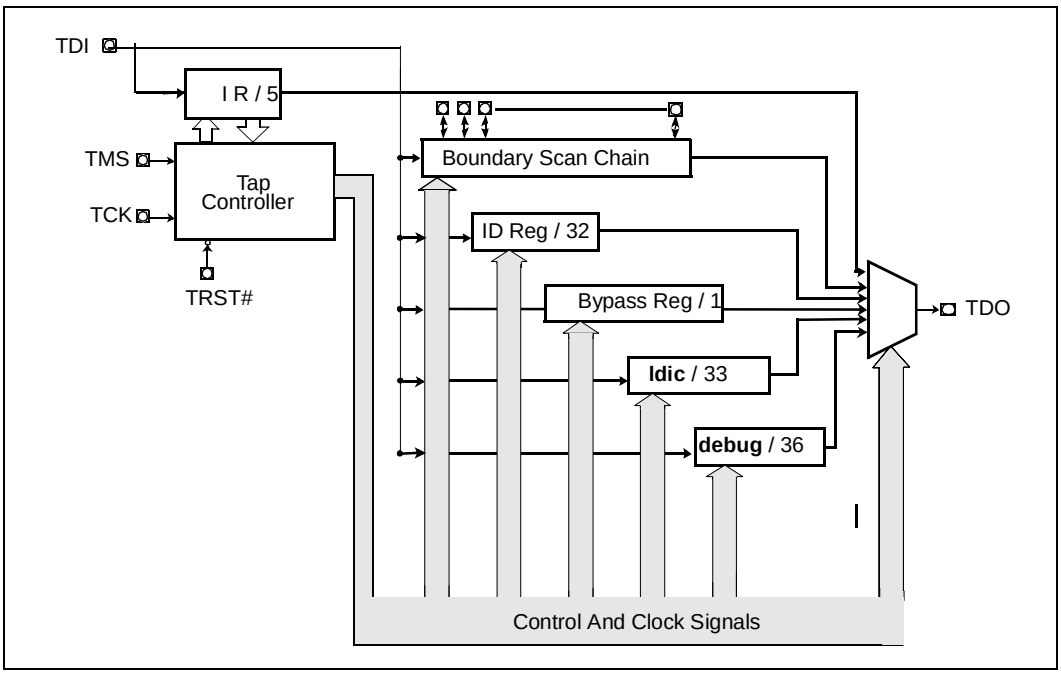

C.2.2 TAP Pins

The Intel® 80200 processor TAP is composed of four input connections (TMS, TCK, TRST# and TDI) and one output connection (TDO). These pins are described in Table C-1. The TAP pins provide access to the instruction register and the test data registers.

Table C-1. TAP Controller Pin Definitions

Pin Name	Mnemonic	Type	Definition
Test Clock	TCK	Input	Clock input for the TAP controller, the instruction register, and the test data registers. The JTAG unit retains its state when TCK is stopped at “0” or “1”.
Test Mode Select	TMS	Input	Controls the operation of the TAP controller. The TMS input is pulled high when not being driven. TMS is sampled on the rising edge of TCK.
Test Data In	TDI	Input	Serial date input to the instruction and test data registers. Data at TDI is sampled on the rising edge of TCK. Like TMS, TDI is pulled high when not being driven. Data shifted from TDI through a register to TDO appears non-inverted at TDO.
Test Data Out	TDO	Output	Used for serial data output. Data at TDO is driven at the falling edge of TCK and provides an inactive (high-Z) state when scanning is not in progress. The non-shift inactive state is provided to support parallel connection of TDO outputs at the board or module level.
Asynchronous Reset	TRST#	Input	Provides asynchronous initialization of the test logic. TRST# is pulled low when not being driven. Assertion of this pin puts the TAP controller in the Test_Logic_Reset (initial) state. An external source drives this signal high for TAP controller operation.

C.2.3 Instruction Register (IR)

The instruction register holds instruction codes shifted through the Test Data Input (TDI) pin. The instruction codes are used to select the specific test operation to be performed and the test data register to be accessed.

The instruction register is a parallel-loadable, master/slave-configured 5-bit wide, serial-shift register with latched outputs. Data is loaded into the IR serially through the TDI pin clocked by the rising edge of TCK when the TAP controller is in the Shift_IR state. The shifted-in instruction becomes active upon latching from the master-stage to the slave-stage in the Update_IR state. At that time the IR outputs along with the TAP finite state machine outputs are decoded to select and control the test data register selected by that instruction. Upon latching, all actions caused by any previous instructions must terminate.

The instruction determines the test to be performed, the test data register to be accessed, or both (Table C-2). The IR is five bits wide. When the IR is selected in the Shift_IR state, the most significant bit is connected to TDI, and the least significant bit is connected to TDO. TDI is shifted into IR on each rising edge of TCK, as long as TMS remains asserted. When the processor enters Capture_IR TAP controller state, fixed parallel data (0001_2) is captured. During Shift_IR, when a new instruction is shifted in through TDI, the value 0001_2 is always shifted out through TDO least significant bit first. This helps identify instructions in a long chain of serial data from several devices.

Upon activation of the TRST# pin, the latched instruction asynchronously changes to the **idcode** instruction. If the TAP controller moved into the Test_Logic_Reset state other than by reset activation, the opcode changes as TDI is shifted, and becomes active on the falling edge of TCK. See Figure C-4 for an example of loading the instruction register.

C.2.3.1. Boundary-Scan Instruction Set

The Intel® 80200 processor supports three mandatory boundary scan instructions (**bypass**, **sample**/**preload** and **extest**). The Intel® 80200 processor also contains seven additional public instructions along with seven Intel® 80200 processor private instructions. Table C-2 lists the Intel® 80200 processor instruction codes and Table C-3 describes each instruction.

Table C-2. JTAG Instruction Set

Instruction Code	Instruction Name	Instruction Code	Instruction Name
00000_2	**extest**	01011_2	private
00001_2	**sample**	01100_2	private
00010_2	dbgrx	01101_2	private
00011_2	private	01110_2	not used
00100_2	**clamp**	01111_2	not used
00101_2	private	10000_2	dbgtx
00110_2	not used	10001_2	not used
00111_2	ldic	10010_2	not used
01000_2	**highz**	10011_2 through 11101_2	not used
01001_2	dcsr	11110_2	**idcode**
01010_2	private	11111_2	bypass

Table C-3. IEEE Instructions

Instruction / Requisite	Opcode	Description
extest IEEE 1149.1 Required	00000_2	**extest** initiates testing of external circuitry, typically board-level interconnects and off chip circuitry. **extest** connects the Boundary-Scan register between TDI and TDO in the Shift_DR state only. When **extest** is selected, all output signal pin values are driven by values shifted into the Boundary-Scan register and may change only on the falling-edge of TCK in the Update_DR state. Also, when **extest** is selected, all system input pin states must be loaded into the Boundary-Scan register on the rising-edge of TCK in the Capture_DR state. Values shifted into input latches in the Boundary-Scan register are never used by the processor's internal logic.
sample IEEE 1149.1 Required	00001_2	**sample**/**preload** performs two functions: • When the TAP controller is in the Capture-DR state, the **sample** instruction occurs on the rising edge of TCK and provides a snapshot of the component's normal operation without interfering with that normal operation. The instruction causes Boundary-Scan register cells associated with outputs to **sample** the value being driven by or to the processor. • When the TAP controller is in the Update-DR state, the **preload** instruction occurs on the falling edge of TCK. This instruction causes the transfer of data held in the Boundary-Scan cells to the slave register cells. Typically the slave latched data is then applied to the system outputs by means of the **extest** instruction.
dbgrx	00010_2	See Chapter 13, "Software Debug".
clamp	00100_2	**clamp** instruction allows the state of the signals driven from Intel® 80200 processor pins to be determined from the boundary-scan register while the Bypass register is selected as the serial path between TDI and TDO. Signals driven from the component pins do not change while the **clamp** instruction is selected.
ldic	00111_2	See Chapter 13, "Software Debug".
highz	01000_2	The execution of **highz** generates a signal that is read on the rising-edge of $\overline{\text{RESET}}$. If this signal is found asserted, the device floats all its output pins. Also, when this instruction is active, the Bypass register is connected between TDI and TDO. This register can be accessed via the JTAG Test-Access Port throughout the device operation. Access to the Bypass register can also be obtained with the **bypass** instruction.
dcsr	01001_2	See Chapter 13, "Software Debug".
idcode IEEE 1149.1 Optional	11110_2	**idcode** is used in conjunction with the device identification register. It connects the identification register between TDI and TDO in the Shift_DR state. When selected, **idcode** parallel-loads the hard-wired identification code (32 bits) on TDO into the identification register on the rising edge of TCK in the Capture_DR state. **NOTE:** The device identification register is not altered by data being shifted in on TDI.
dbgtx	10000_2	See Chapter 13, "Software Debug".
bypass IEEE 1149.1 Required	11111_2	**bypass** instruction selects the Bypass register between TDI and TDO pins while in SHIFT_DR state, effectively bypassing the processor's test logic. 0_2 is captured in the CAPTURE_DR state. While this instruction is in effect, all other test data registers have no effect on the operation of the system. Test data registers with both test and system functionality perform their system functions when this instruction is selected.

C.2.4 TAP Test Data Registers

The Intel® 80200 processor contains a device identification register and two test data registers (Bypass and RUNBIST). Each test data register selected by the TAP controller is connected serially between TDI and TDO. TDI is connected to the test data register's most significant bit. TDO is connected to the least significant bit. Data is shifted one bit position within the register towards TDO on each rising edge of TCK. The following sections describe each of the test data registers. See Figure C-5 for an example of loading the data register.

C.2.4.1. Device Identification Register

The Device Identification register is a 32-bit register containing the manufacturer's identification code, part number code and version code. The identification register is selected only by the **idcode** instruction. When the TAP controller's Test_Logic_Reset state is entered, **idcode** is automatically loaded into the instruction register. The Device Identification register has a fixed parallel input value that is loaded in the Capture_DR state.

The value of this register is shown in Table C-4.

Table C-4. JTAG ID Register Value

Stepping	Value
A0	0x09263013
B0	0x29263013

C.2.4.2. Bypass Register

The required Bypass Register, a one-bit shift register, provides the shortest path between TDI and TDO when either of a **bypass**, **highz** or **clamp** instructions are in effect. This allows rapid movement of test data to and from other components on the board. This path can be selected when no test operation is being performed. While the Bypass register is selected, data is transferred from TDI to TDO without inversion.

Any instruction that does not make use of another test data register may select the Bypass register as its active TDI to TDO path.

C.2.4.3. Boundary-Scan Register

The Boundary-Scan register is a required set of serial-shiftable register cells.

C.2.5 TAP Controller

The TAP controller is a 16-state synchronous finite state machine that controls the sequence of test logic operations. The TAP can be controlled via a bus master. The bus master can be either automatic test equipment or a component (i.e. PLD) that interfaces to the Test Access Port (TAP). The TAP controller changes state only in response to a rising edge of TCK or power-up. The value of the test mode state (TMS) input signal at a rising edge of TCK controls the sequence of state changes. The TAP controller is automatically initialized on powerup. In addition, the TAP controller can be initialized by applying a high signal level on the TMS input for five TCK periods.

Behavior of the TAP controller and other test logic in each controller state is described in the following subsections. For greater detail on the state machine and the public instructions, refer to IEEE 1149.1 Standard Test Access Port and Boundary-Scan Architecture Document.

Figure C-2. TAP Controller State Diagram

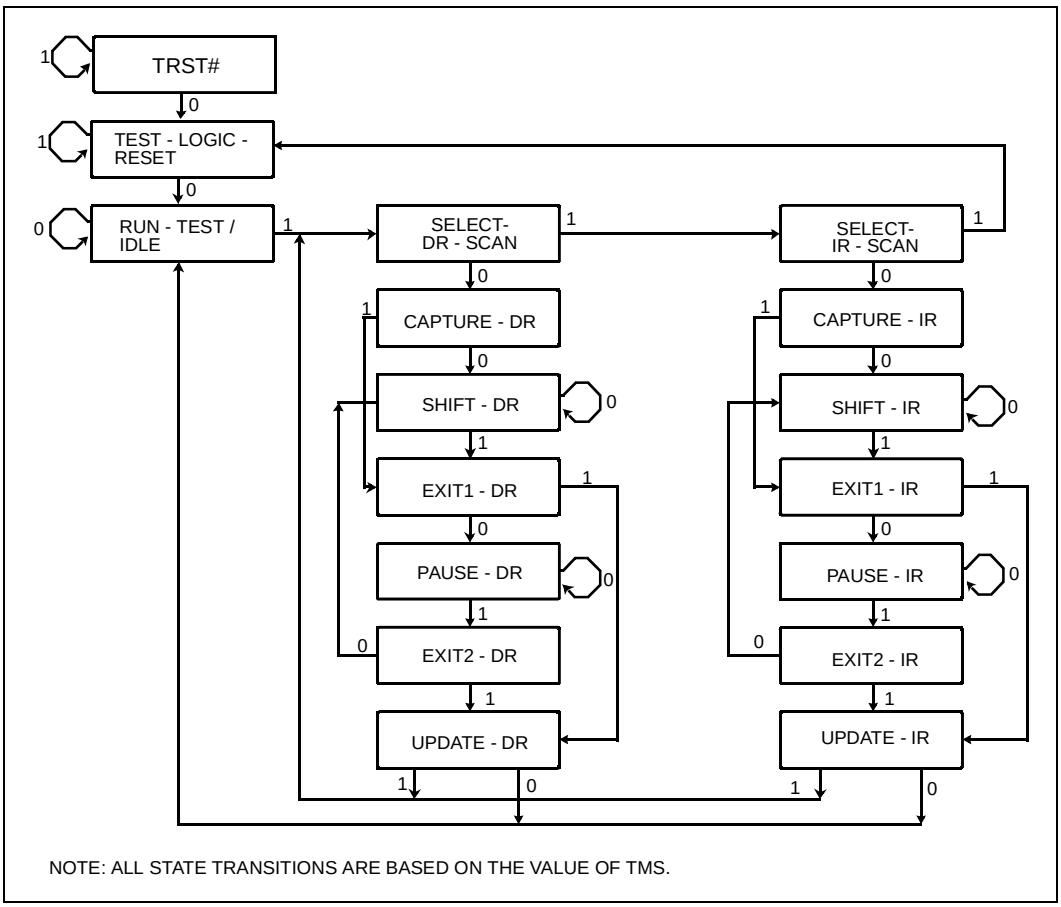

C.2.5.1. Test Logic Reset State

In this state, test logic is disabled to allow normal operation of the Intel® 80200 processor. Test logic is disabled by loading the **idcode** register. No matter what the state of the controller, it enters Test-Logic-Reset state when the TMS input is held high (1) for at least five rising edges of TCK. The controller remains in this state while TMS is high. The TAP controller is also forced to enter this state by enabling TRST#.

If the controller exits the Test-Logic-Reset controller states as a result of an erroneous low signal on the TMS line at the time of a rising edge on TCK (for example, a glitch due to external interference), it returns to the test logic reset state following three rising edges of TCK with the TMS line at the intended high logic level. Test logic operation is such that no disturbance is caused to on-chip system logic operation as the result of such an error.

C.2.5.2. Run-Test/Idle State

The TAP controller enters the Run-Test/Idle state between scan operations. The controller remains in this state as long as TMS is held low. In the Run-Test/Idle state the **runbist** instruction is performed; the result is reported in the RUNBIST register. Instructions that do not call functions generate no activity in the test logic while the controller is in this state. The instruction register and all test data registers retain their current state. When TMS is high on the rising edge of TCK, the controller moves to the Select-DR-Scan state.

C.2.5.3. Select-DR-Scan State

The Select-DR-Scan state is a temporary controller state. The test data registers selected by the current instruction retain their previous state. If TMS is held low on the rising edge of TCK when the controller is in this state, the controller moves into the Capture-DR state and a scan sequence for the selected test data register is initiated. If TMS is held high on the rising edge of TCK, the controller moves into the Select-IR-Scan state.

The instruction does not change while the TAP controller is in this state.

C.2.5.4. Capture-DR State

When the controller is in this state and the current instruction is **sample**/**preload**, the Boundary-Scan register captures input pin data on the rising edge of TCK.Test data registers that do not have parallel input are not changed. Also if the **sample**/**preload** instruction is not selected while in this state, the Boundary-Scan registers retain their previous state.

The instruction does not change while the TAP controller is in this state.

If TMS is high on the rising edge of TCK, the controller enters the Exit1-DR. If TMS is low on the rising edge of TCK, the controller enters the Shift-DR state.

C.2.5.5. Shift-DR State

In this controller state, the test data register, which is connected between TDI and TDO as a result of the current instruction, shifts data one bit position nearer to its serial output on each rising edge of TCK. Test data registers that the current instruction selects but does not place in the serial path, retain their previous value during this state.

The instruction does not change while the TAP controller is in this state.

If TMS is high on the rising edge of TCK, the controller enters the Exit1-DR state. If TMS is low on the rising edge of TCK, the controller remains in the Shift-DR state.

C.2.5.6. Exit1-DR State

This is a temporary controller state. When the TAP controller is in the Exit1-DR state and TMS is held high on the rising edge of TCK, the controller enters the Update-DR state, which terminates the scanning process. If TMS is held low on the rising edge of TCK, the controller enters the Pause-DR state.

The instruction does not change while the TAP controller is in this state. All test data registers selected by the current instruction retain their previous value during this state.

C.2.5.7. Pause-DR State

The Pause-DR state allows the test controller to temporarily halt the shifting of data through the test data register in the serial path between TDI and TDO. The test data register selected by the current instruction retains its previous value during this state. The instruction does not change in this state.

The controller remains in this state as long as TMS is low. When TMS goes high on the rising edge of TCK, the controller moves to the Exit2-DR state.

C.2.5.8. Exit2-DR State

This is a temporary state. If TMS is held high on the rising edge of TCK, the controller enters the Update-DR state, which terminates the scanning process. If TMS is held low on the rising edge of TCK, the controller enters the Shift-DR state.

The instruction does not change while the TAP controller is in this state. All test data registers selected by the current instruction retain their previous value during this state.

C.2.5.9. Update-DR State

The Boundary-Scan register is provided with a latched parallel output. This output prevents changes at the parallel output while data is shifted in response to the **extest**, **sample/preload** instructions. When the Boundary-Scan register is selected while the TAP controller is in the Update-DR state, data is latched onto the Boundary-Scan register's parallel output from the shift-register path on the falling edge of TCK. The data held at the latched parallel output does not change unless the controller is in this state.

While the TAP controller is in this state, all of the test data register's shift-register bit positions selected by the current instruction retain their previous values.

The instruction does not change while the TAP controller is in this state.

When the TAP controller is in this state and TMS is held high on the rising edge of TCK, the controller enters the Select-DR-Scan state. If TMS is held low on the rising edge of TCK, the controller enters the Run-Test/Idle state.

C.2.5.10. Select-IR Scan State

This is a temporary controller state. The test data registers selected by the current instruction retain their previous state. In this state, if TMS is held low on the rising edge of TCK, the controller moves into the Capture-IR state and a scan sequence for the instruction register is initiated. If TMS is held high on the rising edge of TCK, the controller moves to the Test-Logic-Reset state.

The instruction does not change in this state.

C.2.5.11. Capture-IR State

When the controller is in the Capture-IR state, the shift register contained in the instruction register loads the fixed value 0001_2 on the rising edge of TCK.

The test data register selected by the current instruction retains its previous value during this state. The instruction does not change in this state. While in this state, holding TMS high on the rising edge of TCK causes the controller to enter the Exit1-IR state. If TMS is held low on the rising edge of TCK, the controller enters the Shift-IR state.

C.2.5.12. Shift-IR State

When the controller is in this state, the shift register contained in the instruction register is connected between TDI and TDO and shifts data one bit position nearer to its serial output on each rising edge of TCK. The test data register selected by the current instruction retains its previous value during this state. The instruction does not change.

If TMS is held high on the rising edge of TCK, the controller enters the Exit1-IR state. If TMS is held low on the rising edge of TCK, the controller remains in the Shift-IR state.

C.2.5.13. Exit1-IR State

This is a temporary state. If TMS is held high on the rising edge of TCK, the controller enters the Update-IR state, which terminates the scanning process. If TMS is held low on the rising edge of TCK, the controller enters the Pause-IR state.

The test data register selected by the current instruction retains its previous value during this state.

The instruction does not change and the instruction register retains its state.

C.2.5.14. Pause-IR State

The Pause-IR state allows the test controller to temporarily halt the shifting of data through the instruction register. The test data registers selected by the current instruction retain their previous values during this state.

The instruction does not change and the instruction register retains its state.

The controller remains in this state as long as TMS is held low. When TMS goes high on the rising edges of TCK, the controller moves to the Exit2-IR state.

C.2.5.15. Exit2-IR State

This is a temporary state. If TMS is held high on the rising edge of TCK, the controller enters the Update-IR state, which terminates the scanning process. If TMS is held low on the rising edge of TCK, the controller enters the Shift-IR state.

This test data register selected by the current instruction retains its previous value during this state. The instruction does not change and the instruction register retains its state.

C.2.5.16. Update-IR State

The instruction shifted into the instruction register is latched onto the parallel output from the shift-register path on the falling edge of TCK. Once latched, the new instruction becomes the current instruction. Test data registers selected by the current instruction retain their previous values.

If TMS is held high on the rising edge of TCK, the controller enters the Select-DR-Scan state. If TMS is held low on the rising edge of TCK, the controller enters the Run-Test/Idle state.

C.2.5.17. Boundary-Scan Example

In the example that follows, two command actions are described. The example starts in the reset state, a new instruction is loaded and executed. See Figure C-3 for a JTAG example. The steps are:

1. Load the **sample**/**preload** instruction into the Instruction Register:
   a. Select the Instruction register scan.
   b. Use the Shift-IR state four times to read the least through most significant instruction bits into the instruction register (we do not care that the old instruction is being shifted out of the TDO pin).
   c. Enter the Update-IR state to make the instruction take effect.
   d. Exit the Instruction register.
2. Capture and shift the data onto the TDO pin:
   a. Select the Data register scan state.
   b. Capture the pin information into the n-stage Boundary-Scan register.
   c. Enter and stay in the shift-DR state for n cycles. These TDO values are compared against expected data to determine if component operation and connection are correct. Record the TDO value after each cycle. New serial data enters the boundary-scan register through the TDI pin, while old data is scanned out.
   d. Pass through the Exit1-DR and Update-DR to continue.

This example does not make use of the pause states. Those states would be more useful where we do not control the clock directly. The pause states let the clock tick without affecting the shift registers.

The old instruction was *abcd* in the example. It is known that the original value is the ID code since the example starts from the reset state. Other times it represents the previous opcode. The new instruction opcode is 0001_2 (**sample**/**preload**). All pins are captured into the serial Boundary-Scan register and the values are output to the TDO pin.

The clock signal drawn at the top of the diagram is drawn as a stable symmetrical clock. This is not in practice the most common case. Instead the clocking is usually done by a program writing to a port bit. The TMS and TDI signals are written by software and then the software makes the clock go high. The software typically often lowers the clock input quickly. The program can then read the TDO pin.

Figure C-3. JTAG Example

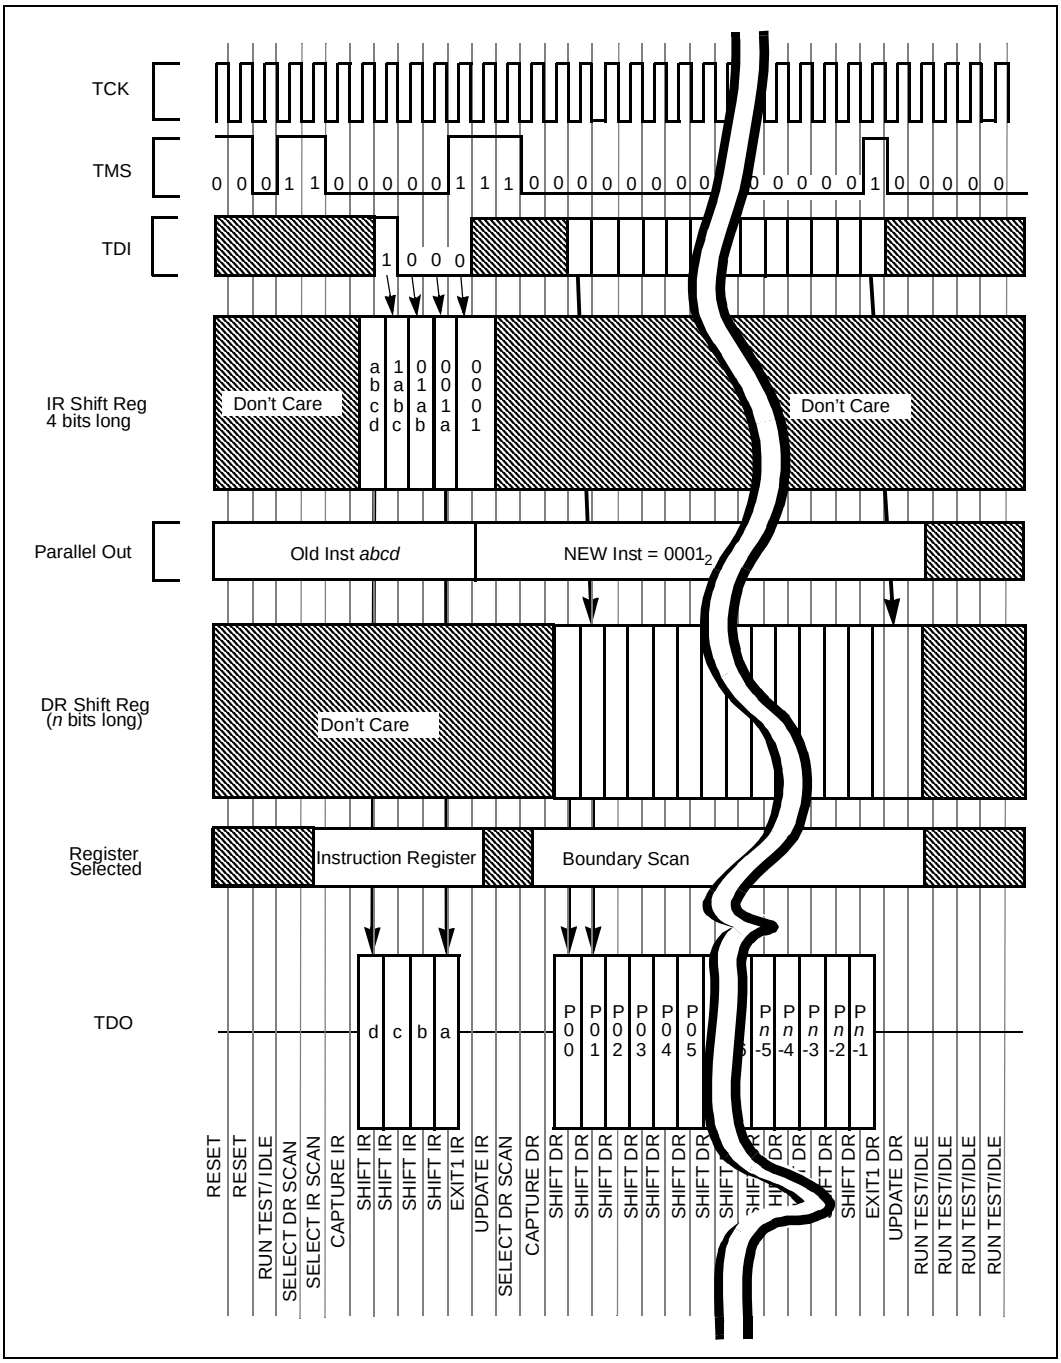

Figure C-4. Timing Diagram Illustrating the Loading of Instruction Register

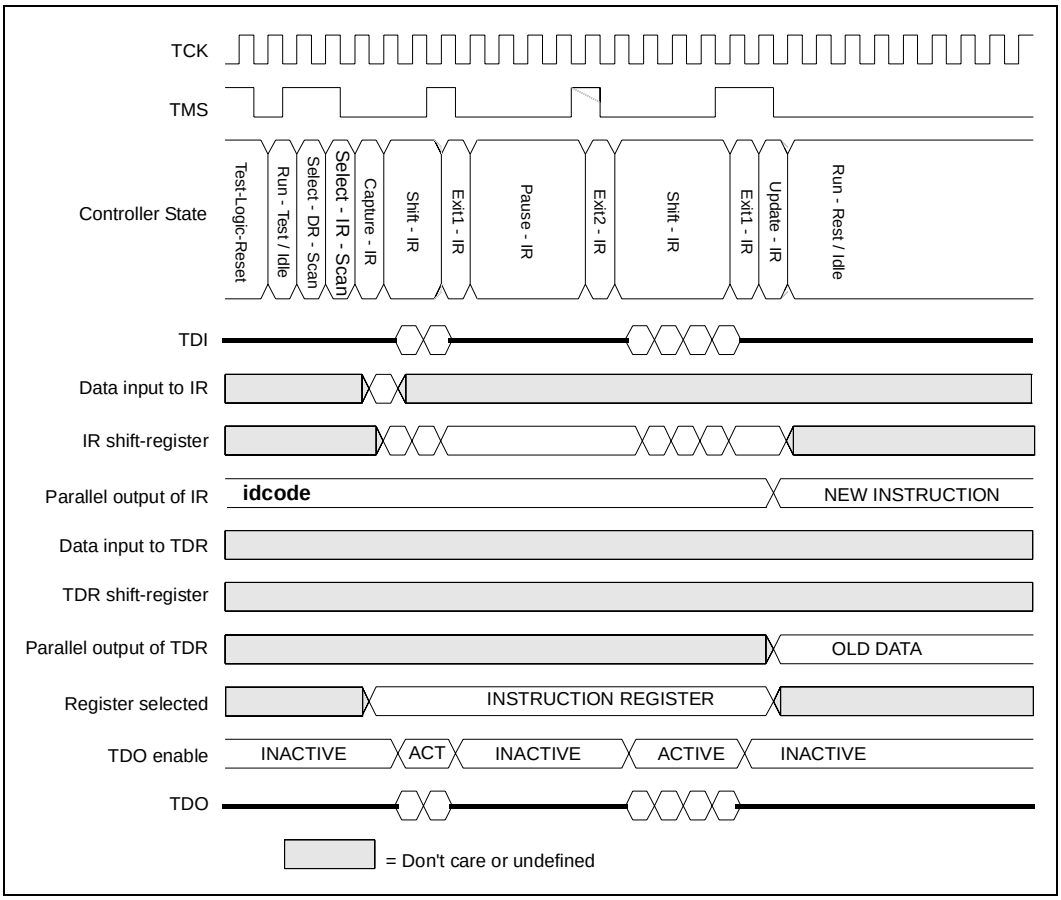
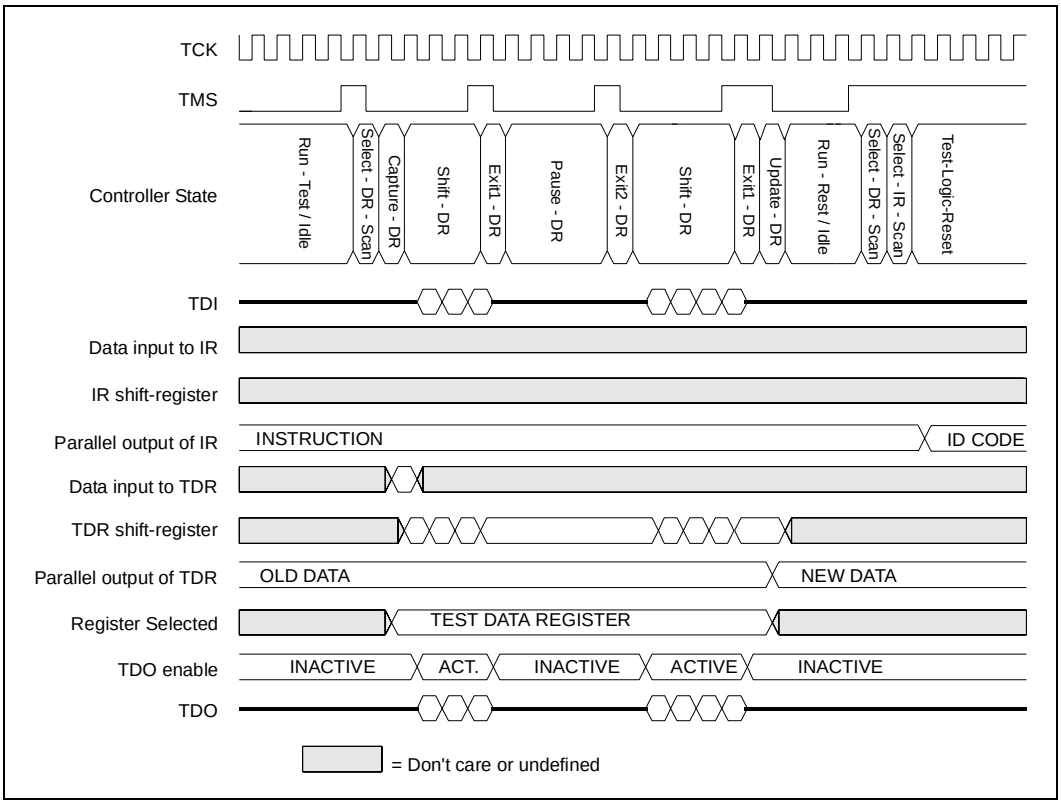
Figure C-5. Timing Diagram Illustrating the Loading of Data Register