

Cortex[™]-R5 and Cortex-R5F

Revision: r1p1

Technical Reference Manual



Cortex-R5 and Cortex-R5F

Technical Reference Manual

Copyright © 2010-2011 ARM. All rights reserved.

Release Information

The following changes have been made to this book.

Change history			
Date	Issue	Confidentiality	Change
03 August 2010	A	Confidential	First release for r0p0
29 October 2010	B	Non-Confidential	First release for r1p0
11 February 2011	C	Non-Confidential	First release for r1p1

Proprietary Notice

Words and logos marked with ® or ™ are registered trademarks or trademarks of ARM® in the EU and other countries, except as otherwise stated below in this proprietary notice. Other brands and names mentioned herein may be the trademarks of their respective owners.

Neither the whole nor any part of the information contained in, or the product described in, this document may be adapted or reproduced in any material form except with the prior written permission of the copyright holder.

The product described in this document is subject to continuous developments and improvements. All particulars of the product and its use contained in this document are given by ARM in good faith. However, all warranties implied or expressed, including but not limited to implied warranties of merchantability, or fitness for purpose, are excluded.

This document is intended only to assist the reader in the use of the product. ARM shall not be liable for any loss or damage arising from the use of any information in this document, or any error or omission in such information, or any incorrect use of the product.

Where the term ARM is used it means “ARM or any of its subsidiaries as appropriate”.

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by ARM and the party that ARM delivered this document to.

Product Status

The information in this document is final, that is for a developed product.

Web Address

<http://www.arm.com>

Contents

Cortex-R5 and Cortex-R5F Technical Reference Manual

	Preface	
	About this book	viii
	Feedback	xii
Chapter 1	Introduction	
	1.1 About the processor	1-2
	1.2 Compliance	1-3
	1.3 Features	1-4
	1.4 Interfaces	1-5
	1.5 Configurable options	1-6
	1.6 Test features	1-12
	1.7 Product documentation, design flow, and architecture	1-13
	1.8 Changes from previous version	1-15
Chapter 2	Functional Description	
	2.1 About the functions	2-2
	2.2 Interfaces	2-10
	2.3 Clocking and resets	2-12
	2.4 Operation	2-18
Chapter 3	Programmers Model	
	3.1 About the programmers model	3-2
	3.2 Modes of operation and execution	3-3
	3.3 Memory model	3-5
	3.4 Coherency	3-6
	3.5 Data structures	3-8

	3.6	Registers	3-9
	3.7	Program status registers	3-12
	3.8	Exceptions	3-17
	3.9	Acceleration of execution environments	3-28
	3.10	Unaligned and mixed-endian data access support	3-29
	3.11	Big-endian instruction support	3-30
Chapter 4	System Control		
	4.1	About system control	4-2
	4.2	Register summary	4-7
	4.3	Register descriptions	4-9
Chapter 5	Prefetch Unit		
	5.1	About the prefetch unit	5-2
	5.2	Branch prediction	5-3
	5.3	Return stack	5-5
	5.4	Controlling instruction prefetch and program flow prediction	5-6
Chapter 6	Events and Performance Monitor		
	6.1	About the events	6-2
	6.2	About the PMU	6-6
	6.3	Performance monitoring registers	6-7
	6.4	Event bus interface	6-20
Chapter 7	Memory Protection Unit		
	7.1	About the MPU	7-2
	7.2	Memory types	7-7
	7.3	Region attributes	7-8
	7.4	MPU interaction with memory system	7-9
	7.5	MPU faults	7-10
	7.6	MPU software-accessible registers	7-11
Chapter 8	Level One Memory System		
	8.1	About the L1 memory system	8-2
	8.2	About the error detection and correction schemes	8-4
	8.3	Fault handling	8-7
	8.4	About the TCMs	8-13
	8.5	About the caches	8-18
	8.6	Internal exclusive monitor	8-34
	8.7	Memory types and L1 memory system behavior	8-35
	8.8	Error detection events	8-36
Chapter 9	Level Two Interface		
	9.1	About the L2 interface	9-2
	9.2	AXI master interface	9-4
	9.3	AXI master interface transfers	9-7
	9.4	AXI slave interface	9-18
	9.5	Enabling or disabling AXI slave accesses	9-21
	9.6	Accessing RAMs using the AXI slave interface	9-21
	9.7	Peripheral interfaces	9-31
	9.8	Accelerator Coherency Port interface	9-48
Chapter 10	Power Control		
	10.1	About power control	10-2
	10.2	Power management	10-3
Chapter 11	FPU Programmers Model		
	11.1	About the FPU programmers model	11-2

	11.2	General-purpose registers	11-4
	11.3	System registers	11-5
	11.4	Modes of operation	11-12
	11.5	Compliance with the IEEE 754 standard	11-13
Chapter 12	Debug		
	12.1	Debug systems	12-2
	12.2	About the debug unit	12-3
	12.3	Debug register interface	12-5
	12.4	Debug register descriptions	12-10
	12.5	Management registers	12-33
	12.6	Debug events	12-40
	12.7	Debug exception	12-42
	12.8	Debug state	12-45
	12.9	Cache debug	12-50
	12.10	External debug interface	12-51
	12.11	Using the debug functionality	12-54
	12.12	Debugging systems with energy management capabilities	12-70
Chapter 13	Integration Test Registers		
	13.1	About Integration Test Registers	13-2
	13.2	Summary of the processor registers used for integration testing	13-3
	13.3	Processor integration testing	13-4
Appendix A	Signal Descriptions		
	A.1	About the processor signal descriptions	A-2
	A.2	Global signals	A-3
	A.3	Configuration signals	A-4
	A.4	Interrupt signals, including VIC interface signals	A-8
	A.5	L2 interface signals	A-9
	A.6	TCM interface signals	A-22
	A.7	Redundant CPU signals	A-25
	A.8	Debug interface signals	A-26
	A.9	ETM interface signals	A-28
	A.10	Test signals	A-29
	A.11	MBIST signals	A-30
	A.12	Validation signals	A-31
	A.13	FPU signals	A-32
	A.14	Split/Lock	A-33
	A.15	Power modes	A-34
Appendix B	Cycle Timings and Interlock Behavior		
	B.1	About cycle timings and interlock behavior	B-3
	B.2	Register interlock examples	B-6
	B.3	Data processing instructions	B-7
	B.4	QADD, QDADD, QSUB, and QDSUB instructions	B-9
	B.5	Media data-processing	B-10
	B.6	Sum of Absolute Differences (SAD)	B-11
	B.7	Multiplies	B-12
	B.8	Divide	B-14
	B.9	Branches	B-15
	B.10	Processor state updating instructions	B-16
	B.11	Single load and store instructions	B-17
	B.12	Load and Store Double instructions	B-19
	B.13	Load and Store Multiple instructions	B-20
	B.14	RFE and SRS instructions	B-23
	B.15	Synchronization instructions	B-24
	B.16	Coprocessor instructions	B-25
	B.17	SVC, BKPT, Undefined, and Prefetch Aborted instructions	B-26

	B.18	Miscellaneous instructions	B-27
	B.19	Floating-point register transfer instructions	B-28
	B.20	Floating-point load/store instructions	B-29
	B.21	Floating-point single-precision data processing instructions	B-31
	B.22	Floating-point double-precision data processing instructions	B-32
	B.23	Dual issue	B-33
Appendix C	ECC Schemes		
	C.1	ECC scheme selection guidelines	C-2
Appendix D	Memory Ordering		
	D.1	Memory ordering	D-2
	D.2	Virtual AXI peripheral interface	D-3
Appendix E	Revisions		
	Glossary		

Preface

This preface introduces the *Cortex-R5 and Cortex-R5F Technical Reference Manual*. It contains the following sections:

- *About this book* on page viii
- *Feedback* on page xii.

About this book

This book is for the Cortex-R5 and Cortex-R5F processors.

Note

- The Cortex-R5F processor is a Cortex-R5 processor that includes the optional *Floating Point Unit* (FPU) extension.
 - In this book, references to the Cortex-R5 processor also apply to the Cortex-R5F processor, unless the context makes it clear that this is not the case.
-

Product revision status

The *mpn* identifier indicates the revision status of the product described in this book, where:

- | | |
|-----------|--|
| rn | Identifies the major revision of the product. |
| pn | Identifies the minor revision or modification status of the product. |

Intended audience

This book is written for system designers, system integrators, and programmers who are designing or programming a *System-on-Chip* (SoC) that uses the Cortex-R5 processor.

Using this book

This book is organized into the following chapters:

Chapter 1 *Introduction*

Read this for an introduction to the processor and descriptions of the major functional blocks.

Chapter 2 *Functional Description*

Read this for a description of the functionality of the product.

Chapter 3 *Programmers Model*

Read this for a description of the processor registers and programming information.

Chapter 4 *System Control*

Read this for a description of the system control coprocessor registers and programming information.

Chapter 5 *Prefetch Unit*

Read this for a description of the functions of the *Prefetch Unit* (PFU), including dynamic branch prediction and the return stack.

Chapter 6 *Events and Performance Monitor*

Read this for a description of the *Performance Monitoring Unit* (PMU) and the event bus.

Chapter 7 *Memory Protection Unit*

Read this for a description of the *Memory Protection Unit* (MPU) and the access permissions process.

Chapter 8 *Level One Memory System*

Read this for a description of the Level One (L1) memory system.

Chapter 9 *Level Two Interface*

Read this for a description of the features of the Level Two (L2) interface not covered in the *AMBA AXI Protocol Specification*.

Chapter 10 *Power Control*

Read this for a description of the power control facilities.

Chapter 11 *FPU Programmers Model*

Read this for a description of the *Floating Point Unit* (FPU) support in the Cortex-R5F processor.

Chapter 12 *Debug*

Read this for a description of the debug support.

Chapter 13 *Integration Test Registers*

Read this for a description of the Integration Test Registers, and of integration testing of the processor with an ETM-R5 trace macrocell.

Appendix A *Signal Descriptions*

Read this for a description of the inputs and outputs of the processor.

Appendix B *Cycle Timings and Interlock Behavior*

Read this for a description of the instruction cycle timing and instruction interlocks.

Appendix C *ECC Schemes*

Read this for a description of how to select the *Error Checking and Correction* (ECC) scheme depending on the *Tightly-Coupled Memory* (TCM) configuration.

Appendix D *Memory Ordering*

Read this for a description of the processor memory ordering and the virtual AXI peripheral interface.

Appendix E *Revisions*

Read this for a description of the technical changes between released issues of this book.

Glossary Read this for definitions of terms used in this book.

Conventions

Conventions that this book can use are described in:

- *Typographical*
- *Timing diagrams* on page x
- *Signals* on page x.

Typographical

The typographical conventions are:


italic Highlights important notes, introduces special terminology, denotes internal cross-references, and citations.

bold	Highlights interface elements, such as menu names. Denotes signal names. Also used for terms in descriptive lists, where appropriate.
<code>monospace</code>	Denotes text that you can enter at the keyboard, such as commands, file and program names, and source code.
<u>monospace</u>	Denotes a permitted abbreviation for a command or option. You can enter the underlined text instead of the full command or option name.
<i>monospace italic</i>	Denotes arguments to monospace text where the argument is to be replaced by a specific value.
monospace bold	Denotes language keywords when used outside example code.
< and >	Enclose replaceable terms for assembler syntax where they appear in code or code fragments. For example: MRC p15, 0 <Rd>, <CRn>, <CRm>, <Opcode_2>

Timing diagrams

The figure named *Key to timing diagram conventions* explains the components used in timing diagrams. Variations, when they occur, have clear labels. You must not assume any timing information that is not explicit in the diagrams.

Shaded bus and signal areas are undefined, so the bus or signal can assume any value within the shaded area at that time. The actual level is unimportant and does not affect normal operation.



Key to timing diagram conventions

Signals

The signal conventions are:

Signal level	The level of an asserted signal depends on whether the signal is active-HIGH or active-LOW. Asserted means: - HIGH for active-HIGH signals - LOW for active-LOW signals.
Lower-case n	At the start or end of a signal name denotes an active-LOW signal.
Lower-case m	At the end of a signal name denotes a value that is 0 or 1, to indicate the CPU to which the signal applies. In a single processor system, m is always 0.

Additional reading

This section lists publications by ARM and by third parties.

See Infocenter, <http://infocenter.arm.com>, for access to ARM documentation.

See on ARM, <http://onarm.com>, for embedded software development resources.

ARM publications

This book contains information that is specific to this product. See the following documents for other relevant information:

- *AMBA® AXI Protocol Specification* (ARM IHI 0022)
- *AMBA 3 APB Protocol Specification* (ARM IHI 0024)
- *AMBA 3 AHB-Lite Protocol Specification* (ARM IHI 0033)
- *ARM Architecture Reference Manual, ARMv7-A and ARMv7-R edition* (ARM DDI 0406)
- *ARM PrimeCell® Vectored Interrupt Controller (PL192) Technical Reference Manual* (ARM DDI 0273)
- *Cortex-R5 and Cortex-R5F Integration Manual* (ARM DIT 0016)
- *Cortex-R5 and Cortex-R5F Configuration and Sign-off Guide* (ARM DII 0255)
- *CoreSight™ DAP-Lite Technical Reference Manual* (ARM DDI 0316)
- *CoreSight ETM-R5 Technical Reference Manual* (ARM DII 0469)
- *RealView™ Compilation Tools Developer Guide* (ARM DUI 0203)
- *Application Note 98, VFP Support Code* (ARM DAI 0098)
- *ARM Synchronization Primitives* (ARM DHT 0008).

Other publications

This section lists relevant documents published by third parties:

- ANSI/IEEE Std 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*
- JEP106M, *Standard Manufacturer's Identification Code*, JEDEC Solid State Technology Association.

Feedback

ARM welcomes feedback on this product and its documentation.

Feedback on this product

If you have any comments or suggestions about this product, contact your supplier and give:

- The product name.
- The product revision or version.
- An explanation with as much information as you can provide. Include symptoms and diagnostic procedures if appropriate.

Feedback on content

If you have comments on content then send an e-mail to errata@arm.com. Give:

- the title
- the number, ARM DDI 0460C
- the page numbers to which your comments apply
- a concise explanation of your comments.

ARM also welcomes general suggestions for additions and improvements.

Chapter 1

Introduction

This chapter introduces the processor and its features. It contains the following sections:

- *About the processor* on page 1-2
- *Compliance* on page 1-3
- *Features* on page 1-4
- *Interfaces* on page 1-5
- *Configurable options* on page 1-6
- *Test features* on page 1-12
- *Product documentation, design flow, and architecture* on page 1-13
- *Changes from previous version* on page 1-15.

1.1 About the processor

The Cortex-R5 processor is a mid-range CPU for use in deeply-embedded, real-time systems. It implements the ARMv7-R architecture, and includes Thumb-2 technology for optimum code density and processing throughput. The pipeline has a single *Arithmetic Logic Unit* (ALU), but implements limited dual-issuing of instructions for efficient utilization of other resources such as the register file. A hardware *Accelerator Coherency Port* (ACP) is provided to reduce the requirement for slow software cache maintenance operations when sharing memory with other masters.

Interrupt latency is kept low by interrupting and restarting load-store multiple instructions, and by use of a dedicated peripheral port that enables low-latency access to an interrupt controller. The processor has *Tightly-Coupled Memory* (TCM) ports for low-latency and deterministic accesses to local RAM, in addition to caches for higher performance to general memory.

Error Checking and Correction (ECC) is used on the Cortex-R5 processor ports and in *Level 1* (L1) memories to provide improved reliability and address safety-critical applications.

Many of the features, including the caches, TCM ports, and ECC are configurable so that a given processor implementation can be tailored to the application for efficient area usage.

Figure 1-1 shows the processor in a typical system.




Figure 1-1 Example Cortex-R5 system

1.2 Compliance

The processor implements the ARMv7-R architecture and ARMv7 debug architecture. In addition, the Cortex-R5F processor implements the VFPv3-D16 architecture. This includes the VFPv3 instruction set.

The Cortex-R5 processor complies with, or implements, the specifications described in:

- *ARM architecture*
- *Trace macrocell*
- *Advanced Microcontroller Bus Architecture.*
- *Debug architecture.*

This TRM complements architecture reference manuals, architecture specifications, protocol specifications, and relevant external standards. It does not duplicate information from these sources.

1.2.1 ARM architecture

The Cortex-R5 processor implements the ARMv7-R architecture profile that includes the following architecture extensions:

- *Advanced Single Instruction Multiple Data (SIMD) architecture extension for integer and floating-point vector operations*
- *Vector Floating-Point version 3 (VFPv3) architecture extension for floating-point computation that is fully compliant with the IEEE 754 standard*
- *Multiprocessing Extensions for multiprocessing functionality.*

See the *ARM Architecture Reference Manual*.

1.2.2 Trace macrocell

The Cortex-R5 processor implements the ETM v3.3 architecture profile. See the *CoreSight ETM-R5 Technical Reference Manual*.

1.2.3 Advanced Microcontroller Bus Architecture

This Cortex-R5 processor complies with the AMBA 3 protocol. See *AMBA AXI Protocol Specification* and *AMBA 3 APB Protocol Specification*.

1.2.4 Debug architecture

The Cortex-A9 processor implements the ARMv7 Debug architecture that includes support for Security Extensions and CoreSight. See the *CoreSight Architecture Specification*.

1.3 Features

The features of the processor include:

- an integer unit implementing the ARMv7-R instruction set
- optional and separately licensable *Floating Point Unit* (FPU) implementing the VFPv3 instruction set, fully or as single-precision only
- dynamic branch prediction with a global history buffer, and a 4-entry return stack
- an L1 memory system with:
 - optional TCM interfaces with optional support for ECC
 - optional Harvard caches with optional support for parity or ECC
 - optional ARMv7-R architecture *Memory Protection Unit* (MPU).
- the ability to implement and use redundant CPU logic for fault detection
- an L2 memory interface:
 - 64-bit master AXI3 interface for accessing memory and shared peripherals
 - optional 64-bit slave AXI3 interface to TCM memories and cache RAM blocks for DMA of instructions or data and online RAM test
 - 32-bit master AXI3 peripheral interface for accessing local peripherals
 - optional 32-bit master AHB peripheral interface for accessing legacy peripherals
 - optional ACP for hardware coherency between peripheral data transfers and data cache.
- a debug interface to a CoreSight *Debug Access Port* (DAP)
- a trace interface to a CoreSight ETM-R5
- a *Performance Monitoring Unit* (PMU)
- low interrupt latency with restartable instructions
- non-maskable interrupt
- a *Vectored Interrupt Controller* (VIC) port
- option to implement two CPUs within a group, sharing one ACP.

1.4 Interfaces

The processor has the following interfaces:

- 64-bit AXI-master interfaces, one per CPU, for instruction fetch and data access
- 32-bit AXI and AHB master interfaces, per CPU, for data accesses, particularly to peripherals
- 64-bit AXI-slave interfaces, one per CPU, for external access to TCMs and cache RAMs
- TCM interfaces, per CPU, for access to local memory containing instructions and data
- ACP pass through interface, comprising AXI master and slave, up to 64 bits wide, providing limited hardware coherency functions
- VIC interfaces, one per CPU, for the connection of a PL192 VIC
- configuration signals for customizing the behavior of the processor, particularly from reset
- interrupt and event outputs providing information about the behavior of the processor to the wider system
- 32-bit APB slave interfaces and various debug handshake signals, one per CPU, for connection to CoreSight components providing debug features
- ETM interfaces, one per CPU, for connection to a CoreSight ETM-R5 providing instruction and data trace
- *Memory Built-In Self Test* (MBIST) interfaces and scan signals, one per CPU, enabling test during manufacture of local RAMs and logic.

All the processor AMBA interfaces conform to one of the following AMBA 3 specifications:

- *AMBA AXI Protocol Specification*
- *AMBA AHB-Lite Protocol Specification*
- *AMBA APB Protocol Specification*.

The debug interfaces are CoreSight compliant, see the *CoreSight Architecture Specification*.

1.5 Configurable options

Table 1-1 shows the features of the processor that can be configured using either build-configuration or pin-configuration. See *Product documentation, design flow, and architecture* on page 1-13 for information about configuration of the processor. Many of these features, if included, can also be enabled and disabled during software configuration. In a twin-CPU configuration, some of the options can be configured separately for each CPU while for other options both CPUs take the same value. Options that permit independent configuration are highlighted with footnotes in Table 1-1 and Table 1-2 on page 1-9.

Table 1-1 Configurable options

Feature	Options	Sub-options	Build-configuration or pin-configuration
Number of CPUs ^a	Single-CPU (no redundancy)	-	Build
	Redundant CPU	-	Build
	Twin-CPU (no redundancy)	-	Build
	Split/Lock	Safety-mode (redundancy) Performance-mode (twin CPU)	Build and pin
Instruction cache	No I-Cache ^b	-	Build
	I-Cache included ^b	No error checking Parity error checking 64-bit ECC error checking	Build
		4KB (4x1KB ways) ^b 8KB (4x2KB ways) ^b 16KB (4x4KB ways) ^b 32KB (4x8KB ways) ^b 64KB (4x16KB ways) ^b	Build
Data cache	No D-Cache ^b	-	Build
	D-Cache included ^b	No error checking ^b Parity error checking 32-bit ECC error checking	Build
		4KB (4x1KB ways) ^b 8KB (4x2KB ways) ^b 16KB (4x4KB ways) ^b 32KB (4x8KB ways) ^b 64KB (4x16KB ways) ^b	Build
ATCM	No ATCM ports	-	Build and pin
	One ATCM port	No error checking 32-bit ECC error checking 64-bit ECC error checking	Build
		4KB, 8KB, 16KB, 32KB, 64KB, 128KB, 256KB, 512KB, 1MB, 2MB, 4MB, or 8MB ^b	Pin

Table 1-1 Configurable options (continued)

Feature	Options	Sub-options	Build-configuration or pin-configuration
BTCM	No BTCM ports	-	Build and pin
	One BTCM port (B0TCM) ^b	No error checking	Build and pin ^c
		32-bit ECC error checking	
		64-bit ECC error checking	
	Two BTCM ports (B0TCM and B1TCM) ^b	4KB, 8KB, 16KB, 32KB, 64KB, 128KB, 256KB, 512KB, 1MB, 2MB, 4MB, or 8MB ^b	Pin
		No error checking	Build and pin ^c
		32-bit ECC error checking	
		64-bit ECC error checking	
Instruction endianness	Little-endian	2x2KB, 2x4KB, 2x8KB, 2x16KB, 2x32KB, 2x64KB, 2x128KB, 2x256KB, 2x512KB, 2x1MB, 2x2MB, or 2x4MB ^b	Pin
		Interleaved on 64-bit granularity in memory ^b	Pin
		Adjacent in memory ^b	
Instruction endianness	Pin-configured	Little-endian Big-endian	Pin
Floating point (VFP)	No FPU ^b	-	Build
	FPU included ^{bd}	Full implementation Single-precision only	
MPU	No MPU ^b	-	Build
	MPU included ^b	12 MPU regions ^b 16 MPU regions ^b	Build
TCM bus parity	No TCM address and control bus parity	-	Build
	TCM address and control bus parity generated	-	
AXI bus ECC/parity on AXI-master, AXI-slave (if included) and ACP (if included)	No AXI bus ECC/parity	-	Build
	AXI bus ECC/parity generated/ checked	-	
Bus ECC/parity on AXI peripheral port and AHB peripheral port (if included)	No peripheral port bus ECC/parity	-	Build
	Peripheral port bus ECC/parity generated/checked	-	
Breakpoints	2-8 breakpoint register pairs	-	Build

Table 1-1 Configurable options (continued)

Feature	Options	Sub-options	Build-configuration or pin-configuration
Watchpoints	1-8 watchpoint registers	-	Build
ATCM at reset	Disabled ^b	-	Pin
	Enabled ^{b e}	Base address 0x0 ^b Base address configured ^b	Build and pin
BTCM at reset	Disabled ^b	-	Pin
	Enabled ^{b e}	Base address configured ^b Base address 0x0 ^{b f}	Build and pin
Peripheral ID RevAnd field	Any 4-bit value	-	Build
AXI slave interface	No AXI-slave ^b	-	Build
	AXI-slave included ^b	-	
TCM Hard Error Cache	No TCM Hard Error Cache	-	Build
	TCM Hard Error Cache included ^g	-	
Non-Maskable FIQ Interrupt	Disabled (FIQ can be masked by software)	-	Pin
	Enabled	-	
Parity type ^h	Odd parity	-	Pin
	Even parity	-	
AXI coherency port (ACP)	No ACP	-	Build
	ACP included	-	
AHB peripheral port	AXI peripheral port only	-	Build
	AXI and AHB peripheral ports	AHB peripheral port region size: 4KB, 8KB, 16KB, 32KB, 64KB, 128KB, 256KB, 512KB, 1MB, 2MB, 4MB, or 8MB. 16MB, 32MB, 64MB, 128MB, 256MB, 512MB, 1GB, 2GB, 4GB ^b	Build and pin
		AHB peripheral port base address: any size-aligned address ^b	
AXI peripheral interface region size	4KB, 8KB, 16KB, 32KB, 64KB, 128KB, 256KB, 512KB, 1MB, 2MB, 4MB, or 8MB. 16MB, 32MB, 64MB, 128MB, 256MB, 512MB, 1GB, 2GB, 4GB ^b	-	Pin
AXI peripheral interface base address	Any size-aligned address ^b	-	Pin

Table 1-1 Configurable options (continued)

Feature	Options	Sub-options	Build-configuration or pin-configuration
Virtual AXI peripheral interface region size	4KB, 8KB, 16KB, 32KB, 64KB, 128KB, 256KB, 512KB, 1MB, 2MB, 4MB, or 8MB. 16MB, 32MB, 64MB, 128MB, 256MB, 512MB, 1GB, 2GB, 4GB ^b	-	Pin
Virtual AXI peripheral interface base address	Any size-aligned address ^b	-	Pin
Cortex-R5 group ID	Any 4-bit value	-	Pin

- a. See *CPU configurations* on page 1-10 for more information.
- b. This option, or some aspects of it, can be configured separately for each CPU on a twin-CPU build.
- c. The error scheme is a build option only. The number of BTCM ports (none, one, two) is set by both build and pin configuration.
- d. Only available with the Cortex-R5F processor.
- e. Only if the relevant TCM port(s) are included.
- f. The BTCM base address must be size aligned, to the total size of B0TCM + B1TCM.
- g. Only if at least one TCM port is included and uses ECC error checking.
- h. Only relevant if one of the caches includes parity checking, or AXI bus ECC or TCM bus parity is included.

Table 1-2 describes the various features that can be pin-configured to be either enabled or disabled at reset. It also shows which CP15 register field provides software configuration of the feature when the processor is out of reset. All of these fields exist in either the SCTLR, or one of the auxiliary control registers.

Table 1-2 Configurable options at reset

Feature	Options	Register field
Exception endianness	Little-endian/big-endian data for exception handling	SCTLR.EE
Exception state	ARM/Thumb state for exception handling	SCTLR.TE
Exception vector table	Base address for exception vectors: 0x00000000/0xFFFF0000 ^a	SCTLR.V
TCM error checking	ATCM ECC check enable ^{ab}	ACTLR.ATCMPCEN
	BTCM ECC check enabled, for B0TCM and B1TCM together ^{ab}	ACTLR.B0TCMPCEN/ ACTLR.B1TCMPCEN
TCM external errors	ATCM external error enable ^a	ACTLR.ATCMECEN
	BTCM external error enable, for B0TCM and B1TCM independently	ACTLR.B0TCMECEN/ ACTLR.B1TCMECEN
TCM load/store-64 (read-modify-write) behavior	ATCM load/store-64 enable ^{ac}	ACTLR2.ATCMRMW
	BTCM load/store-64 enable ^{ac}	ACTLR2.BTCMRMW
AXI peripheral interface	Region enable ^a	PPX.En
AHB peripheral interface ^d	Region enable ^a	PPH.En

- a. This can be configured separately for each CPU on a twin-CPU build.
- b. Can only be enabled if the appropriate TCM is configured with the appropriate error checking scheme, and the appropriate number of ports
- c. Can only be enabled if the appropriate TCM is not configured with 32-bit ECC.
- d. Can only be enabled if the AHB peripheral port is included.

1.5.1 CPU configurations

A Cortex-R5 processor group can consist of either one or two CPUs. The number of CPUs included and the behavior of these CPUs within the group depends on the configuration used. This section describes the CPU arrangements supported and the functionality of each arrangement.

Single CPU

This configuration includes a single CPU.

Twin CPU

This configuration includes two individual and decoupled CPUs, and a single, optional ACP. It offers higher performance than a standard single CPU configuration. Each CPU has its own cache RAMs, debug logic and bus interfaces to the rest of the SoC. There is only one ACP port in the group. Accesses on this port are kept coherent with both CPUs in the group. For more information about ACP coherency, see *Accelerator Coherency Port interface* on page 9-48. The CPUs do not interact within the processor group boundary but might interact elsewhere in the SoC. Contact your system integrator for more information about programming a device that includes a twin-CPU configuration.

You can configure some aspects of the two CPUs separately, for example cache size. See Table 1-1 on page 1-6 for more information about which configuration options can be configured independently.

There is no internal hardware to maintain coherency between the two CPUs in a twin CPU Cortex-R5 group. Loss of coherency occurs if one CPU tries to access dirty data that is in the cache of the other CPU. For example, if CPU0 attempts to transfer a frame of data to CPU1, using a write-back cacheable memory region, then the frame valid bit might miss in the CPU0 cache and be updated in level-2 memory, while some or all of the frame data can hit in the CPU0 cache and not be updated in level-2 memory. This represents a loss of coherency, because CPU1 can detect a valid frame but reads out-of-date frame data. For more information about coherency, see *Coherency* on page 3-6.

Redundant CPU

In this configuration, there is a single functional CPU and an optional ACP. The configuration also includes a second redundant copy of the majority of the CPU logic, and a redundant copy of the ACP logic if an ACP is configured. The redundant logic is driven by the same inputs as the functional logic. In particular, the redundant CPU logic shares the same cache RAMs as the functional CPU. Therefore only one set of cache RAMs is required. The redundant logic operates in lock-step with the CPU, but does not directly affect the processor behavior in any way. The processor outputs to the rest of the system, and the CPU outputs to the cache RAMs, are driven exclusively by the functional CPU.

Comparison logic can be included, during implementation, to compare the outputs of the redundant logic and the functional logic. These comparators can detect a single fault that occurs in either set of logic because of radiation or circuit failure. When used in conjunction with RAM error detection schemes, the system can be protected from faults.

The input signals **DCCMINP[7:0]** and **DCCMINP2[7:0]** and the output signals **DCCMOUT[7:0]** and **DCCMOUT2[7:0]** enable the comparators to communicate with the rest of the SoC.

ARM provides example comparison logic, but you can change this during implementation. If you are implementing a Redundant CPU configuration, contact ARM for more information.

Split/Lock

Two CPUs are included in this configuration. If an ACP is configured, a functional ACP and a redundant copy of the ACP logic is included. The processor group can operate in one of two modes:

- Split mode** Operates as a twin-CPU configuration. Also known as *performance mode*.
- Locked mode** Operates as a redundant CPU configuration. Also known as *safety mode*.

Switching between these modes is only permitted while the processor group is held in power-on reset. The input signals **SLCLAMP** and **SLSPLIT** are provided to enable the system to control the mode of the processor group. For more information about how to effect a change in processor mode, contact your system integrator.

If you are implementing a Split/Lock configuration, contact ARM for more information.

1.6 Test features

The processor is delivered as fully-synthesizable RTL and is a fully-static design. Scan-chains and test wrappers for production test can be inserted into the design by the synthesis tools during implementation. See the relevant reference methodology documentation for more information.

If the AXI-slave interface is included, production test of the processor cache and TCM RAMs can be done through the dedicated, pipelined MBIST interface. This interface shares some of the multiplexing present in the processor design.

In addition, you can use the AXI slave interface to read and write the cache RAMs and TCM. You can use this feature to test the cache RAMs in a running system. This might be required in a safety-critical system. The TCM can be read and written directly by the program running on the processor. You can also use the AXI slave interface for swapping a test program in to the TCMs for the processor to execute. See *Accessing RAMs using the AXI slave interface* on page 9-21 for more information about how to access the RAMs using the AXI slave interface.

1.7 Product documentation, design flow, and architecture

This section describes the Cortex-R5 and Cortex-R5F processor books and how they relate to the design flow in:

- *Documentation*
- *Design flow*.

See *Additional reading* on page xi for more information about the books described in this section. For information about the relevant architectural standards and protocols, see *Compliance* on page 1-3.

1.7.1 Documentation

The Cortex-R5 processor documentation is as follows:

Technical Reference Manual

The *Technical Reference Manual* (TRM) describes the functionality and the effects of functional options on the behavior of the Cortex-R5 processor. It is required at all stages of the design flow. The choices made in the design flow can mean that some behavior described in the TRM is not relevant. If you are programming the Cortex-R5 processor then contact:

- the implementer to determine the build configuration of the implementation
- the integrator to determine the pin configuration of the device that you are using.

Configuration and Sign-off Guide

The *Configuration and Sign-off Guide* (CSG) describes:

- the available build configuration options and related issues in selecting them
- how to configure the *Register Transfer Level* (RTL) source files with the build configuration options
- the processes to sign off the configured design.

The ARM product deliverables include reference scripts and information about using them to implement your design. Reference methodology flows supplied by ARM are example reference implementations. Contact your EDA vendor for EDA tool support.

The CSG is a confidential book that is only available to licensees.

Integration Manual

The *Integration Manual* (IM) describes how to integrate the Cortex-R5 processor into a SoC. It describes the pins that the integrator must tie off to configure the macrocell for the required integration. Some of the integration is affected by the configuration options used when implementing the Cortex-R5 processor.

The IM is a confidential book that is only available to licensees.

1.7.2 Design flow

The Cortex-R5 processor is delivered as synthesizable RTL. Before it can be used in a product, it must go through the following process:

Implementation

The implementer configures and synthesizes the RTL to produce a hard macrocell. This might include integrating RAMs into the design.

Integration The integrator connects the implemented design into a SoC. This includes connecting it to a memory system and peripherals.

Programming

This is the last process. The system programmer develops the software required to configure and initialize the Cortex-R5 processor, and tests the required application software.

Each process can be performed by a different party. Implementation and integration choices affect the behavior and features of the Cortex-R5 processor. The implementer can implement a macrocell that includes some of the SoC components in addition to the Cortex-R5 processor. In this situation, they must perform some of the integration before implementation. The integrator of such a macrocell has fewer integration tasks to perform, and fewer option choices to make.

The operation of the final device depends on:

Build configuration

The implementer chooses the options that affect how the RTL source files are pre-processed. These options usually include or exclude logic that affects one or more of the area, maximum frequency, and features of the resulting macrocell.

For example, the BTCM interface can be configured to have zero, one (B0TCM) or two (B0TCM and B1TCM) ports. If one port is chosen, the logic for the second port is excluded from the macrocell, although the pins remain, and the second port (B1TCM) cannot be used on that macrocell.

Configuration inputs

The integrator configures some features of the Cortex-R5 processor by tying inputs to specific values. These configurations affect the start-up behavior before any software configuration is made. They can also limit the options available to the software.

For example, if the build configuration for the macrocell includes both BTCM ports, the integrator can choose how many ports to actually use, and therefore how many RAMs must be integrated with the macrocell. If the integrator only wishes to use one BTCM port, they can connect RAM to the B0TCM port only, and tie the **ENTCM1IFm** input to zero to indicate that the B1TCM is not available.

Software configuration

The programmer configures the Cortex-R5 processor by programming particular values into registers. This affects the behavior of the Cortex-R5 processor.

For example, the enable bit in the BTCM Region Register controls whether or not memory accesses are performed to the BTCM interface. However, the BTCM cannot, and must not, be enabled if the build configuration does not include any BTCM ports, or if the pin configuration indicates that no RAMs have been integrated onto the BTCM ports.

————— Note —————

This manual refers to *implementation-defined* features that are applicable to build configuration options. Reference to a feature that is *included* means that the appropriate build and pin configuration options are selected. Reference to an *enabled* feature means one that has also been configured by software.

1.8 Changes from previous version

This section describes the differences in functionality between product revisions:

r0p0 First release.

r0p0-r1p0 Functional changes are:

- Adds option for single-precision only floating point support, in addition to existing double-and-single-precision support.
 - Configurable options. See Table 1-1 on page 1-6.
 - FLOAT_PRECISION bit in Build Options 1 register. See *c15, Build Options 1 Register* on page 4-79.
 - VFP instructions undefined in single-precision. See *VFP instructions in a single-precision configuration* on page 11-2.
 - Change to MVFR0 register to indicate double and single-precision support. See Table 11-7 on page 11-10.
- Changes the behavior of the AXI slave port for instruction and data cache accesses. See *Cache RAM access* on page 9-23.
- Adds the V7A&R MP extensions. See:
 - *c0, Multiprocessor Affinity Register* on page 4-18
 - *c0, Cache Level ID Register* on page 4-36
- SCTLr enables SWP and SWPB to be Undefined. See Table 4-24 on page 4-39.
- Adds support for the ARM UDIV and SDIV instructions. See *Instruction Set Attributes Registers* on page 4-27
- Adds ID values for r1p0. See Table 1-3 on page 1-15.

r1p0-r1p1 Functional changes are:

- Adds ID values for r1p1. See Table 1-3 on page 1-15.

Chapter 2

Functional Description

This chapter describes the functionality of the Cortex-R5 processor. It contains the following sections:

- *About the functions* on page 2-2
- *Interfaces* on page 2-10
- *Clocking and resets* on page 2-12
- *Operation* on page 2-18.

2.1 About the functions

Figure 2-1 shows the structure of the processor. Figure 2-2 shows the structure of a CPU within the processor




Figure 2-1 Processor block diagram




Figure 2-2 CPU block diagram

The *PreFetch Unit* (PFU) fetches instructions from the memory system, predicts branches, and passes instructions to the *Data Processing Unit* (DPU). The DPU executes all instructions and uses the *Load/Store Unit* (LSU) for data memory transfers. The PFU and LSU interface to the L1 memory system that contains L1 instruction and data caches and the TCM interfaces. The L1 caches in turn connect to the L2 memory system, and the LSU has a more direct connection to the L2 memory system by means of the peripheral port. The L1 data cache interfaces to the μSCU to perform cache maintenance as required for coherency with ACP transactions.

This section describes the main components of the processor:

- *Data Processing Unit*
- *Load/Store Unit*
- *PreFetch Unit*
- *L1 memory system*
- *L2 AXI interfaces* on page 2-5
- *Dual-redundant core* on page 2-6
- *Split/lock* on page 2-6
- *Hard error features* on page 2-6
- *Debug* on page 2-6
- *System control coprocessor* on page 2-7
- *Interrupt handling* on page 2-7
- *Power management* on page 2-8

2.1.1 Data Processing Unit

The DPU holds most of the program-visible state of the processor, such as general-purpose registers, status registers and control registers. It decodes and executes instructions, operating on data held in the registers in accordance with the ARM architecture. Instructions are fed to the DPU from the PFU through a buffer. The DPU performs instructions that require data to be transferred to or from the memory system by interfacing to the LSU. See Chapter 3 *Programmers Model* for more information.

Floating Point Unit

The *Floating Point Unit* (FPU) is an optional part of the DPU that includes the VFP register file and status registers. It performs floating-point operations on the data held in the VFP register file. See Chapter 11 *FPU Programmers Model* for more information.

2.1.2 Load/Store Unit

The LSU manages all load and store operations, interfacing with the DPU to the TCMs, caches, peripheral ports, and L2 memory interfaces.

2.1.3 PreFetch Unit

The PFU obtains instructions from the instruction cache, the TCMs, or from external memory and predicts the outcome of branches in the instruction stream. See Chapter 5 *Prefetch Unit* for more information.

Branch prediction

The branch predictor is a global type that uses history registers and a 256-entry pattern history table.

Return stack

The PFU includes a 4-entry return stack to accelerate returns from procedure calls.

2.1.4 L1 memory system

The processor L1 memory system includes the following features:

- separate instruction and data caches

- flexible TCM interfaces
- 64-bit datapaths throughout the memory system
- MPU that supports configurable memory region sizes
- export of memory attributes for L2 memory system
- parity or ECC supported on local memories.

For more information of the blocks in the L1 memory system, see:

- *Instruction and data caches*
- *Memory Protection Unit*
- *TCM interfaces*
- *Error correction and detection on page 2-5.*

Instruction and data caches

You can configure the processor to include separate instruction and data caches. The caches have the following features:

- Support for independent configuration of the instruction and data cache sizes between 4KB and 64KB.
- Pseudo-random cache replacement policy.
- 8-word cache line length. Cache lines can be either write-back or write-through, determined by MPU region.
- Ability to disable each cache independently.
- Streaming of sequential data from LDM and LDRD operations, and sequential instruction fetches.
- Critical word first filling of the cache on a cache miss.
- Implementation of all the cache RAM blocks and the associated tag and valid RAM blocks using standard ASIC RAM compilers.

Memory Protection Unit

An optional MPU provides memory attributes for embedded control applications. You can configure the MPU to have eight or twelve regions, each with a minimum resolution of 32 bytes. MPU regions can overlap, and the highest numbered region has the highest priority.

The MPU checks for protection and memory attributes, and some of these can be passed to an external L2 memory system.

For more information, see Chapter 7 *Memory Protection Unit*.

TCM interfaces

Because some applications do not cache well, there are two TCM interfaces that permit connection to configurable memory blocks of *Tightly-Coupled Memory* (ATCM and BTCM). These ensure high-speed access to code or data. As an option, the BTCM can have two memory ports for increased bandwidth.

An ATCM typically holds interrupt or exception code that must be accessed at high speed, without any potential delay resulting from a cache miss.

A BTCM typically holds a block of data for intensive processing, such as audio or video processing.

The TCMs are external to the processor. This provides flexibility in optimizing the TCM subsystem for performance, power, and RAM type. The **INITRAMAm** and **INITRAMBm** pins enable booting from the ATCM or BTCM, respectively. Both the ATCM and BTCM support wait states.

For more information, see Chapter 8 *Level One Memory System*.

Error correction and detection

To increase the tolerance of the system to soft memory faults, you can configure the caches for either:

- parity generation and error correction/detection
- ECC code generation, single-bit error correction, and two-bit error detection.

Similarly, you can configure the TCM interfaces for ECC code generation, single-bit error correction, and two-bit error detection.

For more information, see Chapter 8 *Level One Memory System*.

2.1.5 L2 AXI interfaces

The L2 AXI interfaces enable the L1 memory system to have access to peripherals and to external memory using an AXI master and AXI slave port and the peripheral ports. See Chapter 9 *Level Two Interface* for more information.

AXI master interface

The AXI master interface provides a high bandwidth interface to second level caches, on-chip RAM, peripherals, and interfaces to external memory. It consists of a single AXI port with a 64-bit read channel and a 64-bit write channel for instruction and data fetches.

The AXI master can run at the same frequency as the processor, or at a lower synchronous frequency. If asynchronous clocking is required an external asynchronous AXI slice is required.

AXI slave interface

The AXI slave interface enables AXI masters, including the AXI master port of the processor, to access data and instruction cache RAMs and TCMs on the AXI system bus. You can use this for DMA into and out of the TCM RAMs and for software test of the TCM and cache RAMs.

The slave interface can run at the same frequency as the processor or at a lower, synchronous frequency. If asynchronous clocking is required an external asynchronous AXI slice is required.

Bits in the Auxiliary Control Register and Slave Port Control Register can control access to the AXI slave. Access to the TCM RAMs can be granted to any master, to only privileged masters, or completely disabled. Access to the cache RAMs can be separately controlled in a similar way.

Peripheral interfaces

The peripheral interfaces provide low latency interfaces to on-chip RAM and peripherals for non-cached data accesses only. They consist of a single 32-bit AXI port, accessed through two interfaces, and a single 32-bit AHB port accessed through a single interface.

The peripheral interfaces can run at the same frequency as the processor, or at a lower synchronous frequency. If asynchronous clocking is required, an external asynchronous AXI slice and asynchronous AHB bridge is required.

2.1.6 Dual-redundant core

The processor can be implemented with a second, redundant copy of most of the logic. This second core shares the input pins and the cache RAMs of the master core, so only one set of cache RAMs is required. The master core drives the output pins and the cache RAMs.

Comparison logic can be included during implementation that compares the outputs of the redundant core with those of the master core. If a fault occurs in the logic of either core, because of radiation or circuit failure, this is detected by the comparison logic. Used in conjunction with the RAM error detection schemes, this can help protect the system from faults. The inputs **DCCMINP[7:0]** and **DCCMINP2[7:0]** and the outputs **DCCMOUT[7:0]** and **DCCMOUT2[7:0]** enable the comparison logic inside the processor to communicate with the rest of the system.

ARM provides example comparison logic, but you can change this during implementation. If you are implementing a processor with dual-redundant cores, contact ARM for more information. If you are integrating a Cortex-R5 macrocell with dual-redundant cores, contact the implementer for more information.

2.1.7 Split/lock

The Cortex-R5 processor can be configured so that it can be switched, under reset, between a twin-CPU performance mode and a dual-redundant safety mode. This feature imposes extra constraints on the software usage model. Contact ARM for information on how it can be used.

2.1.8 Hard error features

The error correction features of the processor are targeted at soft errors. The processor contains features that enable it to recover from a limited set of hard errors. Contact ARM for information on hard error effects and these features.

2.1.9 Debug

Each CPU has a CoreSight compliant *Advanced Peripheral Bus version 3* (APBv3) debug interface. This permits system access to debug resources, for example, the setting of watchpoints and breakpoints.

The processor provides extensive support for real-time debug and performance profiling.

The following sections give an overview of debug:

- *System performance monitoring*
- *ETM interface*
- *Real-time debug facilities* on page 2-7.

System performance monitoring

This is a group of counters that you can configure to monitor the operation of the processor and memory system. For more information, see *About the PMU* on page 6-6.

ETM interface

The *Embedded Trace Macrocell* (ETM) interface enables you to connect an external ETM unit to the processor for real-time code tracing of the core in an embedded system.

The ETM interface collects various processor signals and drives these signals from the processor. The interface is unidirectional and runs at the full speed of the processor. The ETM interface connects directly to the external ETM unit without any additional glue logic. You can disable the ETM interface for power saving. For more information, see the *CoreSight ETM-R5 Technical Reference Manual*.

Real-time debug facilities

Each CPU contains an EmbeddedICE logic unit to provide real-time debug facilities. It has:

- up to eight breakpoints
- up to eight watchpoints
- a *Debug Communications Channel* (DCC).

Note

The number of breakpoints and watchpoints is configured during implementation, see *Configurable options* on page 1-6.

The EmbeddedICE logic monitors the internal address and data buses. You access the EmbeddedICE logic through the memory-mapped APB interface.

The processor implements the ARMv7 Debug architecture, including the extensions of the architecture to support CoreSight.

See Chapter 12 *Debug* for more information on debug.

The EmbeddedICE logic supports two modes of debug operation:

Halt mode On a debug event, such as a breakpoint or watchpoint, the debug logic stops the processor and forces it into debug state. This enables you to examine the internal state of the processor, and the external state of the system, independently from other system activity. When the debugging process completes, the processor and system state are restored, and normal program execution resumes.

Monitor debug mode

On a debug event, the processor generates a debug exception instead of entering debug state, as in halt mode. The exception entry enables a debug monitor program to debug the processor while enabling critical interrupt service routines to operate on the processor. The debug monitor program can communicate with the debug host over the DCC or any other communications interface in the system.

2.1.10 System control coprocessor

The system control coprocessor provides configuration and control of the memory system and its associated functionality. Other system-level operations, such as cache maintenance operations, are also managed through the system control coprocessor.

For more information, see *System control and configuration* on page 4-2.

2.1.11 Interrupt handling

Interrupt handling in the processor is compatible with previous ARM architectures, but has several additional features to improve interrupt performance for real-time applications.

VIC port

The core has a dedicated port that enables an external interrupt controller, such as the ARM PrimeCell *Vectored Interrupt Controller* (VIC), to supply a vector address along with an *Interrupt Request* (IRQ) signal. This provides faster interrupt entry, but you can disable it for compatibility with earlier interrupt controllers.

Note

If you do not have a VIC in your design, you must ensure the **nIRQm** and **nFIQm** signals are asserted, held LOW, and remain LOW until the exception handler clears them.

Low interrupt latency

On receipt of an interrupt, the processor abandons any pending restartable memory operations. Restartable memory operations are the multiword transfer instructions LDM, LDRD, STRD, STM, PUSH, and POP that can access Normal memory.

To minimize the interrupt latency, ARM recommends that you do not perform:

- multiple accesses to areas of memory marked as Device or Strongly Ordered
- SWP operations to slow areas of memory.

Exception processing

The ARMv7-R architecture contains exception processing instructions to reduce interrupt handler entry and exit time:

SRS	Save return state to a specified stack frame.
RFE	Return from exception using data from the stack.
CPS	Change processor state, such as interrupt mask setting and clearing, and mode changes.

2.1.12 Power management

The processor includes several microarchitectural features to reduce energy consumption:

- Accurate branch and return prediction, reducing the number of incorrect instruction fetch and decode operations.
- The caches use sequential access information to reduce the number of accesses to the tag RAMs and to unmatched data RAMs.
- Extensive use of gated clocks and gates to disable inputs to unused functional blocks. Because of this, only the logic actively in use to perform a calculation consumes any dynamic power.

Each CPU supports four levels of power management:

Run mode	This mode is the normal mode of operation where all of the functionality of the CPU is available.
Standby mode	This mode disables most of the clocks of the CPU, while keeping the device powered up. This reduces the power drawn to the static leakage current and the minimal clock power overhead required to enable the device to wake up from the Standby mode.

- Dormant mode** The processor can be implemented in such a way as to support Dormant mode. Dormant mode is a power saving mode in which the CPU logic, but not the TCM and cache RAMs, is powered down. The CPU state, apart from the cache and TCM state, is stored to memory before entry into Dormant mode, and restored after exit.
- Contact ARM for more information on preparing the Cortex-R5 processor to support Dormant mode.
- Shutdown mode** This mode has the entire CPU powered down. All state, including cache and TCM state, must be saved externally. After power-up, the assertion of reset returns the CPU to the run state.

For more information on the power management features, see Chapter 10 *Power Control*.

2.2 Interfaces

The processor has the following interfaces for external access:

- *AXI master interface*
- *Peripheral interfaces*
- *AXI slave interface*
- *TCM interfaces*
- *ACP interface*
- *Interrupt and VIC interface*
- *Configuration interface*
- *Interrupt and event outputs* on page 2-11
- *APB Debug interface* on page 2-11
- *ETM interface* on page 2-11
- *Test interface* on page 2-11.

2.2.1 AXI master interface

AXI master interface on page 9-4 describes the AXI master interface. *AXI master port* on page A-9 and *AXI master port error detection signals* on page A-11 describe the associated signals. The *AMBA AXI Protocol Specification* describes the AXI protocol.

2.2.2 Peripheral interfaces

Peripheral interfaces on page 9-31 describes the peripheral interfaces. *AXI peripheral port* on page A-18 to *AHB peripheral port error detection signals* on page A-21 describe the associated signals. The *AMBA AXI Protocol Specification* and the *AMBA 3 AHB-Lite Protocol Specification* describe the AXI and AHB-Lite protocols respectively.

2.2.3 AXI slave interface

AXI slave interface on page 9-18 describes the AXI slave interface. *AXI slave port* on page A-12 and *AXI slave port error detection signals* on page A-14 describe the associated signals. The *AMBA AXI Protocol Specification* describes the AXI protocol.

2.2.4 TCM interfaces

About the TCMs on page 8-13 describes the TCM interfaces. *TCM interface signals* on page A-22 describes the associated signals.

2.2.5 ACP interface

Accelerator Coherency Port interface on page 9-48 describes the ACP interface. *ACP slave port* on page A-15 to *ACP master port error detection signals* on page A-17 describe the associated signals. The *AMBA AXI Protocol Specification* describes the AXI protocol.

2.2.6 Interrupt and VIC interface

Interrupts on page 3-19 describes the interrupts. *Interrupt signals, including VIC interface signals* on page A-8 describes the associated signals.

2.2.7 Configuration interface

Configuration signals on page A-4 describes the configuration signals.

2.2.8 Interrupt and event outputs

Chapter 6 *Events and Performance Monitor* describes events and the interrupts they can generate. *Exceptions* on page 11-14 describes the FPU exception outputs. *Interrupt signals, including VIC interface signals* on page A-8, *ETM interface signals* on page A-28, *Validation signals* on page A-31, and *FPU signals* on page A-32 describe the associated signals.

2.2.9 APB Debug interface

AMBA APBv3 is used for debugging purposes. CoreSight is the ARM architecture for multi-processor trace and debug. CoreSight defines what debug and trace components are required and how they are connected. See the *CoreSight Architecture Specification for more information*. *Debug interface signals* on page A-26 describes the debug APB interface signals.

Note

The APB debug interface can also connect to a DAP-Lite. For more information on the DAP-Lite, see the *CoreSight DAP-Lite Technical Reference Manual*.

2.2.10 ETM interface

You can connect an ETM-R5 to the processor through the ETM interface. The ETM-R5 provides instruction and data trace for the processor. The *CoreSight ETM-R5 Technical Reference Manual* describes how the ETM-R5 connects to the processor.

The ETM interface includes these signals:

- an instruction interface
- a data interface
- an event interface
- other connections to the ETM.

ETM interface signals on page A-28 describes the associated signals. *Event bus interface* on page 6-20 describes the event bus.

2.2.11 Test interface

The test interface provides support for test during manufacture of the processor using *Memory Built-In Self Test* (MBIST). *MBIST signals* on page A-30 describes the test interface signals.

2.3 Clocking and resets

Before you can run application software on the processor, it must be reset and initialized, including loading the appropriate software-configuration. This section describes the signals for clocking and resetting the processor. It contains the following sections:

- *Resets*
- *Reset modes* on page 2-13
- *Clocking* on page 2-16.

See *Initialization* on page 2-18 for information on software initialization.

2.3.1 Resets

Each Cortex-R5 CPU has the following inputs:

nRESETm	Main CPU reset. Resets the non-debug CPU logic.
DBGRESETmn	CPU debug reset. Resets core-domain debug logic. This includes breakpoints, watchpoints and the DCC registers.
PRESETDBGmn	CPU debug reset. Resets debug-domain debug logic and the APB interface of the CPU.

Note

- For more information about the split between core-domain and debug-domain logic, see the *ARM Architecture Reference Manual*.
 - Cortex-R5 implements separate core and debug domains with the minimal architected set of debug domain registers.
-

The Cortex-R5 processor group, containing one or two CPUs, has the following resets:

ACPRESETn	ACP reset. Resets the ACP logic and both the ACP slave and master AXI interfaces.
nSYSPORESET	Power-on reset. Resets the entire processor group including all implemented CPUs, debug logic and ACP. See <i>Effects of resets on debug registers</i> on page 12-8.

The following input is related to the reset functionality:

nCPUHALTm	This signal, when asserted, stops the CPU from fetching instructions out of reset.
------------------	--

All these signals are active-LOW and are suitably synchronized within the processor. You must take care when generating these reset signals, for example, to ensure that they are glitch-free.

2.3.2 Reset modes

The reset signals in the processor enable you to reset different parts of the design independently. Table 2-1 shows the reset signals, and the combinations and possible applications that you can use them in.

Table 2-1 Reset modes

Reset mode	nRESETm	DBG RESETmn	PRESET DBGmn	ACP RESETn	nSYSPORESET	nCPU HALTm	Application
Power-on reset	0/x	x	x	x	0	x	Power-up reset, full-system reset. Hard or cold reset.
CPU reset	0	1	1	x	1	x	Watchdog reset, soft reset or warm reset. Debug logic remains active to permit debugging through reset.
CPU power-up reset	0	0	1	x	1	x	Reset of CPU, on wake-up from dormant or shutdown modes.
Debug reset	x	0	0	x	1	x	Debugger and debug system reset.
ACP reset	x	x	x	0	1	x	Coherent peripheral reset.
Normal	1	1	1	1	1	1	Normal run mode.
Halt	1	1	1	1	1	0	Halt mode with CPU not fetching instructions, provided normal mode has not been entered since last reset

All reset signals are synchronized within the processor. You do not have to synchronize either edge of any of the reset signals. Unless otherwise stated, whenever **nRESETm** is asserted, it must be held asserted for at least four **CLKIN** cycles to ensure correct reset operation.

———— Note ————

If you are implementing either a dual-redundant core or a Split/Lock configuration, contact ARM for additional reset requirements.

This section of the manual describes:

- *Power-on reset*
- *CPU reset* on page 2-14
- *Normal operation* on page 2-15
- *Halt operation* on page 2-15.

Power-on reset

You must apply power-on or cold reset to the processor when power is first applied to the system. A power-on reset must consist of one of the following:

- Assert **nSYSPORESET** and keep it asserted for at least four **CLKIN** cycles. See Figure 2-3 on page 2-14.




Figure 2-3 Power-on Reset

- Assert **nSYSPORESET** and **nRESETm** together, holding **nRESETm** asserted for at least four **CLKIN** cycles. See Figure 2-4.




Figure 2-4 Power-on reset

The processor implements synchronizers for **nSYSPORESET**. You do not have to synchronize either edge of **nSYSPORESET**.

After applying power-up reset to the processor, you must initialize various registers. See *Initialization* on page 2-18 for more information.

CPU reset

A CPU or warm reset initializes the majority of the CPU logic, excluding the ACP and debug logic. Typically, you use CPU reset to reset a system that has been operating for some time, for example when a watchdog timer expires. The processor debug logic remains active, to permit debugging of the reset handling software.

You can safely reset either or both of the CPUs independently of the ACP.

If you are implementing a twin-CPU configuration you must ensure that a given CPU is quiescent before resetting it independently of the other CPU. A CPU is quiescent when all of the following are true:

- either **nWFPIPESTOPPEDm** or **nWFPIPESTOPPEDm** is LOW
- all transactions to the CPU from the system have completed
- the system cannot issue new stimulus to the CPU.

CPU power-up reset

You must apply a CPU power-up reset when the processor wakes up from either dormant or shutdown mode. A CPU power-up reset must consist of the following:

- Assert **nRESETm** and **DBGRESETnm** together and keep them asserted for at least four **CLKIN** cycles on wake-up from dormant or shutdown mode.
- Assert **nRESETm** only and keep it asserted for at least four **CLKIN** cycles on wake-up from emulated dormant or emulated shutdown mode. The processor debug logic is kept active to permit debugging of the wake-up software.

After applying power-up reset to a CPU, you must initialize various registers. See *Initialization* on page 2-18 for more information.

Debug Reset

A debug reset initializes all the debug and non-debug logic of the processor, excluding the ACP logic. This reset causes a debugger to lose connection to the processor, and therefore ARM recommends you should apply it only on request by the debugger, for example on detection of a fatal error condition.

ACP reset

An ACP reset resets the internal ACP logic and the ACP master and slave AXI ports. You can use ACP reset when the peripheral connected to the ACP port is reset. You must not assert ACP reset independently of the CPU resets, unless the ACP is quiescent. The ACP is quiescent when both of the following are true:

- **ACPIDLE** is asserted
- the system cannot issue new transactions to the ACP.

Normal operation

During normal operation, neither processor reset nor power-on reset is asserted. If the EmbeddedICE-RT logic is not used, the value of **PRESETDBGmn** does not matter.

Halt operation

When **nCPUHALTm** is asserted, and **nSYSPORESET** and **nRESETm** are deasserted, the CPU is out of reset, but the PFU is inhibited from fetching instructions. When the CPU is halted in this way, you can, for example, use the AXI slave interface to store instructions in the TCMs using DMA. You can then deassert **nCPUHALTm** and the PFU starts fetching the preloaded instructions from TCMs. When the CPU has started to fetch, **nCPUHALTm** must not be asserted again except when the CPU is reset.

Independent resets

When the Cortex-R5 processor is configured with an ACP, you can reset the CPU or CPUs independently of the ACP. In a twin-CPU configuration it is possible to reset the CPUs independently of each other. Each CPU and the ACP has its own AMBA ports, and in a typical system some or all of these are ultimately connected to the same bus infrastructure. In such a system, to preserve ongoing transactions from other masters, the bus infrastructure is not normally reset when only one of the CPUs or the ACP is reset. To avoid loss of synchronization between bus infrastructure that is not reset and logic that is reset, you must ensure that the logic is quiescent before reset is applied to it. If reset is applied to the bus infrastructure at the same time as the connected logic, the logic does not have to be quiescent.

A CPU is quiescent when:

- **nWFPIPESTOPPEDm** or **nWFPIPESTOPPEDm** is asserted
- all transactions to the CPU from the system have completed
- the system can send no new stimulus to the CPU.

The ACP is quiescent when:

- **ACPIDLE** is asserted
- the system can send no new transactions to the ACP.

2.3.3 Clocking

The processor has a single clock input, **CLKIN**, that is used for the CPU or both CPUs in a twin-CPU configuration. The same clock is used for the ACP ports and logic, and the debug-APB interfaces.

The clock can be stopped indefinitely without loss of state.

The additional clock input, **CLKIN2**, is related to the dual-redundant core functionality, if included. If you are integrating a Cortex-R5 processor with dual-redundant core, contact the implementer of that macrocell for information about how to connect the clock inputs.

This section describes:

- *AMBA interface clocking*
- *Clock gating.*

AMBA interface clocking

The AXI master, AXI slave, ACP, debug-APB, and AXI and AHB peripheral ports must be connected to the AMBA systems that are synchronous to the processor clock, **CLKIN**, even if this might be at a lower frequency. This means that every rising edge on the AMBA system clock must be synchronous to a rising edge on **CLKIN**.

The AXI master interface clock enable signal **ACLKENMm**, the AXI slave interface clock enable signal **ACLKENSs**, ACP clock enable **ACLKENC**, debug-APB block enable **PCLKENDBGm**, and AHB and AXI peripheral port clock enables **ACLKENP** and **HCLKENP** respectively must be asserted on every **CLKIN** rising edge for which there is a simultaneous rising edge on the AXI system clock.

Figure 2-5 shows an example in which the processor is clocked at 400MHz (**CLKIN**), while the AXI system connected to the AXI master interface is clocked at 200MHz (**ACLKM**). The **ACLKENMm** clock indicates the relationship between the two clocks.




Figure 2-5 AXI interface clocking

If the AMBA system connected to an interface is clocked at the same frequency as the processor, then the corresponding clock enable signal must be tied HIGH.

Clock gating

In Standby Mode the CPU can gate its own clock to save power. See Chapter 10 *Power Control* for more information about Standby Mode. You can use the **nCLKSTOPPEDm** output to gate the clock to the TCMs when the CPU is gating its own clock in Standby mode. If you do, you must design the logic so that the TCM clock starts running within three cycles of **nCLKSTOPPEDm** going HIGH.

Figure 2-6 on page 2-17 shows an example of an ATCM access occurring immediately after CPU0 exits Standby Mode. **nCLKSTOPPED0** indicates when the CPU internal clock, shown as **CPU_CLK0**, has been restarted. The clock to the ATCM, shown as **ATCM_CLK0**, has been gated off in Standby Mode and is restarted by the third cycle in order to permit the ATCM to

respond to the access that CPU0 presents by asserting **ATCEN00**. This example shows the worst-case, that is, the earliest TCM access that the CPU can generate after exiting Standby Mode.




Figure 2-6 Standby, wake-up

2.4 Operation

When you power-up the Cortex-R5 processor, you must first reset it, as described in *Clocking and resets* on page 2-12. When it is out of reset, and no longer halted, it starts to fetch and execute instructions from an address and according to the instruction set as described in *Reset* on page 3-19. The processor initially fetches instructions from, and transfers data to and from either the TCM interfaces or the level-2 memory interfaces.

The processor also responds to stimulus received on its interfaces, for example interrupts, or transactions received on the AXI slave interface.

2.4.1 Initialization

When the processor has started executing, but before you can run application software on the processor, it must be initialized, including loading the appropriate software-configuration. This section describes the steps that the software must take to initialize the processor after reset.

Most of the architectural registers in the processor, such as r0-r14, and s0-s31 and d0-d15 when floating-point is included, are not reset. Because of this, you must initialize these for all modes before they are used, using an immediate-MOV instruction, or a PC-relative load instruction. The *Current Program Status Register* (CPSR) is given a known value on reset. This is described in the *ARM Architecture Reference Manual*. The reset values for the CP15 registers are described along with the registers in Chapter 4 *System Control*.

In addition, before you run the application, you might want to:

- program particular values into various registers, for example, stack pointers
- enable various processor features, for example, error correction
- program particular values into memory, for example, the TCMs.

Other initialization requirements are described in:

- *MPU*
- *FPU*
- *Caches* on page 2-19
- *TCM* on page 2-19.

MPU

If the processor has been built with an MPU, before you can use it you must:

- program and enable at least one of the regions
- enable the MPU in the SCTLr.

See *c6, MPU memory region programming registers* on page 4-53. Do not enable the MPU unless at least one MPU region is programmed and active. If the MPU is enabled, before using the TCM interfaces you must program MPU regions to cover the TCM regions to give access permissions to them.

FPU

If the processor has been built with a *Floating Point Unit* (FPU) you must enable it before VFP instructions can be executed:

- enable access to the FPU in the coprocessor access control register, see *c1, Coprocessor Access Control Register* on page 4-47
- enable the FPU by setting the EN-bit in the FPExc register, see *Floating-Point Exception Register, FPExc* on page 11-8.

Note

Floating-point logic is only available with the Cortex-R5F processor.

Caches

If the processor has been built with instruction or data caches, these must be invalidated before they are enabled, otherwise Unpredictable behavior can occur. See *Cache operations* on page 4-59.

If you are using an error checking scheme in the cache, you must enable this by programming the Auxiliary Control Register before invalidating the cache, to ensure that the correct error code or parity bits are calculated when the cache is invalidated. See *c1, Auxiliary Control Register* on page 4-41. An invalidate all operation never reports any ECC or parity errors.

If you are using the ACP, you must perform the data cache invalidation before initiating coherent ACP transactions. Until then, you must not depend on the coherency maintenance information signals.

TCM

The processor does not initialize the TCM RAMs. It is not essential to initialize all the memory attached to the TCM interface but ARM recommends that you do. In addition, the main application might require you to preload instructions or data into the TCM. This section describes various ways that you can perform data preloading. You can also configure the processor to use the TCMs from reset.

Preloading TCMs

You can write data to the TCMs using either store instructions or the AXI slave interface. Depending on the method you choose, you might require:

- particular hardware on the SoC that you are using
- boot code
- a debugger connected to the processor.

Methods to preload TCMs include:

Memory copy with running boot code

The boot code includes a memory copy routine that reads data from a ROM, and writes it into the appropriate TCM. You must enable the TCM to do this, and it might be necessary to give the TCM one base address while the copy is occurring, and a different base address when the application is being run.

Copy data from the debug communications channel

The boot code includes a routine to read data from the *Debug Communications Channel* (DCC) and write it into the TCM. The debug host feeds the data for this operation into the DCC by writing to the appropriate registers on the processor APB debug port.

Execute code in debug halt state

The processor is put into debug halt state by the debug host, that then feeds instructions into the processor through the *Instruction Transfer Register* (DBGITR). The processor executes these instructions, that replace the boot code in either of the previous two methods.

DMA into TCM

The SoC includes a *Direct Memory Access* (DMA) device that reads data from a ROM, and writes it to the TCMs through the AXI slave interface.

Write to TCM directly from debugger

A *Debug Access Port* (DAP) in the system is used to generate AMBA transactions to write data into the TCMs through the AXI slave interface. This DAP is controlled from the debug host through a JTAG chain.

Preloading TCMs with ECC

The error codes in the TCM RAM, if configured with an error scheme, are not initialized by the processor. Before a RAM location is read with ECC checking enabled, the error codes must be initialized. To calculate the error code correctly, the logic must have all the data in the data chunk that those bits protect. Therefore, when the TCM is being initialized, the writes must be of the same width and aligned to the data chunk that the error scheme protects.

You can initialize the TCM RAM with error checking turned on or off, according to the following rules. See *c1, Auxiliary Control Register* on page 4-41. The error code written to the TCM are valid for the data provided, even if the error checking is turned off.

If the slave port is used, write transactions must be used that write to the TCM memory as follows:

- If the error scheme is 32-bit ECC, the write transaction must start at a 32-bit aligned addresses and write a continuous block of memory, containing a multiple of 4 bytes. All bytes in the block must be written, that is, have their byte lane strobe asserted.
- If the error scheme is 64-bit ECC, the write transaction must start at a 64-bit aligned addresses and write a continuous block of memory, containing a multiple of 8 bytes. All bytes in the block must be written, that is, have their byte lane strobe asserted.

If initialization is done by running code on the processor, this is best done by a loop of stores that write to the whole of the TCM memory as follows:

- If the scheme is 32-bit ECC, use *Store Word* (STR), *Store Two Words* (STRD), or *Store Multiple Words* (STM) instructions to 32-bit aligned addresses.
- If the scheme is 64-bit ECC, use STRD or STM that has an even number of registers in the register list, with a 64-bit aligned starting address.

Note

You can use the alignment-checking features of the processor to ensure that memory accesses are 32-bit aligned, but there is no checking for 64-bit alignment. If you are using STRD or STM, an alignment fault is generated if the address is not 32-bit aligned. For the same behavior with STR instructions, enable strict-alignment-checking by setting the A-bit in the SCTL. See *c1, System Control Register* on page 4-38.

If the error scheme is 64-bit ECC, a simpler way to initialize the TCM is:

- Ensure error checking is off.
- Turn on 64-bit store behavior using CP15. See *c15, Secondary Auxiliary Control Register* on page 4-44.
- Write to the TCM using any store instructions, or any AXI write transactions. The processor performs read-modify-write accesses to ensure that all writes are to 64-bit aligned quantities, even though error checking is turned off.

Note

You can enable error checking and 64-bit store behavior on a per-TCM interface basis. References in this section, to these controls relate to whichever TCM is being initialized.

Using TCMs from reset

The processor can be pin-configured to enable the TCM interfaces from reset, and to select the address at which each TCM appears from reset. See *TCM initialization* on page 8-15 for more information. This enables you to configure the processor to boot from TCM but, to do this, the TCM must first be preloaded with the boot code. The **nCPUHALTm** pin can be asserted while the processor is in reset to stop the processor from fetching and executing instructions after coming out of reset. While the processor is halted in this way, the TCMs can be preloaded with the appropriate data. When the **nCPUHALTm** pin is deasserted, the processor starts fetching instructions from the reset vector address in the normal way.

Note

When **nCPUHALTm** has been deasserted to start the processor fetching, **nCPUHALTm** must not be asserted again except when the processor is under processor or power-on reset, that is, **nRESETm** asserted. The processor does not halt if the **nCPUHALTm** pin is asserted while the processor is running.

Peripheral Interfaces

The memory regions used by the peripheral interfaces are fixed during integration. Before you access any peripherals that are in those regions, and attached to the peripheral ports, you must enable the peripheral interfaces. The AXI peripheral interface and the AHB peripheral interface can be enabled from reset by tying **INITPPXm** and **INITPPHm** HIGH respectively. If they are not enabled at reset your software must enable them by writing to the appropriate CP15 region register. See *Peripheral interface region registers* on page 4-84. The virtual AXI peripheral interface can only be enabled by software.

Note

The virtual peripheral interface region is a sub-region of the AXI peripheral interface region. If the AXI peripheral interface is enabled, but the virtual AXI peripheral interface is not, then all accesses to this region of memory use the AXI peripheral port. Enabling the virtual AXI peripheral interface affects only the ordering and ID behavior of the transactions, not the physical port that they use.

Chapter 3

Programmers Model

This chapter describes the processor registers and provides an overview for programming the processor. It contains the following sections:

- *About the programmers model* on page 3-2
- *Modes of operation and execution* on page 3-3
- *Memory model* on page 3-5
- *Coherency* on page 3-6
- *Data structures* on page 3-8
- *Registers* on page 3-9
- *Program status registers* on page 3-12
- *Exceptions* on page 3-17
- *Acceleration of execution environments* on page 3-28
- *Unaligned and mixed-endian data access support* on page 3-29
- *Big-endian instruction support* on page 3-30.

3.1 About the programmers model

The processor implements the ARMv7-R architecture that provides:

- the 32-bit ARM instruction set
- the Thumb-2 technology introduced in ARMv6T2, that extends the Thumb instruction set to a variable-length instruction set, that supports both 16-bit and 32-bit instructions.

For more information on the ARM and Thumb instruction sets, see the *ARM Architecture Reference Manual*. This chapter describes some of the main features of the architecture but, for a complete description, see the *ARM Architecture Reference Manual*.

This chapter also makes reference to older versions of the ARM architecture that the processor does not implement. These references are included to contrast the behavior of the Cortex-R5 processor with other processors you might have used that implement an older version of the architecture.

3.2 Modes of operation and execution

This section describes:

- *Instruction set states*
- *Modes of operation*

3.2.1 Instruction set states

The processor has two instruction set states:

ARM state	The processor executes 32-bit, word-aligned ARM instructions in this state.
Thumb state	The processor executes 32-bit and 16-bit halfword-aligned Thumb instructions in this state.

———— Note ————

Transition between ARM state and Thumb state does not affect the processor mode or the register contents.

Switching state

The instruction set state of the processor can be switched between ARM state and Thumb state:

- Using the BX and BLX instructions, by a load to the PC, or with a data-processing instruction that does not set flags, with the PC as the destination register. Switching state is described in the *ARM Architecture Reference Manual*.

———— Note ————

When the BXJ instruction is used the processor invokes the BX instruction.

- Automatically on an exception. You can write an exception handler routine in ARM or Thumb code. For more information, see *Exceptions* on page 3-17.

Interworking ARM and Thumb state

The processor enables you to mix ARM and Thumb code. For more information about interworking ARM and Thumb, see the *RealView Compilation Tools Developer Guide*.

3.2.2 Modes of operation

In each state there are seven modes of operation:

- *User* (USR) mode is the usual mode for the execution of ARM or Thumb programs. It is used for executing most application programs.
- *Fast interrupt* (FIQ) mode is entered on taking a fast interrupt.
- *Interrupt* (IRQ) mode is entered on taking a normal interrupt.
- *Supervisor* (SVC) mode is a protected mode for the operating system and is entered on taking a *Supervisor Call* (SVC), formerly SWI.
- *Abort* (ABT) mode is entered after a data or instruction abort.
- *System* (SYS) mode is a privileged user mode for the operating system.
- *Undefined* (UND) mode is entered when an Undefined Instruction exception occurs.

Modes other than User mode are collectively known as Privileged modes. Privileged modes are used to service interrupts or exceptions, or access protected resources.

3.3 Memory model

The processor views memory as a linear collection of bytes numbered in ascending order from zero. For example, bytes 0-3 hold the first stored word, and bytes 4-7 hold the second stored word.

The processor can treat words of data in memory as being stored in either:

- *Byte-invariant big-endian format*
- *Little-endian format.*

Additionally, the processor supports mixed-endian and unaligned data accesses. For more information, see the *ARM Architecture Reference Manual*.

3.3.1 Byte-invariant big-endian format

In byte-invariant big-endian (BE-8) format, the processor stores the most significant byte of a word at the lowest-numbered byte, and the least significant byte at the highest-numbered byte. Figure 3-1 shows byte-invariant big-endian (BE-8) format.




Figure 3-1 Byte-invariant big-endian (BE-8) format

3.3.2 Little-endian format

In little-endian format, the lowest-numbered byte in a word is the least significant byte of the word and the highest-numbered byte is the most significant. Figure 3-2 shows little-endian format.




Figure 3-2 Little-endian format

3.4 Coherency

In a system with multiple bus masters, memory coherency problems can occur when the masters access the same memory locations, if one or more of the masters has an associated cache. For example if there are two masters, A and B, each with its own level-one cache, the following problems can occur:

- Master A writes to a location in level-2 memory, but write-back caching is used so that the new value resides in the cache belonging to master A. Subsequent reads of the same address, by master B, access the old value held in the level-2 memory. This might continue indefinitely.
- Master B reads a location in level-2 memory and caches the value read in its own cache. Master A writes to the same location, as above except that the write data propagates to the level-2 memory. If master B reads the same location it gets the old value held in its cache, rather than the new value that master A wrote.

In a twin-CPU configuration of the Cortex-R5 processor, each CPU can have its own level-1 cache. The Cortex-R5 processor might also be integrated into a system with other bus masters. In both cases the coherency problems can occur.

There are a number of solutions to these problems, including:

Data is not shared

If the two masters never access the same data, there can be no coherency issues.

Data that is to be shared between masters is not cached

In the Cortex-R5 processor, data that is in a shared region is never cached in the level-1 caches, even if the region is also cacheable. However, if a Cortex-R5 CPU is connected to a level-2 cache, then data in a shared region might be cached in its level-2 cache, leading to coherency problems, depending on how the level-2 cache is configured. See *Region attributes* on page 7-8 for information about setting memory region attributes.

Data that is to be shared between masters is only cached in coherent caches

If all the bus masters use the same level-2 cache, and do not cache the data in their level-1 cache, then the data stored in the level-2 cache is coherent.

Software coherency

Cache maintenance operations can be used to manipulate the caches so that shared data is visible to other bus masters. In the first example, after master A writes into its cache data that is to be shared by master B, it must also clean the appropriate cache locations to ensure that the level-2 memory has been updated. In the second example, after master A writes data to the level-2 memory, it must cause master B to invalidate the appropriate cache locations in its cache so that master B reads the new value from level-2 memory.

The requirement for cache-clean operations can be avoided by using write-through caching, but invalidate operations are always required. In all cases, barrier operations are required to ensure that the level-2 memory updates have taken place before the cache maintenance operations are performed. Cortex-R5 cache maintenance operations are described in *Cache operations* on page 4-59.

Hardware coherency

Coherency logic, associated with the masters and their caches, performs the appropriate cache manipulation operations to ensure coherency of data that is shared between the masters. ARM multi-processing (MP) technology provides

hardware coherency between multiple CPUs and their associated caches within a cluster, for data that is in a shared memory region. A twin-CPU Cortex-R5 group is not an MP-cluster. No hardware coherency is provided between the two CPUs, see *CPU configurations* on page 1-10 for more information. The Cortex-R5 processor does provide hardware coherency with an external master in limited situations using the ACP. See *Accelerator Coherency Port interface* on page 9-48 for more information.

3.5 Data structures

The processor supports these data types:

- doubleword, 64-bit
- word, 32-bit
- halfword, 16-bit
- byte, 8-bit.

Note

- When any of these types are described as unsigned, the N-bit data value represents a non-negative integer in the range 0 to $+2^N-1$, using normal binary format.
 - When any of these types are described as signed, the N-bit data value represents an integer in the range -2^{N-1} to $+2^{N-1}-1$, using two's complement format.
-

For best performance you must align these data types in memory as follows:

- doubleword quantities aligned to 8-byte boundaries, *doubleword aligned*
- word quantities aligned to 4-byte boundaries, *word aligned*
- halfword quantities aligned to 2-byte boundaries *halfword aligned*
- byte quantities can be placed on any byte boundary.

The processor supports mixed-endian and unaligned access. For more information, see *Unaligned and mixed-endian data access support* on page 3-29.

Note

You cannot use LDRD, LDM, STRD, or STM instructions to access 32-bit quantities if they are not 32-bit aligned.

3.6 Registers

The processor has a total of 37 program registers:

- 31 general-purpose 32-bit registers
- six 32-bit status registers.

These registers are not all accessible at the same time. The processor state and operating mode determine the registers that are available to the programmer.

3.6.1 The register set

In the processor the same register set is used in both the ARM and Thumb states. Sixteen general registers and one or two status registers are accessible at any time. In Privileged modes, alternative mode-specific banked registers become available. Figure 3-3 on page 3-11 shows the registers that are available in each mode.

The register set contains 16 directly-accessible registers, R0-R15. Another register, the *Current Program Status Register* (CPSR), contains condition code flags, status bits, and current mode bits. Registers R0-R12 are general-purpose registers that hold either data or address values. Registers R13, R14, R15, and the CPSR have these special functions:

Stack pointer Software normally uses register R13 as a *Stack Pointer* (SP). The SRS and RFE instructions use Register R13.

Link Register Register R14 is used as the subroutine *Link Register* (LR). Register R14 receives the return address when a *Branch with Link* (BL or BLX) instruction is executed.

You can use R14 as a general-purpose register at all other times. The corresponding banked registers R14_svc, R14_irq, R14_fiq, R14_abt, and R14_und similarly hold the return values when interrupts and exceptions are taken, or when BL or BLX instructions are executed within interrupt or exception routines.

Program Counter Register R15 holds the PC:

- in ARM state this is word-aligned
- in Thumb state this is halfword-aligned.

———— Note ————

There are special cases for reading R15:

- reading the address of the current instruction plus, either:
 - 4 in Thumb state
 - 8 in ARM state.
- reading 0x00000000 (zero).

There are special cases for writing R15:

- causing a branch to the address that was written to R15
- ignoring the value that was written to R15
- writing bits [31:28] of the value that was written to R15 to the condition flags in the CPSR, and ignoring bits [27:0] (used for the MRC instruction only).

You must not assume any of these special cases unless it is explicitly stated in the instruction description. Instead, you must treat instructions with register fields equal to R15 as Unpredictable.

For more information, see the *ARM Architecture Reference Manual*.

In Privileged modes, another register, the *Saved Program Status Register* (SPSR), is accessible. This contains the condition code flags, status bits, and current mode bits saved as a result of the exception that caused entry to the current mode.

Banked registers have a mode identifier that indicates which mode they relate to. Table 3-1 lists these identifiers.

Table 3-1 Register mode identifiers

Mode	Mode identifier
User	usr ^a
Fast interrupt	fiq
Interrupt	irq
Supervisor	svc
Abort	abt
System	usr ^a
Undefined	und

a. The *usr* identifier is usually omitted from register names. It is only used in descriptions where the User or System mode register is specifically accessed from another operating mode.

FIQ mode has seven banked registers mapped to R8–R14 (R8_fiq–R14_fiq). As a result, many FIQ handlers do not have to save any registers.

The Supervisor, Abort, IRQ, and Undefined modes each have alternative mode-specific registers mapped to R13 and R14, permitting a private stack pointer and link register for each mode.

Figure 3-3 on page 3-11 shows the register set, and those registers that are banked.

General registers and program counter

System and User	FIQ	Supervisor	Abort	IRQ	Undefined
R0	R0	R0	R0	R0	R0
R1	R1	R1	R1	R1	R1
R2	R2	R2	R2	R2	R2
R3	R3	R3	R3	R3	R3
R4	R4	R4	R4	R4	R4
R5	R5	R5	R5	R5	R5
R6	R6	R6	R6	R6	R6
R7	R7	R7	R7	R7	R7
R8	 R8_fiq	R8	R8	R8	R8
R9	 R9_fiq	R9	R9	R9	R9
R10	 R10_fiq	R10	R10	R10	R10
R11	 R11_fiq	R11	R11	R11	R11
R12	 R12_fiq	R12	R12	R12	R12
R13	 R13_fiq	 R13_svc	 R13_abt	 R13_irq	 R13_und
R14	 R14_fiq	 R14_svc	 R14_abt	 R14_irq	 R14_und
R15	R15 (PC)	R15 (PC)	R15 (PC)	R15 (PC)	R15 (PC)

Program status registers

CPSR	CPSR	CPSR	CPSR	CPSR	CPSR
	 SPSR_fiq	 SPSR_svc	 SPSR_abt	 SPSR_irq	 SPSR_und

 = banked register

Figure 3-3 Register organization

Note

For 16-bit Thumb instructions, the high registers, R8–R15, are not part of the standard register set. You can use special variants of the MOV instruction to transfer a value from a low register, in the range R0–R7, to a high register, and from a high register to a low register. The CMP instruction enables you to compare high register values with low register values. The ADD instruction enables you to add high register values to low register values. For more information, see the *ARM Architecture Reference Manual*.

3.7 Program status registers

The processor contains one **CPSR** and five **SPSRs** for exception handlers to use. The program status registers:

- hold information about the most recently performed ALU operation
- control the enabling and disabling of interrupts
- set the processor operating mode.

Figure 3-4 shows the bit arrangement in the status registers.




Figure 3-4 Program status register

The following sections explain the meanings of these bits:

- *The N, Z, C, and V bits*
- *The Q bit on page 3-13*
- *The IT bits on page 3-13*
- *The J bit on page 3-14*
- *The DNM bits on page 3-14*
- *The GE bits on page 3-14*
- *The E bit on page 3-15*
- *The A bit on page 3-15*
- *The I and F bits on page 3-15*
- *The T bit on page 3-15*
- *The M bits on page 3-15*
- *Modification of PSR bits by MSR instructions on page 3-16.*

3.7.1 The N, Z, C, and V bits

The N, Z, C, and V bits are the condition code flags. You can optionally set them with arithmetic and logical operations, and also with MSR instructions and MRC instructions to R15. The processor tests these flags in accordance with an instruction's condition code to determine whether to execute that instruction.

In ARM state, most instructions can execute conditionally on the state of the N, Z, C, and V bits. The exceptions are:

- BKPT
- CPS
- LDC2
- MCR2
- MCRR2

- MRC2
- MRRC2
- PLD
- RFE
- SETEND
- SRS
- STC2.

In Thumb state, the processor can only execute the Branch instruction conditionally. Other instructions can be made conditional by placing them in the *If-Then* (IT) block. For more information about conditional execution in Thumb state, see the *ARM Architecture Reference Manual*.

3.7.2 The Q bit

Certain multiply and fractional arithmetic instructions can set the Sticky Overflow, Q, flag:

- QADD
- QDADD
- QSUB
- QDSUB
- SMLAD
- SMLAxy
- SMLAWy
- SMLSD
- SMUAD
- SSAT
- SSAT16
- USAT
- USAT16.

The Q flag is sticky in that, when an instruction sets it, this bit remains set until an MSR instruction writing to the CPSR explicitly clears it. Instructions cannot execute conditionally on the status of the Q flag.

To determine the status of the Q flag you must read the PSR into a register and extract the Q flag from this. For information of how the Q flag is set and cleared, see individual instruction definitions in the *ARM Architecture Reference Manual*.

3.7.3 The IT bits

IT[7:5] encodes the base condition code for the current IT block, if any. It contains b000 when no IT block is active.

IT[4:0] encodes the number of instructions that are to be conditionally executed, and whether the condition for each is the base condition code or the inverse of the base condition code. It contains b00000 when no IT block is active.

When an IT instruction is executed, these bits are set according to the condition in the instruction, and the *Then* and *Else* (T and E) parameters in the instruction. During execution of an IT block, IT[4:0] is shifted to:

- reduce the number of instructions to be conditionally executed by one
- move the next bit into position to form the least significant bit of the condition code.

For more information on the operation of the IT execution state bits, see the *ARM Architecture Reference Manual*.

3.7.4 The J bit

The J bit in the CPSR returns 0 when read.

Note

You cannot use an MSR to change the J bit in the CPSR.

3.7.5 The DNM bits

Software must not modify the *Do Not Modify* (DNM) bits. These bits are:

- Readable, to preserve the state of the processor, for example, during process context switches.
- Writable, to enable the processor to restore its state. To maintain compatibility with future ARM processors, and as good practice, use a read-modify-write strategy when you change the CPSR.

3.7.6 The GE bits

Some of the SIMD instructions set GE[3:0] as greater-than-or-equal bits for individual halfwords or bytes of the result, as Table 3-2 shows.

Table 3-2 GE[3:0] settings

	GE[3]	GE[2]	GE[1]	GE[0]
Instruction	A op B greater than or equal to C	A op B greater than or equal to C	A op B greater than or equal to C	A op B greater than or equal to C
Signed				
SADD16	$[31:16] + [31:16] \geq 0$	$[31:16] + [31:16] \geq 0$	$[15:0] + [15:0] \geq 0$	$[15:0] + [15:0] \geq 0$
SSUB16	$[31:16] - [31:16] \geq 0$	$[31:16] - [31:16] \geq 0$	$[15:0] - [15:0] \geq 0$	$[15:0] - [15:0] \geq 0$
SADDSUBX	$[31:16] + [15:0] \geq 0$	$[31:16] + [15:0] \geq 0$	$[15:0] - [31:16] \geq 0$	$[15:0] - [31:16] \geq 0$
SSUBADDX	$[31:16] - [15:0] \geq 0$	$[31:16] - [15:0] \geq 0$	$[15:0] + [31:16] \geq 0$	$[15:0] + [31:16] \geq 0$
SADD8	$[31:24] + [31:24] \geq 0$	$[23:16] + [23:16] \geq 0$	$[15:8] + [15:8] \geq 0$	$[7:0] + [7:0] \geq 0$
SSUB8	$[31:24] - [31:24] \geq 0$	$[23:16] - [23:16] \geq 0$	$[15:8] - [15:8] \geq 0$	$[7:0] - [7:0] \geq 0$
Unsigned				
UADD16	$[31:16] + [31:16] \geq 2^{16}$	$[31:16] + [31:16] \geq 2^{16}$	$[15:0] + [15:0] \geq 2^{16}$	$[15:0] + [15:0] \geq 2^{16}$
USUB16	$[31:16] - [31:16] \geq 0$	$[31:16] - [31:16] \geq 0$	$[15:0] - [15:0] \geq 0$	$[15:0] - [15:0] \geq 0$
UADDSUBX	$[31:16] + [15:0] \geq 2^{16}$	$[31:16] + [15:0] \geq 2^{16}$	$[15:0] - [31:16] \geq 0$	$[15:0] - [31:16] \geq 0$
USUBADDX	$[31:16] - [15:0] \geq 0$	$[31:16] - [15:0] \geq 0$	$[15:0] + [31:16] \geq 2^{16}$	$[15:0] + [31:16] \geq 2^{16}$
UADD8	$[31:24] + [31:24] \geq 2^8$	$[23:16] + [23:16] \geq 2^8$	$[15:8] + [15:8] \geq 2^8$	$[7:0] + [7:0] \geq 2^8$
USUB8	$[31:24] - [31:24] \geq 0$	$[23:16] - [23:16] \geq 0$	$[15:8] - [15:8] \geq 0$	$[7:0] - [7:0] \geq 0$

Note

GE bit is 1 if $A \text{ op } B \geq C$, otherwise 0.

The SEL instruction uses GE[3:0] to select which source register supplies each byte of its result. See the *ARM Architecture Reference Manual* for more information.

3.7.7 The E bit

ARM and Thumb instructions are provided to set and clear the E bit. The E bit controls load/store endianness. See the *ARM Architecture Reference Manual* for information on where the E bit is used.

3.7.8 The A bit

The A bit is set automatically by certain exceptions and is written by privileged software. It disables asynchronous Data Aborts. For more information on how to use the A bit, see *Asynchronous abort masking* on page 3-24.

3.7.9 The I and F bits

The I and F bits are the interrupt disable bits:

- when the I bit is set, IRQ interrupts are disabled
- when the F bit is set, FIQ interrupts are disabled.

Software can use MSR, CPS, MOVS pc, SUBS pc, LDM ..., {..pc}^, or RFE instructions to change the values of the I and F bits. They are also set automatically by some exceptions.

When NMFI are enabled, updates to the F bit are restricted. For more information see *Non-maskable fast interrupts* on page 3-20.

3.7.10 The T bit

The T bit reflects the instruction set state:

- when the T bit is set, the processor executes in Thumb state
- when the T bit is clear, the processor executes in ARM state.

Note

Never use an MSR instruction to force a change to the state of the T bit in the CPSR. The processor ignores any attempt to modify the T bit using an MSR instruction.

3.7.11 The M bits

M[4:0] are the mode bits. These bits determine the processor operating mode as Table 3-3 shows.

Table 3-3 PSR ode bit values

M[4:0]	Mode
b10000	User
b10001	FIQ
b10010	IRQ

Table 3-3 PSR ode bit values (continued)

M[4:0]	Mode
b10011	Supervisor
b10111	Abort
b11011	Undefined
b11111	System

Note

- In Privileged mode an illegal value programmed into M[4:0] causes the processor to enter System mode.
- In User mode M[4:0] can be read. Writes to M[4:0] are ignored.

3.7.12 Modification of PSR bits by MSR instructions

In the ARMv7-R architecture each CPSR bit falls into one of these categories:

- Bits that are freely modifiable from any mode, either directly by MSR instructions or by other instructions whose side-effects include writing the specific bit or writing the entire CPSR.

Bits in Figure 3-4 on page 3-12 that are in this category are N, Z, C, V, Q, GE[3:0], and E.

- Bits that an MSR instruction must never modify, and so must only be written as a side-effect of another instruction. If an MSR instruction tries to modify these bits, the results are architecturally Unpredictable. In the processor these bits are not affected.

The bits in Figure 3-4 on page 3-12 that are in this category are the execution state bits [26:24], [15:10], and [5].

- Bits that can only be modified from Privileged modes, and that instructions completely protect from modification while the processor is in User mode. Entering a processor exception is the only way to modify these bits while the processor is in User mode, as described in *Exceptions* on page 3-17.

Bits in Figure 3-4 on page 3-12 that are in this category are A, I, F, and M[4:0].

3.8 Exceptions

Exceptions are taken whenever the normal flow of a program must temporarily halt, for example, to service an interrupt from a peripheral. Before attempting to handle an exception, the processor preserves the critical parts of the current processor state so that the original program can resume when the handler routine has finished.

This section provides information of the processor exception handling:

- *Exception entry and exit summary*
- *Reset* on page 3-19
- *Interrupts* on page 3-19
- *Aborts* on page 3-23
- *Supervisor call instruction* on page 3-25
- *Undefined Instruction* on page 3-26
- *Breakpoint instruction* on page 3-26
- *Exception vectors* on page 3-27.

Note

When the processor is in debug halt state, and an exception occurs, it is handled differently to normal. See *Exceptions in debug state* on page 12-48 for more information

3.8.1 Exception entry and exit summary

Table 3-4 summarizes the PC value preserved in the relevant R14 on exception entry, and the recommended instruction for exiting the exception handler.

Table 3-4 Exception entry and exit

Exception or entry	Recommended return instruction	Previous state		Notes
		ARMR14_x	Thumb R14_x	
SVC ^a	MOVS PC, R14_svc	IA + 4	IA + 2	Where the IA is the address of the SVC or Undefined Instruction.
UNDEF	Varies ^b	IA + 4	IA + 2	
PABT	SUBS PC, R14_abt, #4	IA + 4	IA + 4	Where the IA is the address of instruction that had the Prefetch Abort.
FIQ	SUBS PC, R14_fiq, #4	IA + 4	IA + 4	Where the IA is the address of the instruction that was not executed because the FIQ or IRQ took priority.
IRQ	SUBS PC, R14_irq, #4	IA + 4	IA + 4	
DABT	SUBS PC, R14_abt, #8	IA + 8	IA + 8	Where the IA is the address of the Load or Store instruction that generated the Data Abort.
RESET	NA	-	-	The value saved in R14_svc on reset is Unpredictable.
BKPT	SUBS PC, R14_abt, #4	IA + 4	IA + 4	Software breakpoint.

a. Formerly SWI.

b. The return instruction you must use after an Undefined Instruction exception has been handled depends on whether you want to retry the undefined instruction or not and, if not, on the size of the Undefined instruction.

Taking an exception

When taking an exception the processor:

1. Preserves the address of the next instruction in the appropriate R14(LR). When the exception is taken from:
 - ARM state**
The processor writes the address of the instruction into the LR, offset by a value (current IA + 4 or IA + 8 depending on the exception) that causes the program to resume from the correct place on return.
 - Thumb state**
The processor writes the address of the instruction into the LR, offset by a value (current IA + 2, IA + 4 or IA + 8 depending on the exception) that causes the program to resume from the correct place on return.
2. Copies the CPSR into the appropriate SPSR. Depending on the exception type, the processor might modify the IT execution state bits of the CPSR prior to this operation to facilitate a return from the exception.
3. Forces the CPSR mode bits to a value that depends on the exception and clears the IT execution state bits in the CPSR.
4. Sets the E bit based on the state of the EE bit in the SCTLR, see *c1, System Control Register* on page 4-38.
5. The T bit is set based on the state of the TE bit in the SCTLR, see *c1, System Control Register* on page 4-38.
6. Forces the PC to fetch the next instruction from the relevant exception vector.

The processor can also set the interrupt disable flags to prevent otherwise unmanageable nesting of exceptions.

Leaving an exception

When an exception has completed, the exception handler must move the LR, minus an offset, to the PC. The offset varies according to the type of exception, as Table 3-4 on page 3-17 shows.

Typically the return instruction is an arithmetic or logical operation with the S bit set and Rd = R15, so the processor copies the SPSR back to the CPSR. Alternatively, an LDM $\dots, \{..pc\}^A$ or RFE instruction can perform a similar operation if the return state has been pushed onto a stack.

———— Note ————

The action of restoring the CPSR from the SPSR:

- Automatically restores the T, E, A, I, and F bits to the value they held immediately prior to the exception.
- Normally resets the IT execution state bits to the values held immediately prior to the exception. If the exception handler wants to return to the following instruction, these bits might require to be manually advanced to avoid applying the incorrect condition codes to that instruction. For more information about the IT instruction elements and Undefined instructions, and an example of the exception handler code, see the *ARM Architecture Reference Manual*.

Because SVC handlers are always expected to return after the SVC instruction, the IT execution state bits are automatically advanced when an exception is taken prior to copying the CPSR into the SPSR.

3.8.2 Reset

When the **nRESETm** signal is driven LOW a reset occurs, and the processor abandons the executing instruction.

When **nRESETm** and **nCPUHALTm** are driven HIGH again the processor:

1. Forces CPSR M[4:0] to b10011 (Supervisor mode) and sets the A, I, and F bits in the CPSR. The E bit is set based on the state of the **CFGEE** pin. Other bits in the CPSR are indeterminate.
2. Forces the PC to fetch the next instruction from the reset vector address.
3. Reverts to ARM state or Thumb state depending on the state of the **TEINIT** pin, and resumes execution.

After reset, all register values except the PC and CPSR are indeterminate.

See *Resets* on page 2-12 for more information on the reset behavior for the processor.

3.8.3 Interrupts

The processor has two interrupt inputs, for normal interrupts (**nIRQm**) and fast interrupts (**nFIQm**). Each interrupt pin, when asserted and not masked, causes the processor to take the appropriate type of interrupt exception. See *Exceptions* on page 3-17 for more information. The CPSR.F and CPSR.I bits control masking of fast and normal interrupts respectively.

A number of features exist to improve the interrupt latency, that is, the time taken between the assertion of the interrupt input and the execution of the interrupt handler. By default, the processor uses the *Low Interrupt Latency* (LIL) behaviors introduced in version 6 and later of the ARM architecture. The processor also has a port for connection of a *Vectored Interrupt Controller* (VIC), and supports *Non-Maskable Fast Interrupts* (NMFI).

The following subsections describe interrupts:

- *Interrupt request*
- *Fast interrupt request* on page 3-20
- *Non-maskable fast interrupts* on page 3-20
- *Low interrupt latency* on page 3-20
- *Interrupt controller* on page 3-21.

Interrupt request

The IRQ exception is a normal interrupt caused by a LOW level on the **nIRQm** input. An IRQ has a lower priority than an FIQ, and is masked on entry to an FIQ sequence. You must ensure that the **nIRQm** input is held LOW until the processor acknowledges the interrupt request, either from the VIC interface or the software handler.

Irrespective of whether the exception is taken from ARM state or Thumb state, an IRQ handler returns from the interrupt by executing:

```
SUBS PC, R14_irq, #4
```

You can disable IRQ exceptions within a Privileged mode by setting the CPSR.I bit to b1. See *Program status registers* on page 3-12. IRQ interrupts are automatically disabled when an IRQ occurs, by setting the CPSR.I bit. You can use nested interrupts but it is up to you to save any corruptible registers and to re-enable IRQs by clearing the CPSR.I bit.

Fast interrupt request

The *Fast Interrupt Request* (FIQ) reduces the execution time of the exception handler relative to a normal interrupt. FIQ mode has eight private registers to reduce, or even remove the requirement for register saving (minimizing the overhead of context switching).

An FIQ is externally generated by taking the **nFIQm** input signal LOW. You must ensure that the **nFIQm** input is held LOW until the processor acknowledges the interrupt request from the software handler.

Irrespective of whether exception entry is from ARM state or Thumb state, an FIQ handler returns from the interrupt by executing:

```
SUBS PC, R14_fiq, #4
```

If *Non-Maskable Fast Interrupts* (NMFI) are not enabled, you can mask FIQ exceptions by setting the CPSR.F bit to b1. For more information see:

- *Program status registers* on page 3-12
- *Non-maskable fast interrupts*.

FIQ and IRQ interrupts are automatically masked by setting the CPSR.F and CPSR.I bits when an FIQ occurs. You can use nested interrupts but it is up to you to save any corruptible registers and to re-enable interrupts.

Non-maskable fast interrupts

When NMFI behavior is enabled, FIQ interrupts cannot be masked by software. Enabling NMFI behavior ensures that when the FIQ mask, that is, the CPSR.F bit, has been cleared by the reset handler, fast interrupts are always taken as quickly as possible, except during handling of a fast interrupt. This makes the fast interrupt suitable for signaling critical events. NMFI behavior is controlled by a configuration input signal **CFGNMFI_m**, that is asserted HIGH to enable NMFI operation. There is no software control of NMFI.

Software can detect whether NMFI operation is enabled by reading the NMFI bit of the SCTL_R:

NMFI == 0 Software can mask FIQs by setting the CPSR.F bit to b1.

NMFI == 1 Software cannot mask FIQs.

For more information see *c1, System Control Register* on page 4-38.

When the NMFI bit in the SCTL_R is b1:

- an instruction writing b0 to the CPSR.F bit clears it to b0
- an instruction writing b1 to the CPSR.F bit leaves it unchanged
- the CPSR.F bit can be set to b1 only by an FIQ or reset exception entry.

Low interrupt latency

Low Interrupt Latency (LIL) is a set of behaviors that reduce the interrupt latency for the processor, and is enabled by default. That is, the FI bit [21] in the SCTL_R is Read-as-One.

LIL behavior enables accesses to Normal memory, including multiword accesses and external accesses, to be abandoned part-way through execution so that the processor can react to a pending interrupt faster than would otherwise be the case. When an instruction is abandoned in this way, the processor behaves as if the instruction was not executed at all. If, after handling the interrupt, the interrupt handler returns to the program in the normal way using instruction `SUBS pc, r14, #4`, the abandoned instruction is re-executed. This means that some of the memory accesses generated by the instruction are performed twice.

Memory that is marked as Strongly Ordered or Device type is typically sensitive to the number of reads or writes performed. Because of this, instructions that access Strongly Ordered or Device memory are never abandoned when they have started accessing memory. These instructions always complete either all or none of their memory accesses. The same is true of all accesses to the AXI peripheral port, regardless of the memory type. Therefore, to minimize the interrupt latency, you must avoid the use of multiword load/store instructions to memory locations that are marked as Strongly Ordered or Device or are in the AXI or virtual AXI peripheral interface.

Interrupt controller

The processor includes a VIC port for connection of a *Vectored Interrupt Controller* (VIC). An interrupt controller is a peripheral that handles multiple interrupt sources. Features usually found in an interrupt controller are:

- multiple interrupt request inputs, one for each interrupt source, and one or more amalgamated interrupt request outputs to the processor
- the ability to mask out particular interrupt requests
- prioritization of interrupt sources for interrupt nesting.

In a system with an interrupt controller with these features, software is still required to:

- determine from the interrupt controller which interrupt source is requesting service
- determine where the service routine for that interrupt source is loaded
- mask or clear that interrupt source, before re-enabling processor interrupts to permit another interrupt to be taken.

A VIC does all these in hardware to reduce the interrupt latency. It supplies the starting address of the service routine corresponding to the highest priority asserted interrupt source directly to the processor. When the processor has accepted this address, it masks the interrupt so that the processor can re-enable interrupts without clearing the source. The PL192 VIC is an *Advanced Microcontroller Bus Architecture* (AMBA) compliant, *System-on-Chip* (SoC) peripheral that is developed, tested, and licensed by ARM.

You can use the VIC port to connect a PL192 VIC to the processor. See the *ARM PrimeCell Vectored Interrupt Controller (PL192) Technical Reference Manual* for more information about the PL192 VIC. You can enable the VIC port by setting the VE bit in the SCTLR. When the VIC port is enabled and an IRQ occurs, the processor performs an handshake over the VIC interface to obtain the address of the handling routine for the IRQ.

Interrupt entry flowchart

Figure 3-5 on page 3-22 is a flowchart for processor interrupt recognition. It shows all the necessary decisions and actions for a complete interrupt entry.




Figure 3-5 Interrupt entry sequence

For information on the I and F bits that Figure 3-5 shows, see *Program status registers* on page 3-12. For information on the V and VE bits that Figure 3-5 shows, see *c1, System Control Register* on page 4-38.

3.8.4 Aborts

When the processor memory system cannot complete a memory access successfully, an abort is generated. Aborts can occur for a number of reasons, for example:

- a permission fault indicated by the MPU
- an error response to a transaction on the AMBA memory bus
- an error detected in the data by the ECC checking logic.

An error occurring on an instruction fetch generates a *prefetch abort*. Errors occurring on data accesses generate *data aborts*. Aborts are also categorized as being either *synchronous*, previously known as precise, or *asynchronous*, previously known as imprecise.

When a prefetch or data abort occurs, the processor takes the appropriate type of exception. See *Exception entry and exit summary* on page 3-17 for more information. Additional information about the type of abort is stored in registers, and signaled as events. See *Fault handling* on page 8-7 for more information about the types of fault that can cause an abort and the information that the processor provides about these faults.

Prefetch aborts

When a *Prefetch Abort* (PABT) occurs, the processor marks the prefetched instruction as invalid, but does not take the exception until the instruction is to be executed. If the instruction is not executed, for example because a branch occurs while it is in the pipeline, the abort does not take place.

All prefetch aborts are synchronous.

Data aborts

An error occurring on a data memory access can generate a data abort. If the instruction generating the memory access is not executed, for example, because it fails its condition codes, or is interrupted, the data abort does not take place.

A *Data Abort* (DABT) can be either synchronous or asynchronous, depending on the type of fault that caused it.

The Cortex-R5 processor implements the *base restored Data Abort model*, as opposed to a *base updated Data Abort model*.

With the *base restored Data Abort model*, when a Data Abort exception occurs during the execution of a memory access instruction, the processor hardware always restores the base register to the value it contained before the instruction was executed. For more information, see the *ARM Architecture Reference Manual*.

Synchronous aborts

A synchronous abort, also known as a precise abort, is one for which the exception is guaranteed to be taken on the instruction that generated the aborting memory access. The abort handler can use the value in the Link Register (r14_abt) to determine which instruction generated the abort, and the value in the Saved Program Status Register (SPSR_abt) to determine the state of the processor when the abort occurred.

Asynchronous aborts

An asynchronous abort, also known as an imprecise abort, is one for which the exception is taken on a later instruction than the instruction that generated the aborting memory access. The abort handler cannot determine which instruction generated the abort, or the state of the processor when the abort occurred. Therefore, asynchronous aborts are normally fatal.

Asynchronous aborts can be generated by store instructions to Normal-type or Device-type memory. When the store instruction is committed, the data is normally written into a buffer that holds the data until the memory system has sufficient bandwidth to perform the write access. This gives read accesses higher priority. The write data can be held in the buffer for a long period, during which many other instructions can complete. If an error occurs when the write is finally performed, this generates an asynchronous abort.

Asynchronous abort masking

The nature of asynchronous aborts means that they can occur while the processor is handling a different abort. If an asynchronous abort generates a new exception in such a situation, the `r14_abt` and `SPSR_abt` values are overwritten. If this occurs before the data is pushed to the stack in memory, the state information about the first abort is lost. To prevent this from happening, the CPSR contains a mask bit to indicate that an asynchronous abort cannot be accepted, the A-bit. When the A-bit is set, any asynchronous abort that occurs is held pending by the processor until the A-bit is cleared, when the exception is actually taken. The A-bit is automatically set when abort, IRQ or FIQ exceptions are taken, and on reset. You must only clear the A-bit in an abort handler after the state information has either been stacked to memory, or is no longer required.

Only one pending asynchronous abort of each asynchronous abort type is supported. The processor supports the following pending asynchronous aborts:

- AXI-master port external error.
If a subsequent external error is signaled while another one is pending, the later one is ignored and only one abort is taken.
- One TCM write external error for each TCM port.
- Cache write parity or ECC error.
If a subsequent cache parity or ECC error is signaled while another one is pending, the later one is normally ignored and only one abort is taken. However, if the pending error was correctable, and the later one is not correctable, the pending error is ignored, and one abort is taken for the error that cannot be corrected.
- AXI peripheral port external error from either main or virtual interface access.
If a subsequent AXI peripheral port error is signalled while another one is pending, the later one is ignored and only one abort is taken.
- AHB peripheral port external error.

Memory barriers

When a store instruction, or series of instructions has been executed to normal-type or device-type memory, it is sometimes necessary to determine whether any errors occurred because of these instructions. Because most of these errors are reported asynchronously, they might not generate an abort exception until some time after the instructions are executed. To ensure that all possible errors have been reported, you must execute a DSB instruction. Abort exceptions are only taken because of these errors if they are not masked, that is, the CPSR A-bit is clear. If the A-bit is set, the aborts are held pending.

Aborts in Strongly Ordered and Device memory

When a memory access generates an abort, the instruction generating that access is abandoned, even if it has not completed all its memory accesses, and the abort exception is taken. The abort handler can then do one of the following:

- fix the error and return to the instruction that was abandoned, to re-execute it
- perform the appropriate data transfers on behalf of the aborted instruction and return to the instruction after the abandoned instruction
- treat the error as fatal and terminate the process.

If the abort handler returns to the abandoned instruction, some of the memory accesses generated are repeated. The effect is that multiword load/store instructions can access the same memory location twice. The first access occurs before the abort is detected, and the second when the instruction is restarted.

In Strongly Ordered or Device type memory, repeating memory accesses might have unacceptable side-effects. Therefore, if the abort handler can fix the error and re-execute the aborted instruction, you must ensure that for all memory errors on multiword load/store instructions, either:

- all side effects of repeating accesses are inconsequential
- the error must either occur on the first word accessed or not at all.

The instructions that this rule applies to are:

- All forms of ARM instructions LDM, and LDRD, all forms of STM, STRD including VFP variants, and unaligned LDR, STR, LDRH, and STRH
- Thumb instructions LDMIA, LDRD, SDRD, PUSH, POP, and STMIA including VFP variants, and unaligned LDR, STR, LDRH, and STRH.

Abort handler

If you configure the processor with parity or ECC on the caches or the TCMs, and the abort handler is in one of these memories, then it is possible for a parity or ECC error to occur in the abort handler. If the error is not recoverable, then a synchronous abort occurs and the processor loops until the next interrupt. The LR and SPSR values for the original abort are also lost. Therefore, you must construct software that ensures that no synchronous aborts occur when in the abort handler. This means the abort handler must be in external memory and not cached.

3.8.5 Supervisor call instruction

You can use the *SuperVisor Call* (SVC) instruction (formerly SWI) to enter Supervisor mode, usually to request a particular supervisor function. The SVC handler reads the opcode to extract the SVC function number. A SVC handler returns by executing the following instruction, irrespective of the processor operating state:

```
MOVS PC, R14_svc
```

This action restores the PC and CPSR, and returns to the instruction following the SVC.

IRQs are disabled when a software interrupt occurs.

The processor modifies the IT execution state bits on exception entry so that the values that the processor writes into the SPSR are correct for the instruction following the SVC. This means that the SVC handler does not have to perform any special action to accommodate the IT instruction. For more information on the IT instruction, see the *ARM Architecture Reference Manual*.

3.8.6 Undefined Instruction

The processor takes the Undefined Instruction exception when:

- a double-precision VFP operation is attempted when only single-precision support is implemented
- a VFP operation is attempted when the VFP is not enabled

Software can use this mechanism to extend the ARM instruction set by emulating Undefined coprocessor instructions. Undefined Instruction exceptions also occur when a UDIV or SDIV instruction is executed, the value in Rm is zero, and the DZ bit in the SCTLR is set.

If the handler is required to return after the instruction that caused the Undefined Instruction exception, it must:

- Advance the IT execution state bits in the SPSR before restoring SPSR to CPSR. This is so that the correct condition codes are applied to the next instruction on return. The pseudo-code for advancing the IT bits is:

```
Mask = SPSR[11,10,26,25];
if (Mask != 0) {
    Mask = Mask << 1;
    SPSR[12,11,10,26,25] = Mask;
}
if (Mask[3:0] == 0) {
    SPSR[15:12] = 0;
}
```

- Obtain the instruction that caused the Undefined Instruction exception and return correctly after it. Exception handlers must also be aware of the potential for both 16-bit and 32-bit instructions in Thumb state.

After testing the SPSR and determining the instruction was executed in Thumb state, the Undefined handler must use the following pseudo-code or equivalent to obtain this information:

```
addr = R14_undef - 2
instr = Memory[addr,2]
if (instr >> 11) > 28 { /* 32-bit instruction */
    instr = (instr << 16) | Memory[addr+2,2]
    if (emulating, so return after instruction wanted) {
        R14_undef += 2 //
    } //
}
```

After this, instr holds the instruction (in the range 0x0000-0xE7FF for a 16-bit instruction, 0xE8000000-0xFFFFFFFF for a 32-bit instruction), and the exception can be returned from using a MOVs PC, R14 to return after it.

IRQs are disabled when an Undefined Instruction trap occurs. For more information about Undefined instructions, see the *ARM Architecture Reference Manual*.

3.8.7 Breakpoint instruction

A breakpoint (BKPT) instruction operates as though the instruction causes a Prefetch Abort.

A breakpoint instruction does not cause the processor to take the Prefetch Abort exception until the instruction is to be executed. If the instruction is not executed, for example because a branch occurs while it is in the pipeline, the breakpoint does not take place.

After dealing with the breakpoint, the handler executes the following instruction irrespective of the processor operating state:

```
SUBS PC, R14_abt, #4
```

This action restores both the PC and the CPSR, and retries the breakpointed instruction.

Note

If the ETM-R5 is configured into Halt debug-mode, a breakpoint instruction causes the processor to enter debug state. See *Halting debug-mode debugging* on page 12-3.

3.8.8 Exception vectors

You can configure the location of the exception vector addresses by setting the V bit in CP15 c1 System Control Register to enable HIVECS, as Table 3-5 shows.

Table 3-5 Configuration of exception vector address locations

Value of V bit	Exception vector base location
0	0x00000000
1 (HIVECS)	0xFFFF0000

Table 3-6 shows the exception vector addresses and entry conditions for the different exception types.

Table 3-6 Exception vectors

Exception	Offset from vector base	Mode on entry	A bit on entry	F bit on entry	I bit on entry
Reset	0x00	Supervisor	Set	Set	Set
Undefined Instruction	0x04	Undefined	Unchanged	Unchanged	Set
Software interrupt	0x08	Supervisor	Unchanged	Unchanged	Set
Abort (prefetch)	0x0C	Abort	Set	Unchanged	Set
Abort (data)	0x10	Abort	Set	Unchanged	Set
IRQ	0x18	IRQ	Set	Unchanged	Set
FIQ	0x1C	FIQ	Set	Set	Set

3.9 Acceleration of execution environments

Because the ARMv7-R architecture requires Jazelle® software compatibility, three Jazelle registers are implemented in the processor.

Table 3-7 shows the Jazelle register instruction summary and the response to the instructions.

Table 3-7 Jazelle register instruction summary

Register	Instruction	Response
Jazelle ID	MRC p14, 7, <Rd>, c0, c0, 0	Read as zero
	MCR p14, 7, <Rd>, c0, c0, 0	Ignore writes
Jazelle main configuration	MRC p14, 7, <Rd>, c2, c0, 0	Read as zero
	MCR p14, 7, <Rd>, c2, c0, 0	Ignore writes
Jazelle OS control	MRC p14, 7, <Rd>, c1, c0, 0	Read as zero
	MCR p14, 7, <Rd>, c1, c0, 0	Ignore writes

Note

Because no hardware acceleration is present in the processor, when the BXJ instruction is used, the BX instruction is invoked.

3.10 Unaligned and mixed-endian data access support

The processor supports unaligned memory accesses. Unaligned memory accesses was introduced with ARMv6. Bit [22] of c1, Control Register is always 1.

The processor supports byte-invariant big-endianness BE-8 and little-endianness LE. The processor does not support word-invariant big-endianness BE-32. Bit [7] of c1, Control Register is always 0.

For more information on unaligned and mixed-endian data access support, see the *ARM Architecture Reference Manual*.

3.11 Big-endian instruction support

The processor supports little-endian or big-endian instruction format, and is dependent on the setting of the **CFGIE** pin. This is reflected in bit [31] of the SCTLR. For more information, see *c1, System Control Register* on page 4-38.

Note

The facility to use big-endian or little-endian instruction format is an implementation option, and you can therefore remove it in specific implementations. If this facility is not present, the **CFGIE** pin is still reflected in the SCTLR but the instruction format is always little-endian. The Build Options Register indicates whether the processor has been built with instruction endianness control. See *Build Options Registers* on page 4-79.

Chapter 4

System Control

This chapter describes the system control registers, their structure, operation, and how to use them. It contains the following sections:

- *About system control* on page 4-2
- *Register summary* on page 4-7
- *Register descriptions* on page 4-9.

4.1 About system control

This section gives an overview of the system control coprocessor. For more information of the registers in the system control coprocessor, see *Register descriptions* on page 4-9.

The system control coprocessor, CP15, controls and provides status information for the functions implemented in the processor. The main functions of the system control coprocessor are:

- overall system control and configuration
- cache configuration and management
- *Memory Protection Unit* (MPU) configuration and management
- system performance monitoring.

The system control coprocessor does not exist in a distinct physical block of logic.

4.1.1 System control and configuration

The system control and configuration registers provide overall management of:

- memory functionality
- interrupt behavior
- exception handling
- program flow prediction
- coprocessor access rights for CP0-CP13, including the VFP, CP10-11.

The system control and configuration registers also provide the processor ID and information on configured options.

The system control and configuration registers consist of 18 read-only registers and seven read/write registers. Figure 4-1 shows the arrangement of registers in this functional group.




Figure 4-1 System control and configuration registers

Some of the functionality depends on how you set external signals at reset.

4.1.2 MPU control and configuration

The MPU control and configuration registers:

- control program access to memory
- designate areas of memory as either:
 - Normal, Non-cacheable
 - Normal, Cacheable
 - Device
 - Strongly Ordered.
- detect MPU faults and external aborts.

The MPU control and configuration registers consist of one read-only register and 11 read/write registers. Figure 4-2 shows the arrangement of registers in this functional group.




Figure 4-2 MPU control and configuration registers

4.1.3 Cache control and configuration

The cache control and configuration registers:

- provide information on the size and architecture of the instruction and data caches
- control cache maintenance operations that include clean and invalidate caches, drain and flush buffers, and address translation
- override cache behavior during debug or interruptible cache operations.

The cache control and configuration registers consist of three read-only registers, one read/write register, and a number of write-only registers. Figure 4-3 on page 4-4 shows the arrangement of the registers in this functional group.




Figure 4-3 Cache control and configuration registers

4.1.4 Interface control and configuration

The interface control and configuration registers:

- indicate the size, number and status of the TCM regions
- define and enable TCM regions.
- indicate the size and address of the peripheral interface regions
- enable the peripheral interface regions
- control AXI-slave interface permissions

The interface control and configuration registers consist of two read-only registers and six read/write registers. Figure 4-4 shows the arrangement of registers.




Figure 4-4 TCM control and configuration registers

4.1.5 System performance monitor

The performance monitor registers:

- control the monitoring operation
- count events.

The system performance monitor consists of 12 read/write registers. Figure 4-5 on page 4-5 shows the arrangement of registers in this functional group.




Figure 4-5 System performance monitor registers

System performance monitoring counts system events, such as cache misses, pipeline stalls, and other related features to enable system developers to profile the performance of their systems. It can generate interrupts when the number of events reaches a given value.

For more information on the programmers model of the performance counters, see the *ARM Architecture Reference Manual*. See Chapter 6 *Events and Performance Monitor* for more information on the registers.

4.1.6 System validation

The system validation registers extend the use of the system performance monitor registers to provide some functions for validation. You must not use them for other purposes. The system validation registers schedule and clear:

- resets
- interrupts
- fast interrupts
- external debug requests.

The system validation registers consist of nine read/write registers and one write-only register. Figure 4-6 shows the arrangement of registers.




Figure 4-6 System validation registers

You can only change the cache size to a size supported by the cache RAMs implemented in your design.

4.2 Register summary

The system control coprocessor is a set of registers that you can write to and read from. Some of the registers permit more than one type of operation. The functional groups for the registers are:

- *System control and configuration* on page 4-2
- *MPU control and configuration* on page 4-3
- *Cache control and configuration* on page 4-3
- *Interface control and configuration* on page 4-4
- *System performance monitor* on page 4-4
- *System validation* on page 4-5.

Table 4-1 shows the overall functionality for the system control coprocessor, provided through the registers. The registers are listed in their functional groups.

Table 4-2 on page 4-9 lists the registers in the system control processor, in register order, and gives the reset value for each register.

Table 4-1 System control coprocessor register functions

Function	Register/operation	Reference to description
System identification, control and configuration	Control	<i>c1, System Control Register</i> on page 4-38
	Auxiliary control	<i>c1, Auxiliary Control Register</i> on page 4-41
	Coprocessor Access Control	<i>c1, Coprocessor Access Control Register</i> on page 4-47
	Secondary Auxiliary Control Register	<i>c15, Secondary Auxiliary Control Register</i> on page 4-44
	Main ID ^a	<i>c0, Main ID Register</i> on page 4-14
	Auxiliary ID Register	<i>c0, Auxiliary ID Register</i> on page 4-37
	Product Feature IDs	<i>The Processor Feature Registers</i> on page 4-19 <i>c0, Debug Feature Register 0</i> on page 4-21 <i>c0, Auxiliary Feature Register 0</i> on page 4-22 <i>Memory Model Feature Registers</i> on page 4-22 <i>Instruction Set Attributes Registers</i> on page 4-27
	Multiprocessor ID	<i>c0, Multiprocessor Affinity Register</i> on page 4-18
	Context ID	<i>c13, Context ID Register</i> on page 4-66
	FCSE PID	<i>c13, FCSE PID Register</i> on page 4-66
	Pin Options Register	<i>Pin Options Register</i> on page 4-83
	Build Options Registers	<i>c15, Build Options 1 Register</i> on page 4-79
		<i>c15, Build Options 2 Register</i> on page 4-80
Software compatibility	Thread And Process ID	<i>c13, Thread and Process ID Registers</i> on page 4-67

Table 4-1 System control coprocessor register functions (continued)

Function	Register/operation	Reference to description
MPU control and configuration	Data Fault Status	<i>c5, Data Fault Status Register on page 4-49</i>
	Auxiliary Fault Status	<i>c5, Auxiliary Fault Status Registers on page 4-51</i>
	Instruction Fault Status	<i>c5, Instruction Fault Status Register on page 4-50</i>
	Instruction Fault Address	<i>c6, Instruction Fault Address Register on page 4-53</i>
	Data Fault Address	<i>c6, Data Fault Address Register on page 4-53</i>
	MPU Type	<i>c0, MPU Type Register on page 4-17</i>
	Region Base Address	<i>c6, MPU Region Base Address Registers on page 4-54</i>
	Region Size and Enable	<i>c6, MPU Region Size and Enable Registers on page 4-55</i>
	Region Access Control	<i>c6, MPU Region Access Control Registers on page 4-56</i>
	Memory Region Number	<i>c6, MPU Region Number Register on page 4-59</i>
	Correctable Fault Location Register	<i>Correctable Fault Location Register on page 4-77</i>
Cache control and configuration	Cache Type	<i>c0, Cache Type Register on page 4-15</i>
	Current Cache Size Identification	<i>c0, Cache Size ID Register on page 4-34</i>
	Current Cache Level	<i>c0, Cache Level ID Register on page 4-36</i>
	Cache Size Selection	<i>c0, Cache Size Selection Register on page 4-37</i>
	c7, Cache Operations	<i>Cache operations on page 4-59</i>
	c15, Invalidate all data cache	
Interface control and configuration	TCM Status	<i>c0, TCM Type Register on page 4-16</i>
	Region	<i>c9, BTCM Region Register on page 4-63</i> <i>c9, ATCM Region Register on page 4-64</i> <i>c9, TCM Selection Register on page 4-65</i>
	Slave Port Control	<i>c11, Slave Port Control Register on page 4-65</i>
	Peripheral Port Region Registers.	<i>Peripheral interface region registers on page 4-84</i>
System performance monitoring	Performance monitoring	<i>Chapter 6 Events and Performance Monitor</i>
Validation	System validation	<i>Validation Registers on page 4-68</i>

a. Known as the ID Code Register on previous designs. Returns the device ID code.

4.3 Register descriptions

This section describes all of the registers in the system control coprocessor. The section presents a summary of the registers and descriptions in register order of CRn, Opcode_1, CRm, Opcode_2.

For more information on using the system control coprocessor and the general method of how to access CP15 registers, see the *ARM Architecture Reference Manual*.

4.3.1 Register allocation

Table 4-2 shows a summary of address allocation and reset values for the registers in the system control coprocessor where:

- CRn is the register number within CP15
- Op1 is the Opcode_1 value for the register
- CRm is the operational register
- Op2 is the Opcode_2 value for the register.

Table 4-2 Summary of CP15 registers and operations

CRn	Op1	CRm	Op2	Register or operation	Type	Reset value	Page
c0	0	c0	{0, 3, 6-7}	Main ID	Read-only	0x41xFC15x ^a	page 4-14
			1	Cache Type	Read-only	0x8003C003	page 4-15
			2	TCM Type	Read-only	0x00010001	page 4-16
			4	MPU Type	Read-only	_b	page 4-17
			5	Multiprocessor Affinity	Read-only	_d	page 4-18
	1	c1	0	Processor Feature 0	Read-only	0x00000131	page 4-19
			1	Processor Feature 1	Read-only	0x00000001	page 4-20
			2	Debug Feature 0	Read-only	0x00010400	page 4-21
			3	Auxiliary Feature 0	Read-only	0x00000000	page 4-22
			4	Memory Model Feature 0	Read-only	0x00210030	page 4-22
			5	Memory Model Feature 1	Read-only	0x00000000	page 4-23
			6	Memory Model Feature 2	Read-only	0x01200000	page 4-24
			7	Memory Model Feature 3	Read-only	0x00000211	page 4-26
	2	c2	0	Instruction Set Attributes 0	Read-only	0x01101111	page 4-27
			1	Instruction Set Attributes 1	Read-only	0x13112111	page 4-28
			2	Instruction Set Attributes 2	Read-only	0x21232131	page 4-30
			3	Instruction Set Attributes 3	Read-only	0x01112131	page 4-31
			4	Instruction Set Attributes 4	Read-only	0x00010142	page 4-33
			5	Instruction Set Attributes 5	Read-only	0x00000000	page 4-34
			6-7	Reserved, <i>Read As Zero</i> (RAZ)	Read-only	0x00000000	page 4-34
	c3-c7		0-7	Reserved, RAZ	Read-only	0x00000000	-

Table 4-2 Summary of CP15 registers and operations (continued)

CRn	Op1	CRm	Op2	Register or operation	Type	Reset value	Page		
	1	c8-c15	0-7	Undefined	-	-	-		
		c0	0	Current Cache Size ID	Read-only	_cd	page 4-34		
			1	Current Cache Level ID	Read-only	_c	page 4-36		
			2-6	Undefined	-	-	-		
			7	Auxiliary ID	Read-only	0x00000000	page 4-37		
	c1-c15	0-7	Undefined	-	-	-			
	2	c0	0	Cache Size Selection	Read/write	Unpredictable	page 4-37		
	c1	0	c0	0	System Control	Read/write	_d	page 4-38	
				1	Auxiliary Control	Read/write	_d	page 4-41	
				2	Coprocessor Access	Read/write	0x00000000	page 4-47	
3-7				Undefined	-	-	-		
c1-c15		0-7							
c2-c4	0	c0-c15	0-7						
c5	0	c0	0	Data Fault Status	Read/write	Unpredictable	page 4-49		
			1	Instruction Fault Status	Read/write	Unpredictable	page 4-50		
			2-7	Undefined	-	-	-		
		c1	0	Auxiliary Data Fault Status	Read/write	Unpredictable	page 4-51		
		c1	1	Auxiliary Instruction Fault Status	Read/write	Unpredictable	page 4-51		
			2-7	Undefined	-	-	-		
		c2-c15	0-7						
		c6	0	c0	0	Data Fault Address	Read/write	Unpredictable	page 4-53
					1	Undefined	-	-	-
					2	Instruction Fault Address	Read/write	Unpredictable	page 4-53
3-7	Undefined				-	-	-		
c1	0			MPU Region Base Address	Read/write	0x00000000	page 4-54		
	1			Undefined	-	-	-		
	2			MPU Region Size and Enable	Read/write	0x00000000	page 4-55		
	3			Undefined	-	-	-		
	4			MPU Region Access Control	Read/write	0x00000000	page 4-56		
	5-7			Undefined	-	-	-		
	c2			0	MPU Memory Region Number	Read/write	0x00000000	page 4-59	
				1-7	Undefined	-	-	-	

Table 4-2 Summary of CP15 registers and operations (continued)

CRn	Op1	CRm	Op2	Register or operation	Type	Reset value	Page
		c3-c15	1-7				
c7	0	c0	0-3	Undefined	-	-	-
			4	NOP, previously Wait For Interrupt	Write-only	-	page 4-59
			5-7	Undefined	-	-	-
		c1-c4	0-7				
		c5	0	Invalidate entire instruction cache	Write-only	-	page 4-61
		c5	1	Invalidate instruction cache line by address to Point-of-Unification.	Write-only	-	page 4-61
			2-3	Undefined	-	-	-
			4	Instruction Synchronization Barrier	Write-only	-	page 4-61
			5	Undefined	-	-	-
			6	Invalidate entire branch predictor array (NOP)	Write-only	-	page 4-61
			7	Invalidate address from branch predictor array (NOP)	Write-only	-	page 4-61
		c6	0	Undefined	-	-	-
			1	Invalidate data cache line by physical address	Write-only	-	page 4-61
			2	Invalidate data cache line by Set/Way	Write-only	-	page 4-61
			3-7	Undefined	-	-	-
		c7-9	0-7				
		c10	0				
			1	Clean data cache line by physical address	Write-only	-	page 4-61
			2	Clean data cache line by Set/Way	Write-only	-	page 4-61
			3	Undefined	-	-	-
			4	Data Synchronization Barrier	Write-only	-	page 4-62
			5	Data Memory Barrier	Write-only	-	page 4-62
			6-7	Undefined	-	-	-
		c11	0				
		c11	1	Clean data cache line by physical address to Point-of-Unification	Write-only	-	page 4-61
			2-7	Undefined	-	-	-
		c12	0-7				

Table 4-2 Summary of CP15 registers and operations (continued)

CRn	Op1	CRm	Op2	Register or operation	Type	Reset value	Page
		c13	0				
			1	NOP	Write-only	-	-
			2-7	Undefined	-	-	-
		c14	0				
			1	Clean and invalidate data cache line by physical address to Point-of-Unification	Write-only	-	page 4-61
		c14	2	Clean and invalidate data cache line by Set/Way	Write-only	-	page 4-61
			3-7	Undefined	-	-	-
		c15	0-7				
c8	0	c0-c15	0-7	Undefined	-	-	-
c9	0	c0	0-7	Undefined	-	-	-
		c1	0	BTCM Region	Read/write	..d	page 4-63
			1	ATCM Region	Read/write	..d	page 4-63
			2-7	Undefined	-	-	-
		c2	0	TCM selection	Read/write	0x00000000	page 4-65
			1-7	Undefined	-	-	-
		c12	0-7				
			0	Performance Monitor Control	Read/write	0x41151800	page 6-7
			1	Count Enable Set	Read/write	Unpredictable	page 6-8
			2	Count Enable Clear	Read/write	Unpredictable	page 6-9
			3	Overflow Flag Status	Read/write	Unpredictable	page 6-11
			4	Software Increment	Write-only	-	page 6-12
		c12	5	Performance Counter Selection	Read/write	Unpredictable	page 6-12
			6-7	Undefined	-	-	-
		c13	0	Cycle Count	Read/write	0x00000000	page 6-13
			1	Event Select	Read/write	Unpredictable	page 6-14
			2	Performance Monitor Count	Read/write	0x00000000	page 6-16
			3-7	Undefined	-	-	-
		c14	0	User Enable	Read/write	0x00000000	page 6-16
			1	Interrupt Enable Set	Read/write	Unpredictable	page 6-17
		c14	2	Interrupt Enable Clear	Read/write	Unpredictable	page 6-18
			3-7	Undefined	-	-	-

Table 4-2 Summary of CP15 registers and operations (continued)

CRn	Op1	CRm	Op2	Register or operation	Type	Reset value	Page
		c15	0-7				
c10	0	c0-c15	0-7	Undefined	-	-	-
c11	0	c0	0	Slave Port Control	Read/write	0x00000000	page 4-65
		c0	1-7	Undefined	-	-	-
		c1-c15	0-7				
c12	0	c0-c15	0-7				
c13	0	c0	0	FCSE PID	RAZ, ignore writes	0x00000000	page 4-66
			1	Context ID	Read/write	0x00000000	page 4-66
			2	User read/write Thread and Process ID	Read/write	0x00000000	page 4-67
			3	User Read-only Thread and Process ID	Read/write	0x00000000	page 4-67
			4	Privileged Only Thread and Process ID	Read/write	0x00000000	page 4-67
			5-7	Undefined	-	-	-
		c1-c15	0-7	Undefined	-	-	-
c14	0	c0-c15	0-7				
c15	0	c0	0	Secondary Auxiliary Control	Read/write	..d	page 4-44
			1	Normal AXI Peripheral Interface Region	Read/write	..d	page 4-84
			2	Virtual AXI Peripheral Interface Region	Read/write	0	page 4-84
			3	AHB Peripheral Interface Region	Read/write	..d	page 4-84
			4-7	Undefined	-	-	-
		c1	0	nVAL IRQ Enable Set	Read/write	Unpredictable	page 4-68
			1	nVAL FIQ Enable Set	Read/write	Unpredictable	page 4-69
			2	nVAL Reset Enable Set	Read/write	Unpredictable	page 4-70
			3	nVAL Debug Request Enable Set	Read/write	Unpredictable	page 4-71
			4	nVAL IRQ Enable Clear	Read/write	Unpredictable	page 4-72
			5	nVAL FIQ Enable Clear	Read/write	Unpredictable	page 4-73
			6	nVAL Reset Enable Clear	Read/write	Unpredictable	page 4-74
			7	nVAL Debug Request Enable Clear	Read/write	Unpredictable	page 4-75
		c2	0	Build Options 1	Read-only	..d	page 4-79
			1	Build Options 2	Read-only	..d	page 4-80
			2-6	Undefined	-	-	-
			7	Pin Options	Read-only	..d	page 4-83

Table 4-2 Summary of CP15 registers and operations (continued)

CRn	Op1	CRm	Op2	Register or operation	Type	Reset value	Page
		c3	0	Correctable Fault Location	Read/write	Unpredictable	page 4-77
		c3	1-7	Undefined	-	-	-
		c4	0-7				
		c5	0	Invalidate all data cache	Write-only	-	page 4-61
			1-7	Undefined	-	-	-
		c6-c13	0-7				
		c14	0	Cache Size Override	Write-only	-	page 4-76
			1-7	Undefined	-	-	-
		c15	0-7				

- a. The value of bits [23:20,3:0] of the MIDR depend on product revision. See the register description for more information.
b. Reset value depends on number of MPU regions.
c. Reset value depends on which caches are implemented and their sizes.
d. See register description for more information.

4.3.2 c0, Main ID Register

The MIDR characteristics are:

Purpose Returns the device ID code that contains information about the processor

Usage constraints The MIDR is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-3 on page 4-15.

Figure 4-7 shows the shows the MIDR bit assignments.

31	24	23	20	19	16	15					4	3	0
Implementor				Variant		Architecture		Primary part number				Revision	

Figure 4-7 MIDR bit assignments

Table 4-3 shows the MIDR bit assignments.

Table 4-3 MIDR bit assignments

Bits	Name	Function
[31:24]	Implementer	Indicates implementer. 0x41 = ARM Limited.
[23:20]	Variant	Identifies the major revision of the processor. This is the major revision number <i>n</i> in the <i>rn</i> part of the <i>rnpn</i> description of the product revision status.
[19:16]	Architecture	Indicates the architecture version. 0xF = see feature registers.
[15:4]	Primary part number	Indicates processor part number. 0xC15 = Cortex-R5.
[3:0]	Revision	Identifies the minor revision of the processor. This is the minor revision number <i>n</i> in the <i>pn</i> part of the <i>rnpn</i> description of the product revision status.

Note

If an MRC instruction is executed with CRn = c0, Opcode_1 = 0, CRm = c0, and an Opcode_2 value corresponding to an unimplemented or reserved ID register, the system control coprocessor returns the value of the MIDR.

To access the MIDR, read CP15 with:

MRC p15, 0, <Rd>, c0, c0, 0 ; Read MIDR

For more information on the processor features, see *The Processor Feature Registers* on page 4-19.

4.3.3 c0, Cache Type Register

The CTR characteristics are:

Purpose	Determines the instruction and data minimum line length in bytes, to enable a range of addresses to be invalidated.
Usage constraints	The CTR is: - a read-only register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-4 on page 4-16.

Figure 4-8 shows the CTR bit assignments.

31	28	27	24	23	20	19	16	15	14	13							4	3	0
Reserved		CWG		ERG		DMinLine		1	1	Reserved						IMinLine			

Figure 4-8 CTR bit assignments

Table 4-4 shows the CTR bit assignments.

Table 4-4 CTR bit assignments

Bits	Name	Function
[31:28]	-	Always b1000.
[27:24]	CWG	Cache Write-back Granule 0x0 = No information provided. See maximum cache line size in <i>c0</i> , <i>Cache Size ID Register</i> on page 4-34.
[23:20]	ERG	Exclusives Reservation Granule 0x0 = No information provided.
[19:16]	DMinLine	Indicates log2 of the number of words in the smallest cache line of the data and unified caches controlled by the processor: 0x3 = Eight words in an L1 data cache line.
[15:14]	-	Always 0x3.
[13: 4]	-	Always 0x000.
[3: 0]	IMinLine	Indicates log2 of the number of words in the smallest cache line of the instruction caches controlled by the processor: 0x3 - Eight words in an L1 instruction cache line.

To access the CTR, read CP15 with:

MRC p15, 0, <Rd>, c0, c0, 1 ; Read CTR

4.3.4 c0, TCM Type Register

The TCMTR characteristics are:

Purpose Informs the processor of the number of ATCMs and BTCMs in the system

Usage constraints The TCMTR is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-5 on page 4-17.

Figure 4-9 shows the TCMTR bit assignments.

31	30	29	28				19	18	16	15						3	2	0
0	0	0		Reserved				BTCM			Reserved					ATCM		

Figure 4-9 TCMTR bit assignments

Table 4-5 shows the TCMTR bit assignments.

Table 4-5 TCMTR bit assignments

Bits	Name	Function
[31:29]	-	Always 0, indicating v6 format TCMTR.
[28:19]	-	SBZ.
[18:16]	BTCM	Specifies the number of BTCMs implemented. This is always set to b001 because the processor has one BTCM.
[15:3]	-	SBZ.
[2:0]	ATCM	Specifies the number of ATCMs implemented. Always set to b001. The processor has one ATCM.

To access the TCMTR, read CP15 with:

MRC p15, 0, <Rd>, c0, c0, 2 ; Returns TCMTR

Note

- The ATCM and BTCM fields in the TCMTR occupy the same space respectively as the ITCM and DTCM fields as defined by the ARM architecture. These fields, and the corresponding TCM interfaces, can be considered equivalent to those defined in the architecture.
- The ARM architecture requires only the ITCM to be accessible from both instruction and data sides. In the Cortex-R5 processor, both ATCM and BTCM are accessible from both instruction and data sides.

4.3.5 c0, MPU Type Register

The MPUIR characteristics are:

Purpose	Holds the value for the number of instruction and data memory regions implemented in the processor.
Usage constraints	The MPUIR is: • a read-only register • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-6 on page 4-18.

Figure 4-10 shows the MPUIR bit assignments.

31	16	8	7	1	0
Reserved				DRegion	Reserved
					S

Figure 4-10 MPUIR bit assignments

Table 4-6 shows the MPUIR bit assignments.

Table 4-6 MPUIR bit assignments

Bits	Name	Function
[31:16]	-	SBZ.
[15:8]	DRegion	Specifies the number of unified MPU regions. Set to 0, 12, or 16 data MPU regions.
[7:1]	-	SBZ.
[0]	S	Specifies the type of MPU regions, unified or separate, in the processor. Always set to 0, the processor has unified memory regions.

To access the MPUIR, read CP15 with:

MRC p15, 0, <Rd>, c0, c0, 4 ; Read MPUIR

4.3.6 c0, Multiprocessor Affinity Register

The MPIDR characteristics are:

Purpose	Enables CPUs to be recognized and characterized within a twin-CPU system.
Usage constraints	The MPIDR is: • a read-only register • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-7.

Figure 4-11 shows the MPIDR bit assignments.




Figure 4-11 MPIDR bit assignments

Table 4-7 shows the MPIDR bit assignments.

Table 4-7 MPIDR bit assignments

Bits	Name	Description
[31:30]	-	Multiprocessing extensions: 0b00 = no multiprocessing extensions, applies to Cortex-R5, r0p0 0b11 = processor is part of a uniprocessor system, applies to Cortex-R5, from r1p0.
[29:24]	-	SBZ.

Table 4-7 MPIDR bit assignments (continued)

Bits	Name	Description
[23:16]	Aff2	0x00.
[15:8]	Aff1	Processor groups within a system. Read GROUPID input.
[7:0]	Aff0	Processors within a group: 0x0 = CPU0 0x1 = CPU1, if implemented.

To access the MPIDR, read CP15 with:

MRC p15, 0, <Rt>, c0, c0, 5 ; Read MPIDR

4.3.7 The Processor Feature Registers

There are two Processor Feature Registers, PFR0 and PFR1. This section describes:

- *c0, Processor Feature Register 0*
- *c0, Processor Feature Register 1* on page 4-20.

c0, Processor Feature Register 0

The PFR0 characteristics are:

Purpose	Provides information about the execution state support and programmers model for the processor.
Usage constraints	PFR0 is: • a read-only register • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-8.

Figure 4-12 shows the PFR0 bit assignments.

31				16	15	12	11	8	7	4	3	0
Reserved					State3	State2		State1		State0		

Figure 4-12 PFR0 bit assignments

Table 4-8 shows the PFR0 bit assignments.

Table 4-8 PFR0 bit assignments

Bits	Name	Function
[31:16]	-	SBZ.
[15:12]	State3	Indicates support for Thumb Execution Environment (ThumbEE). 0x0 = no support.

Table 4-8 PFR0 bit assignments (continued)

Bits	Name	Function
[11:8]	State2	Indicates support for acceleration of execution environments in hardware or software. 0x1 = the processor supports acceleration of execution environments in software.
[7:4]	State1	Indicates type of Thumb encoding that the processor supports. 0x3 = the processor supports Thumb encoding with all Thumb instructions.
[3:0]	State0	Indicates support for ARM instruction set. 0x1 = the processor supports ARM instructions.

To access PFR0 read CP15 with:

MRC p15, 0, <Rd>, c0, c1, 0 ; Read PFR0

c0, Processor Feature Register 1

The PFR1 characteristics are:

Purpose	Provides information about the execution state support and programmers model for the processor.
Usage constraints	PFR1 is: - a read-only register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-9.

Figure 4-13 shows the PFR1 bit assignments.


**Figure 4-13 PFR1 bit assignments**

Table 4-9 shows the PFR1 bit assignments.

Table 4-9 PFR1 bit assignments

Bits	Name	Function
[31:12]	-	SBZ.

Table 4-9 PFR1 bit assignments (continued)

Bits	Name	Function
[11:8]	Microcontroller programmers model	Indicates support for Microcontroller programmers model: 0x0 = no support.
[7:4]	Security extension	Indicates support for Security Extensions architecture: 0x0 = no support.
[3:0]	ARMv4 programmers model	Indicates support for standard ARMv4 programmers model: 0x1 = the processor supports the ARMv4 model.

To access the PFR1 read CP15 with:

MRC p15, 0, <Rd>, c0, c1, 1 ; Read PFR1

4.3.8 c0, Debug Feature Register 0

The ID_DFR0 characteristics are:

Purpose Provides information about the debug system for the processor.

Usage constraints ID_DFR0 is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-10.

Figure 4-14 shows the ID_DFR0 bit assignments.


**Figure 4-14 ID_DFR0 bit assignments**

Table 4-10 shows the ID_DFR0 bit assignments.

Table 4-10 ID_DFR0 bit assignments

Bits	Name	Function
[31:24]	-	SBZ.
[23:20]	Microcontroller Debug model - memory mapped	Indicates support for the microcontroller debug model - memory mapped: 0x0 = no support.
[19:16]	Trace debug model - memory mapped	Indicates support for the trace debug model - memory mapped: 0x1 = trace supported, memory mapped access.

Table 4-10 ID_DFR0 bit assignments (continued)

Bits	Name	Function
[15:12]	Trace debug model - coprocessor	Indicates support for the trace debug model - coprocessor: 0x0 = no support.
[11:8]	Core debug model - memory mapped	Indicates the type of embedded processor debug model that the processor supports: 0x4 = ARMv7 based model - memory mapped.
[7:4]	Secure debug model	Indicates the type of secure debug model that the processor supports: 0x0 = no support.
[3:0]	Core debug model - coprocessor	Indicates the type of applications processor debug model that the processor supports: 0x0 = no support.

To access the ID_DFR0 read CP15 with:

```
MRC p15, 0, <Rd>, c0, c1, 2 ; Read ID_DFR0
```

4.3.9 c0, Auxiliary Feature Register 0

The ID_AFR0 characteristics are:

Purpose Provides additional information about the features of the processor.

Usage constraints The ID_AFR0 is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes In this processor, the ID_AFR0 reads as 0x00000000.

To access the ID_AFR0 read CP15 with:

```
MRC p15, 0, <Rd>, c0, c1, 3 ; Read ID_AFR0
```

4.3.10 Memory Model Feature Registers

There are four Memory Model Feature Registers, MMFR0 to MMFR3. They are described in the following subsections:

- *c0, Memory Model Feature Register 0*
- *c0, Memory Model Feature Register 1* on page 4-23
- *c0, Memory Model Feature Register 2* on page 4-24
- *c0, Memory Model Feature Register 3, MMFR3* on page 4-26.

c0, Memory Model Feature Register 0

The ID_MMFR0 characteristics are:

Purpose The ID_MMFR0 provides information about the memory model, memory management, and cache support operations of the processor.

Usage constraints The ID_MMFR0 is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-11.

Figure 4-15 shows the ID_MMFR0 bit assignments.

31	28	27	24	23	20	19	16	15	12	11	8	7	4	3	0
Innermost shareability			FCSE		Auxiliary Registers		TCM support		Shareability levels		Outermost shareability		PMSA		VMSA

Figure 4-15 ID_MMFR0 bit assignments

Table 4-11 shows the ID_MMFR0 bit assignments.

Table 4-11 ID_MMFR0 bit assignments

Bits	Name	Function
[31:28]	Innermost shareability	Indicates the innermost shareability domain implemented. RAZ/Unknown because only one shareability domain is implemented, see [15:12].
[27:24]	FCSE	Indicates support for <i>Fast Context Switch Extension</i> (FCSE). 0x0 = no support.
[23:20]	Auxiliary Registers	Indicates support for the auxiliary registers. 0x2 = the processor supports the Auxiliary Instruction and Data Fault Status Registers (AIFSR and ADFSR) and the Auxiliary Control Register.
[19:16]	TCM support	Indicates support for TCM and associated DMA. 0x1 = implementation defined.
[15:12]	Shareability levels	Indicates the number of shareability levels implemented. 0x0 = one level of shareability implemented
[11:8]	Outermost shareability	Indicates the outermost shareability domain implemented. 0x0 = implemented as non-cacheable
[7:4]	PMSA	Indicates support for <i>Physical Memory System Architecture</i> (PMSA). 0x3 = the processor supports PMSAv7 (subsection support).
[3:0]	VMSA	Indicates support for <i>Virtual Memory System Architecture</i> (VMSA). 0x0 = no support.

To access the ID_MMFR0 read CP15 with:

```
MRC p15, 0, <Rd>, c0, c1, 4 ; Read ID_MMFR0
```

c0, Memory Model Feature Register 1

The ID_MMFR1 Register characteristics are:

Purpose	Provides information about the memory model, memory management, and cache support of the processor.
Usage constraints	The ID_MMFR1 is: - a read-only register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-12 on page 4-24.

Figure 4-16 shows the ID_MMFR1 bit assignments.




Figure 4-16 ID_MMFR1 bit assignments

Table 4-12 shows the ID_MMFR1 bit assignments.

Table 4-12 ID_MMFR1 bit assignments

Bits	Name	Function
[31:28]	Branch predictor	Indicates Branch Predictor management requirements. 0x0 = no MMU present.
[27:24]	L1 test clean operations	Indicates support for test and clean operations on data cache, Harvard or unified architecture. 0x0 = no support.
[23:20]	L1 cache maintenance operations (unified)	Indicates support for L1 cache, entire cache maintenance operations, unified architecture. 0x0 = no support.
[19:16]	L1 cache maintenance operations (Harvard)	Indicates support for L1 cache, entire cache maintenance operations, Harvard architecture. 0x0 = no support.
[15:12]	L1 cache line maintenance operations - Set and Way (unified)	Indicates support for L1 cache line maintenance operations by Set and Way, unified architecture. 0x0 = no support.
[11:8]	L1 cache line maintenance operations - Set and Way (Harvard)	Indicates support for L1 cache line maintenance operations by Set and Way, Harvard architecture. 0x0 = no support.
[7:4]	L1 cache line maintenance operations - MVA (unified)	Indicates support for L1 cache line maintenance operations by address, unified architecture. 0x0 = no support.
[3:0]	L1 cache line maintenance operations - MVA (Harvard)	Indicates support for L1 cache line maintenance operations by address, Harvard architecture. 0x0 = no support.

To access the ID_MMFR1 read CP15 with:

MRC p15, 0, <Rd>, c0, c1, 5 ; Read ID_MMFR1

c0, Memory Model Feature Register 2

The ID_MMFR2 characteristics are:

Purpose The ID_MMFR2 provides information about the memory model, memory management, and cache support operations of the processor.

- Usage constraints** The ID_MMFR2 is:
- a read-only register
 - accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-13.

Figure 4-17 shows the ID_MMFR2 bit assignments.




Figure 4-17 ID_MMFR2 bit assignments

Table 4-13 shows the ID_MMFR2 bit assignments.

Table 4-13 ID_MMFR2 bit assignments

Bits	Name	Function
[31:28]	Hardware access flag	Indicates support for Hardware Access Flag. 0x0 = no support.
[27:24]	WFI	Indicates support for Wait-For-Interrupt stalling. 0x1 = the processor supports Wait-For-Interrupt.
[23:20]	Memory barrier	Indicates support for memory barrier operations. 0x2, = he processor supports: • DSB (formerly DWB) • ISB (formerly Prefetch Flush) • DMB.
[19:16]	TLB maintenance operations (unified)	Indicates support for TLB maintenance operations, unified architecture. 0x0 = no support.
[15:12]	TLB maintenance operations (Harvard)	Indicates support for TLB maintenance operations, Harvard architecture. 0x0 = no support.
[11:8]	L1 cache maintenance range operations (Harvard)	Indicates support for cache maintenance range operations, Harvard architecture. 0x0 = no support.
[7:4]	L1 background prefetch cache operations	Indicates support for background prefetch cache range operations, Harvard architecture. 0x0 = no support.
[3:0]	L1 foreground prefetch cache operations	Indicates support for foreground prefetch cache range operations, Harvard architecture. 0x0 = no support.

To access the ID_MMFR2 read CP15 with:

MRC p15, 0, <Rd>, c0, c1, 6 ; Read ID_MMFR2

c0, Memory Model Feature Register 3, MMFR3

The ID_MMFR3 characteristics are:

- Purpose** Provides information about the two cache line maintenance operations for the processor.
- Usage constraints** The ID_MMFR3 is:
- a read-only register
 - accessible in Privileged mode only.
- Configurations** Available in all processor configurations.
- Attributes** See Table 4-14.

Figure 4-18 shows the ID_MMFR3 bit assignments.




Figure 4-18 ID_MMFR3 bit assignments

Table 4-14 shows the ID_MMFR3 bit assignments.

Table 4-14 ID_MMFR3 bit assignments

Bits	Name	Function
[31:28]	Supersection support	RAZ because this is a PMSA implementation.
[27:24]	-	SBZ
[23:20]	Coherent walk	RAZ because this is a PMSA implementation.
[19:16]	-	SBZ
[15:12]	Maintenance broadcast	Indicates whether cache maintenance operations are broadcast. 0x0 = cache maintenance operations only affect local structures.

Table 4-14 ID_MMFR3 bit assignments (continued)

Bits	Name	Function
[11:8]	Branch predictor maintenance operations	Indicates support for branch predictor maintenance operations in systems with hierarchical cache maintenance operations. 0x2 = supports invalidate entire branch predictor array and invalidate branch predictor by MVA ^a .
[7:4]	Hierarchical cache maintenance operations by Set and Way	Indicates support for hierarchical cache maintenance operations by Set and Way. 0x1 = the processor supports invalidate cache, clean and invalidate, and clean by Set and Way.
[3:0]	Hierarchical cache maintenance operations by MVA	Indicates support for hierarchical cache maintenance operations by address. 0x1 = the processor supports: • Invalidate data cache by address • Clean data cache by address • Clean and invalidate data cache by address • Invalidate instruction cache by address • Invalidate all instruction cache entries.

a. Both of these operations are NOP on Cortex-R5.

To access the ID_MMFR3 read CP15 with:

MRC p15, 0, <Rd>, c0, c1, 7 ; Read ID_MMFR3

4.3.11 Instruction Set Attributes Registers

There are eight Instruction Set Attributes Registers, ID_ISAR0 to ID_ISAR7, but three of these are unused. This section describes:

- *c0, Instruction Set Attributes Register 0*
- *c0, Instruction Set Attributes Register 1* on page 4-28
- *c0, Instruction Set Attributes Register 2* on page 4-30
- *c0, Instruction Set Attributes Register 3* on page 4-31
- *c0, Instruction Set Attributes Register 4* on page 4-33
- *c0, Instruction Set Attributes Register 5* on page 4-34.
- *c0, Instruction Set Attributes Registers 6-7* on page 4-34.

c0, Instruction Set Attributes Register 0

The ID_ISAR0 characteristics are:

Purpose	Provides information about the instruction set that the processor supports, beyond the basic set.
Usage constraints	The ID_ISAR0 is: • a read-only register • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-15 on page 4-28.

Figure 4-19 on page 4-28 shows the ID_ISAR0 bit assignments.




Figure 4-19 ID_ISAR0 bit assignments

Table 4-15 shows the ID_ISAR0 bit assignments.

Table 4-15 ID_ISAR0 bit assignments

Bits	Name	Function
[31:28]	-	SBZ
[27:24]	Divide instructions	Indicates support for divide instructions. 0x1 = Support for UDIV and SDIV in the Thumb ISA. Applies to Cortex-R5, r0p0. 0x2 = Support for UDIV and SDIV in the ARM and Thumb ISA. Applies from Cortex-R5, r1p0.
[23:20]	Debug instructions	Indicates support for debug instructions. 0x1 = the processor supports BKPT.
[19:16]	Coprocessor instructions	Indicates support for coprocessor instructions other than separately attributed feature registers, such as CP15 registers and VFP. 0x0 = no support.
[15:12]	Compare and branch instructions	Indicates support for combined compare and branch instructions. 0x1 = the processor supports combined compare and branch instructions, CBNZ and CBZ.
[11:8]	Bitfield instructions	Indicates support for bitfield instructions. 0x1 = the processor supports bitfield instructions, BFC, BFI, SBFX, and UBFX.
[7:4]	Bit counting instructions	Indicates support for bit counting instructions. 0x1 = the processor supports CLZ.
[3:0]	Atomic instructions	Indicates support for atomic load and store instructions. 0x1 = the processor supports SWP and SWPB.

To access the ID_ISAR0, read CP15 with:

MRC p15, 0, <Rd>, c0, c2, 0 ; Read ID_ISAR0

c0, Instruction Set Attributes Register 1

The ID_ISAR1 characteristics are:

Purpose Provides information about the instruction set that the processor supports beyond the basic set.

Usage constraints The ID_ISAR1 is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-16.

Figure 4-20 shows the ID_ISAR1 bit assignments.




Figure 4-20 ID_ISAR1 bit assignments

Table 4-16 shows the ID_ISAR1 bit assignments.

Table 4-16 ID_ISAR1 bit assignments

Bits	Name	Function
[31:28]	Jazelle instructions	Indicates support for Jazelle instructions. 0x1 = the processor supports: • BXJ instruction • J bit in PSRs. For more information see <i>Program status registers</i> on page 3-12 and <i>Acceleration of execution environments</i> on page 3-28.
[27:24]	Interworking instructions	Indicates support for interworking instructions. 0x3 = the processor supports: • BX, and T bit in PSRs • BLX, and PC loads have BX behavior. • Data-processing instructions in the ARM instruction set with the PC as the destination and the S bit clear have BX-like behavior.
[23:20]	Immediate instructions	Indicates support for immediate instructions. 0x1 = the processor supports: • the MOVT instruction • MOV instruction encodings with 16-bit immediates • Thumb ADD and SUB instructions with 12-bit immediates.
[19:16]	ITE instructions	Indicates support for if then instructions. 0x1 = the processor supports IT instructions.
[15:12]	Extend instructions	Indicates support for sign or zero extend instructions. 0x2 = the processor supports: • SXTB, SXTB16, SXTH, UXTB, UXTB16, and UXTH • SXTAB, SXTAB16, SXTAH, UXTAB, UXTAB16, and UXTAH.

Table 4-16 ID_ISAR1 bit assignments (continued)

Bits	Name	Function
[11:8]	Exception 2 instructions	Indicates support for exception 2 instructions. 0x1 = the processor supports RFE, SRS, and CPS.
[7:4]	Exception 1 instructions	Indicates support for exception 1 instructions. 0x1 = the processor supports LDM (exception return), LDM (user registers), and STM (user registers).
[3:0]	Endian instructions	Indicates support for endianness control instructions. 0x1 = the processor supports SETEND and E bit in PSRs.

To access the ID_ISAR1 read CP15 with:

MRC p15, 0, <Rd>, c0, c2, 1 ; Read ID_ISAR1

c0, Instruction Set Attributes Register 2

The ID_ISAR2 is:

- a read-only register
- accessible in Privileged mode only.

The ID_ISAR2 characteristics are:

Purpose The ID_ISAR2 provides information about the instruction set that the processor supports beyond the basic set.

Usage constraints The ID_ISAR2 is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-17 on page 4-31.

Figure 4-21 shows the ID_ISAR2 bit assignments.




Figure 4-21 ID_ISAR2 bit assignments

Table 4-17 shows the ID_ISAR2 bit assignments.

Table 4-17 ID_ISAR2 bit assignments

Bits	Name	Function
[31:28]	Reversal instructions	Indicates support for reversal instructions. 0x2 = the processor supports REV, REV16, REVSH, and RBIT.
[27:24]	PSR instructions	Indicates support for PSR instructions. 0x1 = the processor supports MRS and MSR, and the exception return forms of data-processing instructions.
[23:20]	Unsigned multiply instructions	Indicates support for advanced unsigned multiply instructions. 0x2 = the processor supports: • UMULL and UMLAL • UMAAL.
[19:16]	Signed multiply instructions	Indicates support for advanced signed multiply instructions. 0x3 = the processor supports: • SMULL and SMLAL • SMLABB, SMLABT, SMLALBB, SMLALBT, SMLALTB, SMLALTT, SMLATB, SMLATT, SMLAWB, SMLAWT, SMULBB, SMULBT, SMULTB, SMULTT, SMULWB, SMULWT, and Q flag in PSRs • SMLAD, SMLADX, SMLALD, SMLALDX, SMLSD, SMLS DX, SMLS LD, SMLS LD, SMMLA, SMMLAR, SMMLS, SMMLSR, SMPUL, SMPULR, SMUAD, SMUADX, SMUSD, and SMUSD.
[15:12]	Multiply instructions	Indicates support for multiply instructions. 0x2 = the processor supports MUL, MLA, and MLS.
[11:8]	Interruptible instructions	Indicates support for multi-access interruptible instructions. 0x1 = the processor supports restartable LDM and STM.
[7:4]	Memory hint instructions	Indicates support for memory hint instructions. 0x3 = the processor supports PLD and PLI. Applies to Cortex-R5, r0p0 0x4 = the processor supports PLD, PLI and PLDW. Applies from Cortex-R5, r1p0
[3:0]	Load/store instructions	Indicates support for additional load and store instructions. 0x1 = the processor supports LDRD and STRD.

To access the ID_ISAR2 read CP15 with:

MRC p15, 0, <Rd>, c0, c2, 2 ; Read ID_ISAR2

c0, Instruction Set Attributes Register 3

The ID_ISAR3 characteristics are:

Purpose	Provides information about the instruction set that the processor supports beyond the basic set.
Usage constraints	The ID_ISAR3 is: • a read-only registers • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-18 on page 4-32.

Figure 4-22 on page 4-32 shows the ID_ISAR3 bit assignments.




Figure 4-22 ID_ISAR3 bit assignments

Table 4-18 shows the ID_ISAR3 bit assignments.

Table 4-18 ID_ISAR3 bit assignments

Bits	Name	Function
[31:28]	ThumbEE extension	Indicates support for ThumbEE Execution Environment extension. 0x0 = no support.
[27:24]	True NOP instructions	Indicates support for true NOP instructions. 0x1 = the processor supports NOP16, NOP32 and various NOP compatible hints in both the ARM and Thumb instruction sets.
[23:20]	Thumb copy instructions	Indicates support for Thumb copy instructions. 0x1 = the processor supports Thumb MOV(3) low register ⇒ low register.
[19:16]	Table branch instructions	Indicates support for table branch instructions. 0x1 = the processor supports table branch instructions, TBB and TBH.
[15:12]	Synchronization primitive instructions	Indicates support for synchronization primitive instructions. 0x2 = the processor supports: • LDREX and STREX • LDREXB, LDREXH, LDREXD, STREXB, STREXH, STREXD, and CLREX.
[11:8]	SVC instructions	Indicates support for SVC (formerly SWI) instructions. 0x1 = the processor supports SVC.
[7:4]	SIMD instructions	Indicates support for <i>Single Instruction Multiple Data</i> (SIMD) instructions. 0x3 = the processor supports: PKHBT, PKHTB, QADD16, QADD8, QASX, QSUB16, QSUB8, QSAX, SADD16, SADD8, SASX, SEL, SHADD16, SHADD8, SHASX, SHSUB16, SHSUB8, SHSAX, SSAT, SSAT16, SSUB16, SSUB8, SSAX, SXTAB16, SXTB16, UADD16, UADD8, UASX, UHADD16, UHADD8, UASX, UHSUB16, UHSUB8, USAX, UQADD16, UQADD8, UQASX, UQSUB16, UQSUB8, UQSAX, USAD8, USADA8, USAT, USAT16, USUB16, USUB8, USAX, UXTAB16, UXTB16, and the GE[3:0] bits in the PSRs.
[3:0]	Saturate instructions	Indicates support for saturate instructions. 0x1 = the processor supports QADD, QDADD, QDSUB, QSUB and Q flag in PSRs.

To access the ID_ISAR3 read CP15 with:

MRC p15, 0, <Rd>, c0, c2, 3 ; Read ID_ISAR3

c0, Instruction Set Attributes Register 4

The ID_ISAR4 characteristics are:

Purpose	Provides information about the instruction set that the processor supports beyond the basic set.
Usage constraints	The ID_ISAR4 is: - a read-only register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-19.

Figure 4-23 shows the ID_ISAR4 bit assignments.




Figure 4-23 ID_ISAR4 bit assignments

Table 4-19 shows the ID_ISAR4 bit assignments.

Table 4-19 ID_ISAR4 bit assignments

Bits	Name	Function
[31:28]	SWP_frac	RAZ because SWP/SWPB instruction support is indicated in ID_ISAR0.
[27:24]	PSR_M_instrs	Indicates support for M-profile instructions for modifying the PSRs. 0x0 = no support.
[23:20]	Exclusive instructions	Indicates support for Exclusive instructions. 0x0 = Only supports synchronization primitive instructions as indicated by bits [15:12] in the ISAR3 register. See <i>c0, Instruction Set Attributes Register 3</i> on page 4-31 for more information.
[19:16]	Barrier instructions	Indicates support for Barrier instructions. 0x1 = the processor supports DMB, DSB, and ISB instructions.
[15:12]	SMC instructions	Indicates support for <i>Secure Monitor Call</i> (SMC) (formerly SMI) instructions. 0x0 = no support.
[11:8]	Write-back instructions	Indicates support for write-back instructions. 0x1 = supports all the writeback addressing modes defined in ARMv7.
[7:4]	With shift instructions	Indicates support for with-shift instructions. 0x4 = the processor supports: - the full range of constant shift options, on load/store and other instructions - register-controlled shift options.
[3:0]	Unprivileged instructions	Indicates support for Unprivileged instructions. 0x2 = the processor supports LDR{SB B SH H}T and STR{B H}T.

To access the ID_ISAR4 read CP15 with:

MRC p15, 0, <Rd>, c0, c2, 4 ; Read ID_ISAR4

c0, Instruction Set Attributes Register 5

The ID_ISAR5 characteristics are:

- Purpose** Provides additional information about the properties of the processor.
- Usage constraints** ID_ISAR5 is:
- a read-only register
 - accessible in Privileged mode only.
- Configurations** Available in all processor configurations.
- Attributes** In the processor, ID_ISAR5 is read as 0x00000000.

To access the ID_ISAR5, read CP15 with:

MRC p15, 0, <Rd>, c0, c2, 5 ; Read Instruction Set Attribute Register 5

c0, Instruction Set Attributes Registers 6-7

ID_ISAR6 and ID_ISAR7 are not implemented, and their positions in the register map are Reserved. They correspond to CP15 accesses with:

MRC p15, 0, <Rd>, c0, c2, 6 ; Read ID_ISAR6

MRC p15, 0, <Rd>, c0, c2, 7 ; Read ID_ISAR7

These registers are read-only, and are accessible in Privileged mode only.

4.3.12 c0, Cache Size ID Register

The CCSIDR Register characteristics are:

- Purpose** Provides information about the size and behavior of the instruction or data cache. Architecturally, there can be up to eight levels of cache, containing instruction, data, or unified caches. This processor contains L1 instruction and data caches only. The CSSELR determines which CCSIDR to select, see *c0, Cache Size Selection Register* on page 4-37.
- Usage constraints** The CCSIDR is:
- a read-only register
 - accessible in Privileged mode only.
- Configurations** Available in all processor configurations.
- Attributes** See Table 4-20 on page 4-35.

Figure 4-24 shows the CCSIDR bit assignments.

31	30	29	28	27					13	12					2	0
W	W	R	W		NumSets					Associativity					Line Size	
T	B	A	A													

Figure 4-24 CCSIDR bit assignments

Table 4-20 shows the CCSIDR bit assignments.

Table 4-20 CCSIDR bit assignments

Bits	Name	Function
[31]	WT	Indicates support available for write-through: 1 = write-through support available ^a
[30]	WB	Indicates support available for write-back: 1 = write-back support available ^a
[29]	RA	Indicates support available for read allocation: 1 = read allocation support available ^a
[28]	WA	Indicates support available for write allocation: 1 = write allocation support available ^a
[27:13]	NumSets	Indicates the number of sets as (number of sets) - 1 ^a
[12:3]	Associativity	Indicates the number of ways as (number of ways) - 1 ^a
[2:0]	LineSize	Indicates the number of words in each cache line ^a

a. See Table 4-21 for valid bit field encodings.

The LineSize field is encoded as 2 less than $\log_2$ of the number of words in the cache line. For example, a value of 0x0 indicates there are four words in a cache line, that is the minimum size for the cache. A value of 0x1 indicates there are eight words in a cache line.

Table 4-21 shows the individual bit field and complete register encodings for the CCSIDR. Use this to match the cache size and level of cache set by the *Current Cache Size Selection Register* (CSSR). See *c0, Cache Size Selection Register* on page 4-37.

Table 4-21 Bit field and register encodings for CCSIDR

Size	Complete register encoding	Register bit field encoding						
		WT	WB	RA	WA	NumSets	Associativity	LineSize
4KB	0xF003E019	1	1	1	1	0x001F	0x3	0x1
8KB	0xF007E019	1	1	1	1	0x003F		
16KB	0xF00FE019	1	1	1	1	0x007F		
32KB	0xF01FE019	1	1	1	1	0x00FF		
64KB	0xF03FE019	1	1	1	1	0x01FF		

To access the CCSIDR read CP15 with:

MRC p15, 1, <Rd>, c0, c0, 0 ; Read CCSIDR

4.3.13 c0, Cache Level ID Register

The CLIDR Register characteristics are:

Purpose	- Indicates the cache levels that are implemented. Architecturally, there can be a different number of cache levels on the instruction and data side. - Captures the point-of-coherency. - Captures the point-of-unification.
Usage constraints	The CLIDR is: - a read-only register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-22.

Figure 4-25 shows the CLIDR bit assignments.




Figure 4-25 CLIDR Register bit assignments

Table 4-22 shows the CLIDR bit assignments.

Table 4-22 CLIDR Register bit assignments

Bits	Name	Function
[31:30]	-	SBZ
[29:27]	LoU	Level of Unification. 0b001 = level 2, if either cache is implemented 0b000 = level 1, if neither instruction nor data cache is implemented.
[26:24]	LoC	Level of Coherency. 0b001 = level 2, if either cache is implemented 0b000 = level 1, if neither instruction nor data cache is implemented.
[23:21]	LoUIS	Level of Unification Inner Shareable 0b000 = MP extensions are not implemented. 0b001 = Level 2
[20:18]	CL 7	0b000 = no cache at CL 7.
[17:15]	CL 6	0b000 = no cache at CL 6.
[14:12]	CL 5	0b000 = no cache at CL 5.
[11:9]	CL 4	0b000 = no cache at CL 4.
[8:6]	CL 3	0b000 = no cache at CL 3.
[5:3]	CL 2	0b000 = no cache at CL 2.

Table 4-22 CLIDR Register bit assignments (continued)

Bits	Name	Function
[2]	CL 1	RAZ. Indicates no unified cache at CL1.
[1]	CL 1	0b001 = data cache is implemented 0b000 = no data cache is implemented.
[0]	CL 1	0b001 = an instruction cache is implemented 0b000 = no instruction cache is implemented.

To access the CLIDR, read CP15 with:

MRC p15, 1, <Rd>, c0, c0, 1 ; Read CLIDR

4.3.14 c0, Auxiliary ID Register

The AIDR is:

- a read-only register
- accessible in Privileged mode only.

The AIDR characteristics are:

Purpose Provides additional information about the processor.

Usage constraints The AIDR is:

- a read-only register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes In this processor, the AIDR reads as 0x00000000.

To access the AIDR read CP15 with:

MRC p15, 1, <Rd>, c0, c0, 7 ; Read AIDR

4.3.15 c0, Cache Size Selection Register

The CSSELR characteristics are:

Purpose Holds the value that the processor uses to select the CSSELR to use.

Usage constraints The CSSELR is:

- a read/write register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-23 on page 4-38.

Figure 4-26 on page 4-38 shows the CSSELR bit assignments.




Figure 4-26 CSSELR bit assignments

Table 4-23 shows the CSSELR bit assignments.

Table 4-23 CSSELR bit assignments

Bits	Name	Function
[31: 4]	-	SBZ.
[3:1]	Level	Identifies which cache level to select. b000 = Level 1 cache This field is read only, writes are ignored.
[0]	InD	Identifies instruction or data cache to use. 1 = instruction 0 = data.

To access the CCSIDRs read or write CP15 with:

MRC p15, 2, <Rd>, c0, c0, 0 ; Read CSSELR
MCR p15, 2, <Rd>, c0, c0, 0 ; Write CSSELR

4.3.16 c1, System Control Register

The SCTLr characteristics are:

Purpose	Provides control and configuration information for: - memory alignment, endianness, protection, and fault behavior - MPU and cache enables and cache replacement strategy - interrupts and the behavior of interrupt latency - the location for exception vectors - program flow prediction.
Usage constraints	The SCTLr is: - a read/write register - accessible in Privileged mode only. - Attempts to read or write the SCTLr from User mode result in an Undefined Instruction exception.
Configurations	Available in all processor configurations.
Attributes	See Table 4-24 on page 4-39.

Figure 4-27 on page 4-39 shows the SCTLr bit assignments.




Figure 4-27 SCTLR bit assignments

Table 4-24 shows the SCTLR bit assignments.

Table 4-24 SCTLR bit assignments

Bits	Name	Function
[31]	IE	Identifies little or big instruction endianness in use: 0 = little-endianness 1 = big-endianness. The primary input CFGIE defines the value. This bit is read-only.
[30]	TE	Thumb exception enable: 0 = enable ARM exception generation 1 = enable Thumb exception generation. The primary input TEINIT defines the reset value.
[29]	AFE	Access Flag Enable. On the processor this bit is SBZ.
[28]	TRE	TEX Remap Enable. On the processor this bit is SBZ.
[27]	NMFI	NMFI, non-maskable fast interrupt enable: 0 = Software can disable FIQs 1 = Software cannot disable FIQs. This bit is read-only. The configuration input CFGNMFI defines its value.
[26]	-	SBZ.
[25]	EE	Determines how the E bit in the CPSR is set on an exception: 0 = CPSR E bit is set to 0 on an exception 1 = CPSR E bit is set to 1 on an exception. The primary input CFGEE defines the reset value.
[24]	VE	Configures vectored interrupt: 0 = exception vector address for IRQ is 0x00000018 or 0xFFFF0018. See V bit. 1 = VIC controller provides handler address for IRQ. The reset value of this bit is 0.
[23:22]	-	SBO.
[21]	FI	Fast Interrupts enable. On the processor Fast Interrupts are always enabled. This bit is SBO.
[20]	-	SBZ.

Table 4-24 SCTLR bit assignments (continued)

Bits	Name	Function
[19]	DZ	Divide by zero: 0 = do not generate an Undefined Instruction exception 1 = generate an Undefined Instruction exception. The reset value of this bit is 0.
[18]	-	SBO.
[17]	BR	MPU background region enable.
[16]	-	SBO.
[15]	-	SBZ.
[14]	RR	Round-robin bit, controls replacement strategy for instruction and data caches: 0 = random replacement strategy 1 = round-robin replacement strategy. The reset value of this bit is 0. The processor always uses a random replacement strategy, regardless of the state of this bit.
[13]	V	Determines the location of exception vectors: 0 = normal exception vectors selected, address range = 0x00000000-0x0000001C 1 = high exception vectors (HIVECS) selected, address range = 0xFFFF0000-0xFFFF001C. The primary input VINITHIm defines the reset value.
[12]	I	Enables L1 instruction cache: 0 = instruction caching disabled. This is the reset value. 1 = instruction caching enabled. If no instruction cache is implemented, then this bit is SBZ.
[11]	Z	Branch prediction enable bit. The processor supports branch prediction. This bit is SBO. The ACTLR can control branch prediction, see <i>c1</i> , <i>Auxiliary Control Register</i> on page 4-41.
[10]	SW	Enables SWP and SWPB instructions 0 = SWP and SWPB are Undefined 1 = SWP and SWPB are executed with full locking support on the bus The reset value of this bit is 0. ^a
[9:7]	-	SBZ.
[6:3]	-	SBO.
[2]	C	Enables L1 data cache: 0 = data caching disabled. This is the reset value. 1 = data caching enabled. If no data cache is implemented, then this bit is SBZ.
[1]	A	Enables strict alignment of data to detect alignment faults in data accesses: 0 = strict alignment fault checking disabled. This is the reset value. 1 = strict alignment fault checking enabled.
[0]	M	Enables the MPU: 0 = MPU disabled. This is the reset value. 1 = MPU enabled. If no MPU is implemented, this bit is SBZ.

- a. Unless explicitly enabled, SWP and SWPB are Undefined

To use the SCTL, ARM recommends that you use a read-modify-write technique. To access the SCTL, read or write CP15 with:

MRC p15, 0, <Rd>, c1, c0, 0 ; Read SCTL
MCR p15, 0, <Rd>, c1, c0, 0 ; Write SCTL

Attempts to read or write the SCTL from User mode results in an Undefined Instruction exception.

4.3.17 c1, Auxiliary Control Register

The ACTLR characteristics are:

Purpose

Controls:

- branch prediction
- performance features
- error and parity logic.

Usage constraints

The ACTLR is:

- A read/write register.
- Accessible in Privileged mode only.
- ARM recommends that any instruction that changes bits [31:28] or [7] is followed by an ISB instruction to ensure that the changes have taken effect before any dependent instructions are executed.

Configurations

Available in all processor configurations.

Attributes

See Table 4-25 on page 4-42.

Figure 4-28 shows the ACTLR bit assignments.




Figure 4-28 ACTLR bit assignments

Table 4-25 shows the ACTLR bit assignments.

Table 4-25 ACTLR bit assignments

Bits	Name	Function
[31]	DICDI ^a	Case C dual issue control: 0 = Enabled. This is the reset value. 1 = Disabled.
[30]	DIB2DI ^a	Case B2 dual issue control: 0 = Enabled. This is the reset value. 1 = Disabled.
[29]	DIB1DI ^a	Case B1 dual issue control: 0 = Enabled. This is the reset value. 1 = Disabled.
[28]	DIADI ^a	Case A dual issue control: 0 = Enabled. This is the reset value. 1 = Disabled.
[27]	B1TCMPCEN	B1TCM ECC check enable: 0 = Disabled 1 = Enabled. The primary input PARECCENRAMm[2] ^b defines the reset value. If the B1TCM is configured with ECC, you must always set this bit to the same value as B0TCMPCEN .
[26]	B0TCMPCEN	B0TCM ECC check enable: 0 = Disabled 1 = Enabled. The primary input PARECCENRAMm[1] ^b defines the reset value. If the B0TCM is configured with ECC, you must always set this bit to the same value as B1TCMPCEN .
[25]	ATCMPCEN	ATCM ECC check enable: 0 = Disabled 1 = Enabled. The primary input PARECCENRAMm[0] ^b defines the reset value.
[24]	AXISCEN	AXI slave cache RAM access enable: 0 = Disabled. This is the reset value. 1 = Enabled. Note When AXI slave cache access is enabled, the caches are disabled and the processor cannot run any cache maintenance operations. If the processor attempts a cache maintenance operation, an Undefined Instruction exception is taken.
[23]	AXISCUEN	AXI slave cache RAM non-privileged access enable: 0 = Disabled. This is the reset value. 1 = Enabled.
[22]	DILSM	Disable <i>Low Interrupt Latency</i> (LIL) on load/store multiples: 0 = Enable LIL on load/store multiples. This is the reset value. 1 = Disable LIL on all load/store multiples.

Table 4-25 ACTLR bit assignments (continued)

Bits	Name	Function
[21]	DEOLP	Disable end of loop prediction: 0 = Enable loop prediction. This is the reset value. 1 = Disable loop prediction.
[20]	DBHE	Disable <i>Branch History</i> (BH) extension: 0 = Enable the extension. This is the reset value. 1 = Disable the extension.
[19]	FRCDIS	Fetch rate control disable: 0 = Normal fetch rate control operation. This is the reset value. 1 = Fetch rate control disabled.
[18]	-	SBZ.
[17]	RSDIS	Return stack disable: 0 = Normal return stack operation. This is the reset value. 1 = Return stack disabled.
[16:15]	BP	This field controls the branch prediction policy: b00 = Normal operation. This is the reset value. b01 = Branch always taken and history table updates disabled. b10 = Branch always not taken and history table updates disabled. b11 = Reserved. Behavior is Unpredictable if this field is set to b11.
[14]	DBWR	Disable write burst in the AXI master: 0 = Normal operation. This is the reset value. 1 = Disable write burst optimization.
[13]	DLFO	Disable linefill optimization in the AXI master: 0 = Normal operation. This is the reset value. 1 = Limits the number of outstanding data linefills to two.
[12]	ERPEG ^c	Enable random parity error generation: 0 = Random parity error generation disabled. This is the reset value. 1 = Enable random parity error generation in the cache RAMs. Note This bit controls error generation logic during system validation. A synthesized ASIC typically does not have such models and this bit is therefore redundant for ASICs.
[11]	DNCH	Disable data forwarding for Non-cacheable accesses in the AXI master: 0 = Normal operation. This is the reset value. 1 = Disable data forwarding for Non-cacheable accesses.
[10]	FORA	Force outer read allocate (ORA) for outer write allocate (OWA) regions: 0 = No forcing of ORA. This is the reset value. 1 = ORA forced for OWA regions.
[9]	FWT	Force write-through (WT) for write-back (WB) regions: 0 = No forcing of WT. This is the reset value. 1 = WT forced for WB regions.

Table 4-25 ACTLR bit assignments (continued)

Bits	Name	Function
[8]	FDSnS	Force D-side to not-shared when MPU is off: 0 = Normal operation. This is the reset value. 1 = D-side normal Non-cacheable forced to Non-shared when MPU is off.
[7]	sMOV	sMOV of a divide does not complete out of order. No other instruction is issued until the divide is finished. 0 = Normal operation. This is the reset value. 1 = sMOV out of order disabled.
[6]	DILS	Disable low interrupt latency on all load/store instructions. 0 = Enable LIL on all load/store instructions. This is the reset value. 1 = Disable LIL on all load/store instructions.
[5:3]	CEC	Cache error control for cache parity and ECC errors. See Table 8-2 on page 8-21 and Table 8-3 on page 8-22 for more information about how these bits are used. The reset value is b100.
[2]	B1TCMECEN	B1TCM external error enable: 0 = Disabled 1 = Enabled. The primary input ERRENRAMm[2] defines the reset value.
[1]	B0TCMECEN	B0TCM external error enable: 0 = Disabled 1 = Enabled. The primary input ERRENRAMm[1] defines the reset value.
[0]	ATCMECEN	ATCM external error enable: 0 = Disabled 1 = Enabled. The primary input ERRENRAMm[0] defines the reset value.

- a. See *Dual issue* on page B-33
b. See *Configuration signals* on page A-4.
c. This bit is only supported if parity error generation is implemented in your design.

To access the ACTLR, read or write CP15 with:

MRC p15, 0, <Rd>, c1, c0, 1 ; Read ACTLR
MCR p15, 0, <Rd>, c1, c0, 1 ; Write ACTLR

4.3.18 c15, Secondary Auxiliary Control Register

The Secondary Auxiliary Control Register characteristics are:

Purpose	Controls: • branch prediction • performance features • error and parity logic.
Usage constraints	The Secondary Auxiliary Control Register is: • A read/write register. • Accessible in Privileged mode only.

- ARM recommends that any instruction that changes bits [20:16] is followed by an ISB instruction to ensure that the changes have taken effect before any dependent instructions are executed.

Configurations Available in all processor configurations.

Attributes See Table 4-26.

Figure 4-29 shows the Secondary Auxiliary Control Register bit assignments.




Figure 4-29 Secondary Auxiliary Control Register bit assignments

Table 4-26 shows the Secondary Auxiliary Control Register bit assignments.

Table 4-26 Secondary Auxiliary Control Register bit assignments

Bits	Name	Function
[31:23]	-	SBZ.
[22]	DCHE	Disable hard-error support in the caches. ^a 0 = Enabled. The cache logic recovers from some hard errors. 1 = Disabled. Most hard errors in the caches are fatal. This is the reset value. See <i>Hard errors</i> on page 8-5 for more information.
[21]	DR2B ^b	Enable random 2-bit error generation in cache RAMs. This bit has no effect unless ECC is configured, see <i>Configurable options</i> on page 1-6. 0 = Disabled. This is the reset value. 1 = Enabled. Note This bit controls error generation logic during system validation. A synthesized ASIC typically does not have such models and this bit is therefore redundant for ASICs.
[20]	DF6DI	F6 dual issue control. ^c 0 = Enabled. This is the reset value. 1 = Disabled.
[19]	DF2DI	F2_Id/F2_st/F2D dual issue control. ^c 0 = Enabled. This is the reset value. 1 = Disabled.

Table 4-26 Secondary Auxiliary Control Register bit assignments (continued)

Bits	Name	Function
[18]	DDI	F1/F3/F4dual issue control. ^c 0 = Enabled. This is the reset value. 1 = Disabled.
[17]	DOODPPF	Out-of-order double-precision floating point instruction control. ^c 0 = Enabled. This is the reset value. 1 = Disabled.
[16]	DOOFMACS	Out-of-order FMACS control. ^c 0 = Enabled. This is the reset value. 1 = Disabled.
[15:14]	-	SBZ.
[13]	IXC	Floating-point inexact exception output mask. ^c 0 = Mask floating-point inexact exception output. The output FPIXCm is forced to zero. This is the reset value. 1 = Propagate floating point inexact exception flag FPSCR.IXC to output FPIXCm .
[12]	OFC	Floating-point overflow exception output mask. ^c 0 = Mask floating-point overflow exception output. The output FPOFCm is forced to zero. This is the reset value. 1 = Propagate floating-point overflow exception flag FPSCR.OFC to output FPOFCm .
[11]	UFC	Floating-point underflow exception output mask. ^c 0 = Mask floating-point underflow exception output. The output FPUFCm is forced to zero. This is the reset value. 1 = Propagate floating-point underflow exception flag FPSCR.UFC to output FPUFCm .
[10]	IOC	Floating-point invalid operation exception output mask. ^c 0 = Mask floating-point invalid operation exception output. The output FPIOCm is forced to zero. This is the reset value. 1 = Propagate floating-point invalid operation exception flag FPSCR.IOC to output FPIOCm .
[9]	DZC	Floating-point divide-by-zero exception output mask. ^c 0 = Mask floating-point divide-by-zero exception output. The output FPDZCm is forced to zero. This is the reset value. 1 = Propagate floating-point divide-by-zero exception flag FPSCR.DZC to output FPDZCm .
[8]	IDC	Floating-point input denormal exception output mask. ^c 0 = Mask floating-point input denormal exception output. The output FPIDCm is forced to zero. This is the reset value. 1 = Propagate floating-point input denormal exception flag FPSCR.IDC to output FPIDCm .
[7:4]	-	SBZ.
[3]	BTCMECC	Correction for internal ECC logic on BTCM ports. ^d 0 = Enabled. This is the reset value. 1 = Disabled.

Table 4-26 Secondary Auxiliary Control Register bit assignments (continued)

Bits	Name	Function
[2]	ATCMECC	Correction for internal ECC logic on ATCM port. ^d 0 = Enabled. This is the reset value. 1 = Disabled.
[1]	BTCMRMW	Enables 64-bit stores for the BTCMs. When enabled, the processor uses read-modify-write to ensure that all reads and writes presented on the BTCM ports are 64 bits wide. ^e 0 = Disabled 1 = Enabled. The primary input RMWENRAMm[1] defines the reset value.
[0]	ATCMRMW	Enables 64-bit stores for the ATCM. When enabled, the processor uses read-modify-write to ensure that all reads and writes presented on the ATCM port are 64 bits wide. ^e 0 = Disabled 1 = Enabled. The primary input RMWENRAMm[0] defines the reset value.

- This bit is RAZ if both caches have neither ECC nor parity.
- This bit is only supported if parity error generation is implemented in your design.
- This bit has no effect unless the *Floating Point Unit* (FPU) has been configured, see *Configurable options* on page 1-6.
- This bit has no effect unless TCM ECC logic has been configured for the respective TCM interface, see *Configurable options* on page 1-6.
- This feature is not available when the TCM interface has been built with 32-bit ECC.

To access the Secondary Auxiliary Control Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c0, 0 ; Read Secondary Auxiliary Control Register
MCR p15, 0, <Rd>, c15, c0, 0 ; Write Secondary Auxiliary Control Register

4.3.19 c1, Coprocessor Access Control Register

The CPACR characteristics are:

Purpose	Sets access rights for coprocessors.
Usage constraints	The CPACR is: - A read/write register. - Accessible in Privileged mode only. - Because this processor does not support coprocessors CP0 through CP9, CP12, and CP13, bits [27:24] and [19:0] in this register are read-as-zero and ignore writes. - CPACR has no effect on access to CP14, the debug control coprocessor, or CP15, the system control coprocessor. The only other coprocessor that the Cortex-R5F CPU includes is the FPU, CP10, and CP11. This register enables software to determine if the FPU exists in the CPU.
Configurations	Available in all processor configurations.
Attributes	See Table 4-27 on page 4-48.

Figure 4-30 on page 4-48 shows the CPACR bit assignments.




Figure 4-30 CPACR bit assignments

Table 4-27 shows the CPACR bit assignments.

Table 4-27 CPACR bit assignments

Bits	Name	Function
[31]	ASEDIS	Advanced-SIMD disable. Read only. 0 = FPU is not configured 1 = FPU is configured, Advanced SIMD is not available.
[30]	D32DIS	D16-D31 disable. Read only. 0 = FPU is not configured 1 = FPU is configured, VFP registers D16-D32 are not available.
[29:28]	-	Read as Zero.
[27:26]	cp13	Read as Zero.
[25:24]	cp12	
[23:22]	cp11	Defines access permissions for the FPU. If the FPU is not included for this processor, these bits are RAZ/WI.
[21:22]	cp10	If the FPU is included, both cp10 and cp11 must be programmed to the same value: b00 = Access denied. Attempts to access generates an Undefined Instruction exception. This is the reset value. b01 = Privileged mode access only b10 = Reserved b11 = Privileged and User mode access.
[19:18]	cp9	Read as Zero.
[17:16]	cp8	
[15:14]	cp7	
[13:12]	cp6	
[11:10]	cp5	
[9:8]	cp4	
[7:6]	cp3	
[5:4]	cp2	
[3:2]	cp1	
[1:0]	cp0	

To access the CPACR, read or write CP15 with:

MRC p15, 0, <Rd>, c1, c0, 2 ; Read CPACR

MCR p15, 0, <Rd>, c1, c0, 2 ; Write CPACR

4.3.20 Fault Status and Address Registers

The processor reports the status and address of faults that occur during its operation. For both data and instruction faults there are two *Fault Status Registers* (FSRs) and one *Fault Address Register* (FAR).

Fields within the Data and Instruction FSRs indicate the priority and source of a fault and the validity of the address in the corresponding FAR. Table 4-28 shows this encoding for the FSRs.

Table 4-28 Fault Status Register encodings

Priority	Sources	FSR [10,3:0]	FAR
Highest	Alignment	0b00001	Valid
	Background	0b00000	Valid
	Permission	0b01101	Valid
	Synchronous External Abort	0b01000	Valid
	Asynchronous External Abort	0b10110	Unpredictable
	Synchronous Parity or ECC Error	0b11001	Valid
	Asynchronous Parity or ECC Error	0b11000	Unpredictable
Lowest	Debug Event	0b00010	Unchanged

All other encodings for these FSR bits are Reserved.

c5, Data Fault Status Register

The DFSR is:

- a read/write register
- accessible in Privileged mode only.

The DFSR characteristics are:

Purpose Holds status information regarding the source of the last data abort.

Usage constraints The DFSR is:

- a read/write register
- accessible in Privileged mode only.

Configurations Available in all processor configurations.

Attributes See Table 4-29 on page 4-50.

Figure 4-31 shows the DFSR bit assignments.




Figure 4-31 DFSR bit assignments

Table 4-29 shows the DFSR bit assignments.

Table 4-29 DFSR bit assignments

Bits	Name	Function
[31:13]	-	SBZ.
[12]	SD	Distinguishes between an AXI Decode or Slave error on an external abort. This bit is only valid for external aborts. For all other aborts types of abort, this bit is set to zero: 0 = AXI Decode error (DECERR), or AHB error, caused the abort 1 = AXI Slave error (SLVERR), or unsupported exclusive access, for example exclusive access using the AHB peripheral port, caused the abort.
[11]	RW	Indicates whether a read or write access caused an abort: 0 = read access caused the abort 1 = write access caused the abort.
[10] ^a	S	Part of the Status field.
[9:8]	-	Always read as 0. Writes ignored.
[7:4]	Domain	SBZ. This is because domains are not implemented in this processor.
[3:0] ^a	Status	Indicates the type of fault generated. To determine the data fault, you must use bit [12] and bit [10] in conjunction with bits [3:0].

a. For more information on how these bits are used in reporting faults, see Table 4-28 on page 4-49.

To use the DFSR read or write CP15 with:

MRC p15, 0, <Rd>, c5, c0, 0 ; Read DFSR
MCR p15, 0, <Rd>, c5, c0, 0 ; Write DFSR

c5, Instruction Fault Status Register

The IFSR characteristics are:

- Purpose** Holds status information regarding the source of the last instruction abort.
- Usage constraints** The IFSR is:
- a read/write register
 - accessible in Privileged mode only.
- Configurations** Available in all processor configurations.
- Attributes** See Table 4-30 on page 4-51.

Figure 4-32 shows the IFSR bit assignments.




Figure 4-32 IFSR bit assignments

Table 4-30 shows the IFSR bit assignments.

Table 4-30 IFSR bit assignments

Bits	Name	Function
[31:13]	-	SBZ.
[12]	SD	Distinguishes between an AXI Decode or Slave error on an external abort. This bit is only valid for external aborts. For all other aborts types of abort, this bit is set to zero: 0 = AXI Decode error (DECERR) caused the abort 1 = AXI Slave error (SLVERR) caused the abort.
[11]	-	SBZ.
[10] ^a	S	Part of the Status field.
[9:8]	-	SBZ.
[7:4]	Domain	SBZ. This is because domains are not implemented in this processor.
[3:0] ^a	Status	Indicates the type of fault generated. To determine the instruction fault, bit [12] and bit [10] must be used in conjunction with bits [3:0].

a. For more information on how these bits are used in reporting faults, see Table 4-28 on page 4-49.

To access the IFSR read or write CP15 with:

MRC p15, 0, <Rd>, c5, c0, 1 ; Read IFSR
MCR p15, 0, <Rd>, c5, c0, 1 ; Write IFSR

c5, Auxiliary Fault Status Registers

There are two auxiliary fault status registers:

- the *Auxiliary Data Fault Status Register* (ADFSR)
- the *Auxiliary Instruction Fault Status Register* (AIFSR).

The auxiliary fault status registers characteristics are:

Purpose Provide additional information about data and instruction parity, ECC, and external TCM errors.

Usage constraints The auxiliary fault status registers are:

- Read/write registers.
- Accessible in Privileged mode only.
- The contents of an auxiliary fault status register are only valid when the corresponding Data or Instruction Fault Status Register indicates that a parity or ECC error has occurred. At other times the contents of the auxiliary fault status registers are Unpredictable.

Configurations Available in all processor configurations.

Attributes See Table 4-31 on page 4-52.

Figure 4-33 on page 4-52 shows the auxiliary fault status registers bit assignments.


**Figure 4-33 Auxiliary fault status registers bit assignments**

Table 4-31 shows the auxiliary fault status registers bit assignments.

Table 4-31 ADFSR and AIFSR bit assignments

Bits	Name	Function
[31:28]	-	SBZ.
[27:24]	CacheWay ^a	The value returned in this field indicates the cache way or ways in which the error occurred.
[23:22]	Side	The value returned in this field indicates the source of the error. See Table 4-32 for the encodings.
[21]	Recoverable error	The value returned in this field indicates if the error is recoverable. 0 = Unrecoverable error. 1 = Recoverable error. This includes all correctable parity/ECC errors and recoverable TCM external errors.
[20]	SideExt	The value returned in this field indicates the source of the error. See Table 4-32 for the encodings.
[19:14]	-	SBZ.
[13:5]	Index ^b	This field returns the index value for the access giving the error.
[4:0]	-	SBZ.

a. This field is only valid for data cache store parity/ECC errors, otherwise it is Unpredictable.

b. This field is only valid for data cache store parity/ECC errors. On the AIFSR, and for TCM accesses, this field SBZ.

Table 4-32 shows the encodings for the SideExt and Side bits.

Table 4-32 SideExt and Side bit encodings

Bit values		Meaning
SideExt	Side	
0	00	Cache/AXIM
0	01	ATCM
0	10	BTCM
0	11	Reserved
1	00	
1	01	AXI peripheral port, including virtual interface
1	10	AHB peripheral port
1	11	Reserved

To access the auxiliary fault status registers, read or write CP15 with:

MRC p15, 0, <Rd>, c5, c1, 0 ; Read Auxiliary Data Fault Status Register

MCR p15, 0, <Rd>, c5, c1, 0 ; Write Auxiliary Data Fault Status Register

MRC p15, 0, <Rd>, c5, c1, 1 ; Read Auxiliary Instruction Fault Status Register
 MCR p15, 0, <Rd>, c5, c1, 1 ; Write Auxiliary Instruction Fault Status Register

c6, Data Fault Address Register

The DFAR characteristics are:

Purpose	Holds the address of the fault when a synchronous abort occurs.
Usage constraints	The DFAR is: • a read/write register • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	The DFAR bits [31:0] contain the address where the synchronous abort occurred.

To access the DFAR read or write CP15 with:

MRC p15, 0, <Rd>, c6, c0, 0 ; Read DFAR
 MCR p15, 0, <Rd>, c6, c0, 0 ; Write DFAR

A write to this register sets the DFAR to the value of the data written. This is useful for a debugger to restore the value of the DFAR.

The processor also updates the DFAR on debug exception entry because of watchpoints. See *Effect of debug exceptions on CP15 registers and DBGWFEAR* on page 12-43 for more information.

c6, Instruction Fault Address Register

The IFAR characteristics are:

Purpose	Holds the address of the instruction that caused a prefetch abort.
Usage constraints	The IFAR is: • a read/write register • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	The IFAR bits [31:0] contain the Instruction Fault address.

To access the IFAR read or write CP15 with:

MRC p15, 0, <Rd>, c6, c0, 2 ; Read IFAR
 MCR p15, 0, <Rd>, c6, c0, 2 ; Write IFAR

A write to this register sets the IFAR to the value of the data written. This is useful for a debugger to restore the value of the IFAR.

4.3.21 c6, MPU memory region programming registers

The MPU memory region programming registers program the MPU regions.

There is one register that specifies which one of the sets of region registers is to be accessed. See *c6, MPU Region Number Register* on page 4-59. Each region has its own registers to specify:

- region base address

- region size and enable
- region access control.

You can implement the processor with 12 or 16 regions, or without an MPU entirely. If you implement the processor without an MPU, then there are no regions and no region programming registers.

Note

- When the MPU is enabled:
 - The MPU determines the access permissions for all accesses to memory, including the TCMs. Therefore, you must ensure that the memory regions in the MPU are programmed to cover the complete TCM address space with the appropriate access permissions. You must define at least one of the regions in the MPU.
 - An access to an undefined area of memory normally generates a background fault.
- For the TCM space the processor uses the access permissions but ignores the region attributes from MPU.

CP15, c9 sets the location of the TCM base address. For more information see *c9, BTCM Region Register* on page 4-63 and *c9, ATCM Region Register* on page 4-64.

c6, MPU Region Base Address Registers

The MPU Region Base Address Register characteristics are:

Purpose	Describes the base address of the region specified by the Memory Region Number Register.
Usage constraints	The MPU Region Base Address Registers are: • 32-bit read/write registers • accessible in Privileged mode only. • The region base address must always align to the region size.
Configurations	Use these registers if the processor is configured with an MPU.
Attributes	See Table 4-33.

Figure 4-34 shows the MPU Region Base Address Registers bit assignments.




Figure 4-34 MPU Region Base Address Registers bit assignments

Table 4-33 shows the MPU Region Base Address Registers bit assignments.

Table 4-33 MPU Region Base Address Registers bit assignments

Bits	Name	Function
[31:5]	Base address	Defines bits [31:5] of the base address of a region
[4:0]	-	SBZ

To access an MPU Region Base Address Register, read or write CP15 with:

MRC p15, 0, <Rd>, c6, c1, 0 ; Read MPU Region Base Address Register
MCR p15, 0, <Rd>, c6, c1, 0 ; Write MPU Region Base Address Register

c6, MPU Region Size and Enable Registers

The MPU Region Size and Enable Register characteristics are:

- Purpose**
- Specifies the size of the region specified by the Memory Region Number Register.
 - Identifies the address ranges that are used for a particular region.
 - Enables or disables the region, and its sub-regions, specified by the Memory Region Number Register.
- Usage constraints** The MPU Region Size and Enable Registers are:
- 32-bit read/write registers
 - accessible in Privileged mode only.
- Configurations** Use these registers if the processor is configured with an MPU.
- Attributes** See Table 4-34.

Figure 4-35 shows the MPU Region Size and Enable Registers bit assignments.




Figure 4-35 MPU Region Size and Enable Registers bit assignments

Table 4-34 shows the MPU Region Size and Enable Registers bit assignments.

Table 4-34 Region Size Register bit assignments

Bits	Name	Function
[31:16]	-	SBZ.
[15:8]	Sub-region disable	Each bit position represents a sub-region, 0-7 ^a . Bit [8] corresponds to sub-region 0 ... Bit [15] corresponds to sub-region 7 The meaning of each bit is: 0 = address range is part of this region 1 = address range is not part of this region.

Table 4-34 Region Size Register bit assignments (continued)

Bits	Name	Function																																						
	-	SBZ.																																						
[5:1]	Region size	Defines the region size: <table><tr><td>b01100 = 8KB</td><td>b10110 = 8MB</td></tr><tr><td>b00000 - b00011=Unpredictable</td><td>b01101 = 16KB</td></tr><tr><td>b00100 = 32 bytes</td><td>b01110 = 32KB</td></tr><tr><td>b00101 = 64 bytes</td><td>b01111 = 64KB</td></tr><tr><td>b00110 = 128 bytes</td><td>b10000 = 128KB</td></tr><tr><td>b00111 = 256 bytes</td><td>b10001 = 256KB</td></tr><tr><td>b01000 = 512 bytes</td><td>b10010 = 512KB</td></tr><tr><td>b01001 = 1KB</td><td>b10011 = 1MB</td></tr><tr><td>b01010 = 2KB</td><td>b10100 = 2MB</td></tr><tr><td>b01011 = 4KB</td><td>b10101 = 4MB</td></tr><tr><td></td><td>b11011 = 16MB</td></tr><tr><td></td><td>b11000 = 32MB</td></tr><tr><td></td><td>b11001 = 64MB</td></tr><tr><td></td><td>b11010 = 128MB</td></tr><tr><td></td><td>b11011 = 256MB</td></tr><tr><td></td><td>b11100 = 512MB</td></tr><tr><td></td><td>b11101 = 1GB</td></tr><tr><td></td><td>b11110 = 2GB</td></tr><tr><td></td><td>b11111 = 4GB.</td></tr></table>	b01100 = 8KB	b10110 = 8MB	b00000 - b00011=Unpredictable	b01101 = 16KB	b00100 = 32 bytes	b01110 = 32KB	b00101 = 64 bytes	b01111 = 64KB	b00110 = 128 bytes	b10000 = 128KB	b00111 = 256 bytes	b10001 = 256KB	b01000 = 512 bytes	b10010 = 512KB	b01001 = 1KB	b10011 = 1MB	b01010 = 2KB	b10100 = 2MB	b01011 = 4KB	b10101 = 4MB		b11011 = 16MB		b11000 = 32MB		b11001 = 64MB		b11010 = 128MB		b11011 = 256MB		b11100 = 512MB		b11101 = 1GB		b11110 = 2GB		b11111 = 4GB.
b01100 = 8KB	b10110 = 8MB																																							
b00000 - b00011=Unpredictable	b01101 = 16KB																																							
b00100 = 32 bytes	b01110 = 32KB																																							
b00101 = 64 bytes	b01111 = 64KB																																							
b00110 = 128 bytes	b10000 = 128KB																																							
b00111 = 256 bytes	b10001 = 256KB																																							
b01000 = 512 bytes	b10010 = 512KB																																							
b01001 = 1KB	b10011 = 1MB																																							
b01010 = 2KB	b10100 = 2MB																																							
b01011 = 4KB	b10101 = 4MB																																							
	b11011 = 16MB																																							
	b11000 = 32MB																																							
	b11001 = 64MB																																							
	b11010 = 128MB																																							
	b11011 = 256MB																																							
	b11100 = 512MB																																							
	b11101 = 1GB																																							
	b11110 = 2GB																																							
	b11111 = 4GB.																																							
[0]	Enable	Enables or disables a memory region: 0 = Memory region disabled. Memory regions are disabled on reset. 1 = Memory region enabled. A memory region must be enabled before it is used.																																						

- a. Sub-region 0 covers the least significant addresses in the region, while sub-region 7 covers the most significant addresses in the region. For more information, see *Subregions* on page 7-3.

To access an MPU Region Size and Enable Register, read or write CP15 with:

MRC p15, 0, <Rd>, c6, c1, 2 ; Read Data MPU Region Size and Enable Register
MCR p15, 0, <Rd>, c6, c1, 2 ; Write Data MPU Region Size and Enable Register

Writing a region size that is outside the range results in Unpredictable behavior.

c6, MPU Region Access Control Registers

The MPU Region Access Control Register characteristics are:

Purpose	Holds the region attributes and access permissions for the region specified by the Memory Region Number Register.
Usage constraints	The MPU Region Access Control Registers are: - read/write registers - accessible in Privileged mode only.
Configurations	Use these registers if the processor is configured with an MPU.
Attributes	See Table 4-35 on page 4-57.

Figure 4-36 shows the MPU Region Access Control Registers bit assignments.


**Figure 4-36 MPU Region Access Control Register bit assignments**

Table 4-35 shows the MPU Region Access Control Registers bit assignments.

Table 4-35 MPU Region Access Control Register bit assignments

Bits	Name	Function
[31:13]	-	SBZ.
[12]	XN	Execute Never. Determines if a region of memory is executable: 0 = all instruction fetches enabled 1 = no instruction fetches enabled.
[11]	-	Reserved.
[10:8]	AP	Access permission. Defines the data access permissions. For more information on AP bit values, see Table 4-38 on page 4-58.
[7:6]	-	SBZ.
[5:3]	TEX	Type extension. Defines the type extension attribute ^a .
[2]	S	Share. Determines if the memory region is Shared or Non-shared: 0 = Non-shared. 1 = Shared. This bit only applies to Normal, not Device or Strongly Ordered memory.
[1]	C	C bit ^a :
[0]	B	B bit ^a :

a. For more information on this region attribute, see Table 4-36.

Table 4-36 shows the encoding for the TEX[2:0], C, and B regions.

Table 4-36 TEX[2:0], C, and B encodings

TEX[2:0]	C	B	Description	Memory Type	Shareable?
000	0	0	Strongly-ordered.	Strongly-ordered	Shareable
000	0	1	Shareable Device.	Device	Shareable
000	1	0	Outer and Inner write-through, no write-allocate.	Normal	S bit ^a
000	1	1	Outer and Inner write-back, no write-allocate.	Normal	S bit ^a
001	0	0	Outer and Inner Non-cacheable.	Normal	S bit ^a
001	0	1	Reserved.	-	-
001	1	0			
001	1	1	Outer and Inner write-back, write-allocate.	Normal	S bit ^a
010	0	0	Non-shareable Device.	Device	Non-shareable
010	0	1	Reserved.	-	-
010	1	X	Reserved.	-	-
011	X	X	Reserved.	-	-
1BB	A	A	Cacheable memory: AA ^b = Inner policy BB ^b = Outer policy	Normal	S bit ^a

- a. Region is Shareable if $S == 1$, and Non-shareable if $S == 0$.
- b. Table 4-37 shows the encoding for these bits.

When $TEX[2] == 1$, the memory region is Cacheable memory, and the rest of the encoding defines the Inner and Outer cache policies:

TEX[1:0] defines the Outer cache policy

C,B defines the Inner cache policy

The same encoding is used for the Outer and Inner cache policies. Table 4-37 shows the encoding.

Table 4-37 Inner and Outer cache policy encoding

Memory attribute encoding	Cache policy
00	Non-cacheable
01	Write-back, write-allocate
10	Write-through, no write-allocate
11	Write-back, no write-allocate

Table 4-38 shows the AP bit values that determine the permissions for Privileged and User data access.

Table 4-38 Access data permission bit encoding

AP bit values	Privileged permissions	User permissions	Description
b000	No access	No access	All accesses generate a permission fault
b001	Read/write	No access	Privileged access only
b010	Read/write	Read-only	Writes in User mode generate permission faults
b011	Read/write	Read/write	Full access
b100	UNP	UNP	Reserved
b101	Read-only	No access	Privileged read-only
b110	Read-only	Read-only	Privileged/User read-only
b111	UNP	UNP	Reserved

To access the MPU Region Access Control Registers read or write CP15 with:

MRC p15, 0, <Rd>, c6, c1, 4 ; Read MPU Region Access Control Register

MCR p15, 0, <Rd>, c6, c1, 4 ; Write MPU Region Access Control Register

To execute instructions in User and Privileged modes:

- the region must have read access as defined by the AP bits
- the XN bit must be set to 0.

Copyright © 2010-2011 ARM. All rights reserved.
Non-Confidential

4-59

- Set and Way:
 - clean
 - invalidate
 - clean and invalidate.
- Address, usually labelled MVA for Modified Virtual Address, but on this processor all addresses are identical:
 - clean
 - invalidate
 - clean and invalidate.

In addition, the maintenance operations use these definitions:

Point of Coherency (PoC)

A point where all instruction and data walks are transparent to any processor in the system.

Point of Unification (PoU)

A point where instruction and data become unified and self-modifying code can function.

Figure 4-38 shows the arrangement of the functions in this group that operate with the MCR and MRC instructions.

———— Note ————

The following operations, as Figure 4-38 shows, are implemented as No Operation, NOP, on the processor:

- Wait For Interrupt, CRm= c0, Opcode_2 = 4
- Invalidate all branch predictors Inner Shareable, CRm= c1, Opcode_2 = 6
- Invalidate Entire Branch Predictor Array, CRm= c5, Opcode_2 = 6
- Invalidate Branch Predictor Array Line using MVA, CRm= c5, Opcode_2 = 7

The *Wait For Interrupt* (WFI) instruction provides the Wait For Interrupt function. For more information see the *ARM Architecture Reference Manual*.




Figure 4-38 Cache operations

In addition to the register c7 cache management functions in this processor, an *Invalidate all data caches* operation is provided as a c15 operation. For convenience, this c15 operation is also described in this section.

Note

- Writing c7 with a combination of CRm and Opcode_2 not listed in Figure 4-38 on page 4-60 results in an Undefined Instruction exception.
 - In this processor, reading from c7 causes an Undefined Instruction exception.
 - All accesses to c7 can only be executed in a Privileged mode of operation, except for the Instruction Synchronization Barrier, Data Synchronization Barrier, and Data Memory Barrier operations. These can be performed in User mode. Attempting to execute a Privileged instruction in User mode results in an Undefined Instruction exception.
 - This processor does not contain an address-based branch predictor array.
-

Invalidate and clean operations

The terms that describe the invalidate, clean, and prefetch operations are defined in the *ARM Architecture Reference Manual*.

You can perform invalidate and clean operations on:

- single cache lines
- entire caches.

Set and Way format

Figure 4-39 shows the Set and Way bit assignments.




Figure 4-39 c7 Set and Way bit assignments

Table 4-40 shows the Set and Way bit assignments.

Table 4-40 c7 Set and Way bit assignments

Bits	Name	Function
[31:30]	Way	Indicates the cache way to invalidate or clean.
[29:S+5]	-	SBZ.
[S+4:5]	Set	Indicates the cache set to invalidate or clean. Because the cache sizes are configurable, the width of the Set field is unique to the cache size. See Table 4-41 on page 4-62.
[4:0]	-	SBZ.

Table 4-41 shows the cache sizes and the resultant bit range for Set.

Table 4-41 Widths of the set field for L1 cache sizes

Size	Set
4KB	[9:5]
8KB	[10:5]
16KB	[11:5]
32KB	[12:5]
64KB	[13:5]

See *c0, Cache Type Register* on page 4-15 for more information on cache sizes.

Address format

Figure 4-40 shows the invalidate and clean operations bit assignments.




Figure 4-40 Invalidate and clean operations bit assignments

Table 4-42 shows the invalidate and clean operations bit assignments.

Table 4-42 Invalidate and clean operations bit assignments

Bits	Name	Function
[31:5]	Address	Specifies the address to invalidate or clean
[4:0]	Reserved	SBZ

Data Synchronization Barrier operation

The purpose of the Data Synchronization Barrier operation is to ensure that all outstanding explicit memory transactions complete before any following instructions begin. This ensures that data in memory is up to date before the processor executes any more instructions.

The Data Synchronization Barrier Register is:

- a write-only operation
- accessible in both User and Privileged mode.

To access the Data Synchronization Barrier operation, write CP15 with:

`MCR p15, 0, <Rd>, c7, c10, 4 ; Data Synchronization Barrier operation`

For more information about memory barriers, see the *ARM Architecture Reference Manual*.

Data Memory Barrier operation

The purpose of the Data Memory Barrier operation is to ensure that all outstanding explicit memory transactions complete before any following explicit memory transactions begin. This ensures that data in memory is up to date before any memory transaction that depends on it.

The Data Memory Barrier operation is:

- write-only
- accessible in User and Privileged mode.

To access the Data Memory Barrier operation write CP15 with:

MCR p15, 0, <Rd>, c7, c10, 5 ; Data Memory Barrier Operation

For more information about memory barriers, see the *ARM Architecture Reference Manual*.

4.3.23 c9, BTCM Region Register

The BTCM Region Register characteristics are:

- Purpose**
- Holds the base address and size of the BTCM.
 - Determines if the BTCM is enabled.
- Usage constraints** The BTCM Region Register is:
- a read/write register
 - accessible in Privileged mode only.
- Configurations** Available in all processor configurations.
- Attributes** See Table 4-43.

Figure 4-41 shows the BTCM Region Register bit assignments.




Figure 4-41 BTCM Region Register bit assignments

Table 4-43 shows the BTCM Region Register bit assignments.

Table 4-43 BTCM Region Register bit assignments

Bits	Name	Function
[31:12]	Base address	Base address. Defines the base address of the BTCM. The base address must be aligned to the size of the BTCM. Any bits in the range $[(\log_2(\text{RAMSize})-1):12]$ are ignored. At reset, if LOCZRAMAm is set to: 0 =The initial base address is 0x0. 1 =The initial base address is implementation-defined. See <i>Configurable options</i> on page 1-6.
[11:7]	-	UNP on reads, SBZ on writes.

Bits	Name	Function															
[6:2]	Size	Size. Indicates the size of the BTCM on reads. On writes this field is ignored. See <i>About the TCMs</i> on page 8-13. <table> <tr> <td>b00000 = 0KB, no TCM</td><td>b00110 = 32KB</td><td>b01010 = 512kB</td></tr> <tr> <td>b00011 = 4KB</td><td>b00111 = 64KB</td><td>b01011 = 1MB</td></tr> <tr> <td>b00100 = 8KB</td><td>b01000 = 128KB</td><td>b01100 = 2MB</td></tr> <tr> <td>b00101 = 16KB</td><td>b01001 = 256KB</td><td>b01101 = 4MB</td></tr> <tr> <td></td><td></td><td>b01110 = 8MB</td></tr> </table>	b00000 = 0KB, no TCM	b00110 = 32KB	b01010 = 512kB	b00011 = 4KB	b00111 = 64KB	b01011 = 1MB	b00100 = 8KB	b01000 = 128KB	b01100 = 2MB	b00101 = 16KB	b01001 = 256KB	b01101 = 4MB			b01110 = 8MB
b00000 = 0KB, no TCM	b00110 = 32KB	b01010 = 512kB															
b00011 = 4KB	b00111 = 64KB	b01011 = 1MB															
b00100 = 8KB	b01000 = 128KB	b01100 = 2MB															
b00101 = 16KB	b01001 = 256KB	b01101 = 4MB															
		b01110 = 8MB															
[1]	-	SBZ.															
[0]	Enable	Enables or disables the BTCM. 0 = Disabled 1 = Enabled. The reset value of this field is determined by the INITRAMBm input pin. This bit is RAZ if the processor has been implemented or integrated without a BTCM.															

```
MRC p15, 0, <Rd>, c9, c1, 0 ; Read BTCM Region Register
MCR p15, 0, <Rd>, c9, c1, 0 ; Write BTCM Region Register
```

Figure 4-42 shows the ATCM Region Register bit assignments.



Table 4-44 shows the ATCM Region Register bit assignments.

Table 4-44 ATCM Region Register bit assignments

Bits	Name	Function
[31:12]	Base address	Base address. Defines the base address of the ATCM. The base address must be aligned to the size of the ATCM. Any bits in the range $[(\log_2(\text{RAMSize})-1):12]$ are ignored. At reset, if LOCZRAMAm is set to: 0 = The initial base address is implementation-defined. See <i>Configurable options</i> on page 1-6 1 = The initial base address is 0x0.
[11:7]	-	UNP on reads, SBZ on writes.
[6:2]	Size	Size. Indicates the size of the ATCM on reads. On writes this field is ignored. See <i>About the TCMs</i> on page 8-13. b00000 = 0KB, no TCM b00110 = 32KB b01010 = 512kB b00011 = 4KB b00111 = 64KB b01011 = 1MB b00100 = 8KB b01000 = 128KB b01100 = 2MB b00101 = 16KB b01001 = 256KB b01101 = 4MB b01110 = 8MB.
[1]	-	SBZ.
[0]	Enable	Enables or disables the ATCM. 0 = Disabled. 1 = Enabled. The reset value of this field is determined by the INITRAMAm input pin. This bit is RAZ if the processor has been implemented or integrated without an ATCM.

To access the ATCM Region Register, read or write CP15 with:

MRC p15, 0, <Rd>, c9, c1, 1 ; Read ATCM Region Register
MCR p15, 0, <Rd>, c9, c1, 1 ; Write ATCM Region Register

4.3.25 c9, TCM Selection Register

The TCM Selection Register determines the TCM region register that the processor writes to. The processor only supports one TCM region for each TCM interface, and the TCM Selection Register Reads-As-Zero and ignores writes. It is only accessible in Privileged mode.

4.3.26 c11, Slave Port Control Register

The Slave Port Control Register characteristics are:

Purpose	- Enables or disables TCM access to the AXI slave port in Privileged or User mode. - Enables access to the cache RAMs through the AXI slave port. See <i>c1, Auxiliary Control Register</i> on page 4-41.
Usage constraints	The Slave Port Control Register is: - a read/write register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-45 on page 4-66.

Figure 4-43 on page 4-66 shows the Slave Port Control Register bit assignments.




Figure 4-43 Slave Port Control Register bit assignments

Table 4-45 shows the Slave Port Control Register bit assignments

Table 4-45 Slave Port Control Register bit assignments

Bits	Name	Function
[31:2]	-	RAZ/UNP.
[1]	Privileged access	Defines level of access for TCM accesses: 0 = Non-privileged and privileged access, reset value 1 = Privileged access only.
[0]	AXI slave enable	Enables or disables the AXI slave port for TCM accesses: 0 = Enables AXI slave port, reset value 1 = Disables AXI slave port.

To access the Slave Port Control Register, read or write CP15 with:

MRC p15, 0, <Rd>, c11, c0, 0 ; Read Slave Port Control Register

MCR p15, 0, <Rd>, c11, c0, 0 ; Write Slave Port Control Register

4.3.27 c13, FCSE PID Register

This processor does not support *Fast Context Switch Extension* (FCSE).

The FCSE *Process Identifier* (PID) Register is accessible in Privileged mode only. This register reads as zero and ignores writes.

4.3.28 c13, Context ID Register

The CONTEXTIDR characteristics are:

Purpose	- Holds a process <i>IDentification</i> (ID) value for the running process. - The <i>Embedded Trace Macrocell</i> (ETM) and the debug logic use this register. The ETM can broadcast its value to indicate the process that is running. You must program each process with a unique number. - Enables process dependent breakpoints and instructions.
Usage constraints	The CONTEXTIDR is: - a read/write register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	The CONTEXTIDR, bits [31:0] contain the process ID number.

To use the CONTEXTIDR, read or write CP15 with:

MRC p15, 0, <Rd>, c13, c0, 1 ; Read CONTEXTIDR

MCR p15, 0, <Rd>, c13, c0, 1 ; Write CONTEXTIDR

4.3.29 c13, Thread and Process ID Registers

The Thread and Process ID Registers provide locations to store the IDs of software threads and processes for *Operating System* (OS) management purposes.

The Thread and Process ID Registers are:

- three read/write registers:
 - User read/write Thread and Process ID Register
 - User read-only Thread and Process ID Register
 - Privileged-only Thread and Process ID Register.
- each accessible in different modes:
 - The User read/write register can be read and written in User and Privileged modes.
 - The User read-only register can only be read in User mode, but can be read and written in Privileged modes.
 - The Privileged-only register can be read and written in Privileged modes only.

To access the Thread and Process ID registers, read or write CP15 with:

```
MRC p15, 0, <Rd>, c13, c0, 2 ; Read User read/write Thread and Proc. ID Register
MCR p15, 0, <Rd>, c13, c0, 2 ; Write User read/write Thread and Proc. ID Register
MRC p15, 0, <Rd>, c13, c0, 3 ; Read User Read Only Thread and Proc. ID Register
MCR p15, 0, <Rd>, c13, c0, 3 ; Write User Read Only Thread and Proc. ID Register
MRC p15, 0, <Rd>, c13, c0, 4 ; Read Privileged Only Thread and Proc. ID Register
MCR p15, 0, <Rd>, c13, c0, 4 ; Write Privileged Only Thread and Proc. ID Register
```

Reading or writing the Thread and Process ID registers has no effect on processor state or operation. These registers provide OS support, and the OS must manage them.

You must clear the contents of all Thread and Process ID registers on process switches to prevent data leaking from one process to another. This is important to ensure the security of data. The reset value of these registers is 0.

4.3.30 Validation Registers

The processor implements a set of validation registers. This section describes:

- *c15, nVAL IRQ Enable Set Register*
- *c15, nVAL FIQ Enable Set Register* on page 4-69
- *c15, nVAL Reset Enable Set Register* on page 4-70
- *c15, VAL Debug Request Enable Set Register* on page 4-71
- *c15, VAL IRQ Enable Clear Register* on page 4-72
- *c15, nVAL FIQ Enable Clear Register* on page 4-73
- *c15, nVAL Reset Enable Clear Register* on page 4-74
- *c15, VAL Debug Request Enable Clear Register* on page 4-75
- *c15, Cache Size Override Register* on page 4-76.

c15, nVAL IRQ Enable Set Register

The nVAL IRQ Enable Set Register characteristics are:

Purpose	Enables any of the PMXEVCNTR Registers, PMXEVCNTR0-PMXEVCNTR2, and CCNT, to generate an interrupt request on overflow. If enabled, the interrupt request is signaled by nVALIRQm being asserted LOW.
Usage constraints	The nVAL IRQ Enable Set Register is: • A read/write register. • Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see <i>c9, User Enable Register</i> on page 6-16.
Configurations	Available in all processor configurations.
Attributes	See Table 4-46.

Figure 4-44 shows the nVAL IRQ Enable Set Register bit assignments.

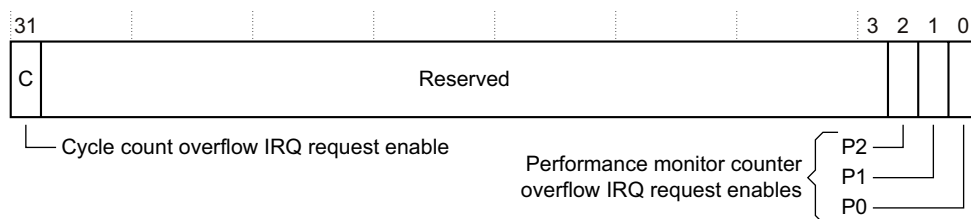


Figure 4-44 nVAL IRQ Enable Set Register bit assignments

Table 4-46 shows the nVAL IRQ Enable Set Register bit assignments.

Table 4-46 nVAL IRQ Enable Set Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow IRQ request
[30: 3]	Reserved	UNP or SBZP

Table 4-46 nVAL IRQ Enable Set Register bit assignments (continued)

Bits	Name	Function
[2]	P2	PMXEVNTR2 overflow IRQ request
[1]	P1	PMXEVNTR1 overflow IRQ request
[0]	P0	PMXEVNTR0 overflow IRQ request

To access the nVAL IRQ Enable Set Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 0 ; Read nVAL IRQ Enable Set Register

MCR p15, 0, <Rd>, c15, c1, 0 ; Write nVAL IRQ Enable Set Register

On reads, this register returns the current setting. On writes, interrupt requests can be enabled by writing a 1 to the appropriate bits. If an interrupt request has been enabled it is disabled by writing to the nVAL IRQ Enable Clear Register, see *c15, VAL IRQ Enable Clear Register* on page 4-72.

If one or more of the IRQ request fields (P2, P1, P0, and C) is enabled, and the corresponding counter overflows, then an IRQ request is indicated by **nVALIRQm** being asserted LOW. This signal might be passed to a system interrupt controller.

c15, nVAL FIQ Enable Set Register

The nVAL FIQ Enable Set Register are:

- Purpose** Enables any of the PMXEVNTR Registers, PMXEVNTR0-PMXEVNTR2, and CCNT, to generate an fast interrupt request on overflow. If enabled, the interrupt request is signaled by **nVALFIQm** being asserted LOW.
- Usage constraints** The nVAL FIQ Enable Set Register is:
- A read/write register.
 - Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see *c9, User Enable Register* on page 6-16.
- Configurations** Available in all processor configurations.
- Attributes** See Table 4-47 on page 4-70.

Figure 4-45 shows the nVAL FIQ Enable Set Register bit assignments.

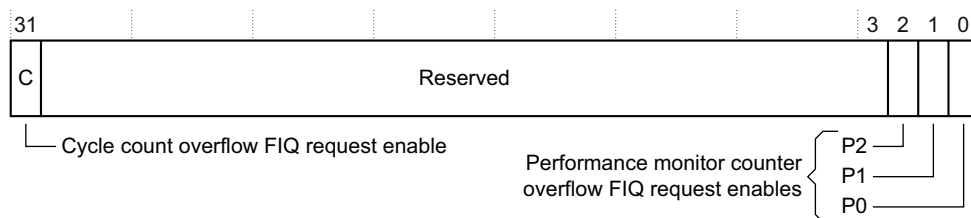


Figure 4-45 nVAL FIQ Enable Set Register bit assignments

Table 4-47 shows the nVAL FIQ Enable Set Register bit assignments

Table 4-47 nVAL FIQ Enable Set Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow FIQ request
[30:3]	Reserved	UNP or SBZP
[2]	P2	PMXEVNTR2 overflow FIQ request
[1]	P1	PMXEVNTR1 overflow FIQ request
[0]	P0	PMXEVNTR0 overflow FIQ request

To access the FIQ Enable Set Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 1 ; Read FIQ Enable Set Register

MCR p15, 0, <Rd>, c15, c1, 1 ; Write FIQ Enable Set Register

On reads, this register returns the current setting. On writes, interrupt requests can be enabled by writing a 1 to the appropriate bits. If an interrupt request has been enabled it is disabled by writing to the FIQ Enable Clear Register, see *c15, nVAL FIQ Enable Clear Register* on page 4-73.

If one or more of the FIQ request fields (P2, P1, P0, and C) is enabled, and the corresponding counter overflows, then an FIQ request is indicated by **nVALFIQm** being asserted LOW. This signal can be passed to a system interrupt controller.

c15, nVAL Reset Enable Set Register

The nVAL Reset Enable Set Register is:

- A read/write register.
- Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see *c9, User Enable Register* on page 6-16.

The nVAL Reset Enable Set Register characteristics are:

Purpose	Enables any of the PMXEVNTR Registers, PMXEVNTR0-PMXEVNTR2, and CCNT, to generate a reset request on overflow. If enabled, the reset request is signaled by nVALRESETm being asserted LOW.
Usage constraints	The nVAL Reset Enable Set Register is: • A read/write register. • Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see <i>c9, User Enable Register</i> on page 6-16.
Configurations	Available in all processor configurations.
Attributes	See Table 4-48 on page 4-71.

Figure 4-46 on page 4-71 shows the nVAL Reset Enable Set Register bit assignments.

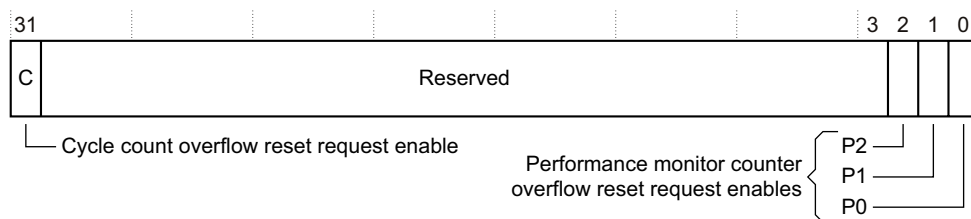


Figure 4-46 nVAL Reset Enable Set Register bit assignments

Table 4-48 shows the nVAL Reset Enable Set Register bit assignments.

Table 4-48 nVAL Reset Enable Set Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow reset request
[30:3]	-	UNP or SBZP
[2]	P2	PMXEVCNTR2 overflow reset request
[1]	P1	PMXEVCNTR1 overflow reset request
[0]	P0	PMXEVCNTR0 overflow reset request

To access the nVAL Reset Enable Set Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 2 ; Read nVAL Reset Enable Set Register
MCR p15, 0, <Rd>, c15, c1, 2 ; Write nVAL Reset Enable Set Register

On reads, this register returns the current setting. On writes, reset requests can be enabled by writing a 1 to the appropriate bits. If a reset request has been enabled, it is disabled by writing to the nVAL Reset Enable Clear Register. See *c15, nVAL Reset Enable Clear Register* on page 4-74.

If one or more of the reset request fields (P2, P1, P0, and C) is enabled, and the corresponding counter overflows, then a reset request is indicated by **nVALRESETm** being asserted LOW. This signal can be passed to a system reset controller.

c15, VAL Debug Request Enable Set Register

The VAL Debug Request Enable Set Register characteristics are:

Purpose	Enables any of the PMXEVCNTR Registers, PMXEVCNTR0-PMXEVCNTR2, and CCNT, to generate a debug request on overflow. If enabled, the debug request is signaled by VALEDBGRQm being asserted HIGH.
Usage constraints	The VAL Debug Request Enable Set Register is: • A read/write register. • Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see <i>c9, User Enable Register</i> on page 6-16.
Configurations	Available in all processor configurations.
Attributes	See Table 4-49 on page 4-72.

Figure 4-47 on page 4-72 shows the VAL Debug Request Enable Set Register bit assignments.

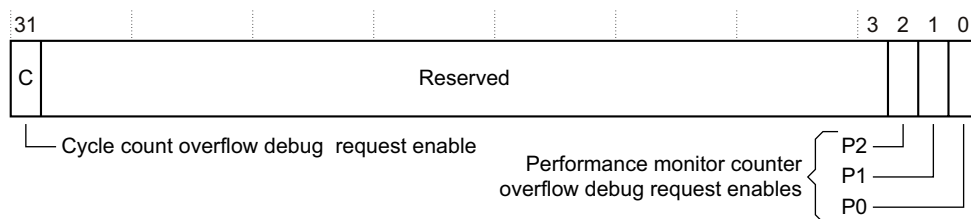


Figure 4-47 VAL Debug Request Enable Set Register bit assignments

Table 4-49 shows the VAL Debug Request Enable Set Register bit assignments.

Table 4-49 VAL Debug Request Enable Set Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow debug request
[30:3]	-	UNP or SBZP
[2]	P2	PMXEVNTR2 overflow debug request
[1]	P1	PMXEVNTR1 overflow debug request
[0]	P0	PMXEVNTR0 overflow debug request

To access the VAL Debug Request Enable Set Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 3 ; Read VAL Debug Request Enable Set Register
MCR p15, 0, <Rd>, c15, c1, 3 ; Write VAL Debug Request Enable Set Register

On reads, this register returns the current setting. On writes, debug requests can be enabled by writing a 1 to the appropriate bits. If a debug request has been enabled, it is disabled by writing to the VAL Debug Request Enable Clear Register. See *c15, VAL Debug Request Enable Clear Register* on page 4-75.

If one or more of the reset request fields (P2, P1, P0, and C) is enabled, and the corresponding counter overflows, then a debug reset request is indicated by **VALEDBGROm** being asserted HIGH. This signal can be passed to an external debugger.

c15, VAL IRQ Enable Clear Register

The VAL IRQ Enable Clear Register characteristics are:

Purpose	Disables overflow IRQ requests from any of the PMXEVNTR Registers, PMXEVNTR0-PMXEVNTR2, and CCNT, for which they have been enabled.
Usage constraints	The VAL IRQ Enable Clear Register is: • A read/write register. • Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see <i>c9, User Enable Register</i> on page 6-16.
Configurations	Available in all processor configurations.
Attributes	See Table 4-50 on page 4-73.

Figure 4-48 on page 4-73 shows the VAL IRQ Enable Clear Register bit assignments.

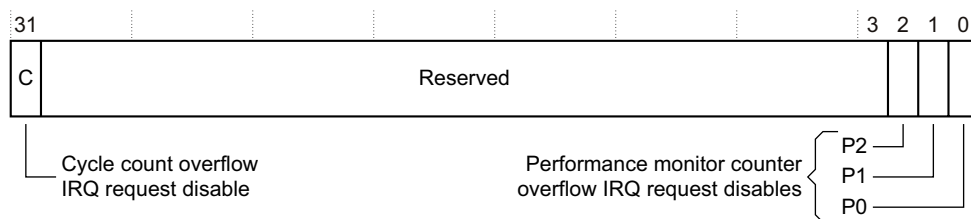


Figure 4-48 VAL IRQ Enable Clear Register bit assignments

Table 4-50 shows the VAL IRQ Enable Clear Register bit assignments.

Table 4-50 VAL IRQ Enable Clear Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow IRQ request
[30:3]	-	UNP or SBZP
[2]	P2	PMXEVCNTR2 overflow IRQ request
[1]	P1	PMXEVCNTR1 overflow IRQ request
[0]	P0	PMXEVCNTR0 overflow IRQ request

To access the VAL IRQ Enable Clear Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 4 ; Read VAL IRQ Enable Clear Register
MCR p15, 0, <Rd>, c15, c1, 4 ; Write VAL IRQ Enable Clear Register

On reads, this register returns the current setting. On writes, overflow interrupt requests that are currently enabled can be disabled by writing a 1 to the appropriate bits.

For more information of how to enable IRQ requests on counter overflows, and how the requests are signaled, see *c15, nVAL IRQ Enable Set Register* on page 4-68.

c15, nVAL FIQ Enable Clear Register

The nVAL FIQ Enable Clear Register characteristics are:

Purpose	Disables overflow FIQ requests from any of the PMXEVCNTR Registers, PMXEVCNTR0-PMXEVCNTR2, and CCNT, that are enabled.
Usage constraints	The nVAL FIQ Enable Clear Register is: • A read/write register. • Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see <i>c9, User Enable Register</i> on page 6-16.
Configurations	Available in all processor configurations.
Attributes	See Table 4-51 on page 4-74.

Figure 4-49 on page 4-74 shows the nVAL FIQ Enable Clear Register bit assignments.




Figure 4-49 nVAL FIQ Enable Clear Register bit assignments

Table 4-51 shows the nVAL FIQ Enable Clear Register bit assignments

Table 4-51 nVAL FIQ Enable Clear Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow FIQ request
[30:3]	-	UNP or SBZP
[2]	P2	PMXEVNTR2 overflow FIQ request
[1]	P1	PMXEVNTR1 overflow FIQ request
[0]	P0	PMXEVNTR0 overflow FIQ request

To access the FIQ Enable Clear Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 5 ; Read FIQ Enable Clear Register
MCR p15, 0, <Rd>, c15, c1, 5 ; Write FIQ Enable Clear Register

On reads, this register returns the current setting. On writes, overflow interrupt requests that are enabled can be disabled by writing a 1 to the appropriate bits.

For information on how to enable FIQ requests on counter overflows, and how the requests are signaled, see *c15, nVAL FIQ Enable Set Register* on page 4-69.

c15, nVAL Reset Enable Clear Register

The nVAL Reset Enable Clear Register characteristics are:

Purpose	Disables overflow reset requests from any of the PMXEVNTR Registers, PMXEVNTR0-PMXEVNTR2, and CCNT, that are enabled.
Usage constraints	The nVAL Reset Enable Clear Register is: • A read/write register. • Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see <i>c9, User Enable Register</i> on page 6-16.
Configurations	Available in all processor configurations.
Attributes	See Table 4-52 on page 4-75.

Figure 4-50 on page 4-75 shows the nVAL Reset Enable Clear Register bit assignments.

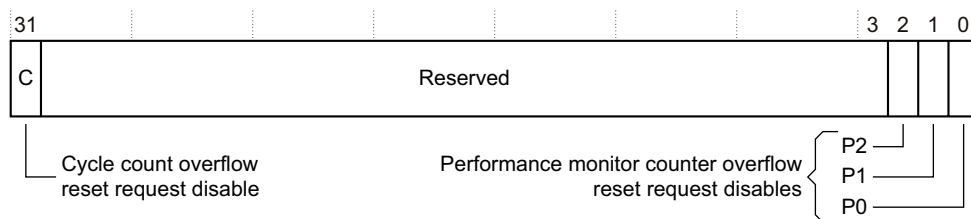


Figure 4-50 nVAL Reset Enable Clear Register bit assignments

Table 4-52 shows the nVAL Reset Enable Clear Register bit assignments.

Table 4-52 nVAL Reset Enable Clear Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow reset request
[30:3]	-	UNP or SBZP
[2]	P2	PMXEVCNTR2 overflow reset request
[1]	P1	PMXEVCNTR1 overflow reset request
[0]	P0	PMXEVCNTR0 overflow reset request

To access the nVAL Reset Enable Clear Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 6 ; Read nVAL Reset Enable Clear Register
MCR p15, 0, <Rd>, c15, c1, 6 ; Write nVAL Reset Enable Clear Register

On reads, this register returns the current setting. On writes, overflow reset requests that are enabled can be disabled by writing a 1 to the appropriate bits.

For more information of how to enable reset requests on counter overflows, and how the requests are signaled, see *c15, nVAL Reset Enable Set Register* on page 4-70.

c15, VAL Debug Request Enable Clear Register

The VAL Debug Request Enable Clear Register characteristics are:

Purpose	Disables overflow debug requests from any of the PMXEVCNTR Registers, PMXEVCNTR0-PMXEVCNTR2, and CCNT, that are enabled.
Usage constraints	The VAL Debug Request Enable Clear Register is: • A read/write register. • Always accessible in Privileged mode. The PMUSERENR Register determines access in User mode, see <i>c9, User Enable Register</i> on page 6-16.
Configurations	Available in all processor configurations.
Attributes	See Table 4-53 on page 4-76.

Figure 4-51 on page 4-76 shows the VAL Debug Request Enable Clear Register bit assignments.

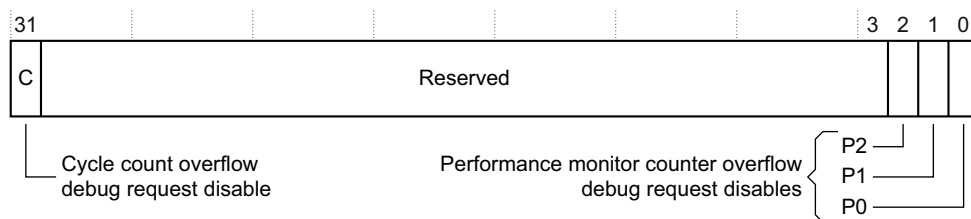


Figure 4-51 VAL Debug Request Enable Clear Register bit assignments

Table 4-53 shows the VAL Debug Request Enable Clear Register bit assignments.

Table 4-53 VAL Debug Request Enable Clear Register bit assignments

Bits	Name	Function
[31]	C	CCNT overflow debug request
[30:3]	-	UNP or SBZP
[2]	P2	PMXEVCNTR2 overflow debug request
[1]	P1	PMXEVCNTR1 overflow debug request
[0]	P0	PMXEVCNTR0 overflow debug request

To access the VAL Debug Request Enable Clear Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c1, 7 ; Read VAL Debug Request Enable Clear Register
MCR p15, 0, <Rd>, c15, c1, 7 ; Write VAL Debug Request Enable Clear Register

On reads, this register returns the current setting. On writes, overflow debug requests that are enabled can be disabled by writing a 1 to the appropriate bits.

For more information of how to enable debug requests on counter overflows, and how the requests are signaled, see *c15, VAL Debug Request Enable Set Register* on page 4-71.

c15, Cache Size Override Register

The Cache Size Override Register characteristics are:

Purpose	Overwrites the caches size fields in the main register. This enables you to choose a smaller instruction and data cache size than is implemented.
Usage constraints	The Cache Size Override Register is: - a write-only register - only accessible in Privileged mode.
Configurations	Available in all processor configurations.
Attributes	See Table 4-54 on page 4-77.

Figure 4-52 shows the Cache Size Override Register bit assignments.




Figure 4-52 Cache Size Override Register bit assignments

Table 4-54 shows the Cache Size Override Register bit assignments.

Table 4-54 Cache Size Override Register bit assignments

Bits	Name	Function
[31:8]	-	SBZ.
[7:4]	Dcache	Defines the data cache size. See Table 4-55.
[3:0]	Icache	Defines the instruction cache size. See Table 4-55.

Table 4-55 shows the encodings for the instruction and data cache sizes.

Table 4-55 Instruction and data cache size encodings

Encoding	Cache size
b0000	4kB
b0001	8kB
b0011	16kB
b0111	32kB
b1111	64kB

To access the Cache Size Override Register, write CP15 with:

MCR p15, 0, <Rd>, c15, c14, 0 ; Write Cache Size Override Register

————— **Note** —————

The VAL Cache Size Override Register can only be used to select cache sizes for which the appropriate RAM has been integrated. Larger cache sizes require deeper data and tag RAMs, and smaller cache sizes require wider tag RAMs. Therefore, it is unlikely that you can change the cache size using this register except using a simulation model of the cache RAMs. ARM recommends that you read the CCSIDR to check the actual cache sizes after writing to the Cache Size Override Register.

4.3.31 Correctable Fault Location Register

The CFLR characteristics are:

Purpose	Indicates the location of the last correctable error that occurred during cache or TCM operations.
Usage constraints	The CFLR is: • a read/write register • accessible in Privileged mode only. • not updated on: — speculative accesses, for example, an instruction fetch for an instruction that is not executed because of a previous branch. — a TCM external error or external retry request. • updated on: — parity or ECC errors in the instruction cache — single-bit ECC errors in the data cache

- parity or multi-bit errors in the data cache when write-through behavior is forced
- single-bit TCM ECC errors.
- updated by the processor, regardless of whether an abort is taken or an access is retried in response to the error.

Configurations Available in all processor configurations.

Attributes See Table 4-56.

Every correctable error that causes a CFLR update also has an associated event. See Table 6-1 on page 6-2 for the events that are related to CFLR updates. If two correctable errors occur simultaneously, for example an AXI slave error and an LSU or PFU error, the LSU or PFU write takes priority. If multiple errors occur, the value in the CFLR reflects the location of the latest event.

The same register is updated by all correctable errors. You can read bits [25:24] to determine whether the error was from a cache or TCM access.

Figure 4-53 shows the CFLR bit assignments, when it indicates a correctable cache error.




Figure 4-53 Correctable Fault Location Register - cache, bit assignments

Table 4-56 shows the CFLR bit assignments, when it indicates a correctable cache error.

Table 4-56 Correctable Fault Location Register - cache, bit assignments

Bits	Name	Function
[31:30]	-	RAZ.
[29:26]	Way	Indicates the Way of the error.
[25:24]	Side	Indicates the source of the error. For cache errors, this value is always 0b00.
[23:14]	-	RAZ.
[13:5]	Index	Indicates the index of the location where the error occurred.
[4:2]	-	RAZ.
[1:0]	Type	Indicates the type of access that caused the error: 0b00 = Instruction cache 0b01 = Data cache 0b11 = ACP.

Figure 4-54 on page 4-79 shows the CFLR bit assignments, when it indicates a correctable TCM error.


**Figure 4-54 Correctable Fault Location Register - TCM, bit assignments**

Table 4-57 shows the CFLR bit assignments, when it indicates a correctable TCM error.

Table 4-57 Correctable Fault Location Register - TCM, bit assignments

Bits	Name	Function
[31:26]	-	RAZ.
[25:24]	Side	Indicates the source of the error: 0b01 = ATCM 0b10 = BTCM.
[23]	-	RAZ.
[22:3]	Address	Indicates the address in the TCM where the error occurred.
[2]	-	RAZ.
[1:0]	Type	Indicates the type of access that caused the error: 0b00 = Instruction 0b01 = Data 0b10 = AXI slave.

To access the Correctable Fault Location Register, read or write CP15 with:

MRC p15, 0, <Rd>, c15, c3, 0 ; Read CFLR
MCR p15, 0, <Rd>, c15, c3, 0 ; Write CFLR

4.3.32 Build Options Registers

Note

In a twin-CPU system, some options can be configured independently for each CPU. For these options, the Options Register reflects the options for the CPU containing the register. Other options are shared, and the options register contains the same value for both CPUs.

c15, Build Options 1 Register

The Build Options 1 Register characteristics are:

Purpose	Reflects the build configuration options used to build the processor.
Usage constraints	The Build Options 1 Register is: - a read-only register - accessible in Privileged mode only - pin-configuration options are shown in a separate register, see <i>Pin Options Register</i> on page 4-83.
Configurations	Available in all processor configurations.

Attributes See Table 4-58.

Figure 4-55 shows the Build Options 1 Register bit assignments.




Figure 4-55 Build Options 1 Register bit assignments

Table 4-58 shows the Build Options 1 Register bit assignments.

Table 4-58 Build Options 1 Register bit assignments

Bits	Name	Function
[31:12]	TCM_HI_INIT_ADDR	Default high address for the TCM.
[11:2]	-	SBZ.
[1]	FLOAT_PRECISION	Indicates whether double-precision floating point is implemented: 0 = Double-precision FP implemented, or no FPU implemented 1 = No double-precision FP implemented.
[0]	PP_BUS_ECC	Indicates whether the peripheral ports were built with bus-ECC: 0 = bus-ECC not included on peripheral ports 1 = bus-ECC included on peripheral ports.

To access the Build Options 1 Register, read CP15 with:

MRC p15, 0, <Rd>, c15, c2, 0 ; Read Build Options 1 Register

c15, Build Options 2 Register

The Build Options 2 Register characteristics are:

- Purpose** Reflects the build configuration options used to build the processor.
- Usage constraints** The Build Options 2 Register is:
- a read-only register
 - accessible in Privileged mode only.
 - pin-configuration options are shown in a separate register, see *Pin Options Register* on page 4-83.
- Configurations** Available in all processor configurations.
- Attributes** See Table 4-58.

Table 4-59 on page 4-81 shows the bit arrangement for the Build Options 2 Register.


**Figure 4-56 Build Options 2 Register bit assignments**

Table 4-59 shows how the bit values correspond with the Build Options 2 Register.

Table 4-59 Build Options 2 Register bit assignments

Bits	Name	Function
[31]	NUM_CPU	Indicates the number of CPUs: 0 = single CPU 1 = twin CPU.
[30]	LOCK_STEP	Indicates whether the CPU has redundant logic running in lock step for checking purposes: 0 = no redundant logic 1 = redundant logic included.
[29]	NO_ICACHE	Indicates whether the CPU contains instruction cache: 0 = CPU contains instruction cache 1 = CPU does not contain instruction cache.
[28]	NO_DCACHE	Indicates whether the CPU contains data cache: 0 = CPU contains data cache 1 = CPU does not contain data cache.
[27:26]	ATCM_ES	Indicates whether an error scheme is implemented on the ATCM interface: 00 = no error scheme 10 = 32-bit error detection and correction 11 = 64-bit error detection and correction.

Table 4-59 Build Options 2 Register bit assignments (continued)

Bits	Name	Function
[25:24]	BTCM_ES	Indicates whether an error scheme is implemented on the BTCM interface(s): 00 = no error scheme 10 = 32-bit error detection and correction 11 = 64-bit error detection and correction.
[23]	NO_IE	Indicates whether the processor supports big-endian instructions: 0 = processor supports big-endian instructions 1 = processor does not support big-endian instructions.
[22]	NO_FPU	Indicates whether the CPU contains a floating point unit: 0 = CPU contains a floating point unit 1 = CPU does not contain a floating point unit.
[21:20]	MPU_REGIONS	Indicates the number of regions in the included CPU MPU: 0b00 = no regions, the MPU has not been included 0b10 = MPU included, with 12 regions 0b11 = MPU included, with 16 regions.
[19:17]	BREAK_POINTS	Indicates the number of break points implemented in each CPU in the processor, minus 1.
[16:14]	WATCH_POINTS	Indicates the number of watch points implemented in each CPU in the processor, minus 1.
[13]	NO_A_TCM_INF	Indicates whether the CPUs contain ATCM ports 0 = CPUs contain ATCM ports 1 = CPUs do not contain ATCM ports.
[12]	NO_B0_TCM_INF	Indicates whether the CPUs contain B0TCM ports: 0 = CPUs contain B0TCM ports 1 = CPUs do not contain B0TCM ports.
[11]	NO_B1_TCM_INF	Indicates whether the CPUs contain B1TCM ports: 0 = CPUs contain B1TCM ports 1 = CPUs do not contain B1TCM ports.
[10]	TCMBUSPARITY	Indicates whether the processor contains TCM address bus parity logic: 0 = processor does not contain TCM address bus parity logic 1 = processor contains TCM address bus parity logic.
[9]	NO_SLAVE	Indicates whether the CPU contains an AXI slave port: 0 = CPU contains an AXI slave port 1 = CPU does not contain an AXI slave port.
[8:7]	ICACHE_ES	Indicates whether an error scheme is implemented for the instruction cache: 0b00 = no error scheme 0b01 = 8-bit parity error detection 0b11 = 64-bit error detection and correction. If the CPU does not contain an I-Cache, these bits are set to 0b00.
[6:5]	DCACHE_ES	Indicates whether an error scheme is implemented for the data cache: 0b00 = no error scheme 0b01 = 8-bit parity error detection 0b10 = 32-bit error detection and correction. If the CPU does not contain a D-Cache, these bits are set to 0b00.

Table 4-59 Build Options 2 Register bit assignments (continued)

Bits	Name	Function
[4]	NO_HARD_ERROR_CACHE	Indicates whether the processor contains cache for corrected TCM errors: 0 = processor contains TCM error cache 1 = processor does not contain TCM error cache.
[3]	AXI_BUS_ECC	Indicates whether the processor contains AXI bus ECC logic. 0 = processor does not contain AXI bus ECC logic 1 = processor contains AXI bus ECC logic.
[2]	SL	Indicates whether the processor has been built with split/lock logic: 0 = no split/lock logic 1 = split/lock logic included.
[1]	AHB_PP	Indicates whether the CPU contain AHB peripheral interfaces: 0 = CPUs do not have AHB peripheral interfaces 1 = CPUs have AHB peripheral interfaces.
[0]	MICRO_SCU	Indicates whether the processor contains an ACP interface: 0 = processor does not contain ACP logic 1 = processor does contain ACP logic.

To access the Build Options 2 Register, read CP15 with:


MRC p15, 0, <Rd>, c15, c2, 1 ; Read Build Options 2 Register

4.3.33 Pin Options Register

The Pin Options Register characteristics are:

Purpose	Describes the value of any pins that control processor options, that are not visible because they: - are exposed in registers - control the initial value of control registers, and are visible in that way.
Usage constraints	The Pin Options Register is: - a read-only register - accessible in Privileged modes only.
Configurations	Available in all processor configurations.
Attributes	See Table 4-60 on page 4-84.

Figure 4-57 on page 4-84 shows the Pin Options Register bit assignments.



4-84

4.3.34 Peripheral interface region registers

There are three peripheral interface region registers, one for each of the:

- The Peripheral Interface Region Register characteristics are:

Usage constraints The Peripheral Interface Region Registers are:

- Read/write registers.
- Accessible in Privileged mode only.
- The enable bits for the LLPP Normal AXI and AHB peripheral interface region registers are initialized immediately after reset, from the values on the **INITPPXm** and **INITPPHm** pins.
- The LLPP Virtual AXI region register enable resets to zero.

Attributes See Table 4-61 on page 4-85.

Figure 4-58 on page 4-85 shows the Peripheral Interface Region Register bit assignments.

**Figure 4-58 Peripheral Interface Region Register bit assignments**

Table 4-61 shows the Peripheral Interface Region Register bit assignments.

Table 4-61 Peripheral Interface Region Register bit assignments

Bit	Name	Type	Function
[31:12]	BaseAddress	RO	The base address of the interface, given as bits [31:12] of the address of the interface in the memory map. This value is configured during integration.
[11:7]	-	RO	Reserved.
[6:2]	Size	RO	Returns the size of the interface configured during integration: 0b00000 = no PP present 0b00011 = 4KB ... 0b10111 = 4GB
[1]	-	RO	Reserved.
[0]	En	RW	Interface enable bit: 0 = Disabled 1 = Enabled. The reset value of this bit is: - for LLPP Normal AXI, determined by INITPPXm - for LLPP Virtual AXI, always 0 - for AHB peripheral interface, determined by INITPPHm.

To access the Peripheral Interface Region Registers, read CP15 with:

MRC p15, 0, <Rt>, c15, c0, 1; Read LLPP Normal AXI region register
MRC p15, 0, <Rt>, c15, c0, 2; Read LLPP Virtual AXI region register
MRC p15, 0, <Rt>, c15, c0, 3; Read AHB peripheral interface region register

Chapter 5

Prefetch Unit

This chapter describes how the *PreFetch Unit* (PFU), in conjunction with the DPU, uses program flow prediction to locate branches in the instruction stream and the strategies used to determine if a branch is likely to be taken or not. It contains the following sections:

- *About the prefetch unit* on page 5-2
- *Branch prediction* on page 5-3
- *Return stack* on page 5-5
- *Controlling instruction prefetch and program flow prediction* on page 5-6.

5.1 About the prefetch unit

The purpose of the PFU is to:

- perform speculative fetch of instructions ahead of the DPU by predicting the outcome of branch instructions
- format instruction data in a way that aids the DPU in efficient execution.

The PFU fetches instructions from the memory system under the control of the DPU, and the internal coprocessors CP14 and CP15. In ARM state the memory system can supply up to two instructions per cycle. In Thumb state the memory system can supply up to four instructions per cycle.

The PFU buffers up to three instruction data fetches in its FIFO. There is an additional FIFO between the PFU and the DPU that can normally buffer up to eight instructions. This reduces or eliminates stall cycles after a branch instruction. This increases the performance of the processor.

Program flow prediction occurs in the PFU by:

- predicting the outcome of conditional branches using the branch predictor and, for direct branches, calculating their destination address using the offset encoded in the instruction
- predicting the destination of procedure returns using the return stack.

The DPU resolves the program flow predictions that the PFU makes.

The PFU fetches the instruction stream as dictated by:

- the Program Counter
- the branch predictor
- procedure returns signaled by the return stack
- exceptions including aborts and interrupts signaled by the DPU
- correction of mispredicted branches as indicated by the DPU.

The PFU starts instruction fetches at a rate that is determined dynamically using a prediction scheme that aims to ensure that the pipeline is kept fed with instructions, without over-fetching instructions that are not used. Fetching of unused instructions consumes extra power and can impact performance.

5.2 Branch prediction

The PFU normally fetches instructions from sequential addresses. If a branch instruction is fetched, the next instruction to be fetched can only be determined with certainty after the instruction has completed execution at the end of the pipeline in the DPU. If the branch is taken, the next instruction to be executed is not sequential. The sequential instructions that the PFU has fetched while the branch instruction was executing must be flushed from the pipeline and the correct instruction fetched. This has the effect of reducing the performance of the processor.

The PFU can detect branches in the Pd-stage of the pipeline, predict whether or not the branch is taken, and determine or predict the target address for a taken branch. This enables the PFU to start fetching instructions at the destination of a taken branch before the branch has completed execution in the DPU. The branch instruction is still executed in the DPU to determine the accuracy of the prediction. If the branch was mispredicted, the pipeline must be flushed and the correct instruction fetched. In general, more branches are correctly predicted than mispredicted so fewer pipeline flushes occur and the performance of the processor is enhanced.

Two major classes of branch are addressed in the processor prediction scheme:

1. Direct branches, including B, BL, CZB, and BLX immediate, where the target address is a fixed offset, encoded in the instruction, from the program counter. If such an instruction has been fetched, and the program counter is known, predicting the destination of the branch only involves predicting whether the instruction passes or fails its condition code, that is, whether the branch is taken or not taken.
2. Indirect branches such as load and *Branch and eXchange* (BX), instructions that write to the PC, that can be identified as a likely return from a procedure call. Two identifiable cases are:
 - loads to the PC from an address derived from R13
 - BX from R0-R14.

In these cases, if the calling operation can also be identified, the likely return address can be stored in the return stack. Typical calling operations are BL and BLX instructions.

Note

Unconditional instructions of either class of program flow are always executed, and do not affect prediction history. Unconditional return stack operations always affect the return stack.

This section describes:

- *Branch predictor*
- *Incorrect predictions and correction* on page 5-4.

5.2.1 Branch predictor

Branch prediction in the processor is dynamic and is based around a global history prediction scheme. In addition, there is extra logic to handle predictions that thrash and to predict the end of long loops.

The global history scheme is an adaptive predictor that learns the behavior of branches during execution, identifying them based on the historical pattern of behavior of the preceding branches. For each pattern of branch behavior, the history table holds a 2-bit hint value. The 2-bit hint indicates if the next branch must be predicted taken or predicted not-taken based on the behavior of previous branches. The history table contains 256 entries.

For loops beyond a certain number of iterations, the branch history is not large enough to learn the history and predict the loop exit. The PFU includes logic to count the number of iterations (up to 31) of a loop, and thereby predict the not-taken branch that exits the loop. If the number of iterations taken exceeds 31, the loop branch is never predicted as not-taken.

If multiple branch histories index into the same hint value, this can cause thrashing in the history table and reduce accuracy of the branch predictor. Logic in the branch predictor detects these cases and provides some hysteresis for the hint value.

For direct branches, the target address is calculated statically from the instruction encoding and the program counter. For indirect branches, the hint value predicts if the branch is taken or not-taken, and the return stack can sometimes be used to predict the target address. When the destination of a branch cannot be calculated statically, or popped from the return stack, PFU assumes the branch to be not-taken.

The PFU updates the history for each occurrence of a branch when the DPU indicates how the branch was resolved.

5.2.2 Incorrect predictions and correction

The DPU resolves branches that the dynamic branch predictor predicts at the Wr-stage of the pipeline, see Figure 2-1 on page 2-2. A misprediction causes the PFU to flush the pipeline and fetch the correct instruction stream.

5.3 Return stack

The call-return stack predicts procedural returns that are program flow changes such as loads, and branch register. The dynamic branch predictor determines if conditional procedure returns are predicted as taken or not-taken, as described in *Branch prediction* on page 5-3. The return stack predicts the target address for unconditional procedure returns, and conditional procedure returns that have been predicted as taken by the branch predictor.

The return stack consists of a 4-entry circular buffer. When the PFU detects a taken procedure call instruction, the PFU pushes the return address onto the return stack. The instructions that the PFU recognizes as procedure calls are, in both the ARM and Thumb instruction sets:

- BL immediate
- BLX immediate
- BLX Rm.

When the return stack detects a taken return instruction, the PFU issues an instruction fetch from the location at the top of the return stack, and pops the return stack. The instructions that the PFU recognizes as procedure returns are, in both the ARM and Thumb instruction sets:

- LDM Rn{!}, {...,pc}
- POP {...,pc}
- LDMIB Rn{!}, {...,pc}
- LDMDA Rn{!}, {...,pc}
- LDMDB Rn{!}, {...,pc}
- LDR pc, [sp], #4
- BX Rm.

Return stack mispredictions can exist when:

- The prediction that a conditional return passed or failed its condition code is not correct.
- The return address of an unconditional or predicted-taken return is not correct.

The return stack has no underflow or overflow detection. Either scenario is likely to cause a misprediction.

Note

The MOV PC, LR instruction is not decoded and is not predicted as a return.

5.4 Controlling instruction prefetch and program flow prediction

In the Cortex-R5 processor, the Z-bit, bit [11] of the SCTL_R, does not control the program flow prediction. The Z-bit is read-as-one, writes-ignored and instead a number of control bits in the Auxiliary Control Register control the program flow and prefetch features. To disable the program flow prediction, you must disable the return stack and set the branch prediction policy to always not-taken. See *c1, Auxiliary Control Register* on page 4-41.

The fetch rate predictor can be disabled by setting FRCDIC in the Auxiliary Control Register. When the predictor is disabled, the PFU fetches instructions at the fastest rate possible.

The dynamic branch predictor is controlled with the BP field in the Auxiliary Control Register. In normal operation the branch prediction is taken from the global history table. You can also force the prediction to be always taken, or always not-taken. When the prediction is forced to a fixed direction, the processor does not update the global history table, and the historic pattern of branches is frozen. You can also disable the loop prediction logic and the logic for preventing thrashing, by setting DEOLP and DBHE respectively.

You can disable the return stack by setting RSDIS in the Auxiliary Control Register. When disabled, pushes onto the stack caused by call instructions are disabled and the stack pointer is frozen.

Chapter 6

Events and Performance Monitor

This chapter describes the *Performance Monitoring Unit* (PMU) and event bus interface. It contains the following sections:

- *About the events* on page 6-2
- *About the PMU* on page 6-6
- *Performance monitoring registers* on page 6-7
- *Event bus interface* on page 6-20.

6.1 About the events

The processor includes logic to detect various events that can occur, for example, a cache miss. These events provide useful information about the behavior of the processor that you can use when debugging or profiling code.

The events are made visible on an output event bus, **EVNTBUSm**, and can be counted using registers in the *Performance Monitoring Unit* (PMU). See *Event bus interface* on page 6-20 for more information about the event bus, and *About the PMU* on page 6-6 for more information about the PMU. Table 6-1 lists the events that are generated, along with the bit position of each event on the event bus, and the numbers that the PMU uses to refer the events. Event reference numbers that are not listed are Reserved. See *Error detection events* on page 8-36 for more information on the CFLR related events.

Table 6-1 Event bus interface bit functions

EVNTBUSm bit position	Description	CFLR update	Event Ref. Value
-	Software increment. The register is incremented only on writes to the Software Increment Register. See <i>c9, Software Increment Register</i> on page 6-12.	-	0x00
[0]	Instruction cache miss. Each instruction fetch from normal Cacheable memory that causes a refill from the level 2 memory system generates this event. Accesses that do not cause a new cache refill, but are satisfied from refilling data of a previous miss are not counted. Where instruction fetches consist of multiple instructions, these accesses count as single events. CP15 cache maintenance operations do not count as events.	-	0x01
[1]	Data cache miss. Each data read from or write to normal Cacheable memory that causes a refill from the level 2 memory system generates this event. Accesses that do not cause a new cache refill, but are satisfied from refilling data of a previous miss are not counted. Each access to a cache line to normal Cacheable memory that causes a new linefill is counted, including the multiple transactions of an LDM and STM. Write-through writes that hit in the cache do not cause a linefill and so are not counted. CP15 cache maintenance operations do not count as events.	-	0x03
[2]	Data cache access. Each access to a cache line is counted including the multiple transactions of an LDM, STM, or other operations. CP15 cache maintenance operations do not count as events.	-	0x04
[3]	Data Read architecturally executed. This event occurs for every instruction that explicitly reads data, including SWP.	-	0x06
[4]	Data Write architecturally executed. This event occurs for every instruction that explicitly writes data, including SWP.	-	0x07
[5]	Instruction architecturally executed ^a .	-	0x08
[6]	Dual-issued pair of instructions architecturally executed.	-	0x5e
[7]	Exception taken. This event occurs on each exception taken.	-	0x09
[8]	Exception return architecturally executed. This event occurs on every exception return, for example, "RFE, MOVs PC, LDM Rn, {...,PC}^".	-	0x0A
[9]	Change to Context ID executed.	-	0x0B

Table 6-1 Event bus interface bit functions (continued)

EVNTBUSm bit position	Description	CFLR update	Event Ref. Value
[10]	Software change of PC, except by an exception, architecturally executed.	-	0x0C
[11]	B immediate, BL immediate or BLX immediate instruction architecturally executed (taken or not taken).	-	0x0D
[12]	Procedure return architecturally executed, other than exception returns, for example, BZ Rm, "LDM Rn, {...,PC}". MOV PC, LR does not generate this event, because it is not predicted as a return.	-	0x0E
[13]	Unaligned access architecturally executed. This event occurs for each instruction that was to an unaligned address that either triggered an alignment fault, or would have done so if the SCTLRA-bit had been set.	-	0x0F
[14]	Branch mispredicted or not predicted. This event occurs for every pipeline flush caused by a branch.	-	0x10
-	Cycle count.	-	0x11
[15]	Branches or other change in program flow that could have been predicted by the branch prediction resources of the processor.	-	0x12
[16]	Stall because instruction buffer cannot deliver an instruction. This can indicate an instruction cache miss. This event occurs every cycle where the condition is present.	-	0x40
[17]	Stall because of a data dependency between instructions. This event occurs every cycle where the condition is present.	-	0x41
[18]	Data cache write-back. This event occurs once for each line that is written back from the cache.	-	0x42
[19]	External memory request. Examples of this are cache refill, Non-cacheable accesses, write-through writes, cache line evictions (write-back).	-	0x43
[20]	Stall because of LSU being busy. This event takes place each clock cycle where the condition is met. A high incidence of this event indicates the pipeline is often waiting for transactions to complete on the external bus.	-	0x44
[21]	Store buffer was forced to drain completely. Examples of this for Cortex-R5 are DMB, Strongly Ordered memory access, or similar events.	-	0x45
-	The number of cycles FIQ interrupts are disabled.	-	0x46
-	The number of cycles IRQ interrupts are disabled.	-	0x47
-	ETMEXTOUTm[0].	-	0x48
-	ETMEXTOUTm[1].	-	0x49
[22]	Instruction cache tag RAM parity or correctable ECC error.	Yes	0x4A
[23]	Instruction cache data RAM parity or correctable ECC error.	Yes	0x4B
[24]	Data cache tag or dirty RAM parity error or correctable ECC error, from data-side or ACP.	Yes	0x4C

Table 6-1 Event bus interface bit functions (continued)

EVNTBUSm bit position	Description	CFLR update	Event Ref. Value
[25]	Data cache data RAM parity error or correctable ECC error.	Yes	0x4D
[26]	TCM fatal ECC error reported from the prefetch unit.	-	0x4E
[27]	TCM fatal ECC error reported from the load/store unit.	-	0x4F
-	Store buffer merge.	-	0x50
-	LSU stall caused by full store buffer.	-	0x51
-	LSU stall caused by store queue full.	-	0x52
-	Integer divide instruction, SDIV or UDIV, executed.	-	0x53
-	Stall cycle caused by integer divide.	-	0x54
-	PLD instruction that initiates a linefill.	-	0x55
-	PLD instruction that did not initiate a linefill because of a resource shortage.	-	0x56
-	Non-cacheable access on AXI master bus.	-	0x57
[28]	Instruction cache access. This is an analog to event 0x04.	-	0x58
-	Store buffer operation has detected that two slots have data in same cache line but with different attributes.	-	0x59
[29]	Dual issue case A (branch).	-	0x5A
[30]	Dual issue case B1, B2, F2 (load/store), F2D.	-	0x5B
[31]	Dual issue other case.	-	0x5C
[32]	Double-precision floating point arithmetic or conversion instruction executed.	-	0x5D
[33]	Data cache data RAM fatal ECC error.	-	0x60
[34]	Data cache tag/dirty RAM fatal ECC error, from data-side or ACP.	-	0x61
[35]	Processor livelock because of hard errors or exception at exception vector.	-	0x62
[36]	Unused.	-	0x63
[37]	ATCM multi-bit ECC error.	-	0x64
[38]	B0TCM multi-bit ECC error.	-	0x65
[39]	B1TCM multi-bit ECC error.	-	0x66
[40]	ATCM single-bit ECC error.	-	0x67
[41]	B0TCM single-bit ECC error.	-	0x68
[42]	B1TCM single-bit ECC error.	-	0x69
[43]	TCM correctable ECC error reported by load/store unit.	Yes	0x6A
[44]	TCM correctable ECC error reported by prefetch unit.	Yes	0x6B
[45]	TCM fatal ECC error reported by AXI slave interface.	-	0x6C
[46]	TCM correctable ECC error reported by AXI slave interface.	Yes	0x6D

Table 6-1 Event bus interface bit functions (continued)

EVNTBUSm bit position	Description	CFLR update	Event Ref. Value
-	All correctable events ^b , OR of: 0x4A ICache tag 0x4B ICache data 0x4C DCache tag/dirty 0x4D DCache data 0x6A LSU TCM 0x6B PFU TCM 0x6D AXI-S TCM 0x70 bus-ECC	Yes	0x6E
-	All fatal events ^b , OR of: 0x60 DCache tag 0x61 DCache tag/dirty 0x4E PFU TCM 0x4F LSU TCM 0x6C AXI-S TCM 0x71 bus -ECC	-	0x6F
[47]	All correctable bus faults ^b	-	0x70
[48]	All fatal bus faults ^b	-	0x71
[49]	ACP D-Cache access, lookup or invalidate.	-	0x72
[50]	ACP D-Cache invalidate.	-	0x73
-	Cycle count	-	0xFF
[51]-[54]	Unused	-	-

- a. If one of the event counters is configured to count this event, then the counter increases by two when a dual-issued pair of instructions are architecturally executed. The **EVNTBUSm[5]** signal is asserted for one cycle only in the same situation - use **EVNTBUSm[6]** to distinguish this situation.
- b. This event is signalled when any one of a number of events occur. It is formed as a logical OR of the constituent events. This means that if two or more of the constituent events occur at the same time, the composite event is only signaled once and the event counter, if configured to count this event, is only incremented by one.

6.2 About the PMU

The PMU consists of three event counting registers, one cycle counting register and 12 CP15 registers, for controlling and interrogating the counters. The performance monitoring registers are always accessible in Privileged mode. You can use the *User Enable* (PMUSERENR) Register to make all of the performance monitoring registers, except for the *Interrupt Enable Set* (PMINTENSET), and *Interrupt Enable Clear* (PMINTENCLR) Registers, accessible in User mode.

All three event counters are read and written through the same CP15 register. The *Performance Counter Selection* (PMSELR) Register determines which counter is read or written. The three Event Selection registers, one per counter, are read and written through one CP15 register in the same way.

Using the control registers, you can enable or disable each of the event counters individually, and read and reset the overflow flag for each counter. Any or all of the counters can be enabled to assert an interrupt request output, **nPMUIRQm**, on overflow.

When the processor is in Debug halt state:

- the PMU does not count events
- events are not visible on the ETM interface
- the *Cycle CouNT* (PMCCNTR) register is halted.

For more information on Debug halt state see Chapter 12 *Debug*.

The PMU only counts events when non-invasive debug is enabled, that is, when either **DBGENm** or **NIDENm** inputs are asserted. The *Cycle Count* (PMCCNTR) Register is always enabled regardless of whether non-invasive debug is enabled, unless the DP bit of the PMCR register is set. See *c9, Performance Monitor Control Register* on page 6-7.

6.3 Performance monitoring registers

The performance monitoring registers are described in:

- *c9, Performance Monitor Control Register*
- *c9, Count Enable Set Register* on page 6-8
- *c9, Count Enable Clear Register* on page 6-9
- *c9, Overflow Flag Status Register* on page 6-11
- *c9, Software Increment Register* on page 6-12
- *c9, Performance Counter Selection Register* on page 6-12
- *c9, Cycle Count Register* on page 6-13
- *c9, Event Type Selection Register* on page 6-14
- *c9, Event Count Registers* on page 6-16
- *c9, User Enable Register* on page 6-16
- *c9, Interrupt Enable Set Register* on page 6-17
- *c9, Interrupt Enable Clear Register* on page 6-18.

6.3.1 c9, Performance Monitor Control Register

The PMCR Register characteristics are:

Purpose Controls the operation of the three count registers, and the PMCCNTR Register.

Usage constraints The PMCR Register is:

- a read/write register
- accessible in:
 - Privileged mode
 - User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.

Configurations Available in all processor configurations.

Attributes See Table 6-2 on page 6-8.

Figure 6-1 shows the bit assignments.

31		24	23		16	15		11	10		6	5	4	3	2	1	0
IMP				IDCODE				N		Reserved		D P	X	D	C	P	E

Figure 6-1 PMCR Register bit assignments

Table 6-2 shows the bit assignments.

Table 6-2 PMCR Register bit assignments

Bits	Name	Function
[31:24]	IMP	Implementer code: 0x41 = ARM
[23:16]	IDCODE	Identification code: 0x15 = Cortex-R5
[15:11]	N	Specifies the number of counters implemented: 0x3 = three counters implemented
[10: 6]	Reserved	RAZ on reads, <i>Should Be Zero or Preserved</i> (SBZP) on writes
[5]	DP	Disable PMCCNTR when prohibited, that is, when non-invasive debug is not enabled: 0 = Count is enabled in prohibited regions. This is the reset value. 1 = Count is disabled in prohibited regions.
[4]	X	Enable export of the events to the event bus for an external monitoring block, for example the ETM, to trace events: 0 = Export disabled. This is the reset value. 1 = Export enabled.
[3]	D	Cycle count divider: 0 = Counts every processor clock cycle. This is the reset value. 1 = Counts every 64th processor clock cycle.
[2]	C	Cycle counter reset: Write one to this bit to reset the cycle counter, PMCCNTR, to zero. This bit Reads-As-Zero.
[1]	P	Event counter reset: Write one to this bit to reset all event counters to zero. This bit Reads-As-Zero.
[0]	E	Enable: 0 = Disable all counters, including PMCCNTR. This is the reset value. 1 = Enable all counters including PMCCNTR.

The PMCR Register is always accessible in Privileged mode. To access the register, read or write CP15 with:

```
MRC p15, 0, <Rd>, c9, c12, 0 ; Read PMCR Register
MCR p15, 0, <Rd>, c9, c12, 0 ; Write PMCR Register
```

6.3.2 c9, Count Enable Set Register

The PMCNTENSET Register characteristics are:

Purpose Enables the Event Count Registers.

Usage constraints The PMCNTENSET Register is:

- accessible in:
 - Privileged mode
 - User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.

- The values in this register are ignored unless the E bit, bit [0], is set in the PMCR Register, see *c9, Performance Monitor Control Register* on page 6-7.

Configurations Available in all processor configurations.

Attributes See Table 6-3.

Figure 6-2 shows the bit assignments.




Figure 6-2 PMCNTENSET Register bit assignments

Table 6-3 shows the bit assignments.

Table 6-3 PMCNTENSET Register bit assignments

Bits	Name	Function
[31]	C	Cycle counter enable
[30:3]	Reserved	UNP on reads, SBZP on writes
[2]	P2	Counter 2 enable
[1]	P1	Counter 1 enable
[0]	P0	Counter 0 enable

To access the PMCNTENSET Register, read or write CP15 with:

MRC p15, 0, <Rd>, c9, c12, 1 ; Read PMCNTENSET Register
MCR p15, 0, <Rd>, c9, c12, 1 ; Write PMCNTENSET Register

When reading this register, any enable that reads as 0 indicates the corresponding counter is disabled. Any enable that reads as 1 indicates the corresponding counter is enabled.

Writing a 1 to a particular count enable bit enables that counter. Writing a 0 to a count enable bit has no effect. You must use the Count Enable Clear Register to disable the counters. All counters are disabled at reset.

The PMCNTENSET Register retains its value when the enable bit of the PMCR is clear, even though its settings are ignored.

6.3.3 c9, Count Enable Clear Register

The PMCNTENCLR Register characteristics are:

Purpose Disables any of the Event Count Registers.

Usage constraints The PMCNTENCLR Register is:

- accessible in:
 - Privileged mode

— User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.

- Configurations**
- Available in all processor configurations.
- Attributes**
- See Table 6-4.

Figure 6-3 shows the bit assignments.




Figure 6-3 PMCNTENCLR Register bit assignments

Table 6-4 shows the bit assignments.

Table 6-4 PMCNTENCLR Register bit assignments

Bits	Name	Function
[31]	C	Cycle counter disable:
[30:3]	Reserved	UNP on reads, SBZP on writes
[2]	P2	Counter 2 enable
[1]	P1	Counter 1 enable
[0]	P0	Counter 0 enable

To access the PMCNTENCLR Register, read or write CP15 with:

```
MRC p15, 0, <Rd>, c9, c12, 2 ; Read PMCNTENCLR Register
MCR p15, 0, <Rd>, c9, c12, 2 ; Write PMCNTENCLR Register
```

When reading this register, any enable that reads as 0 indicates the corresponding counter is disabled. Any enable that reads as 1 indicates the corresponding counter is enabled.

When writing this register, any enable written with a value of 0 is ignored, that is, not updated. Any enable written with a value of 1 clears the counter enable. You must use the Count Enable Set Register to enable the counters. All counters are disabled at reset.

Writing to bits in this register disables individual counters, and clears the corresponding bits in the PMCNTENSET Register, see *c9, Count Enable Set Register* on page 6-8.

You can use the enable, EN, bit [0] of the PMCR Register to disable all performance counters including PMCCNTR, see *c9, Performance Monitor Control Register* on page 6-7.

The PMCNTENCLR and PMCNTENSET Registers retain their values when the enable bit of the PMCR is clear, even though their settings are ignored. The PMCNTENCLR Register can be used to clear the enabled flags for individual counters even when all counters are disabled in the PMCR Register.

6.3.4 c9, Overflow Flag Status Register

The PMOVSr Register characteristics are:

- Purpose** Indicates if event counters have overflowed. All overflow flags are reset to zero.
- Usage constraints** The PMOVSr Register is accessible in:
- Privileged mode
 - User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.
- Configurations** Available in all processor configurations.
- Attributes** See Table 6-5.

Figure 6-4 shows the bit assignments.




Figure 6-4 PMOVSr Register bit assignments

Table 6-5 shows the bit assignments.

Table 6-5 PMOVSr Register bit assignments

Bits	Name	Function
[31]	Cycle counter overflow	Cycle counter overflow flag
[30:3]	Reserved	UNP on reads, SBZP on writes
[2]	P2	Counter 2 overflow flag
[1]	P1	Counter 1 overflow flag
[0]	P0	Counter 0 overflow flag

To access the PMOVSr Register, read or write CP15 with:

```
MRC p15, 0, <Rd>, c9, c12, 3 ; Read PMOVSr Register
MCR p15, 0, <Rd>, c9, c12, 3 ; Write PMOVSr Register
```

If an overflow flag is set to 1 in the PMOVSr register it remains set until one of the following happens:

- writing 1 to the flag bit in the PMOVSr Register clears the flag
- the processor is reset.

The following operations do *not* clear the overflow flags:

- disabling the overflowed counter in the PMCNTENCLR Register
- disabling all counters in the PMCR Register
- resetting the overflowed counter using the PMCR Register.

6.3.5 c9, Software Increment Register

The PMSWINC Register characteristics are:

- Purpose**
- Increments the count of an Event Count Register.
- Usage constraints**
- The PMSWINC Register is:

•

A write-only register that Reads-As-Zero

•

Accessible in:

—

Privileged mode

—

User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.

•

You must only use the PMSWINC Register to increment Event Count Registers when the counter event is set to 0x00, software count, in the Event Select Register, see *c9, Event Type Selection Register* on page 6-14.

If you attempt to use the PMSWINC Register to increment an Event Count Register when the counter event is set to a value other than 0x00 the result is Unpredictable.

Configurations

Available in all processor configurations.

Attributes

See Table 6-6.

Figure 6-5 shows the bit assignments.

31 3 2 1 0

Reserved

Performance monitor counters software increment bits

P2 P1 P0

Figure 6-5 PMSWINC Register bit assignments

Table 6-6 shows the bit assignments.

Table 6-6 PMSWINC Register bit assignments

Bits	Name	Function
[31:3]	Reserved	RAZ on reads, SBZP on writes
[2]	P2	Increment Counter 2
[1]	P1	Increment Counter 1
[0]	P0	Increment Counter 0

To access the PMSWINC Register, write CP15 with:

MCR p15, 0, <Rd>, c9, c12, 4 ; Write PMSWINC Register

6.3.6 c9, Performance Counter Selection Register

The PMSELR Register characteristics are:

Purpose

•

selects an Event Count Register.

ARM DDI 0460C
ID021511

Copyright © 2010-2011 ARM. All rights reserved.
Non-Confidential

6-12

- determines which count register is accessed or controlled by accesses to the Event Selection Register and the Event Count Register.

Usage constraints The PMSELR Register is:

- A read/write register.
- Accessible in:
 - Privileged mode
 - User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.

Configurations Available in all processor configurations.

Attributes See Table 6-7.

Figure 6-6 shows the bit assignments.




Figure 6-6 PMSELR Register bit assignments

Table 6-7 shows the bit assignments.

Table 6-7 PMSELR Register bit assignments

Bits	Name	Function
[31:5]	Reserved	RAZ on reads, SBZP on writes
[4:0]	SEL	Counter select: b00000 = selects counter 0 b00001 = selects counter 1 b00010 = selects counter 2.

Any values programmed in the PMSELR Register other than those specified in Table 6-7 are Unpredictable.

To access the PMSELR Register, read or write CP15 with:

MRC p15, 0, <Rd>, c9, c12, 5 ; Read PMSELR Register
MCR p15, 0, <Rd>, c9, c12, 5 ; Write PMSELR Register

6.3.7 c9, Cycle Count Register

The PMCCNTR Register characteristics are:

Purpose Counts clock cycles.

Usage constraints The PMCCNTR Register:

- Is a 32-bit read/write register.
- Is accessible in:
 - Privileged mode
 - User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.

- Must be disabled before software can write to it. Any attempt by software to write to this register when enabled is Unpredictable.

Configurations Available in all processor configurations.

To access the PMCCNTR read or write CP15 with:

MRC p15, 0, <Rd>, c9, c13, 0 ; Read PMCCNTR Register
MCR p15, 0, <Rd>, c9, c13, 0 ; Write PMCCNTR Register

6.3.8 c9, Event Type Selection Register

There are three Event Type Select Registers in the processor, PMXEVTYPER0 to PMXEVTYPER2, each corresponding to one of the *Performance Monitor Count* (PMXEVCNTR) Registers, PMXEVCNTR0 to PMXEVCNTR2. The register to be accessed is determined by the value in the PMSELR.

The PMXEVTYPER Register characteristics are:

Purpose Selects the events you want a PMXEVCNTR Register to count.

Usage constraints The PMXEVTYPER Register is:

- A read/write register
- Accessible in:
 - Privileged mode
 - User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register* on page 6-16.

Configurations Available in all processor configurations.

Attributes See Table 6-8.

Figure 6-7 shows the bit assignments.




Figure 6-7 PMXEVTYPERx Register bit assignments

Table 6-8 shows the bit assignments.

Table 6-8 PMXEVTYPERx Register bit functions

Bits	Name	Function
[31:8]	-	RAZ or SBZP.
[7:0]	SEL	Event number selected, see Table 6-1 on page 6-2 for values. The reset value of this field is Unpredictable.

To access the PMXEVTYPERx Register, read or write CP15 with:

MRC p15, 0, <Rd>, c9, c13, 1 ; Read PMXEVTYPERx Register
MCR p15, 0, <Rd>, c9, c13, 1 ; Write PMXEVTYPERx Register

The absolute counts of events recorded might vary because of pipeline effects. This has negligible effect except in cases where the counters are enabled for a very short time.

In addition to the counters within the processor, most of the events that Table 6-1 on page 6-2 shows are available to the ETM unit or other external trace hardware to enable monitoring of the events. For information on how to monitor these events, see the *CoreSight ETM-R5 Technical Reference Manual*.

6.3.9 c9, Event Count Registers

There are three Event Count Registers (PMXEVNTR0-PMXEVNTR2) in the processor. Each PMXEVNTR Register, as selected by the PMSELR Register, counts instances of an event selected by the corresponding PMXEVTYPER Register. The register to be accessed is determined by the value in the PMSELR.

Each PMXVCNTR Register is:

- A 32-bit read/write register.
- Accessible in:
 - Privileged mode
 - User mode only when the PMUSERENR.EN bit is set to 1, see *c9, User Enable Register*.

To access the current Event Count Registers, read or write CP15 with:

```
MRC p15, 0, <Rd>, c9, c13, 2 ; Read current PMNx Register
MCR p15, 0, <Rd>, c9, c13, 2 ; Write current PMNx Register
```

6.3.10 c9, User Enable Register

The PMUSERENR Register characteristics are:

Purpose

Enables User mode to have access to:

- the performance monitor registers, see *Performance monitoring registers* on page 6-7
- the validation registers, see *Validation Registers* on page 4-68.

Usage constraints

The PMUSERENR Register:

- is a read/write register
- is writable only in Privileged mode, readable in any processor mode
- does not provide access to the registers that control interrupt generation.

Configurations

Available in all processor configurations.

Attributes

See Table 6-9 on page 6-17.

Figure 6-8 shows the bit assignments.




Figure 6-8 PMUSERENR Register bit assignments

Table 6-9 shows the bit assignments.

Table 6-9 PMUSERENR Register bit assignments

Bits	Name	Function
[31:1]	Reserved	RAZ or SBZP.
[0]	EN	User mode access to performance monitor and validation registers: 0 = Disabled. This is the reset value. 1 = Enabled.

If the EN bit in the PMUSERENR Register is not set, any attempt to access a performance monitor register or a validation register from User mode causes an Undefined Instruction exception.

———— **Note** ————

For more information on access permissions to the performance monitor registers and validation registers, see the *ARM Architecture Reference Manual*.

To access the PMUSERENR Register, read or write CP15 with:

```
MRC p15, 0, <Rd>, c9, c14, 0 ; Read PMUSERENR Register
MCR p15, 0, <Rd>, c9, c14, 0 ; Write PMUSERENR Register
```

6.3.11 c9, Interrupt Enable Set Register

The PMINTENSET Register characteristics are:

Purpose	Determines if any of the PMXEVNTR Registers, PMXEVNTR0-PMXEVNTR2 and PMCCNTR, generate an interrupt request on overflow.
Usage constraints	The PMINTENSET Register is: • a read/write register • accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 6-10 on page 6-18.

Figure 6-9 shows the bit assignments.




Figure 6-9 PMINTENSET Register bit assignments

Table 6-10 shows the bit assignments.

Table 6-10 PMINTENSET Register bit assignments

Bits	Name	Function
[31]	C	PMCCNTR overflow interrupt
[30:3]	Reserved	UNP on reads, SBZP on write
[2]	P2	PMXEVCNTR2 overflow interrupt
[1]	P1	PMXEVCNTR1 overflow interrupt
[0]	P0	PMXEVCNTR0 overflow interrupt

Reading this register returns the current setting, with a 1 in one of the counter bits indicating that interrupts are enabled for that counter. Writing a 1 to a particular interrupt bit enables interrupt generation on overflow of that counter. Writing a 0 has no effect. You can only disable interrupts by writing to the PMINTENCLR Register.

To access the Interrupt Enable Set Register, read or write CP15 with:

MRC p15, 0, <Rd>, c9, c14, 1 ; Read PMINTENSET Register
MCR p15, 0, <Rd>, c9, c14, 1 ; Write PMINTENSET Register

If this unit generates an interrupt, the processor asserts the pin **nPMUIRQm**. You can route this pin to an external interrupt controller for prioritization and masking. This is the only mechanism that signals this interrupt to the processor.

Note

ARM expects that the Performance Monitor interrupt request signal, **nPMUIRQm**, connects to a system interrupt controller.

6.3.12 c9, Interrupt Enable Clear Register

The PMINTENCLR Register characteristics are:

Purpose	Determines if any of the PMXEVCNTR Registers, PMXEVCNTR0-PMXEVCNTR2 and PMCCNTR, generate an interrupt request on overflow.
Usage constraints	The PMINTENCLR Register is: - a read/write register - accessible in Privileged mode only.
Configurations	Available in all processor configurations.
Attributes	See Table 6-11 on page 6-19.

Figure 6-10 on page 6-19 shows the bit assignments.

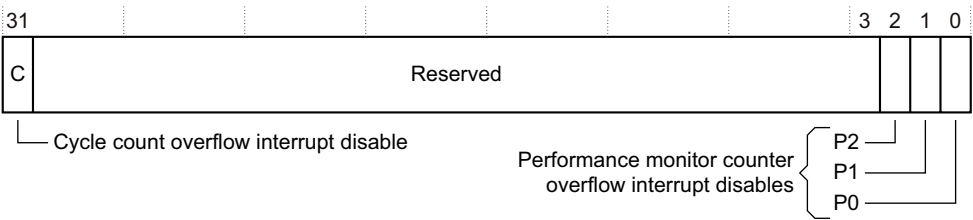


Figure 6-10 PMINTENCLR Register bit assignments

Table 6-11 shows the bit assignments.

Table 6-11 PMINTENCLR Register bit assignments

Bits	Name	Function
[31]	C	PMCCNTR overflow interrupt
[30:3]	Reserved	UNP on reads, SBZP on writes
[2]	P2	PMXEVCNTR2 overflow interrupt
[1]	P1	PMXEVCNTR1 overflow interrupt
[0]	P0	PMXEVCNTR0 overflow interrupt

Reading this register returns the current setting, with a 1 in one of the counter bits indicating that interrupts are enabled for that counter. Writing a 1 to a particular interrupt disable bit disables interrupt generation on overflow of that counter. Writing a 0 has no effect. You can only enable interrupt requests by writing to the PMINTENSET Register.

To access the PMINTENCLR Register, read or write CP15 with:

```
MRC p15, 0, <Rd>, c9, c14, 2 ; Read PMINTENCLR Register
MCR p15, 0, <Rd>, c9, c14, 2 ; Write PMINTENCLR Register
```

6.4 Event bus interface

The event bus, **EVNTBUSm**, is used to signal when an event has occurred. The event bus includes most, but not all, of the events that can be counted by the performance monitoring unit. Each individual event is assigned to an individual bit of this bus, and this bit is asserted for one cycle each time the event occurs.

The event bus only signals events when it is enabled. Set the X bit in the Performance Monitor Control Register to enable the event bus. See *c9, Performance Monitor Control Register* on page 6-7.

See Table 6-1 on page 6-2 to see which bit of the event bus each event is signaled on.

———— Note ————

If an event is being counted in the PMU, the count might not be incremented in exactly the same cycle that the event is signaled on the event bus.

6.4.1 Use of the event bus and counters

The event bus is designed to be connected to the ETM-R5, that enables processor events to trigger tracing for debug purposes. You can also connect it to event counting registers external to the processor, or to an interrupt generator.

Because each **EVNTBUSm** pin is only asserted for one cycle for each occurrence of the event, it is possible to create composite events by ORing various **EVNTBUSm** pins together. A composite event signal like this is asserted when any of the included events occur although, if multiple events occur in the same cycle, the composite event only occurs once.

The processor also has two event input pins, **ETMEXTOUTm[1:0]**. This bus is normally intended for connection to the ETM, and enables the Cortex-R5 performance monitor to count events generated by the ETM. These inputs can alternatively be used for composite events generated external to the processor.

Chapter 7

Memory Protection Unit

This chapter describes the *Memory Protection Unit (MPU)*. It contains the following sections:

- *About the MPU* on page 7-2
- *Memory types* on page 7-7
- *Region attributes* on page 7-8
- *MPU interaction with memory system* on page 7-9
- *MPU faults* on page 7-10
- *MPU software-accessible registers* on page 7-11.

7.1 About the MPU

The MPU works with the L1 memory system to control accesses to and from L1 and external memory. For a full architectural description of the MPU, see the *ARM Architecture Reference Manual*.

The MPU enables you to partition memory into regions and set individual protection attributes for each region. The MPU supports zero, 12, or 16 memory regions.

Note

If the MPU has zero regions, you cannot enable or program the MPU. Attributes are only determined from the default memory map when zero regions are implemented.

Each region is programmed with a base address and size, and the regions can be overlapped to enable efficient programming of the memory map. To support overlapping, the regions are assigned priorities, with region 0 having the lowest priority and region 15 having the highest. The MPU returns access permissions and attributes for the highest priority enabled region where the address hits.

The MPU is programmed using CP15 registers c1 and c6, see *MPU control and configuration* on page 4-3. Memory region control read and write access is permitted only from Privileged modes.

Table 7-1 shows the default memory map.

Table 7-1 Default memory map

Address range	Instruction memory type		Data memory type		Execute Never
	Instruction cache enabled	Instruction cache disabled	Data cache enabled	Data cache disabled	
0xFFFFFFF	Normal	Normal	Strongly Ordered	Strongly Ordered	Instruction execution only permitted if HIVECS is TRUE
0xF000000	Non-cacheable only if HIVECS is TRUE	Non-cacheable only if HIVECS is TRUE			
0xEFFFFFF	-	-	Strongly Ordered	Strongly Ordered	Execute Never
0xC000000					
0xBFFFFFF	-	-	Shared Device	Shared Device	Execute Never
0xA000000					
0x9FFFFFF	-	-	Non-shared Device	Non-shared Device	Execute Never
0x8000000					
0x7FFFFFF	Normal, Cacheable, Non-shared	Normal, Non-cacheable, Non-shared	Normal, Non-cacheable, Shared	Normal, Non-cacheable, Shared	Instruction execution permitted
0x6000000					
0x5FFFFFF	Normal, Cacheable, Non-shared	Normal, Non-cacheable, Non-shared	Normal, WT Cacheable, Non-shared	Normal, Non-cacheable, Shared	Instruction execution permitted
0x4000000					
0x3FFFFFF	Normal, Cacheable, Non-shared	Normal, Non-cacheable, Non-shared	Normal, WBWA Cacheable, Non-shared	Normal, Non-cacheable, Shared	Instruction execution permitted
0x0000000					

This section describes:

- *Memory regions*
- *Overlapping regions* on page 7-4
- *Background regions* on page 7-6
- *TCM regions* on page 7-6
- *Peripheral port regions* on page 7-6.

7.1.1 Memory regions

Before the MPU is enabled, you must program at least one valid protection region. If you do not do this, the processor enters a state that only reset can recover.

When the MPU is disabled, no access permission checks are performed, and memory attributes are assigned according to the default memory map. See Table 7-1 on page 7-2.

For more information on how to enable or disable the MPU, see *MPU interaction with memory system* on page 7-9.

Depending on the implementation, the MPU has a maximum of 12 or 16 regions. Using CP15 register c6 you can specify the following for each region:

- region base address
- region size
- subregion enables
- region attributes
- region access permissions
- region enable.

Region base address

The base address defines the start of the memory region. You must align this to a region-sized boundary. For example, if a region size of 8KB is programmed for a given region, the base address must be a multiple of 8KB.

———— Note ————

If the region is not aligned correctly, this results in Unpredictable behavior.

Region size

The region size is specified as a 5-bit value, encoding a range of values from 32 bytes, a cache-line length, to 4GB. Table 4-34 on page 4-55 shows the encoding.

Subregions

Each region can be split into eight equal sized non-overlapping subregions. An access to a memory address in a disabled subregion does not use the attributes and permissions defined for that region. Instead, it uses the attributes and permissions of a lower priority region or generates a background fault if no other regions overlap at that address. This enables increased protection and memory attribute granularity.

All region sizes between 256 bytes and 4GB support eight subregions. Region sizes below 256 bytes do not support subregions, and the subregion disable field is SBZ/UNP for regions of less than 256 bytes in size.

Region attributes

Each region has a number of attributes associated with it. These control how a memory access is performed when the processor accesses an address that falls within a given region. The attributes are:

- Memory type, one of:
 - Strongly Ordered
 - Device
 - Normal
- Shared or Non-shared
- Non-cacheable
- Write-through Cacheable
- Write-back Cacheable
- Read allocation
- Write allocation.

See *Memory types* on page 7-7 for more information about memory types, and *Region attributes* on page 7-8 for a description of how to assign types and attributes to a region.

Region access permissions

Each region can be given no access, read-only access, or read/write access permissions for Privileged or all modes. In addition, each region can be marked as *eXecute Never* (XN) to prevent instructions being fetched from that region.

For example, if a User mode application attempts to access a *Privileged mode access only* region a permission fault occurs.

The ARM architecture uses constants known as *inline literals* to perform address calculations. The assembler and compiler automatically generate these constants and they are stored inline with the instruction code. To ensure correct operation, only a memory region that has permission for data read access can execute instructions. For more information, see the *ARM Architecture Reference Manual*. For information about how to program access permissions, see Table 4-38 on page 4-58.

Instructions cannot be executed from regions with Device or Strongly-Ordered memory type attributes.

7.1.2 Overlapping regions

You can program the MPU with two or more overlapping regions. For overlapping regions, a fixed priority scheme determines attributes and permissions for memory access to the overlapping region. Attributes and permissions for region 15 take highest priority, those for region 0 take lowest priority. For example:

Region 2	Is 4KB in size, starting from address 0x3000. Privileged mode has full access, and User mode has read-only access.
Region 1	Is 16KB in size, starting from address 0x0000. Both Privileged and User modes have full access.

When the processor performs a data write to address 0x3010 while in User mode, the address falls into both region 1 and region 2, as Figure 7-1 on page 7-5 shows. Because these regions have different permissions, the permissions associated with region 2 are applied. Because User mode is read access only for this region, a permission fault occurs, causing a data abort.




Figure 7-1 Overlapping memory regions

Example of using regions that overlap

You can use overlapping regions for stack protection. For example:

- allocate to region 1 the appropriate size for all stacks
- allocate to region 2 the minimum region size, 32 bytes, and position it at the end of the stack for the current process
- set the region 2 access permissions to No Access.

If the current process overflows the stack it uses, a write access to region 2 by the processor causes the MPU to raise a permission fault.




Figure 7-2 Overlay for stack protection

Example of using subregions

You can use subregions for stack protection. For example:

- Allocate to region 1 the appropriate size for all stacks.
- Set the least-significant subregion disable bit. That is, set the subregion disable field, bits [15:8], of the CP15 MPU Region Size Register to 0x01.

If the current process overflows the stack it uses, a write access by the processor to the disabled subregion causes the MPU to raise a background fault.




Figure 7-3 Overlapping subregion of memory

7.1.3 Background regions

Overlapping regions increase the flexibility of how the regions can be mapped onto physical memory devices in the system. You can also use the overlapping properties to specify a background region. For example, you might have a number of physical memory areas sparsely distributed across the 4GB address space. If a programming error occurs, the processor might issue an address that does not fall into any defined region.

If the address that the processor issues falls outside any of the defined regions and the MPU is enabled, the MPU is hard-wired to abort the access. That is, all accesses for an address that is not mapped to a region in the MPU generate a background fault. You can override this behavior by programming region 0 as a 4GB background region. In this way, if the address does not fall into any of the other 11 regions, the attributes and access permissions you specified for region 0 control the access.

In Privileged modes, you can also override this behavior by setting the BR bit, bit [17], of the SCTL. This causes Privileged accesses that fall outside any of the defined regions to use the default memory map.

7.1.4 TCM regions

Any memory address that you configure to be accessed using a TCM interface is given Normal, Non-shared type attributes, regardless of the attributes of any MPU region that the address also belongs to. Access permissions for an address in a TCM region are preserved from the MPU region that the address also belongs to. For more information, see *About the TCMs* on page 8-13.

7.1.5 Peripheral port regions

Any memory address accessed using one of the peripheral port interfaces is considered to be non-cacheable and eXecute-Never (XN), regardless of the attributes of any MPU region that the address also belongs to. The memory type and other access permissions for such a region are inherited from the MPU region that the address also belongs to. See *Peripheral interface attributes and permissions* on page 9-34.

7.2 Memory types

The ARM architecture defines a set of memory types with characteristics that are suited to particular devices. There are three mutually exclusive memory type attributes:

- Strongly Ordered
- Device
- Normal.

MPU memory regions can each be assigned a memory type attribute. Table 7-2 shows a summary of the memory types.

Table 7-2 Memory attributes summary

Memory type attribute	Shared or Non-shared	Description
Strongly Ordered	-	All memory accesses to Strongly Ordered memory occur in program order. All Strongly Ordered accesses are assumed to be shared.
Device	Shared	For memory-mapped peripherals that several processors share.
	Non-shared	For memory-mapped peripherals that only a single processor uses.
Normal	Shared	For normal memory that is shared between several processors.
	Non-shared	For normal memory that only a single processor uses.

———— **Note** ————

The processor's L1 cache does not cache shared normal regions.

For more information on memory attributes and types, memory barriers, and ordering requirements for memory accesses, see the *ARM Architecture Reference Manual*.

7.2.1 Using memory types

All of the processor interfaces to the external memory system have associated store buffers that help to improve the throughput of accesses to Normal type memory. See *Store buffer* on page 8-18 and *Peripheral interfaces* on page 9-31 for more information. Because of the ordering rules that they must follow, accesses to other types of memory typically have a lower throughput or higher latency than accesses to Normal memory. In particular:

- reads from Device memory must first drain the relevant store buffer of all writes to Device memory and wait for all Device writes to the relevant interface that have been posted onto the bus to complete
- all accesses to Strongly Ordered memory must first drain the store buffer completely and wait for all writes that have been posted onto the buses to complete.

Similarly, when it is accessing Strongly Ordered or Device type memory, the processor's response to interrupts must be modified, and the interrupt response latency is longer. See *Low interrupt latency* on page 3-20 for more information.

To ensure optimum performance, you must understand the architectural semantics of the different memory types. Use Device memory type for appropriate memory regions, typically peripherals, and only use Strongly Ordered memory type for memory regions where it is essential.

7.3 Region attributes

Each region has a number of attributes associated with it. These control how a memory access is performed when the processor accesses an address that falls within a given region. The attributes are:

- Memory type, see *Memory types* on page 7-7, one of:
 - Strongly Ordered
 - Device
 - Normal.
- Shared or Non-shared
- Non-cacheable
- Write-through cacheable
- Write-back cacheable
- Read allocation
- Write allocation.

The Region Access Control Registers use five bits to encode the memory region type. These are the TEX[2:0], C and B bits. Table 4-36 on page 4-57 shows the mapping of these bits to memory region attributes.

Note

In earlier versions of the architecture, the TEX, C, and B bits were known as the Type Extension, Cacheable and Bufferable bits. These names no longer adequately describe the function of the B, C, and TEX bits.

All memory attributes that are Cacheable, write-back or write-through, are also implicitly read-allocate. Table 4-36 on page 4-57 shows which attributes are write-allocate.

In addition, the Region Access Control Registers contain the shared bit, S. This bit only applies to Normal memory, and determines whether the memory region is Shared (1) or Non-shared (0).

When the processor performs a memory access through its AXI bus master interface:

- the Inner attributes are indicated on the **A*INNERMm** signals.
- the Outer attributes are indicated on the **A*CACHEMm** signals.

For the encodings, see Table 9-2 on page 9-6.

Similarly, for memory accesses performed through the AXI peripheral port, the Outer attributes are indicated on the **A*CACHEPm** signals.

For more information on region attributes, see the *ARM Architecture Reference Manual*.

7.4 MPU interaction with memory system

This section describes how to enable and disable the MPU. After you enable or disable the MPU, the pipeline must be flushed using ISB and DSB instructions to ensure that all subsequent instruction fetches and data accesses see the effect of turning on or off the MPU.

Before you enable or disable the MPU you must:

1. Program all relevant CP15 registers. This includes setting up at least one memory region that covers the executing code, and that the attributes and permissions of that region are the same as the attributes and permissions of the region in the default memory map that covers the code, and that the region is executable in Privileged mode.
2. Clean and invalidate the data caches.
3. Disable caches.
4. Invalidate the instruction cache.

The following code is an example of enabling the MPU:

```
MRC p15, 0, R1, c1, c0, 0    ; read CP15 register 1
ORR R1, R1, #0x1
DSB
MCR p15, 0, R1, c1, c0, 0    ; enable MPU
ISB
Fetch from programmed memory map
Fetch from programmed memory map
Fetch from programmed memory map
Fetch from programmed memory map
```

The following code is an example of disabling the MPU:

```
MRC p15, 0, R1, c1, c0, 0    ; read CP15 register 1
BIC R1, R1, #0x1
DSB
MCR p15, 0, R1, c1, c0, 0    ; disable MPU
ISB
Fetch from default memory map
Fetch from default memory map
Fetch from default memory map
Fetch from default memory map
```

Table 7-1 on page 7-2 shows the default memory map.

7.5 MPU faults

The MPU can generate three types of fault:

- *Background fault*
- *Permission fault*
- *Alignment fault.*

When a fault occurs, the memory access or instruction fetch is synchronously aborted, and a prefetch abort or data abort exception is taken as appropriate. No memory accesses are performed on the AXI bus master interface or peripheral ports. For more information about fault handling, see *Fault handling* on page 8-7.

7.5.1 Background fault

A background fault is generated when the MPU is enabled and a memory access is made to an address that is not within an enabled subregion of an MPU region. A background fault does not occur if the background region is enabled and the access is Privileged. See *Background regions* on page 7-6.

7.5.2 Permission fault

A permission fault is generated when a memory access does not meet the requirements of the permissions defined for the memory region that it accesses. See *Region access permissions* on page 7-4.

7.5.3 Alignment fault

An alignment fault is generated if a data access is performed to an address that is not aligned for the size of the access, and strict alignment is required for the access. A number of instructions that access memory, for example, LDM and STC, require strict alignment. See the *ARM Architecture Reference Manual* for more information. In addition, strict alignment can be required for all data accesses by setting the A-bit in the SCTLR. See *c1, System Control Register* on page 4-38.

7.6 MPU software-accessible registers

Figure 4-2 on page 4-3 shows the CP15 registers that control the MPU.

When the MPU is not present, the *c6*, *MPU memory region programming registers* on page 4-53 read as zero and ignore writes in Privileged mode. No Undefined Instruction exceptions are taken.

Chapter 8

Level One Memory System

This chapter describes the processor *Level one* (L1) memory system. It contains the following sections:

- *About the L1 memory system* on page 8-2
- *About the error detection and correction schemes* on page 8-4
- *Fault handling* on page 8-7
- *About the TCMs* on page 8-13
- *About the caches* on page 8-18
- *Internal exclusive monitor* on page 8-34
- *Memory types and L1 memory system behavior* on page 8-35
- *Error detection events* on page 8-36.

8.1 About the L1 memory system

The processor L1 memory system can be configured during implementation and integration. It can consist of:

- separate instruction and data caches
- multiple *Tightly-Coupled Memory* (TCM) areas
- a *Memory Protection Unit* (MPU).

The instruction-side and data-side can each optionally have their own L1 caches. The cache architecture is Harvard, that is, only instructions can be fetched from the I-Cache, and only data can be fetched from the D-Cache. In parallel with each of the caches are two areas of dedicated RAM accessible to both the instruction and data sides. These are regions of TCM. You can implement one TCM using the ATCM interface and up to two TCMs using the BTCM interface. Figure 8-1 on page 8-3 shows this.

Memory accesses, required for fetching instructions and for data transfer instructions, are performed to the appropriate TCM if the address is in an enabled TCM region. Remaining instruction accesses and remaining data accesses that are not in a peripheral interface region are looked up in the appropriate L1 cache if they are cacheable. Accesses that are not serviced by the L1 memory system are passed to the L2 memory system through the AXI-master interface or one of the peripheral interfaces. See Chapter 9 *Level Two Interface* for more information about the L2 memory system.

Each TCM and cache can be configured at implementation time to have an error detection and correction scheme to protect the data stored in the memory from errors. Each TCM interface also has support for logic external to the processor to tell the processor that an error has occurred.

The MPU handles accesses to both the instruction and data sides. The MPU is responsible for protection checking, address access permissions, and memory attributes for all accesses. Some of these attributes can be passed to the L2 memory system through the AXI master or peripheral ports. See Chapter 7 *Memory Protection Unit* for more information about the MPU.

The L1 memory system includes a monitor for exclusive accesses. Exclusive load and store instructions, for example LDREX and STREX, can be used with the appropriate memory monitoring to provide inter-process or inter-processor synchronization and semaphores. See the *ARM Architecture Reference Manual* for more information. The internal monitor can handle some exclusive monitoring internally to the processor, see *Internal exclusive monitor* on page 8-34 for more information.




Figure 8-1 Memory system block diagram

8.2 About the error detection and correction schemes

In silicon devices, stray radiation and other effects can cause the data stored in a RAM to be corrupted. The TCMs and caches on a Cortex-R5 processor can be configured to detect and correct errors that can occur in the RAMs. Extra, redundant data is computed by the processor and stored in the RAMs alongside the real data. When the processor reads data from the RAMs, it checks that the redundant data is consistent with the real data and can either signal an error, or attempt to correct the error.

A number of different error schemes are available, and are described in:

- *Parity*
- *64-bit ECC* on page 8-5
- *32-bit ECC* on page 8-5.

Each has different properties in terms of the number of errors that can be detected, and corrected, and the amount of extra RAM required to store the redundant data. Because different logic is required for each scheme, the scheme must be chosen in the build-configuration, although you can enable or disable, or change the behavior of the error schemes using software-configuration. This section describes the generic properties of each of the schemes. See Appendix C *ECC Schemes* for more information about the advantages and disadvantages of each scheme to the implementer. Operation of the error schemes for the caches is described in *Cache error detection and correction* on page 8-20, and for the TCMs in *TCM internal error detection and correction* on page 8-14.

The error schemes are each described in terms of their operation on a doubleword, 64 bits, of data, because this is the amount of data that the processor L1 memory system can transfer each cycle. The tag and dirty RAMs associated with the caches are different sizes, but the principles are the same. An error is considered to be a single bit of data that has been inverted relative to its correct value.

Figure 8-2 shows the error schemes. The shaded areas represent bits with errors.




Figure 8-2 Error detection and correction schemes

8.2.1 Parity

For each byte, a parity bit is computed and stored with that byte. This requires eight bits of parity, or redundant data per doubleword. With a parity scheme, a single error in a byte or its parity bit can be detected, but not corrected. This means that, provided they are all in different bytes, eight errors can be detected per doubleword. However, if there are two errors in any individual byte, this cannot be detected. Odd or even parity can be used, and this can be pin-configured during integration.

8.2.2 Error checking and correction

The processor supports *Error Checking and Correction* (ECC) schemes for either 64-bits or 32-bits of data, and these have similar properties, although though the size of the data chunk that the ECC scheme applies to is different. For each data chunk, either 32-bits or 64-bits, aligned, a number of redundant code bits are computed and stored with the data. This enables the processor to detect up to two errors in the data chunk or its code bits, and correct any single error in the data chunk or its associated code bits. This is sometimes referred to as a *Single-Error-Correction, Double-Error-Detection* (SEC-DED) ECC scheme.

If there are more than two errors in a data chunk and its associated code bits, they might or might not be detected. The error scheme might interpret such a condition as a single-error and make an unsuccessful attempt at a correction.

64-bit ECC

Eight code bits are computed for each 64 bits of data. The scheme can correct any single error occurring in any doubleword, and detect any two errors occurring in any doubleword.

32-bit ECC

Seven code bits are computed for each 32 bits of data, so 14 bits of redundant data are required for each doubleword. The scheme can correct two errors per doubleword, if they are in different words. Four errors can be detected per doubleword, if there are two in each word.

8.2.3 Read-Modify-Write

The smallest unit of data that the processor can write is a byte. However, both the ECC schemes are computed on data chunks that are larger than this. To write any data to a RAM protected with ECC requires the error code for that data to be recomputed and rewritten. If the entire data chunk is not written, for example, a halfword, 16-bits, is written to address 0x4 of a RAM with a 32-bit error scheme, the error code must be computed partly from the data being written, and partly from data already stored in the RAM. In this example, the halfword in the RAM at address 0x6.

To compute the error code for such a write, the processor must first read data from the RAM, then merge the data to be written with it, to compute the error code, then write the data to the RAM, along with the new error code. This process is referred to as read-modify-write.

8.2.4 Hard errors

The errors described in this chapter are all assumed to be soft errors, that is, one or more bits of the data stored in a RAM chunk are inverted. A new value can still be written to the RAM and read back correctly, unless another soft error occurs in the meantime.

If the error in the memory is a hard error, that is, a physical failure of the RAM circuit so that a bit can never be read or written reliably, the processor might not be able to correct and recover from the error. The processor contains features that enable it to recover from some hard errors. If you are implementing the processor and require these features, contact ARM to discuss the features and your requirements.

8.2.5 Error correction

When a correctable error is detected in data that has been read from a RAM, the processor has various ways of generating the correct data, that follow two schemes:

Correct inline

The error code bits are used to correct the data read from the RAM, and this data is used. This is the simplest way of correcting the data.

Correct-and-retry

The error code bits are used to correct the data, and this data is then written back to the RAM. The processor then repeats the read access by re-executing the instruction that caused the read, and reads the corrected data from the RAM if no more errors have occurred. This takes more clock cycles, at least nine, in the event of an error, but has the side-effect of correcting the data in the RAM so that the errors in the data cannot become worse.

Note

Because RAM errors generally occur infrequently, the extra cycles required to perform correct-and-retry do not have a significant impact on average performance.

The correction method that the processor uses depends on the individual error. The processor uses correct inline error correction when it detects a correctable error on a TCM read made by the AXI-slave interface. The processor uses correct-and-retry correction when it detects a correctable ECC error on a TCM read made by the instruction-side or data-side.

8.3 Fault handling

Faults can occur on instruction fetches for the following reasons:

- MPU background fault
- MPU permission fault
- External AXI slave error (SLVERR)
- External AXI decode error (DECERR)
- Cache parity or ECC error
- TCM ECC error
- TCM external error
- TCM external retry request
- Breakpoints, and vector capture events.

Faults can occur on data accesses for the following reasons:

- MPU background fault
- MPU permission fault
- MPU alignment fault
- External AXI slave error (SLVERR)
- External AXI decode error (DECERR)
- External AHB error
- Cache parity or ECC error
- TCM ECC error
- TCM external error
- TCM external retry request
- Watchpoints.

Fault handling is described in:

- *Faults*
- *Fault status information* on page 8-9
- *Correctable Fault Location Register* on page 8-10
- *Usage models* on page 8-10.

8.3.1 Faults

The classes of fault that can occur are:

- *MPU faults*
- *External faults* on page 8-8
- *Cache and TCM parity and ECC errors* on page 8-8
- *TCM external faults* on page 8-8
- *Debug events* on page 8-9
- *Synchronous and asynchronous aborts* on page 8-9.

MPU faults

The MPU can generate an abort for various reasons. See *MPU faults* on page 7-10 for more information. MPU faults are always synchronous, and take priority over other types of abort. If an MPU fault occurs on an access that is not in the TCM, and is to one of the peripheral ports, is Non-cacheable, or has generated a cache-miss, the AXI/AHB transactions for that access are not performed.

External faults

A memory access performed through the AXI master interface or the AXI peripheral port can generate two different types of error response, a *slave error* (SLVERR) or *decode error* (DECERR). These are known as external errors, because they are generated by the AXI system outside the processor. Synchronous aborts are generated for instruction fetches, data loads, exclusive stores, and data stores to strongly-ordered-type memory. Non-exclusive stores to normal-type or device-type memory generate asynchronous aborts.

Note

- An AXI slave that cannot handle exclusive transactions returns OKAY in response to an exclusive read. This is also treated as an external error, and the processor behaves as if the response was SLVERR.
 - Exclusive doubleword transactions to shared memory on the AXI peripheral port or exclusive transactions to shared memory on the AHB peripheral port are aborted. They are treated as synchronous external errors, and the processor behaves as if the response was SLVERR.
 - An AHB peripheral port slave response of **ERROR** is treated by the processor as a response of SLVERR.
-

Cache and TCM parity and ECC errors

If the processor has been configured with the appropriate build options, it can detect data errors occurring in the cache and TCM RAMs using parity or ECC logic. For more information on cache errors, see *Handling cache parity errors* on page 8-21 and *Handling cache ECC errors* on page 8-22. For more information on TCM errors, see *About the error detection and correction schemes* on page 8-4. Depending on the software configuration of the processor, these errors are either ignored, generate an abort, are automatically corrected without generating an abort, or are corrected and generate an abort. If the processor is in debug-halt-state, an error that is otherwise automatically corrected generates an abort.

Parity and ECC errors can only occur on reads, although these reads might be a side-effect of store instructions. Aborts generated by loads are always synchronous. Aborts generated by store instructions to the TCM are also always synchronous, while those to the cache are always asynchronous. These errors can also occur on some cache-maintenance operations, see *Errors on cache maintenance operations* on page 8-23, and generate asynchronous aborts.

Many of the parity and ECC errors are also signaled by the generation of events. See Chapter 6 *Events and Performance Monitor*. Some of these events are generated when the error is detected, regardless of whether or not an abort is taken. Aborts are only taken when a memory access with an error is committed. Others are signaled when and only when the abort is taken.

Any parity or ECC error that can be corrected by the processor is considered to be a *correctable fault*, regardless of whether or not the processor is configured to correct the fault.

TCM external faults

The TCM port includes signals that can be used to signal an error on a TCM transaction. If enabled, this causes the processor to take the appropriate type of abort for instruction and data accesses, or to generate a SLVERR response to an AXI-slave transaction. Write transactions always generate asynchronous aborts, while read transactions always generate synchronous aborts.

An error signaled on a read transaction can also signal a retry request, that requests that the processor retry the same operation rather than take an exception.

A retry request from the TCM port is considered to be a *recoverable error*. All correctable ECC faults are also considered to be recoverable.

Debug events

The debug logic in the processor can be configured to generate breakpoints or vector capture events on instruction fetches, and watchpoints on data accesses. If the processor is software-configured for monitor-mode debugging, an abort is taken when one of these events occurs, or when a BKPT instruction is executed. For more information, see Chapter 12 *Debug*.

Synchronous and asynchronous aborts

See *Aborts* on page 3-23 for more information about the differences between synchronous and asynchronous aborts.

8.3.2 Fault status information

When an abort occurs, information about the cause of the fault is recorded in a number of registers, depending on the type of abort:

- *Abort exceptions*
- *Synchronous abort exceptions* on page 8-10
- *Asynchronous abort exceptions* on page 8-10.

Abort exceptions

The following registers are updated when any abort exception is taken:

Link Register

The r14_abt register is updated to provide information about the address of the instruction that the exception was taken on, in a similar way to other types of exception. See *Exceptions* on page 3-17 for more information. This information can be used to resume program execution after the abort has been handled.

————— Note —————

When a prefetch abort has occurred, ARM recommends that you do not use the link register value for determining the aborting address, because 32-bit Thumb instructions do not have to be word aligned and can cause an abort on either halfword. This applies even if all of the code in the system does not use the extra 32-bit Thumb instructions introduced in ARMv6T2, because the earlier BL and BLX instructions are both 32 bits long. Use the Fault Address Register instead, as described in this section.

Saved Program Status Register

The SPSR_abt register is updated to record the state and mode of the processor when the exception was taken, in a similar way to other types of exception. See *Exceptions* on page 3-17 for more information.

Fault Status Register

There are two fault status registers, one for prefetch aborts (IFSR) and one for data aborts (DFSR). These record the type of abort that occurred, and whether it occurred on a read or a write. In particular, this enables the abort handler to distinguish between synchronous aborts, asynchronous aborts, and debug events. See *Fault Status and Address Registers* on page 4-49 for more information about the format of this register and the encodings used.

Synchronous abort exceptions

The following registers are updated when a synchronous abort exception is taken:

Fault Address Register

There are two fault address registers, one for prefetch aborts (IFAR) and one for data aborts (DFAR). These indicate the address of the memory access that caused the fault. See *Fault Status and Address Registers* on page 4-49.

Auxiliary Fault Status Register

There are two auxiliary fault status registers, one for prefetch aborts (AIFSR) and one for data aborts (ADFSR). These record additional information about the nature and location of the fault, including whether it was a recoverable error or not, whether it occurred in the cache, AXI-master interface, AXI peripheral port, AHB peripheral port, ATCM or BTCM and, if appropriate, which cache way the error occurred in. The cache index is not recorded on a synchronous abort, because this information can be derived from the fault address. See *Fault Status and Address Registers* on page 4-49.

Asynchronous abort exceptions

The following register is updated when an asynchronous abort exception is taken:

Auxiliary Data Fault Status Register

The ADFSR is updated to indicate whether or not the fault was recoverable, whether it occurred in the cache, AXI-master interface, AXI peripheral port, AHB peripheral port, ATCM or BTCM and, if appropriate, which cache set and way the error occurred in. Because the DFAR is not updated on asynchronous aborts, asynchronous aborts cannot normally be located, except when the error occurred in the cache.

The effect of debug events on these registers is described in *Debug exception* on page 12-42.

8.3.3 Correctable Fault Location Register

Correctable faults are normally automatically corrected by the processor but, depending on the configuration and on the access that generated the fault, an exception might not be generated, and the fault status registers might not be updated. In all cases, information about the location of the fault is recorded in the *Correctable Fault Location Register* (CFLR).

The CFLR also records information about ACP D-Cache lookups that cause a correctable error.

All correctable faults are recorded in the same register, regardless of whether it was an instruction-fetch, a data-access, an AXI-slave access, or an ACP coherency maintenance operation that generated the fault, and whether the fault occurred in the ATCM, BTCM or cache. The CFLR contains information to identify what sort of access generated the fault, and which device it occurred in. See *Correctable Fault Location Register* on page 4-77 for more information about the format of this register. Each time the CFLR is updated, the information already in the CFLR is discarded and therefore the CFLR can only contain information about the most recent correctable fault.

8.3.4 Usage models

This section describes some ways in which errors can be handled in a system. Exactly how you program the processor to handle errors depends on the configuration of your processor and system, and what you are trying to achieve.

If an abort exception is taken, the abort handler reads the information in the link register, SPSR, and fault status registers to determine the type of abort. Some types of abort are fatal to the system, and others can be fixed, and program execution resumed. For example, an MPU background fault might indicate a stack overflow, and be rectified by allocating more stack and reprogramming the MPU to reflect this. Alternatively, an asynchronous external abort might indicate that a software error meant that a store instruction occurred to an unmapped memory address. Such an abort is fatal to the system or process because no information is recorded about the address the error occurred on, or the instruction that caused the error.

Table 8-1 shows which types of abort are typically fatal because either the location of the error is not recorded or the error is unrecoverable. Some aborts that are marked as not fatal might turn out to be fatal in some systems when the cause of the error has been determined. For example, an MPU background fault might indicate a stack overflow, that can be rectified, or it might indicate that, because of a bug, the software has accessed a nonexistent memory location, that can be fatal. These cases can be distinguished by determining the location where the error occurred. If an error is unrecoverable, that is, it is not a correctable parity or ECC error, and it is not a TCM external retry request, it is normally fatal regardless of whether or not the location of the error is recorded. When an abort is taken on an external TCM, parity, or ECC error, the appropriate Auxiliary Fault Status Register records whether the error was recoverable. See *Fault Status and Address Registers* on page 4-49.

Table 8-1 Types of aborts

Type	Conditions	Source	Synchronous	Fatal
MPU fault	Access not permitted by MPU ^a	MPU	Yes	No
Synchronous External	Load using L2 memory interface	AXI, AHB	Yes	No
Asynchronous External	Store to Normal or Device memory using L2 memory interface	AXI, AHB	No	Yes
Synchronous Parity/ECC Cache	Load from cache ^b	Cache	Yes	Maybe ^c
Synchronous ECC TCM	Load/store from/to TCM ^d	TCM	Yes	Maybe ^c
Synchronous TCM external error	Load/store from/to TCM ^e	TCM	Yes	Yes
Asynchronous Parity/ECC Cache	Store to cache or cache maintenance operation ^b	Cache	No	Maybe ^c
Asynchronous TCM external error	Store to TCM ^e	TCM	No	Yes

a. See *MPU faults* on page 7-10 for more information about the types of MPU fault.

b. See *Cache error detection and correction* on page 8-20 for more information about parity/ECC errors from the cache.

c. These types of error can be correctable or uncorrectable. Uncorrectable errors are typically fatal. Correctable errors are automatically corrected by the hardware and might not cause the abort handler to be called. See *Cache error detection and correction* on page 8-20 and *TCM internal error detection and correction* on page 8-14.

d. See *TCM internal error detection and correction* on page 8-14 for more information about ECC errors from the TCM.

e. Aborts generated by external TCM errors are always unrecoverable, and therefore fatal, see *External TCM errors* on page 8-16 for more information about external errors from the TCM.

Correctable errors

In a system in which the processor is configured to automatically correct ECC errors without taking an abort exception, you can still configure it to respond to such errors. Connect the event output or outputs that indicate a correctable error to an interrupt controller. When such an event occurs, the interrupt input to the processor is set, and the processor takes an interrupt exception. When your interrupt handler has identified the source of the interrupt as a correctable error, it can read the CFLR to determine where the ECC error occurred. You can examine this information to identify trends in such errors. By masking the interrupt when necessary, your

software can ensure that when critical code is executing, the processor corrects the error automatically, but delays examining information about the error until after the critical code has completed.

When the processor is in debug halt-state, any correctable error is corrected as appropriate, but the memory access is not repeated to fetch the correct data, therefore the instruction generating the error does not complete successfully. Instead, the sticky synchronous abort flag in the DBGDSCR is set. See *CP14 c1, Debug Status and Control Register* on page 12-14.

8.4 About the TCMs

The processor has two TCM interfaces to support the connection of local memories. The ATCM interface has one TCM port. The BTCM interface can support one or two TCM ports. Each TCM port is a physical connection on the processor that is suitable for connection to SRAM with minimal glue logic. These ports are optimized for low latency memory.

The TCM ports are designed to be connected to RAM, or RAM-like memory, that is, Normal-type memory. The processor can issue speculative read accesses on these interfaces, and interrupt store instructions that have issued some but not all of their write accesses. Therefore, both read and write accesses through the TCM interfaces can be repeated. This means that the TCM ports are generally not suitable for read- or write-sensitive devices such as FIFOs. ROM can be connected to the TCM ports, but normally only if ECC is not used. See *Hard errors* on page 8-5. If the access is speculative, the processor ignores any error or retry signaled on the TCM port.

The TCM ports also have wait and error signals to support slow memories and external error detection and correction. For more information, see *External TCM errors* on page 8-16.

The PFU can read data using the TCM interfaces. The LSU and AXI slave can each read and write data using the TCM interfaces.

Each TCM interface has a dedicated base address that you can place anywhere in the physical address map, and must not be backed by memory implemented externally. The ATCM and BTCM interfaces must have separate base addresses and must not overlap.

This section describes:

- *TCM attributes and permissions*
- *ATCM and BTCM configuration* on page 8-14
- *TCM internal error detection and correction* on page 8-14
- *TCM arbitration* on page 8-15
- *TCM initialization* on page 8-15
- *TCM port protocol* on page 8-16
- *External TCM errors* on page 8-16
- *AXI slave interfaces for TCMs* on page 8-17.

8.4.1 TCM attributes and permissions

Accesses to the TCMs from the LSU and PFU are checked against the MPU for access permission. Memory access attributes and permissions are not exported on this interface. Reads that generate an MPU fault are broadcast on the TCM interface but the abort is taken before the data is used, ensuring protection is maintained.

TCMs always behave as Non-cacheable Non-shared Normal memory, irrespective of the memory type attributes defined in the MPU for a memory region containing addresses held in the TCM. Access permissions for TCM accesses are the same as the permission attributes that the MPU assigns to the same address. See Chapter 7 *Memory Protection Unit* for more information about memory attributes, types, and permissions.

———— Note ————

Any address in an MPU region with device or strongly-ordered memory type attributes is implicitly given *execute-never* (XN) permissions. If such an address is also in a TCM region, XN permissions are applied to TCM accesses to that address. None of the other device or strongly-ordered behaviors apply to an address in a TCM region.

8.4.2 ATCM and BTCM configuration

The TCM interfaces are configured during implementation and integration.

You can configure the ATCM interface to be removed, and not included in the processor design. If implemented, the ATCM can have only a single port.

You can configure the BTCM interface to:

- be removed, and not included in the processor design
- have a single BTCM port
- have two banked BTCM ports, interleaved on either:
 - Bit [3] of the address
 - The most significant bit of the BTCM interface address. This depends on the size of the BTCM.

During implementation, you can configure the ATCM and/or the BTCM to use an error-protection scheme to protect the data stored in the TCM, see *TCM internal error detection and correction*.

The size of each TCM interface is configured during integration. The permissible TCM sizes are:

- 0KB
- 4KB
- 8KB
- 16KB
- 32KB
- 64KB
- 128KB
- 256KB
- 512KB
- 1MB
- 2MB
- 4MB
- 8MB.

If the BTCM interface has two ports, the size of the RAM attached to each port is half the total size for the BTCM interface.

The size of the TCM interfaces is visible to software in the TCM Region Registers, see *c9, BTCM Region Register* on page 4-63 and *c9, ATCM Region Register* on page 4-64. All TCM interface build configuration options can be read from the Build Options Registers, see *c15, Build Options 1 Register* on page 4-79 and *c15, Build Options 2 Register* on page 4-80.

8.4.3 TCM internal error detection and correction

Each TCM interface can be configured with either 32-bit ECC, or 64-bit ECC error schemes. Both the BTCM ports must have the same error scheme. This section describes these error schemes.

If a TCM interface has been built with either 32-bit or 64-bit ECC error checking, you can enable this by setting the appropriate bits in the Auxiliary Control Register. See *c1, Auxiliary Control Register* on page 4-41. On the BTCM interface, ECC checking can only be enabled for both ports or neither port. You can pin-configure the processor to set the enable bits and therefore enable ECC checking on reset, by tying off the **PARECCENRAMm** input as required.

When a fatal error, that is, a 2-bit ECC error, is detected on a TCM read, an error is generated. Instruction and data reads generate the appropriate type of synchronous abort, and the AXI-slave interface returns a SLVERR response to the AXI system.

When a correctable error, that is, a 1-bit ECC error, is detected on a TCM read made by the AXI-slave interface, the processor corrects the data inline before returning to the system.

When a correctable ECC error is detected on a TCM read made by the instruction-side or data-side, the processor normally generates the correct data and writes it back to the TCM. In the meantime, the processor retries the read to fetch the correct instruction or data. By setting the appropriate bits in the Secondary Auxiliary Control Register, you can disable this behavior. See *c15, Secondary Auxiliary Control Register* on page 4-44. Instead of correcting the error in the TCM, the processor generates the appropriate type of synchronous abort.

All ECC code generation and ECC checking must be performed on a complete data chunk, either 32-bits or 64-bits depending on the configuration. If a read access smaller than the data chunk is required, the whole chunk is read. If a write smaller than the data chunk is required, the processor must perform read-modify-write to generate the correct data and ECC code, but it only does this when ECC error checking is enabled. The data read as part of the read-modify-write sequence is checked for ECC errors, and the errors are handled in the same way as for any other TCM read. The ECC code is generated and written to the TCM for every write, regardless of whether error checking is enabled or not, but the code is only correct if the write was of a complete data chunk or if the processor performed read-modify-write to generate the complete data chunk. All data and instruction aborts generated by the ECC logic are indicated in the appropriate FSR as being a synchronous parity error.

8.4.4 TCM arbitration

Each TCM port receives requests from the LSU, PFU, and AXI slave. In most cases, the LSU has the highest priority, followed by the PFU, with the AXI slave having lowest priority.

When a higher-priority device is accessing a TCM port, an access from a lower-priority device must stall.

When either the LSU or the AXI slave interface is performing a read-modify-write operation on a TCM port, various internal data hazards exist for either the AXI-slave interface or the LSU. In these cases, additional stall cycles are generated, beyond those normally required for arbitration. For optimum performance of the processor when configured with ECC, ensure that all write bursts to the TCM from the AXI slave interface write an entire data chunk, that is, 32-bits or 64-bits, naturally aligned, depending on the error scheme.

8.4.5 TCM initialization

You can enable the processor to boot from the ATCM or the BTCM. The **INITRAMAm** and **INITRAMBm** pins, when tied HIGH, enable the ATCM and the BTCM respectively on leaving reset. The **LOCZRAMAm** pin forces one of the TCMs to have its base address at 0x0. If **LOCZRAMAm** is tied HIGH, the initial base address of the ATCM is 0x0, otherwise the initial base address of the BTCM is 0x0. In both cases, the initial base address of the other TCM is implementation-defined, see *Configurable options* on page 1-6.

The ATCM Region Register and BTCM Region Register respectively determine the base address for the ATCM and BTCM. For information on how to read the TCM region registers, see *c9, BTCM Region Register* on page 4-63 or *c9, ATCM Region Register* on page 4-64 as appropriate. For information about pre-loading data into the TCMs, see *TCM* on page 2-19.

8.4.6 TCM port protocol

Each TCM port operates independently to read and write data to and from the memory attached to it. Information about which memory location is to be accessed is passed on the TCM port along with write data and associated error code, if appropriate. In addition, the TCM port provides information about whether the access results from an instruction fetch from the PFU, a data access from the LSU, or a DMA transfer from the AXI slave interface. Each TCM port can also be configured to have an associated parity bit, computed from the address and control signals for that port.

Read data and associated error code or parity bits are read back from the TCM port. In addition, the TCM memory controller can indicate that the processor must wait one or more cycles before reading the response, or signal that an error has occurred and must be either aborted or retried. For more information about TCM errors, see *External TCM errors*.

8.4.7 External TCM errors

Each TCM port has a number of features that support the integration of a TCM RAM with an error checking scheme implemented in the RAM controller logic outside of the processor, that is, by the integrator.

Errors can be signaled to each TCM port if the external error checking scheme detects one and, if enabled, the processor generates an instruction or data abort or an AXI error response as appropriate. On a TCM read from either the instruction-side or data-side, the TCM controller can indicate that the read must be retried instead of generating an abort.

You can enable external errors for each TCM port individually by setting the appropriate bits in the Auxiliary Control Register. See *c1, Auxiliary Control Register* on page 4-41. If external errors are not enabled for a TCM port, the processor ignores any error signaled on that port. You can pin-configure the processor to set the enable bits, and therefore enable external error checking on reset, by tying off the **ERRENRAMm** input as required.

In addition, an external error detection scheme might require that data is read and written in particular sized chunks. The load/store-64 feature, when enabled for a particular TCM interface, causes all loads and stores to the TCM ports to be of 64-bits of data. This feature is also known as *Read-Modify-Write* (RMW), because it causes the processor to generate read-modify-write sequences for any store of less than 64-bits. You can enable RMW behavior for each TCM interface individually by setting the appropriate bits in the Secondary Auxiliary Control Register. See *c1, Auxiliary Control Register* on page 4-41. You can pin-configure the processor to set the enable bits and therefore RMW behavior on reset, by tying off the **RMWENRAMm** input as required.

Note

The load/store-64 feature is not available on any TCM interface that has been configured with 32-bit ECC.

The error inputs on each TCM port can also be used to signal other types of error, for example, when an address accessed is out of range for the RAM attached to the TCM port. Errors signaled on writes from the data-side generate an asynchronous abort. All other aborts generated by external errors are synchronous. The type of abort is shown in the appropriate FSR as either synchronous or asynchronous parity error.

8.4.8 AXI slave interfaces for TCMs

The processor has a 64-bit AXI slave interface that provides access to the TCM interfaces from the AXI bus. This interface is included by default, but can be excluded during configuration of the processor.

You can use the slave interface for access to the TCM memories. This also enables you to construct a system with a consistent view of memory. That is, the TCMs can be available at the same address to the processor and to the system bus.

The AXI slave interface accesses have lower priority than the LSU or PFU accesses.

The MPU does not check accesses from the AXI slave. You can configure the processor to enable privileged or nonprivileged access to the TCM interfaces from the AXI slave port.

The AXI slave interface does not support locked and exclusive accesses. This means that AXI masters, other than the processor, cannot safely use semaphores in the TCMs. Although the Cortex-R5 processor can use semaphores in the TCMs for inter-process synchronization, you must not use the AXI-slave interface to write to TCM semaphores. The processor has no logic to preserve its own exclusivity against such writes.

For more information on the AXI slave interface, see *AXI slave interface* on page 9-18.

8.5 About the caches

The L1 memory system can be configured to include instruction and data caches of varying sizes. You can configure whether each cache controller is included and, if it is, configure the size of each cache independently. The cached instructions or data are fetched from external memory using the AXI-master L2 memory interface. The cache controllers use RAMs that are integrated into the Cortex-R5 processor during implementation.

Any access that is not for a TCM or peripheral port is handled by the appropriate cache controller. If the access is to non-shared Cacheable memory, and the cache is enabled, a lookup is performed in the cache and, if found in the cache, that is, a cache hit, the data is fetched from or written into the cache. When the cache is not enabled and for Non-cacheable or shared memory, the accesses are performed using the AXI-master interface.

Both caches allocate a memory location to a cache line on a cache miss because of a read, that is, all Cacheable locations are *Read-Allocate* (RA). In addition, the data cache can allocate on a write access if the memory location is marked as *Write-Allocate* (WA). When a cache line is allocated, the appropriate memory is fetched into a linefill buffer by the AXI-master interface before being written to the cache. See *Linefill buffers and the AXI master interface* on page 9-5. The linefill buffers always fetch the requested data first, return it, and then fetch the rest of the cache line. This enables the data read to be used by the pipeline without waiting for the linefill to complete and is known as *critical word first* and *non-blocking* behavior. If subsequent instructions require data from the same cache line, this can also be returned when it has been fetched without waiting for the linefill to complete, that is, the caches also support *streaming*. If an error is reported to the AXI-master interface for a linefill, the linefill does not update the cache RAMs, but an abort is only generated if the error was reported on the critical word.

If all the cache lines in a set are valid, to allocate a different address to the cache, the cache controller must evict a line from the cache.

Writes accesses that hit in the cache are written into the cache RAMs. If the memory location is marked as *Write-Through* (WT), the write is also performed on the AXI-master interface, so that the data stored in the RAM remains coherent with the external memory system. If the memory is *Write-Back* (WB), the cache line is marked as dirty, and the write is only performed on the AXI-master interface when the line is evicted. When a dirty cache line is evicted, the data is passed to the Eviction Buffer in the AXI-master interface to be written to the external memory system. See *Eviction buffer* on page 9-6 for more information.

The cache controllers also manage the cache maintenance operations described in *Cache maintenance operations* on page 8-19.

Each cache can also be configured with either parity or ECC error checking schemes. If an error checking scheme is implemented and enabled, then the tags associated with each line, and data read from the cache are checked whenever a lookup is performed in the cache. See *Cache error detection and correction* on page 8-20 for more information.

For more information on the general rules about memory attributes and behavior, see the *ARM Architecture Reference Manual*.

8.5.1 Store buffer

The cache controller includes a store buffer to hold data before it is written to the cache RAMs or passed to the AXI master interface. The store buffer has four entries. Each entry can contain up to 64 bits of data and a 32-bit address. All write requests from the data-side that are not to a TCM or peripheral interface are stored in the store buffer.

Store buffer merging

The store buffer has merging capabilities. If a previous write access has updated an entry, other write accesses on the same line can merge into this entry. Merging is only possible for stores to Normal memory.

Merging is possible between several entries that can be linked together if the data inside the different entries belong to the same cache line.

No merging occurs for writes to Strongly Ordered or Device memory. The processor automatically drains the store buffer as necessary before performing Strongly Ordered accesses or Device reads.

Store buffer behavior

The store buffer directs write requests to the following blocks:

- Cache controller for Cacheable write hits:
The store buffer sends a cache lookup to check that the cache hits in the specified line, and if so, the store buffer merges its data into the cache when the entry is drained.
- AXI master interface:
 - For Non-cacheable stores or write-through Cacheable stores, a write access is performed on the AXI master interface.
 - For write-back, write-allocate stores that miss in the data cache, a linefill is started using either of the two linefill buffers. When the linefill data is returned from the L2 memory system, the data in the store buffer is merged into the linefill buffer to be subsequently written into the cache.

Store buffer draining

A store buffer entry is drained if:

- All bytes in the entry have been written. This might result from merging.
- The entry can be merged into a linefill buffer.
- The entry contains a store to Device or Strongly Ordered memory.
- The entry has been waiting for merge data for too long.

The store buffer is completely drained when:

- An explicit drain request is done for:
 - system control coprocessor cache maintenance operations
 - a DMB or DSB instruction
 - a load or store to Strongly Ordered memory
 - an exclusive load or store to Shared memory
 - a SWP or SWPB to Non-cacheable memory.
- The store buffer is full or likely to become full.

The store buffer is drained of all stores to Device memory before a load is performed from Device memory.

8.5.2 Cache maintenance operations

All cache maintenance operations are done through the system control coprocessor, CP15. The system control coprocessor operations supported for the data cache are:

- Invalidate all

- Invalidate by address (MVA)
- Invalidate by Set/Way combination
- Clean by address (MVA)
- Clean by Set/Way combination
- Clean and Invalidate by address (MVA)
- Clean and Invalidate by Set/Way combination
- *Data Memory Barrier (DMB)* and *Data Synchronization Barrier (DSB)* operations.

The system control coprocessor operations supported for the instruction cache are:

- Invalidate all
- Invalidate by address.

For more information on cache operations, see *Cache operations* on page 4-59.

8.5.3 Cache error detection and correction

This section describes how the processor detects, handles, reports, and corrects cache memory errors. Memory errors detected with parity or ECC have *Fault Status Register (FSR)* values to distinguish them from other abort causes.

This section describes:

- *Error build options*
- *Address decoder faults* on page 8-21
- *Handling cache parity errors* on page 8-21
- *Handling cache ECC errors* on page 8-22
- *Errors on instruction cache read* on page 8-23
- *Errors on data cache read* on page 8-23
- *Errors on data cache write* on page 8-23
- *Errors on evictions* on page 8-23
- *Errors on cache maintenance operations* on page 8-23.

Error build options

The caches can detect and correct errors depending on the build options used in the implementation. The build options for the instruction cache can be different to the data cache.

If the parity build option is enabled, the cache is protected by parity bits. For both the instruction and data cache, the data RAMs include one parity bit per byte of data. The tag RAM contains one parity bit to cover the tag and valid bit.

If the ECC build option is enabled:

- The instruction cache is protected by a 64-bit ECC scheme. The data RAMs include eight bits of ECC code for every 64 bits of data. The tag RAMs include seven bits of ECC code to cover the tag and valid bit.
- The data cache is protected by a 32-bit ECC scheme. The data RAMs include seven bits of ECC code for every 32 bits of data. The tag RAMs include seven bits of ECC code to cover the tag and valid bit. The dirty RAM includes four bits of ECC to cover the dirty bit and the two outer attributes bits of each cache line.

Address decoder faults

The error detection schemes described in this section provide protection against errors that occur in the data stored in the cache RAMs. Each RAM normally includes a decoder that enables access to that data and, if an error occurs in this logic, it is not normally detected by these error detection schemes. The processor includes features that enable it to detect some address decoder faults. If you are implementing the processor and require these features, contact ARM to discuss the features and your requirements.

Handling cache parity errors

Table 8-2 shows the behavior of the processor on a cache parity error, depending on bits [5:3] of the Auxiliary Control Register, see *c1, Auxiliary Control Register* on page 4-41.

Table 8-2 Cache parity error behavior

Value	Behavior
b000	Generate abort on parity errors ^a , force write-through, enable hardware recovery
b001	
b010	
b011	Reserved
b100	Disable parity checking
b101	Do not generate abort on parity errors, force write-through, enable hardware recovery
b110	
b111	Reserved

a. Parity errors caused by ACP coherency maintenance operations do not generate aborts

See *Disabling or enabling error checking* on page 8-32 for information on how to safely change these bits.

Hardware recovery

When parity checking is enabled, hardware recovery is always enabled. Memory marked as write-back write-allocate behaves as write-through. This ensures that cache lines can never be dirty, therefore the error can always be recovered from by invalidating the cache line that contains the parity error. The processor automatically performs this invalidation when an error is detected. The correct data can then be re-read from the L2 memory system.

Parity aborts

If aborts on parity errors are enabled, software is notified of the error by a data abort or prefetch abort. The error is still automatically corrected by the hardware even if an abort is generated.

If abort generation is not enabled, the hardware recovery including the access retry is invisible to software. If required, software can use events and the Correctable Fault Location Register to monitor the errors that are detected and corrected. See *Error detection events* on page 8-36 and *Correctable Fault Location Register* on page 4-77.

Parity errors, caused by ACP coherency maintenance operations, never generate aborts.

Handling cache ECC errors

Table 8-3 shows the behavior of the processor on a cache ECC error, depending on bits [5:3] of the Auxiliary Control Register, see *c1, Auxiliary Control Register* on page 4-41.

Table 8-3 Cache ECC error behavior

Value	Behavior
b000	Generate abort on ECC errors ^a , enable hardware recovery
b001	
b010	Generate abort on ECC errors ^a , force write-through, enable hardware recovery
b011	Reserved
b100	Disable ECC checking
b101	Do not generate abort on ECC errors, enable hardware recovery
b110	Do not generate abort on ECC errors, force write-through, enable hardware recovery
b111	Reserved

a. ECC errors caused by ACP coherency maintenance operations do not generate aborts

See *Disabling or enabling error checking* on page 8-32 for information on how to safely change these bits.

When ECC checking is enabled, hardware recovery is always enabled. When an ECC error is detected, the processor tries to evict the cache line containing the error. If the line is clean, it is invalidated, and the correct data is reloaded from the L2 memory system. If the line is dirty, the eviction writes the dirty data out to the L2 memory system, and in the process it corrects any 1-bit errors. The corrected data is then reloaded from the L2 memory system.

If a 2-bit error is detected in a dirty line, the error is not correctable. If the 2-bit error is in the tag or dirty RAM, no data is written to the L2 memory system. If the 2-bit error is in the data RAM, the cache line is written to the L2 memory system, but the AXI master port **WSTRBMM** signal is LOW for the data that contains the error. If an uncorrectable error is detected, an abort is always generated because data might have been lost. It is expected that such a situation can be fatal to the software process running.

If one of the force write-through settings is enabled, memory marked as write-back write-allocate behaves as write-through. This ensures that cache lines can never be dirty, therefore the error can always be recovered from by invalidating the cache line that contains the ECC error.

You can recover from all detectable errors in the instruction cache, because the instruction cache can never contain dirty data.

ECC aborts

If aborts on ECC errors are enabled, software is notified of the error by a data abort or prefetch abort. The error is still automatically corrected by the hardware even if an abort is generated.

If abort generation is not enabled, the hardware recovery including the access retry of correctable errors is invisible to software. If required, software can use events and the Correctable Fault Location Register to monitor the errors that are detected and corrected. See *Error detection events* on page 8-36 and *Correctable Fault Location Register* on page 4-77.

ECC errors, caused by ACP coherency maintenance operations, never generate aborts.

Errors on instruction cache read

All parity or ECC errors detected on instruction cache reads are correctable. If aborts are enabled, a synchronous prefetch abort exception occurs. The instruction FAR gives the address that caused the error to be detected. The instruction FSR indicates a parity error on a read. The auxiliary FSR indicates that the error was in the cache and which cache Way the error was in.

Errors on data cache read

If parity or ECC aborts are enabled, or an uncorrectable ECC error is detected, a synchronous data abort exception occurs. The data FAR gives the address that caused the error to be detected. The data FSR indicates a synchronous read parity error. The auxiliary FSR indicates that the error was in the cache and which cache Way the error was in.

Errors on data cache write

If parity or ECC aborts are enabled, or an uncorrectable ECC error is detected, an asynchronous data abort exception occurs. Because the abort is asynchronous, the data FAR is Unpredictable. The data FSR indicates an asynchronous write parity error. The auxiliary FSR indicates that the error was in the cache and which cache Way and Index the error was in.

In write-through cache regions the store that caused the error is written to external memory using the L2 memory interface so data is not lost and the error is not fatal.

Errors on evictions

If the cache controller has determined a cache miss has occurred, it might have to do an eviction before a linefill can take place. This can occur on reads, and on writes if write-allocation is enabled for the region. Certain cache maintenance operations also generate evictions. If it is a data-cache line that is dirty, an ECC error might be detected on the line being evicted:

- if the error is correctable, it is corrected inline before the data is written to the external memory using the L2 memory interface
- if there is an uncorrectable error in the tag or dirty RAM, the write is not done and an asynchronous abort occurs
- if there is an uncorrectable error in the data RAM, the AXI master port **WSTRBMm** signal is deasserted for the words with an error, and an asynchronous abort occurs.

An asynchronous abort can also occur on a correctable error depending on the Auxiliary Control Register bits [5:3], see *c1, Auxiliary Control Register* on page 4-41. Any detected error is signaled with the appropriate event.

———— Note ————

When parity checking is enabled, force write-through is always enabled. Therefore the cache lines can never be dirty, and so evictions are not required. Force write-through can also be enabled with ECC checking.

Errors on cache maintenance operations

The following sections describe errors on cache maintenance operations:

- *Invalidate all instruction cache* on page 8-24
- *Invalidate all data cache* on page 8-24
- *Invalidate instruction cache by address* on page 8-24
- *Invalidate data cache by address* on page 8-24

- *Invalidate data cache by set/way*
- *Clean data cache by address*
- *Clean data cache by set/way on page 8-25*
- *Clean and invalidate data cache by address on page 8-25*
- *Clean and invalidate data cache by set/way on page 8-25.*

Invalidate all instruction cache

This operation ignores all errors in the cache and sets all instruction cache entries to invalid regardless of error events. This operation cannot generate an asynchronous abort, and no error events are signaled.

Invalidate all data cache

This operation ignores all errors in the cache and sets all data cache entries to invalid regardless of errors. This operation cannot generate an asynchronous abort and no error events are signaled.

Invalidate instruction cache by address

This operation requires a cache lookup. Any errors found in the set that was looked up are fixed by invalidating that line and, if the address in question is found in the set, it is invalidated.

This operation cannot generate an asynchronous abort. Any detected error is signaled with the appropriate event.

Invalidate data cache by address

This operation requires a cache lookup. Any correctable errors found in the set that was looked up are fixed and, if the address in question is found in the set, it is invalidated.

Any uncorrectable errors cause an asynchronous abort. An asynchronous abort can also be raised on a correctable error if aborts on RAM errors are enabled in the Auxiliary Control Register.

Any detected error is signaled with the appropriate event.

Invalidate data cache by set/way

This operation does not require a cache lookup. It refers to a particular cache line.

The entry at the given set/way is marked as invalid regardless of any errors. This operation cannot generate an asynchronous abort. Any detected error is signaled with the appropriate event.

Clean data cache by address

This operation requires a cache lookup. Any correctable errors found in the set that was looked up are fixed and, if the address in question is found in the set, the instruction carries on with the clean operation. When the tag lookup is done, the dirty RAM is checked.

———— Note ————

When force write-through is enabled, the dirty bit is ignored.

If the tag or dirty RAM has an uncorrectable error, the data is not written to memory.

If the line is dirty, the data is written back to external memory. If the data has an uncorrectable error, the words with the error have their **WSTRBMM** AXI signal deasserted. If there is a correctable error, the line has the error corrected inline before it is written back to memory.

Any uncorrectable errors cause an asynchronous abort. An asynchronous abort can also be raised on a correctable error if aborts on RAM errors are enabled in the Auxiliary Control Register.

Any detected error is signaled with the appropriate event.

Clean data cache by set/way

This operation does not require a cache lookup. It refers to a particular cache line.

The tag and dirty RAMs for the cache line are checked.

Note

When force write-through is enabled, the dirty bit is ignored.

If the tag or dirty RAM has an uncorrectable error, the data is not written to memory.

If the line is dirty, the data is written back to external memory. If the data has an uncorrectable error, the words with the error have their **WSTRBMm** AXI signal deasserted. If there is a correctable error, the line has the error corrected inline before it is written back to memory.

Any uncorrectable errors found cause an asynchronous abort. An asynchronous abort can also be raised on a correctable error if aborts on RAM errors are enabled in the Auxiliary Control Register.

Any detected error is signaled with the appropriate event.

Clean and invalidate data cache by address

This operation requires a cache lookup. Any correctable errors found in the set that was looked up are fixed and, if the address in question is found in the set, the instruction carries on with the clean and invalidate operation. When the tag lookup is done, the dirty RAM is checked.

Note

When force write-through is enabled, the dirty bit is ignored.

If the tag or dirty RAM has an uncorrectable error, the data is not written to memory.

If the line is dirty, the data is written back to external memory. If the data has an uncorrectable error, the words with the error have their **WSTRBMm** AXI signal deasserted. If there is a correctable error, the line has the error corrected inline before it is written back to memory.

Any uncorrectable errors found cause an asynchronous abort. An asynchronous abort can also be raised on a correctable error if aborts on RAM errors are enabled in the Auxiliary Control Register.

Any detected error is signaled with the appropriate event.

Clean and invalidate data cache by set/way

This operation does not require a cache lookup. It refers to a particular cache line.

The tag and dirty RAMs for the cache line are checked.

Note

When force write-through is enabled, the dirty bit is ignored.

If the tag or dirty RAM has an uncorrectable error, the data is not written to memory.

If the line is dirty, the data is written back to external memory. If the data has an uncorrectable error, the words with the error have their **WSTRBMM** AXI signal deasserted. If there is a correctable error, the line has the error corrected inline before it is written back to memory.

Any uncorrectable errors found cause an asynchronous abort. An asynchronous abort can also be raised on a correctable error if aborts on RAM errors are enabled in the Auxiliary Control Register.

Any detected error is signaled with the appropriate event.

Errors on ACP coherency maintenance operations

Coherency maintenance operations are issued to the data cache controller when the ACP processes coherent write transactions. See *Accelerator Coherency Port interface* on page 9-48 for more information on the ACP.

These operations require data cache lookups. Any correctable errors found in the set that was looked up are fixed and, if the address is found in the set and not marked as dirty, it is invalidated.

Any detected error is signaled with the appropriate event.

8.5.4 Cache RAM organization

This section describes RAM organization in the following sections:

- *Tag RAM*
- *Dirty RAM* on page 8-27
- *Data RAM* on page 8-28.

Tag RAM

The tag RAMs consist of four ways of up to 512 lines. The width of the RAM depends on the build options selected, and the size of the cache. The following tables show the tag RAM bits:

- Table 8-4 shows the tag RAM bits when parity is implemented
- Table 8-5 shows the tag RAM bits when ECC is implemented
- Table 8-6 on page 8-27 shows the tag RAM bits when neither parity nor ECC is implemented.

Table 8-4 Tag RAM bit descriptions, with parity

Bit in the tag cache line	Description
Bit [23]	Parity bit
Bit [22]	Valid bit
Bits [21:0]	Tag value

Table 8-5 Tag RAM bit descriptions, with ECC

Bit in the tag cache line	Description
Bits [29:23]	ECC code bits
Bit [22]	Valid bit
Bits [21:0]	Tag value

Table 8-6 Tag RAM bit descriptions, no parity or ECC

Bit in the tag cache line	Description
Bit [22]	Valid bit
Bits [21:0]	Tag value

A cache line is marked as valid by bit [22] of the tag RAM. Each valid bit is associated with a whole cache line, so evictions always occur on the entire line.

Table 8-7 shows the tag RAM cache sizes and associated RAM organization, assuming no parity or ECC. For parity, the width of the tag RAMs must be increased by one bit. For ECC, the width of the tag RAMs must be increased by seven bits.

Table 8-7 Cache sizes and tag RAM organization

Cache size	Tag RAM organization
4KB	4 banks 23 bits 32 lines
8KB	4 banks 22 bits 64 lines
16KB	4 banks 21 bits 128 lines
32KB	4 banks 20 bits 256 lines
64KB	4 banks 19 bits 512 lines

Dirty RAM

For the data cache only, the dirty RAM stores the following information:

- two bits for line outer attributes for evictions
- one line dirty bit
- four ECC code bits if the ECC build option is selected.

The dirty RAM array consists of one bank of up to 512 12-bit lines, 4 ways x 3 bits. If ECC is enabled, the dirty RAM is 28 bits wide. Each line of dirty RAM contains all the information of the four ways for a given index.

Each time a dirty bit is written, the outer bits of the line and, if implemented, the ECC code bits, are also written. The dirty RAM is bit-enabled. Table 8-8 shows the organization of a dirty RAM line.

Table 8-8 Organization of a dirty RAM line

Bit in the dirty cache line	Description
Bits [6:3]	ECC bits, if implemented
Bits [2:1]	Outer attributes that are re-encoded on AWCACHEDMm when an eviction is sent to the AXI bus: 01 = WB, WA 10 = WT 11 = WB, no WA 00 = Non-cacheable.
Bit [0]	Dirty bit

Data RAM

Data RAM is organized as eight banks of 32-bit wide lines, or in the instruction cache as four banks of 64-bit wide lines. This RAM organization means that it is possible to:

- Perform a cache look-up with one RAM access, all banks selected together. This is done for nonsequential read operations. Figure 8-3 shows this.
- Select the appropriate bank RAM for sequential read operations. Figure 8-4 shows this.
- Write a line to the eviction buffer in one cycle, a 256-bit read access.
- Fill a line in one cycle from the linefill buffer, a 256-bit write access.

Figure 8-3 shows a cache look-up being performed on all banks with one RAM access.




Figure 8-3 Nonsequential read operation performed with one RAM access.

Figure 8-4 shows the appropriate bank RAM being selected for a sequential read operation.




Figure 8-4 Sequential read operation performed with one RAM access

The data RAM organization is optimized for 64-bit read operations, because with the same address, two words on the same way can be selected.

Data RAM sizes depend on the build option selected, and are described in:

- *Data RAM sizes without parity or ECC implemented* on page 8-29
- *Data RAM sizes with parity implemented* on page 8-29
- *Data RAM sizes with ECC implemented* on page 8-30.

Data RAM sizes without parity or ECC implemented

Table 8-9 shows the organization for instruction and data caches when neither parity nor ECC is implemented.

Table 8-9 Instruction cache data RAM sizes, no parity or ECC

Cache size	Data RAMs
4KB, 4 1KB ways	4 banks 64 bits 128 lines or 8 banks 32 bits 128 lines
8KB, 4 2KB ways	4 banks 64 bits 256 lines or 8 banks 32 bits 256 lines
16KB, 4 4KB ways	4 banks 64 bits 512 lines or 8 banks 32 bits 512 lines
32KB, 4 8KB ways	4 banks 64 bits 1024 lines or 8 banks 32 bits 1024 lines
64KB, 4 16KB ways	4 banks 64 bits 2048 lines or 8 banks 32 bits 2048 lines

Table 8-10 Data cache data RAM sizes, no parity or ECC

Cache size	Data RAMs
4KB, 4 1KB ways	8 banks 32 bits 128 lines
8KB, 4 2KB ways	8 banks 32 bits 256 lines
16KB, 4 4KB ways	8 banks 32 bits 512 lines
32KB, 4 8KB ways	8 banks 32 bits 1024 lines
64KB, 4 16KB ways	8 banks 32 bits 2048 lines

Data RAM sizes with parity implemented

Table 8-11 shows the organization for instruction and data caches when parity is implemented. For parity error detection, one bit is added per byte, so four bits are added for each RAM bank.

Table 8-11 Instruction cache data RAM sizes, with parity

Cache size	Data RAMs
4KB, 4 1KB ways	4 banks 72 bits 128 lines or 8 banks 36 bits 128 lines
8KB, 4 2KB ways	4 banks 72 bits 256 lines or 8 banks 36 bits 256 lines
16KB, 4 4KB ways	4 banks 72 bits 512 lines or 8 banks 36 bits 512 lines
32KB, 4 8KB ways	4 banks 72 bits 1024 lines or 8 banks 36 bits 1024 lines
64KB, 4 16KB ways	4 banks 72 bits 2048 lines or 8 banks 36 bits 2048 lines

Table 8-12 Data cache data RAM sizes, with parity

Cache size	Data RAMs
4KB, 4 1KB ways	8 banks 36 bits 128 lines
8KB, 4 2KB ways	8 banks 36 bits 256 lines
16KB, 4 4KB ways	8 banks 36 bits 512 lines
32KB, 4 8KB ways	8 banks 36 bits 1024 lines
64KB, 4 16KB ways	8 banks 36 bits 2048 lines

Table 8-13 shows the organization of the data cache RAM bits when parity is implemented.

Table 8-13 Data cache RAM bits, with parity

RAM bits	Description
Bit [35]	Parity bit for byte[31:24]
Bit [34]	Parity bit for byte[23:16]
Bit [33]	Parity bit for byte[15:8]
Bit [32]	Parity bit for byte[7:0]
Bits [31:0]	Data[31:0]

Parity bits are grouped together in bits[35:32] so that data and parity bits are easily differentiated. With this design the parity bit is selected alongside the related data byte, so that when data is updated, the parity bit is also updated.

Data RAM sizes with ECC implemented

Table 8-14 shows the organization for the instruction cache when ECC is implemented. For ECC error detection, eight bits are added per 64 bits, so four bits are added for each RAM bank.

Table 8-14 Instruction cache data RAM sizes with ECC

Cache size	Data RAMs
4KB, 4 1KB ways	4 banks 72 bits 128 lines or 8 banks 36 bits 128 lines
8KB, 4 2KB ways	4 banks 72 bits 256 lines or 8 banks 36 bits 256 lines
16KB, 4 4KB ways	4 banks 72 bits 512 lines or 8 banks 36 bits 512 lines
32KB, 4 8KB ways	4 banks 72 bits 1024 lines or 8 banks 36 bits 1024 lines
64KB, 4 16KB ways	4 banks 72 bits 2048 lines or 8 banks 36 bits 2048 lines

Table 8-15 shows the organization for the data cache when ECC is implemented. For ECC error detection, seven bits are added per 32 bits, so seven bits are added for each RAM bank.

Table 8-15 Data cache data RAM sizes with ECC

Cache size	Data RAMs
4KB, 4 1KB ways	8 banks 39 bits 128 lines
8KB, 4 2KB ways	8 banks 39 bits 256 lines
16KB, 4 4KB ways	8 banks 39 bits 512 lines
32KB, 4 8KB ways	8 banks 39 bits 1024 lines
64KB, 4 16KB ways	8 banks 39 bits 2048 lines

Table 8-16 shows the organization of the data cache RAM bits when ECC is implemented.

Table 8-16 Data cache RAM bits, with ECC

RAM bits	Description
Bits [39:32]	ECC code bits for data [31:0]
Bits [31:0]	Data [31:0]

8.5.5 Cache interaction with memory system

This section describes how to enable or disable the cache RAMs, and to enable or disable error checking. After you enable or disable the instruction cache, you must issue an ISB instruction to flush the pipeline. This ensures that all subsequent instruction fetches see the effect of enabling or disabling the instruction cache.

After reset, you must invalidate each cache before enabling it.

When disabling the data cache, you must clean the entire cache to ensure that any dirty data is flushed to L2 memory.

Before enabling the data cache, you must invalidate the entire data cache if L2 memory might have changed since the cache was disabled.

Before enabling the instruction cache, you must invalidate the entire instruction cache if L2 memory might have changed since the cache was disabled.

See *Enabling or disabling AXI slave accesses* on page 9-21 and *Accessing RAMs using the AXI slave interface* on page 9-21 for information about how to access the cache RAMs using the AXI slave interface.

Disabling or enabling all of the caches

The following code is an example of enabling caches:

```
MRC p15, 0, r1, c1, c0, 0 ; Read System Control Register configuration data
ORR r1, r1, #0x1 <<12    ; instruction cache enable
ORR r1, r1, #0x1 <<2     ; data cache enable
DSB
MCR p15, 0, r0, c15, c5, 0 ; Invalidate entire data cache
MCR p15, 0, r0, c7, c5, 0  ; Invalidate entire instruction cache
MCR p15, 0, r1, c1, c0, 0 ; enabled cache RAMs
ISB
```

The following code is an example of disabling the caches:

```
MRC p15, 0, r1, c1, c0, 0 ; Read System Control Register configuration data
BIC r1, r1, #0x1 <<12    ; instruction cache disable
BIC r1, r1, #0x1 <<2      ; data cache disable
DSB
MCR p15, 0, r1, c1, c0, 0 ; disabled cache RAMs
ISB
; Clean entire data cache. This routine depends on the data cache size. It can be
omitted if it is known that the data cache has no dirty data
```

Disabling or enabling instruction cache

The following code is an example of enabling the instruction cache:

```
MRC p15, 0, r1, c1, c0, 0 ; Read System Control Register configuration data
ORR r1, r1, #0x1 <<12    ; instruction cache enable
MCR p15, 0, r0, c7, c5, 0 ; Invalidate entire instruction cache
MCR p15, 0, r1, c1, c0, 0 ; enabled instruction cache
ISB
```

The following code is an example of disabling the instruction cache:

```
MRC p15, 0, r1, c1, c0, 0 ; Read System Control Register configuration data
BIC r1, r1, #0x1 <<12    ; instruction cache enable
MCR p15, 0, r1, c1, c0, 0 ; disabled instruction cache
ISB
```

Disabling or enabling data cache

The following code is an example of enabling the data cache:

```
MRC p15, 0, r1, c1, c0, 0 ; Read System Control Register configuration data
ORR r1, r1, #0x1 <<2      ; data cache enable
DSB
MCR p15, 0, r0, c15, c5, 0 ; Invalidate entire data cache
MCR p15, 0, r1, c1, c0, 0 ; enabled data cache
```

The following code is an example of disabling the cache RAMs:

```
MRC p15, 0, r1, c1, c0, 0 ; Read System Control Register configuration data
BIC r1, r1, #0x1 <<2      ; data cache disable
DSB
MCR p15, 0, r1, c1, c0, 0 ; disabled data cache
; Clean entire data cache. This routine depends on the data cache size. It can be
omitted if it is known that the data cache has no dirty data.
```

Disabling or enabling error checking

Software must take care when changing the error checking bits in the Auxiliary Control Register. If the bits are changed when the caches contain data, the parity or ECC bits in the caches might not be correct for the new setting, resulting in unexpected errors and data loss. Therefore the bits in the Auxiliary Control Register must only be changed when both caches are turned off and the entire cache must be invalidated after the change.

The following code is the recommended sequence to perform the change:

```
MRC p15, 0, r0, c1, c0, 0 ; Read System Control Register
BIC r0, r0, #0x1 << 2      ; Disable data cache bit
BIC r0, r0, #0x1 << 12     ; Disable instruction cache bit
DSB
MCR p15, 0, r0, c1, c0, 0 ; Write System Control Register
ISB ; Ensures following instructions are not executed from cache
```



```

; Clean entire data cache. This routine depends on the data cache size. It can be
omitted if it is known that the data cache has no dirty data, for example if the cache
has not been enabled yet.
MRC p15, 0, r1, c1, c0, 1 ; Read Auxiliary Control Register
; Change bits 5:3 as required
MCR p15, 0, r1, c1, c0, 1 ; Write Auxiliary Control Register
MCR p15, 0, r0, c15, c5, 0 ; Invalidate entire data cache
MCR p15, 0, r0, c7, c5, 0 ; Invalidate entire instruction cache
MRC p15, 0, r0, c1, c0, 0 ; Read System Control Register
ORR r0, r0, #0x1 << 2 ; Enable data cache bit
ORR r0, r0, #0x1 << 12 ; Enable instruction cache bit
DSB
MCR p15, 0, r0, c1, c0, 0 ; Write System Control Register
ISB

```

8.6 Internal exclusive monitor

The processor L1 memory system has an internal exclusive monitor. This is a two state, open and exclusive, state machine that manages load/store exclusive (LDREXB, LDREXH, LDREX, LDREXD, STREXB, STREXH, STREX and STREXD) accesses and clear exclusive (CLREX) instructions. You can use these instructions, operating in the L1 memory system, to construct semaphores and ensure synchronization between different processes. By adding an external exclusive monitor, you can also use these instructions in the L2 memory system to construct semaphores and ensure synchronization between different processors. See the *ARM Architecture Reference Manual* for more information about how these instructions work.

When a load-exclusive access is performed, the internal exclusive monitor moves to the exclusive state. It moves back to the open state when a store exclusive access or clear exclusive instruction is performed. The internal exclusive monitor holds exclusivity state for an individual Cortex-R5 CPU only. It does not record the address of the memory that a load-exclusive access was performed to and it does not observe accesses from the other CPU in a twin-CPU group. Any store exclusive access performed when the state is open fails. If the state is exclusive, the access passes if it is to non-shared memory but, if it is to shared memory, the access must be performed as an exclusive using the L2 memory interface. Whether the shared store-exclusive access passes or fails depends on the state of an external exclusive monitor that can track accesses made by other processors in the system.

8.7 Memory types and L1 memory system behavior

The behavior of the L1 memory system depends on the type attribute of the memory that is being accessed:

- Only Normal, non-shared memory regions can be cached in the RAMs. Caching only takes place if the appropriate cache is enabled and the memory type is *Cacheable*.
- The store buffer can merge any stores to Normal memory. See *Store buffer* on page 8-18 for more information.
- Only Normal memory is considered restartable, that is, a multi-word transfer can be abandoned part way through because of an interrupt, to be restarted after the interrupt has been handled. See *Interrupts* on page 3-19 for more information about interrupt behavior.
- Only the internal exclusive monitor is used for exclusive accesses to Non-shared memory. Exclusive accesses to shared memory are checked using the internal monitor and also, if necessary, any external monitor, using the L2 memory interface.
- Accesses resulting from SWP and SWPB instructions to Normal, non-shared memory are not marked as locked when performed using the L2 memory interface.

Note

Not all types of exclusive or swap access are permitted to peripheral interface regions. See *Semaphores* on page 9-48.

Table 8-17 summarizes the processor memory types and associated behavior.

Table 8-17 Memory types and associated behavior

Memory type		Can be cached	Merging	Restartable	Internal exclusives	Locked swaps
Normal	Shared	No	Yes	Yes	Partially	Yes
	Non-shared	Yes	Yes	Yes	Yes	No
Device	Shared	No	No	No	Partially	Yes
	Non-shared	No	No	No	Yes	Yes
Strongly Ordered	Shared	No	No	No	Partially	Yes

8.8 Error detection events

The processor generates a number of events related to the internal error detection and correction schemes in the TCMs and caches. For more information, see Table 6-1 on page 6-2. This section describes:

- *TCM error events*
- *Instruction-cache error events*
- *Data-cache error events*
- *Events and the CFLR.*

8.8.1 TCM error events

TCM ECC error events are only signaled for TCM reads, although this includes the read-modify-write sequence performed for some stores. Most errors detected by the ECC logic are signaled twice:

- once on a TCM-centric event
- once on a processor-centric event.

The TCM-centric events consist of two events per TCM port, one for fatal, that is, 2-bit ECC errors and one for correctable, that is, 1-bit ECC errors. These events are generated three clock cycles after the data read cycle. Consequently, these events are sometimes signaled on speculative TCM reads, such as instructions that are prefetched but never executed because of a branch earlier in the instruction sequence.

Note

When an external error is signaled on a TCM access, the TCM-centric events are still generated as appropriate, based on the data returned, as if no external error had been signaled.

The processor-centric TCM events are only signaled for errors in data that would have otherwise been used by the processor. Errors on speculative reads never generate these errors. They consist of fatal and correctable events for:

- the prefetch unit, to signal errors on instruction fetches
- the load/store unit, to signal errors on data accesses
- the AXI slave interface, to signal errors on DMA accesses.

8.8.2 Instruction-cache error events

All parity and ECC errors are correctable in the I-Cache. Therefore there are only two events, to indicate when an error is detected in a read from the tag RAM, or from the data RAM. These events are only signaled for non-speculative instruction fetches and certain cache maintenance operations. See *Cache error detection and correction* on page 8-20.

8.8.3 Data-cache error events

The D-Cache can generate fatal and correctable errors, and therefore has four events, one for each type of error in the data RAM and in the tag or dirty RAMs. These events are only signaled for non-speculative data accesses, cache line evictions, coherency maintenance operations, and certain cache maintenance operations. See *Cache error detection and correction* on page 8-20.

8.8.4 Events and the CFLR

The *Correctable Fault Location Register* (CFLR) records the location of the last correctable error detected on a non-speculative access or coherency maintenance operations. See *Correctable Fault Location Register* on page 4-77 for more information. Every correctable error

that is recorded in the CFLR also generates an event. See Table 6-1 on page 6-2 to see which events are CFLR-related. For correctable cache errors, the CLFR does not record whether the error occurred in the data RAM or tag/dirty RAM. This distinction is only made by the events.

Chapter 9

Level Two Interface

This chapter describes the features of the Level two (L2) interface not covered in the *AMBA AXI Protocol Specification*. It contains the following sections:

- *About the L2 interface* on page 9-2
- *AXI master interface* on page 9-4
- *AXI master interface transfers* on page 9-7
- *AXI slave interface* on page 9-18
- *Enabling or disabling AXI slave accesses* on page 9-21
- *Accessing RAMs using the AXI slave interface* on page 9-21
- *Peripheral interfaces* on page 9-31
- *Accelerator Coherency Port interface* on page 9-48.

9.1 About the L2 interface

This section describes the processor L2 interface. The L2 interface consists of:

- AXI master interface
- AXI slave interface
- three peripheral interfaces
- ACP interface.

The processor is designed for use in larger chip designs using the *Advanced Microcontroller Bus Architecture* (AMBA) AXI and AHB protocols. Instruction fetches and data accesses that the L1 memory system does not service, and peripheral accesses, are performed through the AXI-master interface or one of the peripheral interfaces. See:

- *AXI master interface* on page 9-4 for more information about the AXI master interface
- *AXI peripheral port transfers* on page 9-35 for more information about the AXI peripheral interface
- *AHB peripheral port transfers* on page 9-42 for more information about the AHB peripheral interface.

External AXI masters, that can include the processor itself, can use the AXI slave interface to access the processor RAMs. You can use the AXI slave interface for DMA access into and out of the TCMs or to perform software test of the cache RAMs. See *AXI slave interface* on page 9-18.

The ACP interface enables the Cortex-R5 processor to observe memory transactions that other AXI masters perform, and keep the L1 caches coherent with those transactions. See *Accelerator Coherency Port interface* on page 9-48 for more information about the ACP interface.

You can configure all of the ports associated with the L2 interfaces with bus-ECC. The bus-ECC feature uses additional signals to communicate redundant information, enabling the detection or correction of errors that occur on the bus signals. See *Bus ECC* for more information.

9.1.1 Bus ECC

You can configure a Cortex-R5 processor with bus ECC to protect the integrity of AMBA bus signals. The bus ECC feature of the Low Latency Peripheral Port is configured separately from the other bus interfaces.

Bus ECC uses both parity and *Single Error Correct Double Error Detect* (SEC-DED) *Error Correcting Codes* (ECC). The Cortex-R5 processor computes and checks parity bits as odd or even, depending on the value of the **PARITYLEVEL** primary input, except for AXI handshake signals that have fixed, odd parity.

ECC and parity errors, detected by the Cortex-R5 processor, do not directly cause aborts, exceptions or otherwise affect the CPU operation. Instead, event primary outputs notify the system of correctable or fatal errors. The CPU treats all bus control and response signals as correct, even if parity errors are reported. It is possible that fatal, that is double-bit, ECC errors might cause more data corruption. This can result in the CPU operating on corrupted data, or behaving unpredictably, based on corrupted control or response signals.

Bus ECC functionality checks for errors on every bus transfer the CPU performs. This can include speculative accesses for which data is later discarded. The CPU:

- reports bus faults for all transfers whose data it uses
- never reports bus data faults for transfers where the bus master sees an error response.

AXI Interfaces

The Cortex-R5 processor uses the following scheme to protect AXI signals:

- Fixed, odd parity on channel **VALID** and **READY** signals.
- Parity on address and control payload signals. Each parity bit protects a maximum of eight payload bits.
- SEC-DED ECC to protect read and write data payload.

AHB Interfaces

The Cortex-R5 processor uses the following scheme to protect AHB signals:

- Parity on address and control signals. Each parity bit protects a maximum of eight payload bits
- SEC-DED ECC to protect data payload.

Debug APB Interface

Bus ECC is not available for this interface.

Notifications

The Cortex-R5 bus ECC feature provides the following notifications:

- Correctable errors on read data received by the AXI Master and Peripheral Port, through primary outputs.
- Correctable errors on write data received by the AXI Slave through a primary output.
- Logical address of transfers with correctable errors on master ports, to doubleword granularity.
- Memory chip select and logical address of transfers with correctable errors on the AXI Slave, to doubleword granularity.
- Fatal errors on AXI ports using one primary output bit per channel per port.
- Fatal errors on the AHB Peripheral Port through a primary output.
- A correctable bus fault event in the event bus, **EVNTBUSm**. See *About the events* on page 6-2.
- A fatal bus fault event in the event bus, **EVNTBUSm**. See *About the events* on page 6-2.
- Increments to correctable and fatal bus fault event counters for the *Performance Monitoring Unit* (PMU). See *About the events* on page 6-2.

Concurrent Bus Fault Events

The Cortex-R5 event bus and PMU logic monitors bus fault events on all Cortex-R5 AXI and AHB interfaces simultaneously. It merges bus faults that occur in the same CPU clock cycle, on different bus interfaces. For example, if correctable errors occur on both the AXI master and AXI slave, in the same CPU clock cycle, only one event is logged.

Bus Master Correctable Error Address Reporting

The Cortex-R5 processor has one primary output for reporting the logical address of a transfer with a correctable error, on the AXI master port, or the AXI and AHB Peripheral Ports. Concurrent correctable bus faults on the AXI master and the Peripheral Port cause the address to be reported for the AXI master only. Correctable errors do not occur concurrently on the AHB and AXI Peripheral Ports, see *Peripheral interfaces* on page 9-31 for more information about the Cortex-R5 Peripheral Port.

9.2 AXI master interface

The processor has a single AXI master interface, with one port that is used for:

- I_Cache linefills
- D_Cache linefills and evictions
- *Non-cacheable* (NC) Normal-type memory instruction fetches
- NC Normal-type memory data accesses
- Device and Strongly-ordered type data accesses, normally to peripherals.

The port is 64 bits wide, and conforms to the AXI3 standard as described in the *AMBA AXI Protocol Specification*. Within the AXI standard, the master port uses a number of extension signals to indicate inner memory attributes and, if configured with bus-ECC, parity or ECC information. See *AXI extensions* on page 9-6 for more information about attribute encodings and *Bus ECC* on page 9-2 for more information about bus-ECC.

The master interface can run at the same frequency as the processor or at a lower synchronous frequency. See *AMBA interface clocking* on page 2-16 for more information.

———— Note ————

References in this section to an *AXI slave* refer to the AXI slave in the external system that is connected to the Cortex-R5 AXI master port. This is not necessarily the Cortex-R5 AXI slave port.

The following sections describe the attributes of the AXI master interface, and provide information about the types of burst generated:

- *Identifiers for AXI bus accesses* on page 9-5
- *Write response* on page 9-5
- *Linefill buffers and the AXI master interface* on page 9-5
- *Eviction buffer* on page 9-6
- *AXI extensions* on page 9-6.
- *Memory system implications for AXI accesses* on page 9-7.

Table 9-1 shows the AXI master interface attributes.

Table 9-1 AXI master interface attributes

Attribute	Value	Comments
Write issuing capability	4	Made up of four outstanding writes that can be evictions, single writes, or write bursts. ^a
Read issuing capability	7	Made up of five linefills on the data side, one NC read on the data side, and one read on the instruction side, that can be NC or linefill.
Combined issuing capability	11 ^a	-

Table 9-1 AXI master interface attributes (continued)

Attribute	Value	Comments
Write ID capability	2	-
Write interleave capability	1	The AXI master interface presents all write data in order.
Read ID capability	7	Made up of five linefills on the data side, one NC read on the data side, and one linefill or NC read on the instruction side.

- a. When there are three outstanding write transactions, only data is issued for the fourth. Only three outstanding write addresses are issued.

9.2.1 Identifiers for AXI bus accesses

Accesses on the AXI bus use ID values as follows:

Outstanding write/read access on different IDs

This means, for example, that a *Non-cacheable* (NC) read and linefills can be outstanding on the AXI bus simultaneously as long as the IDs are different.

At the same time, there can be:

- up to seven outstanding reads, each with one of seven different ID values, that consists of:
 - a data side read NC access, RID0
 - an instruction side read NC access or an instruction side read Cacheable access, RID1
 - five outstanding data side linefills on the AXI bus, RID3 - RID7.
- up to two IDs on outstanding writes, that consist of:
 - single or burst NC writes or *write-through* (WT) writes, WID0
 - evictions, WID1.

Outstanding write accesses with the same ID

When the address and data of the first write are both put on AXI bus, another write request with same ID can be sent when the address or data channel is released. For example, the new address can be sent with the same ID, before the target accepts the data of the first write.

———— Note ————

- The AXI master does not generate two outstanding read accesses with the same ID.
- The AXI master does not interleave write data from two different bursts, even if the bursts have different IDs.

9.2.2 Write response

The AXI master requires that the slave does not return a write response until it has received both the write data and the write address.

9.2.3 Linefill buffers and the AXI master interface

On the data side there are two *LineFill Buffers* (LFBs), LFB0 and LFB1. Each request from the data cache controller or from the *STore Buffer* (STB) can be allocated to either LFB0 or LFB1.

On the instruction side, there is one LFB. This is the *Instruction LFB* (ILFB), that treats instruction linefill requests or Non-cacheable instruction reads in the same way.

The linefill buffers:

- get returned data from the AXI bus for linefill requests
- get returned data from the AXI bus for any Non-cacheable LDR or LDMs
- get data from the STB to write as a burst on the AXI bus (LFB0 and LFB1 only).

Single writes do not use LFBs.

The LFBs are 256 bits wide so that an entire cache line can be written to the cache RAMs in one cycle. While the LFB is being filled from L2 memory, its bytes can be merged with write data from the STB.

9.2.4 Eviction buffer

As soon as a linefill is requested, the selected evicted cache line is loaded into the *Eviction Buffer* (EVB). The EVB forwards this information to the AXI bus when possible.

The EVB has a structure of 256 bits for data and 32 bits for the address. See *Cache line write-back (eviction)* on page 9-12 for more information about the AXI transaction generated.

The EVB is removed if cache RAMs are not implemented for the processor.

9.2.5 AXI extensions

The Cortex-R5 AXI master interface uses the **ARCACHEMm**, **AWCACHEMm**, AXI signals and the **ARSHAREMm**, **AWSHAREMm**, **ARINNERMm**, and **AWINNERMm** extension signals to indicate the memory attributes of the transfer, as returned by the MPU. Table 9-2 shows the encodings used for these signals. **ARCACHEMm** and **AWCACHEMm** of the master interface are generated from the memory type and outer region attributes. **ARINNERMm** and **AWINNERMm** are generated from the memory type and inner region attributes. **ARSHAREMm** and **AWSHAREMm** are asserted for transactions to shared memory regions.

Table 9-2 ARCACHEMm, AWCACHEMm, ARINNERMm, and AWINNERMm encodings

Encoding ^a	Meaning
b0000	Strongly Ordered
b0001	Device
b0011	Non-cacheable
b0110	Cacheable, write-through, allocate on reads only
b0111	Cacheable, write-back, allocate on reads only
b1111	Cacheable write-back, allocate on reads and writes

a. All encodings not shown in the table are reserved.

Additional AXI extension signals on all the AXI master channels are used for bus-ECC and parity information.

9.2.6 Memory system implications for AXI accesses

The attributes of the memory being accessed can affect an AXI access. The L1 memory system can cache any Normal memory address that is marked as either:

- Cacheable, write-back, read- and write-allocate, non-shared
- Cacheable, write-through, read-allocate only, non-shared.

However, Device and Strongly Ordered memory is always Non-cacheable. Also, any unaligned access to Device or Strongly Ordered memory generates an alignment fault and therefore does not cause any AXI transfer. This means that the access examples given in this chapter never show unaligned accesses to Device or Strongly Ordered memory.

9.3 AXI master interface transfers

The processor conforms to the AXI3 specification, but it does not generate all the AXI transaction types that the specification permits. This section describes the types of AXI transaction that the Cortex-R5 AXI master does not generate. If you are designing an AXI slave to work only with the Cortex-R5 processor, and there are no other AXI masters in your system, you can take advantage of these restrictions and the interface attributes, described in Table 9-1 on page 9-4, to simplify the slave.

This section also contains tables that show some examples of the types of AXI burst that the processor generates. However, because a particular type of transaction is not shown here does not mean that the processor does not generate such a transaction.

———— Note ————

An AXI slave device connected to the Cortex-R5 AXI master port must be capable of handling every kind of transaction permitted by the AXI specification, except where there is an explicit statement in this chapter that such a transaction is not generated. You must not infer any additional restrictions from the example tables given. Restrictions described here apply to the r0p0 to r1p1 revisions of the processor, but might not be true for future revisions.

Load and store instructions to Non-cacheable memory might not result in an AXI transfer because the data might either be retrieved from, or merged into the internal store data buffers. The exceptions to this are loads or stores to Strongly Ordered or Device memory. These always result in AXI transfers. See *Strongly Ordered and Device transactions* on page 9-8.

Restrictions on AXI transfers on page 9-8 describes restrictions on the type of transfers that the Cortex-R5 AXI master interface generates. If a CPUM exists and is powered up, the buffered write response and read data channel ready signals, **BREADYMm** and **RREADYMm**, are always asserted. They are, however, deasserted when the CPU enters Dormant or Shutdown mode. You must not make any other assumptions about the AXI handshaking signals, except that they conform to the *AMBA AXI Protocol Specification*.

The following sections give examples of transfers generated by the AXI master interface:

- *Restrictions on AXI transfers* on page 9-8
- *Strongly Ordered and Device transactions* on page 9-8
- *Linefills* on page 9-12
- *Cache line write-back (eviction)* on page 9-12
- *Non-cacheable reads* on page 9-12
- *Non-cacheable or write-through writes* on page 9-14
- *AXI transaction splitting* on page 9-15
- *Normal write merging* on page 9-16.

9.3.1 Restrictions on AXI transfers

The Cortex-R5 AXI master interface applies the following restrictions to the AXI transactions it generates:

- A burst never transfers more than 32 bytes.
- The burst length is never more than 8 transfers.
- No transaction ever crosses a 32-byte boundary in memory. See *AXI transaction splitting* on page 9-15.
- FIXED bursts are never used.
- The write address channel always issues INCR type bursts, and never WRAP or FIXED.
- WRAP type read bursts, see *Linefills* on page 9-12:
 - are used only for linefills (reads) of Cacheable Normal non-shared memory
 - always have a size of 64 bits, and a length of 4 transfers
 - always have a start address that is 64-bit aligned.
- If the transfer size is 8 bits or 16 bits then the burst length is always 1 transfer.
- The transfer size is never greater than 64 bits, because it is a 64-bit AXI bus.
- Instruction fetches, identified by **ARPROT[2]**, are always a 64 bit transfer size, and never locked or exclusive.
- Transactions to Device and Strongly Ordered memory are always to addresses that are aligned for the transfer size. See *Strongly Ordered and Device transactions*.
- Exclusive and Locked accesses are always to addresses that are aligned for the transfer size.
- Write data is never interleaved.
- In addition to these restrictions, there are various limitations to the ID values that the AXI master interface uses. See *Identifiers for AXI bus accesses* on page 9-5.

9.3.2 Strongly Ordered and Device transactions

A load or store instruction to or from Strongly Ordered or Device memory always generates AXI transactions of the same size as implied by the instruction. All accesses using LDM, STM, LDRD, or STRD instructions to Strongly Ordered or Device memory occur as 32-bit transfers.

LDRB

Table 9-3 shows the values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for a Non-cacheable LDRB from bytes 0-7 in Strongly Ordered or Device memory.

Table 9-3 Non-cacheable LDRB

Address[2:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x0 (byte 0)	0x00	Incr	8-bit	1 data transfer
0x1 (byte 1)	0x01	Incr	8-bit	1 data transfer
0x2 (byte 2)	0x02	Incr	8-bit	1 data transfer
0x3 (byte 3)	0x03	Incr	8-bit	1 data transfer

Table 9-3 Non-cacheable LDRB (continued)

Address[2:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x4 (byte 4)	0x04	Incr	8-bit	1 data transfer
0x5 (byte 5)	0x05	Incr	8-bit	1 data transfer
0x6 (byte 6)	0x06	Incr	8-bit	1 data transfer
0x7 (byte 7)	0x07	Incr	8-bit	1 data transfer

LDRH

Table 9-4 shows the values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for a Non-cacheable LDRH from halfwords 0-3 in Strongly Ordered or Device memory.

Table 9-4 LDRH from Strongly Ordered or Device memory

Address[2:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x0 (halfword 0)	0x00	Incr	16-bit	1 data transfer
0x2 (halfword 1)	0x02	Incr	16-bit	1 data transfer
0x4 (halfword 2)	0x04	Incr	16-bit	1 data transfer
0x6 (halfword 3)	0x06	Incr	16-bit	1 data transfer

Note

A load of a halfword from Strongly Ordered or Device memory addresses 0x1, 0x3, 0x5, or 0x7 generates an alignment fault.

LDR or LDM that transfers one register

Table 9-5 shows the values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for a Non-cacheable LDR or an LDM that transfers one register, (an LDM1) in Strongly Ordered or Device memory.

Table 9-5 LDR or LDM1 from Strongly Ordered or Device memory

Address[2:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x0 (word 0)	0x00	Incr	32-bit	1 data transfer
0x4 (word 1)	0x04	Incr	32-bit	1 data transfer

Note

A load of a word from Strongly Ordered or Device memory addresses 0x1, 0x2, 0x3, 0x5, 0x6, or 0x7 generates an alignment fault.

LDM that transfers five registers

Table 9-6 shows the values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for a Non-cacheable LDM that transfers five registers (an LDM5) in Strongly Ordered or Device memory.

Table 9-6 LDM5, Strongly Ordered or Device memory

Address[3:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x0 (word 0)	0x00	Incr	32-bit	5 data transfers
0x4 (word 1)	0x04	Incr	32-bit	5 data transfers
0x8 (word 2)	0x08	Incr	32-bit	5 data transfers
0xC (word 3)	0x0C	Incr	32-bit	5 data transfers

Note

A load-multiple from address 0x1, 0x2, 0x3, 0x5, 0x6, 0x7, 0x9, 0xA, 0xB, 0xD, 0xE, or 0xF generates an alignment fault.

STRB

Table 9-7 shows the values of **AWADDRMm**, **AWBURSTMm**, **AWSIZEMm**, and **AWLENMm** for an STRB to Strongly Ordered or Device memory over the AXI master port.

Table 9-7 STRB to Strongly Ordered or Device memory

Address[2:0]	AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x0 (byte 0)	0x00	Incr	8-bit	1 data transfer	b00000001
0x1 (byte 1)	0x01	Incr	8-bit	1 data transfer	b00000010
0x2 (byte 2)	0x02	Incr	8-bit	1 data transfer	b00000100
0x3 (byte 3)	0x03	Incr	8-bit	1 data transfer	b00001000
0x4 (byte 4)	0x04	Incr	8-bit	1 data transfer	b00010000
0x5 (byte 5)	0x05	Incr	8-bit	1 data transfer	b00100000
0x6 (byte 6)	0x06	Incr	8-bit	1 data transfer	b01000000
0x7 (byte 7)	0x07	Incr	8-bit	1 data transfer	b10000000

STRH

Table 9-8 shows the values of **AWADDRMm**, **AWBURSTMm**, **AWSIZEMm**, and **AWLENMm** for an STRH over the AXI master port to Strongly Ordered or Device memory.

Table 9-8 STRH to Strongly Ordered or Device memory

Address[2:0]	AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x0 (halfword 0)	0x00	Incr	16-bit	1 data transfer	b00000011
0x2 (halfword 1)	0x02	Incr	16-bit	1 data transfer	b00001100
0x4 (halfword 2)	0x04	Incr	16-bit	1 data transfer	b00110000
0x6 (halfword 3)	0x06	Incr	16-bit	1 data transfer	b11000000

———— Note ————

A store of a halfword to Strongly Ordered or Device memory addresses 0x1, 0x3, 0x5, or 0x7 generates an alignment fault.

STR or STM of one register

Table 9-9 shows the values of **AWADDRMm**, **AWBURSTMm**, **AWSIZEMm**, and **AWLENMm** for an STR or an STM that transfers one register (an STM1) over the AXI master port to Strongly Ordered or Device memory.

Table 9-9 STR or STM1 to Strongly Ordered or Device memory

Address[2:0]	AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x0 (word0)	0x00	Incr	32-bit	1 data transfer	b00001111
0x4 (word 1)	0x04	Incr	32-bit	1 data transfer	b11110000

———— Note ————

A store of a word to Strongly Ordered or Device memory addresses 0x1, 0x2, 0x3, 0x5, 0x6, or 0x7 generates an alignment fault.

STM of seven registers

Table 9-10 shows the values of **AWADDRMm**, **AWBURSTMm**, **AWSIZEMm**, and **AWLENMm** for an STM that writes seven registers (an STM7) over the AXI master port to Strongly Ordered or Device memory.

Table 9-10 STM7 to Strongly Ordered or Device memory to word 0 or 1

Address[4:0]	AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	First WSTRBMm
0x00 (word 0)	0x00	Incr	32-bit	7 data transfers	b00001111
0x04 (word 1)	0x04	Incr	32-bit	7 data transfers	b11110000

———— Note ————

A store-multiple to address 0x1, 0x2, 0x3, 0x5, 0x6, or 0x7 generates an alignment fault.

9.3.3 Linefills

Loads and instruction fetches from Normal, Cacheable memory that do not hit in the cache generate a cache linefill when the appropriate cache is enabled. Table 9-11 shows the values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for cache linefills.

Table 9-11 Linefill behavior on the AXI interface

Address[4:0] ^a	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x00-0x07	0x00	Wrap	64-bit	4 data transfers
0x08-0x0F	0x08	Wrap	64-bit	4 data transfers
0x10-0x17	0x10	Wrap	64-bit	4 data transfers
0x18-0x1F	0x18	Wrap	64-bit	4 data transfers

a. These are the bottom five bits of the address of the access that cause the linefill, that is, the address of the critical word.

9.3.4 Cache line write-back (eviction)

When a valid and dirty cache line is evicted from the d-cache, a write-back of the data must occur. Table 9-12 shows the values of **AWADDRMm**, **AWBURSTMm**, **AWSIZEMm**, and **AWLENMm** for cache line write-backs, over the AXI master interface.

Table 9-12 Cache line write-back

AWADDRMm[4:0]	AWBURSTMm	AWSIZEMm	AWLENMm
0x00	Incr	64-bit	4 data transfers

9.3.5 Non-cacheable reads

Load instructions accessing Non-cacheable Normal memory generate AXI bursts that are not necessarily the same size or length as the instruction implies. In addition, if the data to be read is contained in the store buffer, the instruction might not generate an AXI read transaction at all.

The tables in this section give examples of the types of AXI transaction that might result from various load instructions, accessing various addresses in Non-cacheable Normal memory. They are provided as examples only, and are not an exhaustive description of the AXI transactions. Depending on the state of the processor, and the timing of the accesses, the actual bursts generated might have a different size and length to the examples shown, even for the same instruction.

Table 9-13 shows possible values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for an LDRH from bytes 0-7 in Non-cacheable Normal memory.

Table 9-13 LDRH from Non-cacheable Normal memory

Address[2:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x0 (byte 0)	0x00	Incr	16-bit	1 data transfer
0x1 (byte 1)	0x00	Incr	32-bit	1 data transfer
0x2 (byte 2)	0x00	Incr	64-bit	1 data transfer
0x3 (byte 3)	0x03	Incr	32-bit	2 data transfers

Table 9-13 LDRH from Non-cacheable Normal memory (continued)

Address[2:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x4 (byte 4)	0x04	Incr	16-bit	1 data transfer
0x5 (byte 5)	0x04	Incr	32-bit	1 data transfer
0x6 (byte 6)	0x06	Incr	16-bit	1 data transfer
0x7 (byte 7)	0x07	Incr	32-bit	2 data transfers

Table 9-14 shows possible values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for a Non-cacheable LDR or an LDM that transfers one register, an LDM1.

Table 9-14 LDR or LDM1 from Non-cacheable Normal memory

Address[2:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x0 (byte 0) (word 0)	0x00	Incr	32-bit	1 data transfer
0x1 (byte 1)	0x01	Incr	64-bit	1 data transfer
0x2 (byte 2)	0x00	Incr	64-bit	1 data transfer
0x3 (byte 3)	0x00	Incr	64-bit	2 data transfers
0x4 (byte 4) (word 1)	0x04	Incr	32-bit	1 data transfer
0x5 (byte 5)	0x05	Incr	32-bit	2 data transfers
0x6 (byte 6)	0x06	Incr	16-bit	1 data transfer
	0x08	Incr	16-bit	1 data transfer
0x7 (byte 7)	0x04	Incr	32-bit	2 data transfers

Table 9-15 show possible values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for a Non-cacheable LDM that transfers five registers (an LDM5).

Table 9-15 LDM5, Non-cacheable Normal memory or cache disabled

Address[4:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x00 (word 0)	0x00	Incr	64-bit	3 data transfers
0x04 (word 1)	0x04	Incr	64-bit	3 data transfers
0x08 (word 2)	0x08	Incr	64-bit	3 data transfers
0x0C (word 3)	0x0C	Incr	64-bit	3 data transfers
0x10 (word 4)	0x10	Incr	64-bit	2 data transfers
	0x00	Incr	32-bit	1 data transfer
0x14 (word 5)	0x14	Incr	64-bit	2 data transfers
	0x00	Incr	64-bit	1 data transfer

Table 9-15 LDM5, Non-cacheable Normal memory or cache disabled (continued)

Address[4:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x18 (word 6)	0x18	Incr	64-bit	1 data transfer
	0x00	Incr	64-bit	2 data transfers
0x1C (word 7)	0x1C	Incr	32-bit	1 data transfer
	0x00	Incr	64-bit	2 data transfers

9.3.6 Non-cacheable or write-through writes

Store instructions to Non-cacheable or write-through Normal memory generate AXI bursts that are not necessarily the same size or length as the instruction implies. The AXI master port asserts byte-lane-strobes, **WSTRBMm[7:0]**, to ensure that only the bytes that were written by the instruction are updated.

The tables in this section give examples of the types of AXI transaction that might result from various store instructions, accessing various addresses in Non-cacheable Normal memory. They are provided as examples only, and are not an exhaustive description of the AXI transactions. Depending on the state of the processor, and the timing of the accesses, the actual bursts generated might have a different size and length to the examples shown, even for the same instruction.

In addition, write operations to Normal memory can be merged to create more complex AXI transactions. See *Normal write merging* on page 9-16 for examples.

Table 9-16 shows possible values of **AWADDRMm**, **AWBURSTMm**, **AWSIZEMm**, and **AWLENMm** for an STRH to Normal memory.

Table 9-16 STRH to Cacheable write-through or Non-cacheable Normal memory

Address[2:0]	AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x0 (byte 0)	0x00	Incr	32-bit	1 data transfer	b00000011
0x1 (byte 1)	0x00	Incr	32-bit	1 data transfer	b00000110
0x2 (byte 2)	0x02	Incr	64-bit	1 data transfer	b00001100
0x3 (byte 3)	0x03	Incr	32-bit	2 data transfers	b00001000 b00010000
0x4 (byte 4)	0x04	Incr	16-bit	1 data transfer	b00110000
0x5 (byte 5)	0x05	Incr	32-bit	1 data transfer	b01100000
0x6 (byte 6)	0x06	Incr	16-bit	1 data transfer	b11000000
0x7 (byte 7)	0x07	Incr	8-bit	1 data transfer	b10000000
	0x08	Incr	8-bit	1 data transfer	b00000001

Table 9-17 shows possible values of **AWADDRMm**, **AWBURSTMm**, **AWSIZEMm**, and **AWLENMm** for an STR or an STM that transfers one register, an STM1, to Normal memory through the AXI master port.

Table 9-17 STR or STM1 to Cacheable write-through or Non-cacheable Normal memory

Address[2:0]	AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x0 (byte 0) (word 0)	0x00	Incr	32-bit	1 data transfer	b00001111
0x1 (byte 1)	0x01	Incr	64-bit	1 data transfer	b00011110
0x2 (byte 2)	0x00	Incr	64-bit	1 data transfer	b00111100
0x3 (byte 3)	0x03	Incr	64-bit	2 data transfers	b01111000 b00000000
0x4 (byte 4) (word 1)	0x04	Incr	32-bit	1 data transfer	b11110000
0x5 (byte 5)	0x05	Incr	32-bit	2 data transfers	b11100000 b00000001
0x6 (byte 6)	0x06	Incr	16-bit	1 data transfer	b11000000
	0x08	Incr	16-bit	1 data transfer	b00000011
0x7 (byte 7)	0x04	Incr	32-bit	2 data transfers	b10000000 b00000111

9.3.7 AXI transaction splitting

The processor splits AXI bursts when it accesses addresses across a cache line boundary, that is, a 32-byte boundary. An instruction that accesses memory across one or two 32-byte boundaries generates two or three AXI bursts respectively. The following examples show this behavior. They are provided as examples only, and are not an exhaustive description of the AXI transactions. Depending on the state of the processor, and the timing of the accesses, the actual bursts generated might have a different size and length to the examples shown, even for the same instruction.

For example, `LDMIA R10, {R0-R5}` loads six words from memory. The number of AXI transactions generated by this instruction depends on the base address, `R10`:

- If all six words are in the same cache line, there is a single AXI transaction. For example, for `LDMIA R10, {R0-R5}` with `R10 = 0x1008`, the interface might generate a burst of three, 64-bit read transfers, as shown in Table 9-18.

Table 9-18 AXI transaction splitting, all six words in same cache line

ARADDRMm	ARBURSTMm	ARsizEMm	ARLENMm
0x1008	Incr	64-bit	3 data transfers

- If the data comes from two cache lines, then there are two AXI transactions. For example, for LDMIA R10, {R0-R5} with R10 = 0x1010, the interface might generate one burst of two 64-bit reads, and one burst of a single 64-bit read, as shown in Table 9-19.

Table 9-19 AXI transaction splitting, data in two cache lines

ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x1010	Incr	64-bit	2 data transfers
0x1020	Incr	64-bit	1 data transfer

Table 9-20 shows possible values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for an LDR or LDM1 to Non-cacheable Normal memory that crosses a cache line boundary.

Table 9-20 Non-cacheable LDR or LDM1 crossing a cache line boundary

Address[4:0]	ARADDRMm	ARBURSTMm	ARSIZEMm	ARLENMm
0x1D (byte 29)	0x1C	Incr	32-bit	1 data transfer
	0x00	Incr	32-bit	1 data transfer
0x1E (byte 30)	0x1E	Incr	16-bit	1 data transfer
	0x00	Incr	64-bit	1 data transfer
0x1F (byte 31)	0x1F	Incr	8-bit	1 data transfer
	0x00	Incr	32-bit	1 data transfer

Table 9-21 shows possible values of **ARADDRMm**, **ARBURSTMm**, **ARSIZEMm**, and **ARLENMm** for an STRH to Non-cacheable Normal memory that crosses a cache line boundary.

Table 9-21 Cacheable write-through or Non-cacheable STRH crossing a cache line boundary

Address[4:0]	AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x1F (byte 31)	0x1F	Incr	8-bit	1 data transfer	b10000000
	0x00	Incr	16-bit	1 data transfer	b00000001

9.3.8 Normal write merging

A store instruction to Non-cacheable, or write-through Normal memory might not result in an AXI transfer because of the merging of store data in the internal buffers.

The STB can detect when it contains more than one write request to the same cache line for write-through Cacheable or Non-cacheable Normal memory. This means it can combine the data from more than one instruction into a single write burst to improve the efficiency of the AXI port. If the AXI master receives several write requests that do not form a single contiguous burst it can choose to output a single burst, with the **WSTRBW** signal low for the bytes that do not have any data.

For write accesses to Normal memory, the STB can perform writes out of order, if there are no address dependencies. It can do this to best use its ability to merge accesses.

The instruction sequence in Example 9-1 on page 9-17 shows the merging of writes.

Example 9-1 Write merging

```

MOV r0, #0x4000
STRH r1, [r0, #0x18]; Store a halfword at 0x4018
STR  r2, [r0, #0xC] ; Store a word at 0x400C
STMIA r0, {r4-r7}    ; Store four words at 0x4000
STRB r3, [r0, #0x1D]; Store a byte at 0x401D

```

If the memory at address 0x4000 is marked as Strongly Ordered or Device type memory, the AXI transactions shown in Table 9-22 are generated.

Table 9-22 AXI transactions for Strongly Ordered or Device type memory

AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x4018	Incr	16-bit	1 data transfer	0b00000011
0x400C	Incr	32-bit	1 data transfer	0b11110000
0x4000	Incr	32-bit	4 data transfers	0b00001111 0b11110000 0b00001111 0b11110000
0x401D	Incr	8-bit	1 data transfer	0b00100000

In Example 9-1, each store instruction produces an AXI burst of the same size as the data written by the instruction.

Table 9-23 shows a possible resulting transaction if the same memory is marked as Non-cacheable Normal, or Cacheable write-through.

Table 9-23 AXI transactions for Non-cacheable Normal or Cacheable write-through memory

AWADDRMm	AWBURSTMm	AWSIZEMm	AWLENMm	WSTRBMm
0x4000	Incr	64-bit	4 data transfers	0b11111111 0b11111111 0b00000000 0b00100011

In this example:

- The store buffer has merged the STRB and STRH writes into one buffer entry, and therefore a single AXI transfer, the fourth in the burst.
- The writes, that occupy three buffer entries, have been merged into a single AXI burst of four transfers.
- The write generated by the STR instruction has not occurred, because it was overwritten by the STM instruction.
- The write transfers have occurred out of order with respect to the original program order.

The transactions shown in Table 9-23 on page 9-17 show this behavior. They are provided as examples only, and are not an exhaustive description of the AXI transactions. Depending on the state of the processor, and the timing of the accesses, the actual bursts generated might have a different size and length to the examples shown, even for the same instruction.

If the same memory is marked as write-back Cacheable, and the addresses are allocated into a cache line, no AXI write transactions occur until the cache line is evicted and performs a write-back transaction. See *Cache line write-back (eviction)* on page 9-12.

9.4 AXI slave interface

The processor has a single AXI slave interface, with one port. The port is 64 bits wide and conforms to the AXI3 standard as described in the *AMBA AXI Protocol Specification*. Within the AXI standard, the slave port uses the extension signals **AWCSELSm** and **ARCSELSm** each as four separate chip select input signals to enable access to:

- BTCM
- ATCM
- instruction cache RAMs
- data cache RAMs.

The external AXI system must generate the chip select signals. The slave interface routes the access to the required RAM.

If the processor is configured with bus-ECC, extension signals are also used for parity and ECC information. See *Bus ECC* on page 9-2 for more information about bus-ECC.

The slave interface can run at the same frequency as the processor or at a lower, synchronous frequency. See *AMBA interface clocking* on page 2-16 for more information. If asynchronous clocking is required, then an external asynchronous AXI register slice is required.

The AXI slave provides access to the TCMs and competes for access to the TCMs with the LSU and PFU. Both the LSU and PFU normally have a higher priority than the AXI slave.

If two BTCM ports are used, you can configure these to interleave in the address map, so any AXI slave access that is denied access to the BTCM on the first cycle of the access gains access on the second cycle when the LSU is using the other port, and can continue in lock-step with the LSU, assuming both are accessing sequential data. Accesses to the ATCM are more likely to encounter a conflict because there is only one port on the interface.

Memory BIST ports are routed through the AXI slave interface logic, to access the RAMs. Memory BIST access is assumed only to occur when no other accesses are taking place, and takes highest priority.

9.4.1 AXI slave interface for cache RAMs

———— Note ————

You must not use the AXI slave to access the cache RAMs at the same time as the ACP. Ensure the ACP is idle before initiating AXI slave transactions to the cache RAMs.

You can use the AXI slave for software testing of the cache RAMs in functional mode. When the AXI slave is enabled to access the RAMs, the processor considers the caches as cache-off, so that the instruction and data requests cannot interact with AXI slave requests. In this state, only AXI slave requests can access the cache RAM and instruction and data requests from the processor are considered as non-cacheable and do not perform any lookup in the caches.

The AXI slave interface accesses each cache RAM individually.

On the instruction cache side the AXI slave can access:

- data cache RAMs, data and parity or ECC code bits
- tag RAMs, tag and valid, and parity or ECC code bits.

On the data cache side, the AXI slave can access:

- data cache RAMs, data and parity or ECC code bits
- tag RAMs, tag and valid, and parity or ECC code bits
- dirty RAM, dirty bit and attributes, and ECC code bits.

A simple decode of four address bits and four way address bits determines which of the data, tag, or dirty RAMs is accessed within the caches. The AXI access is given a SLVERR error response when access to nonexistent cache RAM is indicated.

9.4.2 TCM ECC support

The TCMs can support ECC, as described in *TCM internal error detection and correction* on page 8-14. If a write transaction is issued to the AXI slave, the slave interface calculates the required ECC bits to store to the TCM. If the write data width is smaller than the ECC chunk size then a read-modify-write sequence is automatically performed by the AXI slave.

———— Note ————

It is important to ensure that all writes to TCMs that do not contain the correct ECC bits for their data, such as uninitialized RAMs, are performed with a size of at least the ECC chunk size or with error checking disabled.

If a read transaction is issued to the AXI slave, the slave interface reads the ECC bits and, if error checking is enabled for the appropriate TCM, checks the data for errors. If the interface detects a correctable error, it corrects it inline and returns the correct data on the AXI bus. It does not update the data in the TCM to correct it. If the interface detects an uncorrectable error, it generates a SLVERR error response to the AXI transaction.

9.4.3 External TCM errors

If an error response is given to a TCM access from the AXI slave interface, and external errors are enabled for the appropriate TCM port, the AXI slave returns a SLVERR response to the AXI transaction.

The AXI slave ignores late-error and retry responses from the TCM.

9.4.4 Cache parity and ECC support

When the caches support parity or ECC, the AXI slave interface permits direct read and write access to the parity or ECC code bits. No errors are detected automatically, and on writes the AXI slave does not automatically generate the correct parity or ECC code values.

———— Note ————

The AXI slave interface provides read/write access to the cache RAMs for functional test. It is not suitable for preloading the caches.

9.4.5 AXI slave control

By default, both privileged and non-privileged accesses can be made to the Cortex-R5 TCM RAMs through the AXI slave port. To disable non-privileged accesses, you can set bit [1] in the Slave Port Control Register. You can disable all slave accesses by setting bit [0] of the register. See *c11, Slave Port Control Register* on page 4-65.

Access to the cache RAMs can only be made when bit [24] of the Auxiliary Control Register is set. By default, only privileged accesses can be made to the cache RAMs, but you can enable non-privileged accesses by setting bit [23] of the Auxiliary Control Register. When cache RAM access is enabled, both caches are treated as if they were not enabled. See *c1, Auxiliary Control Register* on page 4-41.

The AXI access is given a SLVERR error response when access is not permitted.

9.4.6 AXI slave characteristics

This section describes the capabilities of the AXI slave interface, and the attributes of its AXI port. You must not make any other assumptions about the behavior of the AXI slave port except that it conforms to the *AMBA AXI 3 Protocol Specification*.

- The AXI slave interface supports merging of data within bursts. When handling an AXI burst of data less than 64-bits wide, the AXI slave interface attempts to perform the minimum number of TCM or cache accesses required to read or write the data. When an ECC error scheme is in use, this sometimes reduces the number of read-modify-write sequences that the AXI slave must perform.
- The AXI slave interface does not support:
 - Security Extensions, all accesses are secure, so **AxPROT**[1] is not used
 - data and instruction transaction signaling, so **AxPROT**[2] is not used
 - memory type and cacheability, so **AxCACHE** is not used
 - atomic accesses. The AXI slave accepts locked transactions but makes no use of the locking information, that is, **AxLOCK**.
- The AXI slave interface has no exclusive access monitor. If there are any exclusive accesses, the AXI slave interface responds with an OKAY response.
- The width of the ID signals for the AXI slave port is 8 bits.

You must avoid building the processor into an AXI system that requires more than 8 bits of ID. The number of bits of ID required by a system can often be reduced by compressing the encoding to remove unused values. The AXI master port does not use all possible values. See *Identifiers for AXI bus accesses* on page 9-5 for more information.

Table 9-24 shows the AXI slave port attributes.

Table 9-24 AXI slave interface attributes

Attribute	Value	Comments
Combined acceptance capability	7	-
Write interleave depth	1	All write data must be presented to the AXI slave interface in order
Read data reorder depth	1	The AXI slave interface returns all read data in order, even if the bursts have different IDs

9.5 Enabling or disabling AXI slave accesses

This section describes how to enable or disable AXI slave accesses to the cache RAMs. When caches are accessible by the AXI slave interface, the caches are considered to be cache-off from the processor. You must ensure that the ACP is idle so that it does not generate any cache RAM accesses. After turning the interface on or off, an ISB instruction must flush the pipeline so that all subsequent instruction fetches return valid data.

The following code is an example of enabling AXI slave accesses to the cache RAMs:

```
MRC p15, 0, R1, c1, c0, 1 ; Read Auxiliary Control Register
ORR R1, R1, #0x1 <<24
; Ensure ACP is idle, that is. cannot access the cache and that no new ACP transactions
; can be generated
DSB
MCR p15, 0, R1, c1, c0, 1 ; enabled AXI slave accesses to the cache RAMs
ISB
; Clean entire data cache. This routine depends on the data cache size. It can be
; omitted if it is known that the data cache has no dirty data
Fetch from uncached memory
Fetch from uncached memory
Fetch from uncached memory
Fetch from uncached memory
```

The following code is an example of disabling AXI slave accesses to the cache RAMs. No cache invalidation is performed because it is assumed that, after accessing the cache RAMs, the AXI slave interface restored the previously valid data to them.

```
MRC p15, 0, R1, c1, c0, 1 ; Read Auxiliary Control Register
BIC R1, R1, #0x1 <<24
DSB
MCR p15, 0, R1, c1, c0, 1 ; disabled AXI slave accesses to the cache RAMs
ISB
; Re-enable ACP transactions
Fetch from cached memory
Fetch from cached memory
Fetch from cached memory
Fetch from cached memory
```

9.6 Accessing RAMs using the AXI slave interface

This section describes how to access the TCM and cache RAMs using the AXI slave interface.

Table 9-25 shows the bits of the **ARCSELSm** or **AWCSELSm** inputs, that determine the target of a transaction. Each signal is a one-hot 4-bit input, with each bit corresponding to a particular RAM or group of RAMs.

Table 9-25 RAM region decode

AxCSELSm bit	One-hot RAM select
[3]	Data cache RAMs
[2]	Instruction cache RAMs
[1]	B0TCM and B1TCM
[0]	ATCM

The remaining addressing information is encoded in **ARADDRSm[22:0]** for reads and **AWADDRSm[22:0]** for writes. The AXI-slave interface does not use the other bits of the address, **ARADDRSm[31:23]** and **AWADDRSm[31:23]**, except for the purposes of bus-ECC. For more information see:

- *TCM RAM access*
- *Cache RAM access* on page 9-23.

Note

Because **AWCSELSm** and **AWADDRSm** are similar to **ARCSELSm** and **ARADDRSm**, the following sections describe their common features as **AxCSELSm** and **AxADDRSm**, noting any differences between them.

9.6.1 TCM RAM access

AxADDRSm[22:3] indicates the address of the doubleword within the TCM that you want to access. If you are accessing a TCM that is smaller than the maximum 8MB, then it is possible to describe an address that is outside of the physical size of the TCM. This is not permitted and results in a SLVERR error response.

Table 9-26 shows the decode of the **AxCSELSm[3:0]** signal, and the state of the address signals for accessing different TCM RAMs. The table also shows the **SLBTCMSBm** configuration input signal that determines which address bit is used to select between the banks of a dual-banked BTCM.

Table 9-27 shows the most significant bit of the address for the different TCM RAM sizes. For split BTCMs, the TCM size is defined to be the total size of both the B0TCM and B1TCM combined. In this situation, the particular BTCM accessed is dependent on either **AxADDRSm[MSB]**, if the input **SLBTCMSBm** is high, or **AxADDRSm[3]** otherwise. For example, if there are split BTCMs and **SLBTCMSBm** is LOW and **AxADDRSm[3]** is HIGH, the access goes to the B1TCM.

Table 9-26 TCM chip-select decode

AxCSELSm[3:0]	BTCM ports	SLBTCMSBm	AxADDRSm[3]	AxADDRSm[MSB]	RAM selected
0001	-	-	-	-	ATCM
0010	1	-	-	-	BTCM
0010	2	0	0	-	B0TCM
0010	2	0	1	-	B1TCM
0010	2	1	-	0	B0TCM
0010	2	1	-	1	B1TCM

Table 9-27 MSB bit for the different TCM RAM sizes

TCM size	AxADDRSm[MSB]
4KB	[11]
8KB	[12]
16KB	[13]
32KB	[14]

Table 9-27 MSB bit for the different TCM RAM sizes (continued)

TCM size	AxADDRSm[MSB]
64KB	[15]
128KB	[16]
256KB	[17]
512KB	[18]
1MB	[19]
2MB	[20]
4MB	[21]
8MB	[22]

An access to the TCM RAMs is given a SLVERR error response if:

- It is outside the physical size of the targeted TCM RAM, that is, bits of **AxADDRSm[22:MSB+1]** are non-zero.
- There is no TCM present. The mapping of bus addresses to **AxCSELSm** and **AxADDRSm** is determined when the processor is integrated. You must understand this mapping to use of the AXI-slave interface within your system.

9.6.2 Cache RAM access

This section contains the following:

- *Memory map when accessing the cache RAMs*
- *D_Cache data RAM single bank accesses* on page 9-27
- *I_Cache Data RAM access* on page 9-28
- *D_Cache data RAM double bank accesses* on page 9-28
- *Tag RAM access* on page 9-30
- *Dirty RAM access* on page 9-30.

Memory map when accessing the cache RAMs

The memory map is divided into 2 regions:

- RAM-Access region
- TRANSFER and AUX register access region

The TRANSFER register enables an AXI master to construct a single RAM access from multiple sub-word accesses, that might be required if the master data width is less than the RAM data width.

The AUX register provides access to Data RAM ECC and parity data, if implemented.

The RAM-Access region initiates all AXI-slave RAM accesses. Reads from this region return data and update the TRANSFER and AUX registers. Writes to this region combine with the data in the TRANSFER and AUX registers, before being committed to the RAM.

Table 9-28 on page 9-24 describes the RAM-Access memory map and Table 9-29 on page 9-25 describes the TRANSFER and AUX memory map.

Any address that is not listed in Table 9-28 or addresses that are explicitly listed as illegal returns a SLVERR.

Any fields marked as RAZ/WI refer to the RAMs. The TRANSFER and AUX registers are not guaranteed to be RAZ/WI.

Table 9-28 RAM-Access space

AxADDRSm bits	Description
[22:19]	Block select: 0000 = single bank data RAM 0001 = tag RAM 0010 = dirty RAM^a 0100 = double bank data RAM^a 1000 = strobed double bank data RAM^a.
[18:15]	Bank select. For D-Cache data RAMs: - Single bank mode. Accesses a single RAM-word: 0001 = Bank 0 or 1 0010 = Bank 2 or 3 0100 = Bank 4 or 5 1000 = Bank 6 or 7 Bit [13] of the address determines which of the two banks is selected for each of these values. 0 = lower numbered bank 1 = higher numbered bank - Double bank mode. Accesses 2 RAM-words from contiguous banks: 0001 = Bank 0 and 1 0010 = Bank 2 and 3 0100 = Bank 4 and 5 1000 = Bank 6 and 7 - Strobed double bank mode. Accesses 2 contiguous banks with byte-strobe support: 0001 = Bank 0 and 1 0010 = Bank 2 and 3 0100 = Bank 4 and 5 1000 = Bank 6 and 7 Strobes used to access the selected banks are derived directly from the WSTRB signals for the access to the RAM-Access space. For I_Cache data RAMs: 0001 = Bank 0 0010 = Bank 1 0100 = Bank 2 1000 = Bank 3 For tag^{bc} and dirty^{ad} RAMs: [15] = Bank 0 [16] = Bank 1 [17] = Bank 2 [18] = Bank 3. For tag-RAM reads, only one-hot encodings are supported. For tag-RAM writes, all combinations are supported, the same data is written to all banks. For dirty RAM accesses, all combinations are supported.

Table 9-28 RAM-Access space (continued)

AxADDRSm bits	Description
[14]	Indicates address space accessed: 0 = RAM-Access
[13:2]	For D_Cache double-bank data RAM accesses: [13:3] = RAM index [2] = bank select For D_Cache single-bank RAM accesses: [13] = bank select [12:2] = RAM index For I_Cache data RAM accesses: [13:3] = RAM index [2] = word select For all other accesses:[13:12] = 0x0[11:3] = RAM index [2] = word select
[1:0]	Byte select

- a. D_Cache only.
- b. For tag-RAM reads, only one-hot encodings are supported.
- c. For tag-RAM writes, all combinations are supported. The same data is written to all banks.
- d. For dirty RAM accesses, all combinations are supported.

Table 9-29 TRANSFER/AUX space

AxADDRSm bits	Description
[22:15]	0x0
[14]	Indicates address space accessed: 1 = TRANSFER/AUX
[13:4]	0x0
[3]	Register accessed: 0 = TRANSFER 1 = AUX
[2]	Word select
[1:0]	Byte select

Only accesses to the RAM-Access space actually perform RAM accesses.

TRANSFER and AUX are intermediate registers that are used by the AXI slave logic to perform RAM accesses.

Note

The physical integration of the RAMs limits the granularity of RAM accesses. This means that:

- A data chunk and its ECC or parity, if implemented, are always updated together.

- It is not possible to access part of a RAM-word unless the RAM-integration guidelines for the processor require that the RAM itself must support this feature.

This requirement exists only for the D_Cache data RAMs and dirty RAMs, that must be implemented by byte-writable RAMs. The AXI slave bus supports the full range of byte-write support to this RAM only.

Writes to the RAM-Access space update TRANSFER with the write data, then use this register, and possibly AUX, to write to the selected RAM. Reads from the RAM-Access space read the RAM contents into TRANSFER, and possibly AUX, and provide the requested portion of the read data from TRANSFER on the AXI interface.

To perform accesses outside these restrictions, you must perform a read-modify-write sequence.

You can also access the TRANSFER and AUX registers directly using the TRANSFER/AUX space. Such accesses do not actually perform RAM accesses. In this way RAM accesses are decoupled from AXI transactions, and a single RAM access can be decomposed into, or composed from, multiple AXI bus accesses. This enables, for example, a master capable only of sub-word accesses to get full access to the RAMs.

All accesses to the TRANSFER and AUX registers are cumulative. This means that data written to the TRANSFER and AUX registers, through direct AXI slave accesses, persists until it is overwritten. Reads from the cache RAMs, occurring as a side effect of AXI slave accesses to the RAM-Access space, also update these registers and overwrite any value previously written. This enables easier read-modify-write (RMW) operation by the master.

The TRANSFER register enables you to transfer data and ECC to the tag and dirty RAMs, and to transfer data to the data RAMs.

The AUX register is used only for transferring ECC to the data RAMs. If neither cache implements parity or ECC, direct accesses to the AUX register return a SLVERR.

For writes, you must ensure that all the data to be written to the selected RAM is initialized, either by prior accesses to TRANSFER/AUX, by the current access to RAM-Access or by a combination of both.

You can perform writes by a variety of sequences involving the RAM-Access space, and possibly also the TRANSFER/AUX space. For example, a write to a data RAM can be done by:

- Multiple writes to the TRANSFER register and AUX register, followed by a single write, with potentially zeroed byte strobes, at the appropriate address to the RAM-Access space
- A single write to the AUX register, if ECC is present, followed by a single write at the appropriate address to the RAM-Access space.

You can perform reads by a similarly varied number of sequences. For example, a read of a data RAM can be done by:

- A single 64-bit read of the RAM-Access space followed by a single 64-bit read of the AUX register
- A byte read of the RAM-Access space followed by several byte-reads to read the rest of the RAM data from the TRANSFER and AUX registers.

The format of the data, for reads and writes, depends on the RAM accessed and the error configuration of the RAM. These formats are described in the following tables. All writes must ensure that the write data is on the correct lane. Reads return data on the lanes described.

D_Cache data RAM single bank accesses

This section applies when you are performing a single bank access.

The location from which data bits are read, or to which they are written, depends on bit [0] of the RAM index. ECC or parity bits are written to, or read from, the lower byte of the AUX register.

Table 9-36 on page 9-29 describes the format of the AUX register for D_Cache data RAM accesses, when ECC is configured.

Table 9-30 Data RAM AUX format, D_Cache, with ECC

Bit	Description
[63:7]	RAZ/WI
[6:0]	ECC32[6:0]

Table 9-37 on page 9-29 describes the format of the AUX register for data RAM, D_Cache, when parity is configured.

Table 9-31 Data RAM AUX format, D_Cache, with parity

Bit	Description
[63:4]	RAZ/WI
[3]	parity for byte 3, data[31:24]
[2]	parity for byte 2, data[23:16]
[1]	parity for byte 1, data[15:8]
[0]	parity for byte 0, data[7:0]

Table 9-38 on page 9-29 describes the format of the AUX register for data RAM, D_Cache, when no error correction is configured.

Table 9-32 Data RAM AUX format, D_Cache, with no error correction

Bit	Description
[63:0]	RAZ/WI

RAM index[0] = 0

Writing The data bits used are the result of the lower word of TRANSFER multiplexed with the lower word of the data sent to the RAM-Access space. The values of **WSTRB** used for this AXI transaction determine which is multiplexed in:

WSTRB=0

data is taken from TRANSFER

WSTRB=1

data is taken from the data bus

Reading The data bits are written to the lower word of TRANSFER, and appear on the lower word of the AXI data bus.

The upper word of TRANSFER is set to zero.

RAM index[0] = 1

- Writing** The data bits used are the result of the upper word of TRANSFER multiplexed with the upper word of the data sent to the RAM-Access space. The values of **WSTRB** used for this AXI transaction determine which is multiplexed in.
- Reading** The data bits are written to the upper word of TRANSFER, and appear on the upper word of the AXI data bus.

I_Cache Data RAM access

Table 9-33 describes the format of the AUX register for data RAM, I-cache, when ECC is configured.

Table 9-33 Data RAM AUX format, I-cache, with ECC

Bit	Description
[63:8]	RAZ/WI
[7:0]	ECC64[7:0] for double word, data[63:0]

Table 9-34 describes the format of the AUX register for data RAM, I-cache, when parity is configured.

Table 9-34 Data RAM AUX format, I-cache, with parity

Bit	Description
[63:8]	RAZ/WI
[7:0]	Parity[7:0] for double word, data[63:0]

Table 9-35 describes the format of the AUX register for data RAM, I-cache, when no error correction is configured.

Table 9-35 Data RAM AUX format, I-cache, with ECC

Bit	Description
[63:0]	RAZ/WI

D_Cache data RAM double bank accesses

This section applies when you are performing a normal, or strobed, double bank access.

Normal accesses read or write all bytes of the doubleword being transferred. Strobed accesses read or write only those bytes specified by the corresponding bit in WSTRB. See Table 9-7 on page 9-10

Table 9-36 describes the format of the AUX register for D_Cache data RAM accesses, when ECC is configured.

Table 9-36 Data RAM AUX format, D_Cache, with ECC

Bit	Description
[63:15]	RAZ/WI
[14:8]	ECC32[6:0] for upper word, data[63:32]
[7]	RAZ/WI
[6:0]	ECC32[6:0] for lower word, data[31:0]

Table 9-37 describes the format of the AUX register for data RAM, D_Cache, when parity is configured.

Table 9-37 Data RAM AUX format, D_Cache, with parity

Bit	Description
[63:12]	RAZ/WI
[11]	parity for byte 3 of upper word, data[63:56]
[10]	parity for byte 2 of upper word, data[55:48]
[9]	parity for byte 1 of upper word, data[47:40]
[8]	parity for byte 0 of upper word, data[39:32]
[7:4]	RAZ/WI
[3]	parity for byte 3 of lower word, data[31:24]
[2]	parity for byte 2 of lower word, data[23:16]
[1]	parity for byte 1 of lower word, data[15:8]
[0]	parity for byte 0 of lower word, data[7:0]

Table 9-38 describes the format of the AUX register for data RAM, D_Cache, when no error correction is configured.

Table 9-38 Data RAM AUX format, D_Cache, with ECC

Bit	Description
[63:0]	RAZ/WI

Tag RAM access

Table 9-39 describes the format of the TRANSFER register for tag RAM, when using ECC.

Table 9-39 Tag RAM TRANSFER ECC format

Bit	Description
[63:30]	RAZ/WI
[29:23]	ECC32 – selected way
[22]	Valid - selected way
[21:0]	Tag - selected way

Table 9-40 describes the format of the TRANSFER register for tag RAM, when using parity.

Table 9-40 Tag RAM TRANSFER parity format

Bit	Description
[63:24]	RAZ/WI
[23]	Parity – selected way
[22]	Valid - selected way
[21:0]	Tag - selected way

Table 9-41 describes the format of the TRANSFER register for tag RAM, when no error correction is configured.

Table 9-41 Tag RAM TRANSFER format, without error correction

Bit	Description
[63:23]	RAZ/WI
[22]	Valid - selected way
[21:0]	Tag - selected way

Dirty RAM access

Table 9-42 describes the format of the TRANSFER register for dirty RAM, when ECC is configured.

Table 9-42 Dirty RAM TRANSFER format, with ECC

Bit	Description
[63:31]	RAZ/WI
[30:27]	ECC32 – way 3
[26:25]	Outer attributes – way 3
[24]	Dirty – way 3
[23]	RAZ/WI

Table 9-42 Dirty RAM TRANSFER format, with ECC (continued)

Bit	Description
[22:19]	ECC32 – way 2
[18:17]	Outer attributes – way 2
[16]	Dirty – way 2
[15]	RAZ/WI
[14:11]	ECC32 – way 1
[10:9]	Outer attributes – way 1
[8]	Dirty – way 1 Dirty – way 1
[7]	RAZ/WI
[6:3]	ECC32 – way 0
[2:1]	Outer attributes – way 0
[0]	Dirty – way 0

Table 9-43 describes the format of the TRANSFER register for dirty RAM, when parity, or no error correction, is configured.

Table 9-43 Dirty RAM TRANSFER format, without ECC

Bit	Description
[63:27]	RAZ/WI
[26:25]	Outer attributes – way 3
[24]	Dirty – way 3
[23:19]	RAZ/WI
[18:17]	Outer attributes – way 2
[16]	Dirty – way 2
[15:11]	RAZ/WI
[10:9]	Outer attributes – way 1
[8]	Dirty – way 1 Dirty – way 1
[7:3]	RAZ/WI
[2:1]	Outer attributes – way 0
[0]	Dirty – way 0

9.7 Peripheral interfaces

The processor has three peripheral interfaces. Accesses to the peripheral interfaces have lower latency, typically to half the latency of accesses to the AXI master interface. The port is used for:

- Device and Strongly-ordered type data accesses, normally to peripherals
- Normal-type memory low bandwidth data accesses, for example mailboxing.

The three peripheral interfaces use two physical ports, a 32-bit wide AXI master port that conforms to the AXI3 standard as described in the *AMBA AXI Protocol Specification* and an optional 32-bit wide AHB-Lite master port that conforms to the AHB-Lite standard as described in the *AMBA AHB Protocol Specification*.

The AXI peripheral port is sub-divided into:

- a virtual interface, referred to as LLPP Virtual AXI or the virtual-AXI peripheral interface
- a non-virtual interface, referred to as LLPP Normal AXI or the AXI peripheral interface.

The LLPP Virtual AXI is independent of the LLPP Normal AXI and the LLPP AHB peripheral interface from an ordering point of view. Accesses to both the AXI peripheral interfaces use the same physical AXI port but have different AXI IDs.

The AXI peripheral port has an address buffer and a data buffer, each of which has three entries. Each entry in the address buffer holds 32 bits of address, and an entry in the data buffer holds 32 bits of data. No merging is possible between the entries of a buffer. The LLPP Normal AXI and LLPP Virtual AXI share the address and data buffer.

The AHB peripheral port has its own address and data buffers. The address buffer has three entries and the data buffer has four entries. Each entry holds 32 bits. No merging is possible between the entries of a buffer.

The maximum number of outstanding write accesses that the processor posts onto the LLPP Virtual AXI is 3 and 15 for the LLPP Normal AXI.

AHB-Lite does not have the ability to do posted and out-of-order transactions, so the AHB peripheral port does not have a separate virtual interface.

Table 9-44 shows the AXI peripheral port attributes.

Table 9-44 AXI peripheral port attributes

Attribute	Value	Comments
Write issuing capability of LLPP Normal AXI	15	15 outstanding writes on (non-virtual) AXI peripheral interface
Write issuing capability of LLPP Virtual AXI	3	3 outstanding writes on virtual AXI peripheral interface
Read issuing capability	1	-
Combined issuing capability	19	Maximum number of posted writes on all AXI peripheral interfaces and a read
Write ID capability	2	-
Write interleave capability	1	The AXI peripheral port presents all write data in order
Read ID capability	2	-

The peripheral ports can run at the same frequency as the processor or at a lower synchronous frequency. See *AMBA interface clocking* on page 2-16 for more information.

In addition, the peripheral ports produce or check parity bits for each AXI or AHB channel. These additional signals are not part of the AXI or AHB specification, though some make use of AXI extension signals.

The following sections describe the attributes of the LLPP interfaces:

- *Peripheral interface configuration* on page 9-33
- *Peripheral interface initialization* on page 9-34
- *Peripheral interface attributes and permissions* on page 9-34

- *Identifiers for AXI peripheral port accesses* on page 9-34
- *Write response* on page 9-34
- *Memory attributes* on page 9-35
- *AXI peripheral port transfers* on page 9-35
- *AHB peripheral port transfers* on page 9-42
- *Semaphores* on page 9-48.

9.7.1 Peripheral interface configuration

The peripheral interfaces are configured during implementation and integration.

You can configure the AHB peripheral port to be removed, and not included in the processor design. The AXI peripheral port is always included and is not optional.

During implementation, you can configure the peripheral ports to use an error-correction scheme to detect and correct signals transferred using the peripheral port buses, see *Bus ECC* on page 9-2.

The size of each peripheral interface is configured during integration. The permissible LLPP Normal AXI, LLPP Virtual AXI, or AHB peripheral interface sizes are:

- 4 KB
- 8 KB
- 16 KB
- 32 KB
- 64 KB
- 128 KB
- 256 KB
- 512 KB
- 1 MB
- 2 MB
- 4 MB
- 8 MB
- 16 MB
- 32 MB
- 64 MB
- 128 MB
- 256 MB
- 512 MB
- 1 GB
- 2 GB
- 4 GB.

The LLPP Virtual AXI is either the same size as the LLPP Normal AXI or a sub-region of it.

The size of the peripheral interfaces is visible to software in the Peripheral Port Region Registers.

9.7.2 Peripheral interface initialization

The LLPP Normal AXI and AHB peripheral interfaces, but not the LLPP Virtual AXI interface, can be enabled from reset by configuring the control pins. Peripheral interface region enables can also be programmed using the System Coprocessor Registers, see *Peripheral interface region registers* on page 4-84. Ensure that peripheral interface region programming is done when the MPU is disabled to prevent unpredictable behavior.

9.7.3 Peripheral interface attributes and permissions

Accesses to the peripheral interfaces from the LSU are checked against the MPU for access permission. Memory access attributes are exported on this interface. Access permissions for peripheral interface accesses are the same as the permission attributes that the MPU assigns to the same address. Instructions cannot be fetched from any of the peripheral interfaces, and therefore they behave as if they have the *eXecute Never* (XN) attribute, regardless of the MPU XN attribute. All instruction fetches from the peripheral interfaces generate a permission fault. See Chapter 7 *Memory Protection Unit* for more information about memory attributes, types, and permissions.

———— Note ————

If a peripheral interface region overlaps with a TCM region then the TCM region gets more priority and the overlapping memory gets the attributes of the TCM region.

The L1 memory system cannot cache any peripheral interface access even if the access is to Normal memory with a Cacheable attribute. Load or store multiple instructions accessing the peripheral port are not performed as long bursts, and are not interruptible-restartable, even when they are in Normal memory. ARM recommends that you do not perform multiples to the peripheral interface regardless of the memory type, because this might impact the interrupt latency.

Any unaligned access to Device or Strongly Ordered memory generates an alignment fault and therefore does not cause any peripheral interface access. This means that the access examples given in this chapter never show unaligned accesses to Device or Strongly Ordered memory.

Also any shared exclusive double to the AXI peripheral port or any shared exclusive to the AHB peripheral port generates an abort and therefore does not cause an access.

9.7.4 Identifiers for AXI peripheral port accesses

Accesses on the AXI peripheral port use ID values as follows:

- ID0 for a read or a write access to the LLPP Normal AXI interface
- ID1 for a read or a write access to the LLPP Virtual AXI interface

9.7.5 Write response

The AXI peripheral port requires that the slave does not return a write response until it has received both the write data and the write address.

9.7.6 Memory attributes

The AXI peripheral port uses the **ARCACHEPm** and **AWCACHEPm** signals to indicate the memory attributes of the transfer, as returned by the MPU. Table 9-45 shows the encoding used for the **ARCACHEPm** and **AWCACHEPm** signals of the master interface. These are generated from the memory type and outer region attributes.

Table 9-45 ARCACHEPm and AWCACHEPm encodings

Encoding ^a	Meaning
b0000	Strongly Ordered
b0001	Device
b0011	Non-cacheable
b0110	Cacheable, write-through, allocate on reads only
b0111	Cacheable, write-back, allocate on reads only
b1111	Cacheable write-back, allocate on reads and writes

a. All encodings not shown in the table are reserved.

9.7.7 AXI peripheral port transfers

The processor conforms to the AXI3 specification, but it does not generate all the AXI transaction types that the specification permits. This section describes the types of AXI transactions that the Cortex-R5 AXI peripheral port does not generate. If you are designing an AXI slave to work only with the Cortex-R5 processor AXI peripheral port, you can take advantage of these restrictions and the interface attributes to simplify the slave.

This section also contains tables that show some examples of the types of AXI burst that the processor generates. However, because a particular type of transaction is not shown here does not mean that the processor does not generate such a transaction.

————— Note —————

An AXI slave device connected to the Cortex-R5 AXI master port must be capable of handling every kind of transaction permitted by the AXI specification, except where there is an explicit statement in this chapter that such a transaction is not generated. You must not infer any additional restrictions from the example tables given.

Restrictions on AXI peripheral transfers on page 9-36 describes restrictions on the type of transfers that the Cortex-R5 AXI peripheral port generates. If a CPUm exists and is powered up, **BREADYPm** and **RREADYPm** are always asserted. They are, however, deasserted when the CPU enters Dormant or Shutdown mode. You must not make any assumptions about the AXI handshaking signals, except that they conform to the *AMBA AXI3 Protocol Specification*.

The following sections give examples of transfers generated by the LLPP AXI interface:

- *Strongly Ordered and Device transactions* on page 9-36
- *Normal reads* on page 9-39
- *Normal Writes* on page 9-41.

Restrictions on AXI peripheral transfers

The Cortex-R5 AXI peripheral port applies the following restrictions to the AXI transactions it generates:

- A burst never transfers more than eight bytes
- The burst length is never more than two transfers
- No transaction ever crosses a 8-byte boundary in memory
- All bursts are incrementing (INCR) bursts
- If the transfer size is 8-bits or 16-bits then the burst length is always one transfer
- The transfer size is never greater than 32 bits
- All transactions are non-secure data accesses
- Transactions to Device and Strongly Ordered memory are always to addresses that are aligned for the transfer size
- Exclusive and Locked accesses are always to addresses that are aligned for the transfer size
- Write data is never interleaved
- ID values can only be 0 or 1 indicating normal AXI or virtual AXI respectively.

Strongly Ordered and Device transactions

A load or store instruction to or from Strongly Ordered or Device memory always generates AXI transactions of the same size as implied by the instruction. All accesses using LDM, STM, LDRD, or STRD instructions to Strongly Ordered or Device memory occur as 32-bit transfers.

LDRB

Table 9-46 shows the values of **ARADDRPm**, **ARBURSTPm**, **ARSIZEPm**, and **ARLENPm** for LDRB from bytes 0-3 in Strongly Ordered or Device memory.

Table 9-46 LDRB transfers

Address[1:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x0 (byte 0)	0x00	Incr	8-bit	1 data transfer
0x1 (byte 1)	0x01	Incr	8-bit	1 data transfer
0x2 (byte 2)	0x02	Incr	8-bit	1 data transfer
0x3 (byte 3)	0x03	Incr	8-bit	1 data transfer

LDRH

Table 9-47 shows the values of **ARADDRPm**, **ARBURSTPm**, **ARSIZEPm**, and **ARLENPm** for LDRH from halfwords 0-1 in Strongly Ordered or Device memory.

Table 9-47 LDRH transfers

Address[1:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x0 (halfword 0)	0x00	Incr	16-bit	1 data transfer
0x2 (halfword 1)	0x02	Incr	16-bit	1 data transfer

Note

A load of a halfword from Strongly Ordered or Device memory addresses 0x1 or 0x3 generates an alignment fault.

LDR or LDM that transfer one register

Table 9-48 shows the values of **ARADDRPm**, **ARBURSTPm**, **ARSIZEPm**, and **ARLENPm** for an LDR or an LDM that transfers one register, an LDM1, in Strongly Ordered or Device memory.

Table 9-48 LDR or LDM transfers

Address[1:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x0 (word 0)	0x00	Incr	32-bit	1 data transfer

Note

A load of a word from Strongly Ordered or Device memory addresses 0x1, 0x2, or 0x3 generates an alignment fault.

LDM that transfers five registers

Table 9-49 shows the values of **ARADDRPm**, **ARBURSTPm**, **ARSIZEPm**, and **ARLENPm** for an LDM that transfers five registers, an LDM5, in Strongly Ordered or Device memory. LDM transfers

Table 9-49 LDM transfers

Address[2:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x0 (word 0)	0x00	Incr	32-bit	2 data transfers
	0x08	Incr	32-bit	2 data transfers
	0x10	Incr	32-bit	1 data transfer
0x4 (word 1)	0x04	Incr	32-bit	1 data transfer
	0x08	Incr	32-bit	2 data transfers
	0x10	Incr	32-bit	2 data transfers

Note

A load-multiple from memory addresses 0x1, 0x2, 0x3, 0x5, 0x6, or 0x7 generates an alignment fault.

STRB

Table 9-50 shows the values of **AWADDRPm**, **AWBURSTPm**, **AWSIZEPm**, and **AWLENPm** for an STRB from bytes 0-3 in Strongly Ordered or Device memory.

Table 9-50 STRB transfers

Address[1:0]	AWADDRPm	AWBURSTPm	AWSIZEPm	AWLENPm	WSTRBPm
0x0 (byte 0)	0x00	Incr	8-bit	1 data transfer	b0001
0x1 (byte 1)	0x01	Incr	8-bit	1 data transfer	b0010
0x2 (byte 2)	0x02	Incr	8-bit	1 data transfer	b0100
0x3 (byte 3)	0x03	Incr	8-bit	1 data transfer	b1000

STRH

Table 9-51 shows the values of **AWADDRPm**, **AWBURSTPm**, **AWSIZEPm**, and **AWLENPm** for an STRB from halfwords 0-1 in Strongly Ordered or Device memory.

Table 9-51 STRH transfers

Address[1:0]	AWADDRPm	AWBURSTPm	AWSIZEPm	AWLENPm	WSTRBPm
0x0 (halfword 0)	0x00	Incr	16-bit	1 data transfer	b0011
0x2 (halfword 1)	0x02	Incr	16-bit	1 data transfer	b1100

Note

A store of a halfword from Strongly Ordered or Device memory addresses 0x1, 0x3, 0x5, or 0x7 generates an alignment fault.

STR or STM of one register

Table 9-52 shows the values of **AWADDRm**, **AWBURSTPm**, **AWSIZEPm**, and **AWLENPm** for an STR or an STM that transfers one register, an STM1, to Strongly Ordered or Device memory.

Table 9-52 STR or STM transfers

Address[1:0]	AWADDRPm	AWBURSTPm	AWSIZEPm	AWLENPm	WSTRBPm
0x0 (word 0)	0x00	Incr	32-bit	1 data transfer	b1111

Note

A store of a word to Strongly Ordered or Device memory addresses 0x1, 0x2, or 0x3 generates an alignment fault.

STM of five registers

Table 9-53 shows the values of **AWADDRm**, **AWBURSTPm**, **AWSIZEPm**, and **AWLENPm** for an STM that writes five registers, an STM5, over the AXI peripheral port to Strongly Ordered or Device memory.

Table 9-53 STM transfers

Address[2:0]	AWADDRPm	AWBURSTPm	AWSIZEPm	AWLENPm	WSTRBPm
0x00 (word 0)	0x00	Incr	32-bit	2 data transfers	b1111 b1111
	0x08	Incr	32-bit	2 data transfers	b1111 b1111
	0x10	Incr	32-bit	1 data transfer	b1111
0x04 (word 1)	0x04	Incr	32-bit	1 data transfer	b1111
	0x08	Incr	32-bit	2 data transfers	b1111 b1111
	0x10	Incr	32-bit	2 data transfers	b1111 b1111

Note

A store-multiple to address 0x1, 0x2, 0x3, 0x5, 0x6, or 0x7 generates an alignment fault.

Normal reads

Load instructions accessing Normal memory generate AXI peripheral port bursts that are always of 32-bit size and not necessarily the same size or length as the instruction implies. The tables in this section give examples of the types of AXI transaction that might result from various load instructions, accessing various addresses in Normal memory. They are provided as examples only, and are not an exhaustive description of the AXI transactions.

Table 9-54 shows possible values of **ARADDRm**, **ARBURSTPm**, **ARSIZEPm**, and **ARLENPm** for an LDRH from bytes 0-7 in Normal memory.

Table 9-54 LDRH transfers

Address[1:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x0 (byte 0)	0x00	Incr	32-bit	1 data transfer
0x1 (byte 1)	0x00	Incr	32-bit	1 data transfer
0x2 (byte 2)	0x00	Incr	32-bit	1 data transfer
0x3 (byte 3)	0x00	Incr	32-bit	2 data transfers
0x4 (byte 4)	0x04	Incr	32-bit	1 data transfer
0x5 (byte 5)	0x04	Incr	32-bit	1 data transfer

Table 9-54 LDRH transfers (continued)

Address[1:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x6 (byte 6)	0x04	Incr	32-bit	1 data transfer
0x7 (byte 7) ^a	0x04	Incr	32-bit	1 data transfer
	0x08	Incr	32-bit	1 data transfer

a. AXI peripheral port transactions do not cross a double word boundary.

Table 9-55 shows possible values of **ARADDRm**, **ARBURSTPm**, **ARSIZEPm**, and **ARLENPm** for an LDR or an LDM that transfers one register, an LDM1, to Normal memory.

Table 9-55 LDR or LDM transfers

Address[1:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x0 (byte 0) (word 0)	0x00	Incr	32-bit	1 data transfer
0x1 (byte 1)	0x00	Incr	32-bit	2 data transfers
0x2 (byte 2)	0x00	Incr	32-bit	2 data transfers
0x3 (byte 3)	0x00	Incr	32-bit	2 data transfers
0x4 (byte 4) (word 1)	0x04	Incr	32-bit	1 data transfer
0x5 (byte 5)	0x04	Incr	32-bit	1 data transfer
	0x08	Incr	32-bit	1 data transfer
0x6 (byte 6)	0x04	Incr	32-bit	1 data transfer
	0x08	Incr	32-bit	1 data transfer
0x7 (byte 7)	0x04	Incr	32-bit	1 data transfer
	0x08	Incr	32-bit	1 data transfer

Table 9-56 shows possible values of **ARADDRm**, **ARBURSTPm**, **ARSIZEPm**, and **ARLENPm** for an LDM that transfers five registers, an LDM5, to Normal memory.

Table 9-56 LDM transfers

Address[1:0]	ARADDRPm	ARBURSTPm	ARSIZEPm	ARLENPm
0x0 (word 0)	0x00	Incr	32-bit	2 data transfers
	0x08	Incr	32-bit	2 data transfers
	0x10	Incr	32-bit	1 data transfer
0x4 (word 1)	0x04	Incr	32-bit	1 data transfer
	0x08	Incr	32-bit	2 data transfers
	0x10	Incr	32-bit	2 data transfers

Normal Writes

Store instructions accessing Normal memory generate AXI peripheral port bursts that are always of 32-bit size and not necessarily the same size or length as the instruction implies. The AXI peripheral port asserts byte-lane strobes, **WSTRBPm[3:0]**, to ensure that only the bytes that were written by the instruction are updated.

The tables in this section give examples of the types of AXI transaction that might result from various store instructions, accessing various addresses in Normal memory. They are provided as examples only, and are not an exhaustive description of the AXI transactions.

Table 9-57 shows the values of **AWADDRPm**, **AWBURSTPm**, **AWSIZEPm**, and **AWLENPm** for an STRH to Normal memory.

Table 9-57 STRH transfers

Address[1:0]	AWADDRPm	AWBURSTPm	AWSIZEPm	AWLENPm	WSTRBPm
0x0 (byte 0)	0x00	Incr	32-bit	1 data transfer	b0011
0x1 (byte 1)	0x00	Incr	32-bit	1 data transfer	b0110
0x2 (byte 2)	0x00	Incr	32-bit	1 data transfer	b1100
0x3 (byte 3)	0x00	Incr	32-bit	2 data transfers	b1000 b0001
0x4 (byte 4)	0x04	Incr	32-bit	1 data transfer	b0011
0x5 (byte 5)	0x04	Incr	32-bit	1 data transfer	b0110
0x6 (byte 6)	0x04	Incr	32-bit	1 data transfer	b1100
0x7 (byte 7)	0x04	Incr	32-bit	1 data transfer	b1000
	0x08	Incr	32-bit	1 data transfer	b0001

Table 9-58 shows the values of **AWADDRPm**, **AWBURSTPm**, **AWSIZEPm**, and **AWLENPm** for an STR or an STM that transfers one register, an STM1, to Normal memory.

Table 9-58 STR or STM transfers

Address[1:0]	AWADDRPm	AWBURSTPm	AWSIZEPm	AWLENPm	WSTRBPm
0x0 (byte 0) (word 0)	0x00	Incr	32-bit	1 data transfer	b1111
0x1 (byte 1)	0x00	Incr	32-bit	2 data transfers	b1110 b0001
0x2 (byte 2)	0x00	Incr	32-bit	2 data transfers	b1100 b0011
0x3 (byte 3)	0x00	Incr	32-bit	2 data transfers	b1000 b0111
0x4 (byte 4) (word 1)	0x04	Incr	32-bit	1 data transfer	b1111
0x5 (byte 5)	0x04	Incr	32-bit	1 data transfer	b1110
	0x08	Incr	32-bit	1 data transfer	b0001

Table 9-58 STR or STM transfers (continued)

Address[1:0]	AWADDRPm	AWBURSTPm	AWSIZEPm	AWLENPm	WSTRBPm
0x6 (byte 6)	0x04	Incr	32-bit	1 data transfer	b1100
	0x08	Incr	32-bit	1 data transfer	b0011
0x7 (byte 7)	0x04	Incr	32-bit	1 data transfer	b1000
	0x08	Incr	32-bit	1 data transfer	b0111

9.7.8 AHB peripheral port transfers

The processor conforms to the AHB-Lite specification, but it does not generate all the AHB transaction types that the specification permits. This section describes the types of AHB transaction that the Cortex-R5 AHB peripheral port does not generate. If you are designing an AHB slave to work only with the Cortex-R5 processor AHB peripheral port, you can take advantage of these restrictions and the interface attributes described in previous sections to simplify the slave.

This section also contains tables that show some of the types of AHB burst that the processor generates. However, because a particular type of transaction is not shown here does not mean that the processor does not generate such a transaction.

———— Note ————

An AHB slave device connected to the Cortex-R5 AHB master port must be capable of handling every kind of transaction permitted by the AHB specification, except where there is an explicit statement in this chapter that such a transaction is not generated. You must not infer any additional restrictions from the example tables given.

Restrictions on AHB peripheral port transfers describes restrictions on the type of transfers that the Cortex-R5 AHB peripheral port generates.

The following sections give examples of transfers generated by the AHB peripheral port:

- *Strongly Ordered and Device transactions* on page 9-43
- *Normal reads* on page 9-46
- *Normal writes* on page 9-47.

Restrictions on AHB peripheral port transfers

The Cortex-R5 AHB peripheral port applies the following restrictions to the AHB transactions it generates:

- A burst never transfers more than eight bytes.
- The burst length is never more than two transfers.
- No transaction ever crosses a 8-byte boundary in memory
- All bursts are either single or 1-beat incrementing bursts, that is, **HBURSTPm[2:0]** is either SINGLE or INCR.
- The transfer type, that is, **HTRANSPm[2:0]** is never BUSY.
- The transfer size is never greater than 32 bits because it is a 32-bit AHB bus.
- If the transfer size is 8 bits or 16 bits then the burst length is always one transfer.

- All transactions are data accesses, that is **HPROTPm[0]** is always 1.
- Transactions to Device and Strongly Ordered memory are always to addresses that are aligned for the transfer size.
- Locked accesses are always to addresses that are aligned for the transfer size.

Strongly Ordered and Device transactions

A load or store instruction, to or from Strongly Ordered or Device memory, always generates AHB transactions of the size implied by the instruction. All accesses using LDM, STM, LDRD or STRD instructions to Strongly Ordered or Device memory occur as 32-bit transfers.

LDRB

Table 9-59 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an LDRB from bytes 0-3 in Strongly Ordered or Device memory.

Table 9-59 LDRB transfers

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (byte 0)	0x00	Single	8-bit
0x1 (byte 1)	0x01	Single	8-bit
0x2 (byte 2)	0x02	Single	8-bit
0x3 (byte 3)	0x03	Single	8-bit

LDRH

Table 9-60 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an LDRH from halfwords 0-1 in Strongly Ordered or Device memory.

Table 9-60 LDRH transfers

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (halfword 0)	0x00	Single	16-bit
0x2 (halfword 1)	0x02	Single	16-bit

———— Note ————

A load of a halfword from Strongly Ordered or Device memory addresses 0x1 or 0x3 generates an alignment fault.

LDR or LDM of one register

Table 9-61 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an LDR or an LDM that transfers one register, an LDM1, in Strongly Ordered or Device memory.

Table 9-61 LDR or LDM of one register

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (word 0)	0x00	Single	32-bit

Note

A load of a word from Strongly Ordered or Device memory addresses 0x1, 0x02, 0x3, 0x5, 0x06, or 0x7 generates an alignment fault.

LDM that transfers five registers

Table 9-62 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an LDM that transfers five registers, an LDM5, in Strongly Ordered or Device memory.

Table 9-62 LDM that transfers five registers

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (word 0)	0x00 0x04	Incr	32-bit
	0x08 0x0C	Incr	32-bit
	0x10	Single	32-bit
0x4 (word 1)	0x04	Single	32-bit
	0x08 0x0C	Incr	32-bit
	0x10 0x14	Incr	32-bit

Note

A load of a word from Strongly Ordered or Device memory addresses 0x1, 0x2, or 0x3 generates an alignment fault.

STRB

Table 9-63 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an STRB from bytes 0-3 in Strongly Ordered or Device memory.

Table 9-63 STRB transfers

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (byte 0)	0x00	Single	8-bit
0x1 (byte 1)	0x01	Single	8-bit
0x2 (byte 2)	0x02	Single	8-bit
0x3 (byte 3)	0x03	Single	8-bit

STRH

Table 9-60 on page 9-43 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an STRH from halfwords 0-1 in Strongly Ordered or Device memory.

Table 9-64 STRH transfers

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (halfword 0)	0x00	Single	16-bit
0x2 (halfword 1)	0x02	Single	16-bit

Note

A store of a halfword to Strongly Ordered or Device memory addresses 0x1 or 0x3 generates an alignment fault.

STR of one register

Table 9-65 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an STR that transfers one register, an STR1, in Strongly Ordered or Device memory.

Table 9-65 STR of one register

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (word 0)	0x00	Single	32-bit

Note

A store of a word to Strongly Ordered or Device memory addresses 0x1, 0x2, or 0x3 generates an alignment fault.

STM of five registers

Table 9-66 shows the values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an STM that transfers five registers, an STM5, over the AHB master port to Strongly Ordered or Device memory.

Table 9-66 STM of five registers

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (word 0)	0x00	Incr	32-bit
	0x04		
	0x08 0x0C	Incr	32-bit
0x4 (word 1)	0x10	Single	32-bit
	0x04	Single	32-bit
	0x08 0x0C		
	0x10 0x14	Incr	32-bit

Note

A store of a word from Strongly Ordered or Device memory addresses 0x1, 0x2, 0x3, 0x5, 0x6, or 0x7 generates an alignment fault.

Normal reads

Load instructions accessing Normal memory generate AHB peripheral port bursts that might not be the same size or length as the instruction implies. The tables in this section give examples of AHB transactions that might result from various load instructions, accessing various addresses in Normal memory. They are examples only, and are not an exhaustive description of the AHB transactions.

LDRH

Table 9-67 shows possible values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an LDRH from bytes 0 to 7 in Normal memory.

Table 9-67 LDRH transfers in Normal memory

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (byte 0)	0x00	Single	16-bit
0x1 (byte 1)	0x01	Single	8-bit
	0x02	Single	8-bit
0x2 (byte 2)	0x02	Single	16-bit
0x3 (byte 3)	0x03	Single	8-bit
	0x04	Single	8-bit
0x4 (byte 4)	0x04	Single	16-bit
0x5 (byte 5)	0x05	Single	8-bit
	0x06	Single	8-bit
0x6 (byte 6)	0x06	Single	16-bit
0x7 (byte 7) ^a	0x07	Single	8-bit
	0x08	Single	8-bit

a. AHB peripheral port transactions do not cross a double word boundary.

LDR

Table 9-68 shows possible values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an LDR from Normal memory.

Table 9-68 LDR transfers in Normal memory

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (byte 0, word 0)	0x00	Single	32-bit

Table 9-68 LDR transfers in Normal memory (continued)

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x1 (byte 1)	0x01	Single	8-bit
	0x02	Single	16-bit
	0x04	Single	8-bit
0x2 (byte 2)	0x02	Single	16-bit
	0x04	Single	16-bit
0x3 (byte 3)	0x03	Single	8-bit
	0x04	Single	16-bit
	0x06	Single	8-bit

Normal writes

Store instructions accessing Normal memory generate AHB peripheral port bursts that might not be the same size or length as the instruction implies. The tables in this section give examples of AHB transactions that might result from various store instructions, accessing various addresses in Normal memory. They are examples only, and are not an exhaustive description of the AHB transactions.

STRH

Table 9-69 shows possible values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an STRH from bytes 0 to 3 in Normal memory.

Table 9-69 STRH transfers in Normal memory

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (byte 0)	0x00	Single	16-bit
0x1 (byte 1)	0x01	Single	8-bit
	0x02	Single	8-bit
0x2 (byte 2)	0x02	Single	16-bit
0x3 (byte 3)	0x03	Single	8-bit
	0x04	Single	8-bit

STR or STM of one register

Table 9-70 shows possible values of **HADDRPm[1:0]**, **HBURSTPm**, and **HSIZEPm** for an STR to Normal memory.

Table 9-70 STR transfers in Normal memory

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x0 (byte 0, word 0)	0x00	Single	32-bit

Table 9-70 STR transfers in Normal memory (continued)

Address[1:0]	HADDRPm[1:0]	HBURSTPm	HSIZEPm
0x1 (byte 1)	0x01	Single	8-bit
	0x02	Single	16-bit
	0x04	Single	8-bit
0x2 (byte 2)	0x02	Single	16-bit
	0x04	Single	16-bit
0x3 (byte 3)	0x03	Single	8-bit
	0x04	Single	16-bit
	0x06	Single	8-bit

9.7.9 Semaphores

The peripheral interfaces use the internal exclusive monitor of the processor L1 memory system to manage load, store and clear exclusive instructions to non-shared memory. The internal monitor checks exclusive accesses to shared memory and also, if necessary, any external monitor using the L2 memory interface. You can use these instructions to construct semaphores and ensure synchronization between different processes or processors. See the *ARM Architecture Reference Manual* for more information about how these instructions work.

Only exclusive instructions to shared memory result in exclusive accesses on the bus. Exclusive accesses to non-shared memory are marked as non-exclusive accesses on the bus.

Exclusive doubles to shared memory on LLPP Normal AXI or LLPP Virtual AXI (LDREXD and STREXD) are aborted. The AHB peripheral port cannot perform any exclusive accesses, so all exclusive accesses to shared memory on the AHB peripheral interface are aborted. The source of an abort because of a shared exclusive double to LLPP Normal AXI or LLPP Virtual AXI, or a shared exclusive to the AHB peripheral interface is encoded in the *Data Fault Status Register* (DFSR) as a Synchronous External AXI Slave Error.

The SWP and SWPB instructions can also be used for memory synchronization. Only swap instructions to shared memory are marked as locked accesses on the bus.

9.8 Accelerator Coherency Port interface

The optional *Accelerator Coherency Port* (ACP) provides memory coherency as introduced in *Coherency* on page 3-6 between each CPU in the Cortex-R5 group and an external master.

The ACP has an AXI slave interface and an AXI master interface:

- the ACP slave interface has one port with only the AW and B channels
- the ACP master interface has one port with only the AW and B channels.

Each port is 64 bits wide, and conforms to the AMBA 3 AXI standard as described in the *AMBA AXI Protocol Specification*.

Within the AXI standard, the ACP slave port uses a number of extension signals to:

- indicate if coherency must be preserved
- give information about coherency maintenance operations
- carry parity information for the bus-ECC feature, if included.

Within the AXI standard, the ACP master port uses a number of extension signals to:

- indicate if coherency must be preserved
- carry parity information for the bus-ECC feature, if included.

See *Bus ECC* on page 9-2 for more information on parity checking and generation in the ACP.

The ACP ports can run at the same frequency as the processor or at a lower synchronous frequency. See *Clocking* on page 2-16 for more information.

The Cortex-R5 ACP memory coherency scheme only provides coherency between an external master connected to the ACP slave port and a CPU with a data cache in the Cortex-R5 group for memory regions configured as inner cacheable write-through in the CPU's MPU. It does not provide coherency for memory regions configured as cacheable write-back.

Note

In a twin-CPU configuration, the ACP maintains memory coherency between the external master and each CPU with a data cache in the Cortex-R5 group, but not between the external master and a CPU without a data cache, or between the two CPUs.

For AXI write transactions going through the ACP and marked as coherent, AW channel sideband signal **AWCOHERENTCS** high, the ACP ensures that there is no cached copy of the data at these addresses in the CPU's data cache when the AXI write completes.

When an AXI write from the external master appears on the ACP slave port's AW channel, the ACP records some information about it and forwards the write transaction to the memory system on the ACP master port's AW channel.

When the memory system sends the write response on the ACP master port's B channel, the ACP records the response and recalls if the transaction was coherent.

If the transaction is not coherent, the ACP forwards the response to the external master on the ACP slave port's B channel.

If the transaction is coherent, the ACP first sends coherency maintenance operations to the CPU's data cache controller for the addresses spanned by the write transaction, and waits until the cache controller has acknowledged that all necessary coherency maintenance operations have been carried out to forward the write response to the ACP slave port's B channel, along with information about the maintenance operations.

Coherency maintenance operations invalidate cache lines when a CPU's data cache holds a copy of data at an address spanned by a coherent external write transaction. However if this cache line is dirty, it is not invalidated and the ACP indicates along with the write response that coherency was not maintained for this transaction.

For each CPU, information on the coherency maintenance operations includes:

- If all addresses were not cached, sideband signal **BMISSCSm**
- If at least one address was cached and potentially dirty in which case coherency has not been maintained, sideband signal **BHITDIRTYCSm**.

If a transaction is not coherent, the ACP always indicates that all addresses were not cached and never indicates that at least one address was cached and potentially dirty.

If a CPU's data cache controller cannot process coherency maintenance requests, because, for example, it is powered down, the ACP always indicates that all addresses were not cached and indicates that at least one address was cached and potentially dirty, only if coherency was not maintained for the write transaction.

Note

- The ACP does not reorder transactions:
 - write address transactions appear on the ACP master port AW channel in the same order as they appeared on the ACP slave port AW channel
 - responses appear on the ACP slave port B channel in the same order as they appeared on the ACP master port B channel.
 - The ACP master port requires that the slave it connects to does not return a write response until it has received both the write data and the write address.
 - You must not use the ACP at the same time as the AXI slave is accessing the cache RAMs. If you use the AXI slave to access the cache RAMs, ensure that it is idle before initiating ACP transactions.
-

The ACP slave interface attributes are described in Table 9-71.

Table 9-71 ACP slave interface attributes

Attribute	Value
Write acceptance capability	4
Write interleave depth	1

The ACP master interface attributes are described in Table 9-72.

Table 9-72 ACP master interface attributes

Attribute	Value
Write issuing capability	2
Write ID capability	4
Write ID width	2

Chapter 10

Power Control

This chapter describes the processor power control functions. It contains the following sections:

- *About power control* on page 10-2
- *Power management* on page 10-3.

10.1 About power control

The features of the processor that improve energy efficiency include:

- branch and return prediction, reducing the number of incorrect instruction fetch and decode operations
- the caches use sequential access information to reduce the number of accesses to the tag RAMs and to unwanted data RAMs.

In the processor, extensive use is also made of gated clocks and gates to disable inputs to unused functional blocks. Only the logic actively in use to perform a calculation consumes any dynamic power.

10.2 Power management

Each CPU in the Cortex-R5 processor supports four different power modes from Run to Shutdown, with decreasing levels of power consumption, but increasing entry and exit costs. The modes are summarized in the following table.

Table 10-1 Power management modes

Mode	CPU clock gated	CPU logic powered	CPU RAMs powered	Exit to Run mode requires
Run	No	Yes	Yes	-
Standby	When idle	Yes	Yes	Pipeline restart
Dormant	Yes	No	Yes	Pipeline restart Restore registers and configuration from memory
Shutdown	Yes	No	No	Pipeline restart Restore registers and configuration from memory Invalidate caches and reinitialize caches and TCMs

If the processor is implemented with twin CPUs, then each CPU can be in a mode independent of the other, provided CPU1 is never in a higher power mode than CPU0 when CPU0 is in Dormant or Shutdown mode. Regardless of the state of the CPUs, the logic for the ACP interfaces and the debug-APB interfaces remain powered up.

A CPU can only enter Dormant or Shutdown modes if it is implemented with the appropriate power gating circuitry and clamp logic, and is integrated into a system with a power controller.

This section describes:

- *Run mode*
- *Standby mode*
- *Dormant mode* on page 10-4
- *Shutdown mode* on page 10-4
- *Power Management Controller* on page 10-5
- *Power mode interaction with ACP* on page 10-5
- *Power mode interaction with debug* on page 10-5.

10.2.1 Run mode

Run mode is the normal mode of operation where all of the functionality of the CPU is available.

10.2.2 Standby mode

Standby mode allows most of the clocks of the device to be disabled, while keeping the design powered up. This reduces the power drawn to the static leakage current, plus a tiny clock power overhead required to enable the device to wake up from the Standby mode.

Entry into Standby mode is performed by executing the *Wait For Interrupt* (WFI) instruction or *Wait For Event* (WFE) instruction. To ensure that the entry into the Standby mode does not affect the memory system on a Cortex-R5 CPU, the WFI and WFE instructions automatically performs a Data Synchronization Barrier operation. This ensures that all explicit memory accesses occur before the WFI or WFE has completed. When this has happened, the CPU stops fetching instructions and asserts **nWFIPIPESTOPPEDm** or **nWFEPIPESTOPPEDm** as appropriate, to indicate that it is in Standby mode.

When the CPU is in Standby mode and it has no outstanding AXI-slave or debug-APB transactions or ACP invalidate requests, then it stops the clock to the majority of its logic. When the CPU clocks are stopped the **nCLKSTOPPEDm** signal is asserted. If the **DBGNOCLKSTOP** input is asserted, the CPU does not stop its clocks or assert **nCLKSTOPPEDm** when in Standby mode.

When the processor is in Standby mode and the AXI slave interface or debug-APB interface receives a transaction or an ACP invalidate request is generated, the processor clocks are temporarily restarted and **nCLKSTOPPEDm** is deasserted to enable it to service the transaction, but it does not return to Run mode.

The CPU exits Standby mode and returns to Run mode in response to a variety of events, depending on whether Standby mode was entered using WFI or WFE.

For WFI, the transition from Standby mode to Run mode is caused by:

- the arrival of an interrupt, whether masked or unmasked
- a debug request, whether debug is enabled or disabled
- a reset.

For WFE, the transition from Standby mode to Run mode is caused by:

- the arrival of an unmasked interrupt
- a debug request, whether debug is enabled or disabled
- an event signalled on the **EVENTIm** input
- a reset.

The debug request can be generated by an externally generated debug request, using the **EDBGRQm** pin on the processor, or from a Debug Halt instruction issued to the processor through the debug *Advanced Peripheral Bus* (APB).

Systems using the VIC interface must ensure that the VIC is not masking any interrupts that are required for restarting the processor when in standby mode.

10.2.3 Dormant mode

In Dormant mode, only the CPU logic, but not the CPU TCM and cache RAMs, is powered down, so that the only power consumption is the static leakage current of the RAMs.

Before entering Dormant mode, you must save the CPU state, except for the cache and TCM state, in memory. When power is restored to the CPU logic, the CPU is returned to Run mode by asserting and deasserting **nRESETm**. You must restore the CPU state as part of the boot process. Because the cache and TCM are not powered down in Dormant mode, you do not have to invalidate or initiate them during boot, and the task can access data in the cache without requiring a cache refill. In Dormant mode, the CPU state, apart from the cache and TCM state, is stored to memory before entry into this mode, and restored after exit. For more information on how to implement and use Dormant mode in your design, contact ARM.

10.2.4 Shutdown mode

In Shutdown mode, the entire CPU is powered down, so that it consumes no power. Before entering Shutdown mode, you must save all the processor state, including any required cache and TCM state in the level-2 memory. This typically includes cleaning the whole data cache. When it is powered up, the CPU is returned to Run mode by asserting and deasserting **nRESETm**. As part of the boot process, you must:

- restore the CPU state if required
- invalidate the caches
- initialize the TCMs as part of the boot process.

10.2.5 Power Management Controller

You can only put the CPU into Dormant mode or Shutdown mode if it is integrated into a system with a memory-mapped *Power Management Controller* (PMXEVNTR). The PMXEVNTR must respond to software running on the CPU to power down the appropriate logic at the right time. The PMXEVNTR must also respond to stimulus from the system, to power up the CPU logic and return it to Run mode.

Both Standby mode and Dormant mode are entered through Standby mode. You must program the PMXEVNTR to indicate which mode you want to enter, then perform the appropriate state-saving operations. After this is done, execute WFI or WFE to enter Standby mode.

When the CPU is in Standby mode, **nWFIPIPESTOPPEDm** or **nWFEPIPESTOPPEDm** is asserted to indicate that the CPU pipeline has quiesced. The PMXEVNTR must also ensure that the system provides no stimulus to the CPU so that the whole CPU is quiesced. For example, no new transactions to the Cortex-R5 AXI-slave interface can be started, and all outstanding transactions must be completed. Only when the CPU is completely quiesced can the PMXEVNTR remove power from the logic. If the system provides stimulus, for example an interrupt to the CPU, after it has entered Standby mode, the CPU might have started to exit Standby mode when the power is removed, that can lead to corruption of the system.

10.2.6 Power mode interaction with ACP

When a CPU is in Standby mode, and a transaction that requires coherency is received by the ACP, the clock for the CPU is restarted, if required, so that coherency maintenance operations can be handled as normal. When the ACP is idle again the clock is gated off again, if appropriate,

When a CPU is in Dormant mode, then its cache contents are live, but it cannot respond to coherency maintenance operations that the ACP generates. For this CPU, for ACP transactions requiring coherency, the coherency maintenance operations information signals indicate that all addresses were not cached, that is, **BMISSCS[m]** is asserted, and indicate that at least one address was cached and potentially dirty, **BHITDIRTYCS[m]**. Because this is usually considered erroneous, ARM recommends that the system is built so that transactions requiring coherency cannot be received by the processor, when one or both of the CPUs are in Dormant mode.

When a CPU is in Shutdown mode, its cache contents are lost and therefore there are no coherency issues with that cache. For this CPU, the coherency maintenance operations information signals indicate that all addresses were not cached, that is, **BMISSCS[m]** is asserted, and do not indicate that at least one address was cached and potentially dirty, that is, **BHITDIRTYCS[m]** is not asserted.

See *Accelerator Coherency Port interface* on page 9-48 for more information about the ACP.

10.2.7 Power mode interaction with debug

When one of the Cortex-R5 CPUs is in Standby mode and a debug-APB access to one of the core registers is received, the clocks for the CPU are restarted, if required, so that the transaction can be serviced as normal. When the transaction is complete, the clock is, gated off again if appropriate.

When a CPU is in Shutdown mode or Dormant mode, the core debug registers, for example, DBGDSCR, are unavailable and an error response is signalled for transactions to these registers. The debug-APB interface and the debug domain registers, for example DIDR, remain available as normal. The power-down status is indicated by the DBGPRSR. See *Device Power-down and Reset Status Register* on page 12-32.

Chapter 11

FPU Programmers Model

This chapter describes the programmers model of the *Floating Point Unit* (FPU). It contains the following sections:

- *About the FPU programmers model* on page 11-2
- *General-purpose registers* on page 11-4
- *System registers* on page 11-5
- *Modes of operation* on page 11-12
- *Compliance with the IEEE 754 standard* on page 11-13.

The Cortex-R5F processor is a Cortex-R5 processor that includes the optional FPU. In this chapter, the generic term *processor* means only the Cortex-R5F processor.

11.1 About the FPU programmers model

The FPU implements the VFPv3-D16 architecture and the Common VFP Sub-Architecture v2. This includes the instruction set of the VFPv3 architecture. See the *ARM Architecture Reference Manual* for information on the VFPv3 instruction set.

11.1.1 FPU functionality

The FPU is an implementation of the *ARM Vector Floating Point v3* architecture, with 16 double-precision registers (VFPv3-D16). It provides floating-point computation functionality that is compliant with the *ANSI/IEEE Std 754-1985, IEEE Standard for Binary Floating-Point Arithmetic*, referred to as the IEEE 754 standard. The FPU supports all data-processing instructions and data types in the VFPv3 architecture as described in the *ARM Architecture Reference Manual*.

The FPU fully supports single-precision and double-precision add, subtract, multiply, divide, multiply and accumulate, and square root operations. It also provides conversions between fixed-point and floating-point data formats, and floating-point constant instructions. The FPU does not support any data processing operations on vectors in hardware. Any data processing instruction that operates on a vector generates an Undefined Instruction exception. The operation can then be emulated in software if necessary.

Optionally, you can configure the FPU to support single-precision only.

Cortex-R5F does not implement either the half-precision conversion or fused-MAC extensions to the VFPv3 architecture.

11.1.2 About the VFPv3-D16 architecture

The VFPv3-D16 architecture only includes 16 double-precision registers. VFPv3 includes 32 double-precision registers by default. An instruction that attempts to access any of the registers D16-D31 generates an Undefined Instruction exception.

11.1.3 VFP instructions in a single-precision configuration

Table 11-1 lists the VFP instructions that are Undefined in a single-precision only configuration. These instructions are <opcode>.<cond>.F64 where opcode is listed in the table:

Table 11-1 Instructions undefined in a single-precision only configuration

Instruction Operation	Opcodes
Vector Multiply Accumulate or Subtract	VMLA, VMLS
Vector Negate Multiply Accumulate or Subtract	VNMLA, VNMLS, VNMUL
Vector Multiply	VMUL
Vector Add	VADD
Vector Subtract	VSUB
Vector Divide	VDIV
Vector Move	VMOV (immediate), VMOV (register)
Vector Absolute	VABS
Vector Negate	VNEG

Table 11-1 Instructions undefined in a single-precision only configuration (continued)

Instruction Operation	Opcodes
Vector Square Root	VSQRT
Vector Compare	VCMP, VCMPE
Vector Convert	VCVT, VCVTR (all supported variants)

Note

The single-precision variants of these instructions (<opcode>.<cond>.F32) execute as normal.

11.2 General-purpose registers

The FPU implements a VFP register bank. This bank is distinct from the ARM register bank.

You can reference the VFP register bank using two explicitly aliased views. Figure 11-1 shows the two views of the register bank and the way the word and doubleword registers overlap.

11.2.1 FPU views of the register bank

In the FPU, you can view the register bank as:

- Sixteen 64-bit doubleword registers, D0-D15.
- Thirty-two 32-bit single-word registers, S0-S31.
- A combination of registers from these views.




Figure 11-1 FPU register bank

The mapping between the registers is as follows:

- S<2n> maps to the least significant half of D<n>
- S<2n+1> maps to the most significant half of D<n>.

For example, you can access the least significant half of the value in D6 by accessing S12, and the most significant half of the elements by accessing S13.

11.3 System registers

The VFPv3 architecture describes the following system registers:

- *Floating-Point System ID Register* on page 11-6
- *Floating-Point Status and Control Register* on page 11-7
- *Floating-Point Exception Register, FPEXC* on page 11-8
- *Media and VFP Feature Registers, MVFR0 and MVFR1* on page 11-9.

Table 11-2 shows the VFP system registers in the Cortex-R5F FPU.

Table 11-2 VFP system registers

Register	VMRS/VMSR <reg> field	Access type	Reset state
Floating-Point System ID Register, FPSID	b0000	Read-only	0x4102315x ^a
Floating-Point Status and Control Register, FPSCR	b0001	Read/write	0x00000000
Floating-Point Exception Register, FPEXC	b1000	Read/write	0x00000000
VFP Feature Register 0, MVFR0	b0111	Read-only	0x10110221
VFP Feature Register 1, MVFR1	b0110	Read-only	0x00000011

a. Bits [3:0] of the FPSID depend on the product revision. See the FPSID register description for more information.

———— Note ————

The FPSID, MVFR0, and MVFR1 Registers are read-only. Attempts to write these registers are ignored.

Table 11-3 shows that a Privileged mode is sometimes required to access a VFP system register. When a Privileged mode is required, an instruction that attempts to access a register in a nonprivileged mode takes the Undefined Instruction exception.

Table 11-3 Accessing VFP system registers

Register	Privileged access		User access	
	FPEXC EN=0	FPEXC EN=1	FPEXC EN=0	FPEXC EN=1
FPSID	Permitted	Permitted	Not permitted	Not permitted
FPSCR	Not permitted	Permitted	Not permitted	Permitted
MVFR0, MVFR1	Permitted	Permitted	Not permitted	Not permitted
FPEXC	Permitted	Permitted	Not permitted	Not permitted

For a VFP system register to be accessible, it must follow the rules in Table 11-3 and the VFP must also be accessible according to the CPACR. See *c1, Coprocessor Access Control Register* on page 4-47 for more information.

Note

All hardware ID information is privileged access only:

FPSID is privileged access only

This is a change in VFPv3 compared to VFPv2.

MVFR registers are privileged access only

User code must issue a system call to determine the features that are supported.

The following sections describe the VFP system registers:

- *Floating-Point System ID Register*
- *Floating-Point Status and Control Register* on page 11-7
- *Floating-Point Exception Register, FPEXC* on page 11-8
- *Media and VFP Feature Registers, MVFR0 and MVFR1* on page 11-9.

11.3.1 Floating-Point System ID Register

The FPSID Register characteristics are:

Purpose	Indicates which VFP implementation is being used.
Usage constraints	The FPSID Register: • is a read-only register • must be accessed in Privileged mode only.
Configurations	Use this register if the device is configured as a Cortex-R5F processor.
Attributes	See Table 11-4.

Figure 11-2 shows the bit assignments.




Figure 11-2 FPSID Register bit assignments

Table 11-4 shows the bit assignments.

Table 11-4 FPSID Register bit assignments

Bits	Name	Function
[31:24]	Implementer	ARM Limited: 0x41 = A
[23]	Hardware or software	0 = hardware implementation
[22:16]	Subarchitecture version	VFP architecture v3 or later with Common VFP subarchitecture v2 ^a : 0x02

Table 11-4 FPSID Register bit assignments (continued)

Bits	Name	Function
[15:8]	Part number	0x31 = Cortex-R5F processor
[7:4]	Variant	0x5 = Cortex-R5F processor
[3:0]	Revision	When the build-configuration includes the floating point unit, this register identifies the revision number of the floating-point unit: 0x0 = r0p0 0x1 = r1p0 0x2 = r1p1

a. For more information about the Common VFP subarchitecture see the *ARM Architecture Reference Manual*.

11.3.2 Floating-Point Status and Control Register

The FPSCR Register characteristics are:

Purpose	Provides all necessary User level control of the floating-point system.
Usage constraints	All bits described as DNM in Figure 11-3 are reserved for future expansion. These bits must be initialized to zeros. To ensure that these bits are not modified, any code other than initialization code must use read-modify-write techniques when writing to FPSCR. Failure to observe this rule can cause Unpredictable results in future systems.
Configurations	Use this register if the device is configured as a Cortex-R5F processor.
Attributes	See Table 11-5.

Figure 11-3 shows the bit assignments.




Figure 11-3 FPSID Register bit assignments

Table 11-5 shows the bit assignments.

Table 11-5 FPSID Register bit assignments

Bits	Name	Function
[31]	N	Set if comparison produces a <i>less than</i> result, resets to zero
[30]	Z	Set if comparison produces an <i>equal</i> result, resets to zero
[29]	C	Set if comparison produces an <i>equal</i> , <i>greater than</i> , or <i>unordered</i> result, resets to zero

Table 11-5 FPSCR Register bit assignments (continued)

Bits	Name	Function
[28]	V	Set if comparison produces an <i>unordered</i> result, resets to zero
[27]	QC	<i>Do Not Modify</i> (DNM)/ <i>Read As Zero</i> (RAZ)
[26]	AHP	DNM/RAZ
[25]	DN	Default NaN mode enable bit: 0 = default NaN mode disabled, this is the reset value 1 = default NaN mode enabled.
[24]	FZ	Flush-to-zero mode enable bit: 0 = flush-to-zero mode disabled, this is the reset value 1 = flush-to-zero mode enabled.
[23:22]	RMODE	Rounding mode control field: b00 = round to nearest (RN) mode, this is the reset value b01 = round towards plus infinity (RP) mode b10 = round towards minus infinity (RM) mode b11 = round towards zero (RZ) mode.
[21:20]	STRIDE	Indicates the vector stride, reset value is 0x0
[19]	-	DNM
[18:16]	LEN	Indicates the vector length, reset value is 0x0
[15]	IDE	RAZ
[14:13]	-	DNM
[12]	IXE	RAZ
[11]	UFE	RAZ
[10]	OFE	RAZ
[9]	DZE	RAZ
[8]	IOE	RAZ
[7]	IDC	Input Subnormal cumulative flag, resets to zero
[6:5]	-	DNM
[4]	IXC	Inexact cumulative flag, resets to zero
[3]	UFC	Underflow cumulative flag, resets to zero
[2]	OFC	Overflow cumulative flag, resets to zero
[1]	DZC	Division by Zero cumulative flag, resets to zero
[0]	IOC	Invalid Operation cumulative flag, resets to zero

11.3.3 Floating-Point Exception Register, FPEXC

The FPEXC Register characteristics are:

Purpose	Provides global enable and disable control of the VFP extension, and indicate how the state of this extension is recorded.
----------------	--

- Usage constraints**
- The FPEXC Register is accessible in Privileged modes only.
 - Clearing EN disables VFP functionality, causing all VFP instructions apart from privileged system register accesses to generate an Undefined Instruction exception.

Configurations Use this register if the device is configured as a Cortex-R5F processor.

Attributes See Table 11-6.

Figure 11-4 shows the bit assignments.




Figure 11-4 FPEXC Register bit assignments

Table 11-6 shows the bit assignments.

Table 11-6 FPEXC Register bit assignments

Bits	Name	Function
[31]	-	RAZ.
[30]	EN	VFP enable bit. Setting EN enables VFP functionality. Reset clears EN.
[29]	DEX	Set when an Undefined Instruction exception is taken because of a vector instruction that would have been executed if the processor supported vectors. This field is cleared when an Undefined Instruction exception is taken for any other reason. Resets to zero. In single-precision only configurations, this bit is not set for any double-precision operations, whether they are vector operations or not.
[28:0]	-	RAZ.

11.3.4 Media and VFP Feature Registers, MVFR0 and MVFR1

The MVFR0 and MVFR1 Register characteristics are:

Purpose Describes the features supported by the FPU.

Usage constraints The MVFR0 and MVFR1 Registers:

- are read-only registers
- are accessible in Privileged modes only.
- ARM recommends that any software attempting to determine the presence or absence of double-precision floating point hardware support uses the MVFR1 register.

Configurations Use this register if the device is configured as a Cortex-R5F processor.

Attributes See Table 11-7 on page 11-10 and Table 11-8 on page 11-10.

Figure 11-5 on page 11-10 shows the MVFR0 Register bit assignments.

31	28	27	24	23	20	19	16	15	12	11	8	7	4	3	0
RM			SV		SR		D		TE		DP		SP		RB

Figure 11-5 MVFR0 Register bit assignments

Table 11-7 shows the MVFR0 Register bit assignments.

Table 11-7 MVFR0 Register bit assignments

Bits	Name	Function
[31:28]	RM	Rounding modes supported: 0x1 = all VFP rounding modes supported.
[27:24]	SV	VFP short vector hardware support: 0x0 = not supported.
[23:20]	SR	VFP hardware square root: 0x1 = supported.
[19:16]	D	VFP hardware divide: 0x1 = supported.
[15:12]	TE	VFP exception trapping: 0x0 = only untrapped exception handling can be selected.
[11:8]	DP	Hardware support for VFP double-precision: 0x0 = no double-precision support present in hardware 0x2 = VFPv3 double-precision HW support present.
[7:4]	SP	Hardware single-precision support: 0x2 = VFPv3 supported.
[3:0]	RB	VFP register bank 16x64-bit register bank support: 0x1 = supported

Figure 11-6 shows the MVFR1 Register bit assignments.

31	28	27	24	23	20	19	16	15	12	11	8	7	4	3	0
Reserved			VFP HPFP		VFP A_SIMD		SP		I		LS		DN		FZ

Figure 11-6 MVFR1 Register bit assignments

Table 11-8 shows the MVFR1 Register bit assignments.

Table 11-8 MVFR1 Register bit assignments

Bits	Name	Function
[31:28]	-	Reserved
[27:24]	VFP HPFP	VFP half-precision conversions: 0x0 = no support.
[23:20]	VFP A_SIMD	Advanced SIMD half-precision conversions: 0x0 = no support.

Table 11-8 MVFR1 Register bit assignments (continued)

Bits	Name	Function
[19:16]	SP	Single-precision floating-point operation support for Advanced SIMD: 0x0 = no support.
[15:12]	I	Integer operation support for Advanced SIMD: 0x0 = no support.
[11:8]	LS	Load and store instruction support for Advanced SIMD: 0x0 = no support.
[7:4]	DN	Indicates whether the VFP hardware supports only Default NaN mode: 0x1 = hardware supports propagation of NaN values in addition to Default NaN mode.
[3:0]	FZ	Indicates whether the VFP hardware supports only Flush-to-Zero mode: 0x1 = hardware supports full denormal arithmetic in addition to Flush-to-Zero mode.

11.4 Modes of operation

The FPU provides three modes of operation to accommodate a variety of applications:

- *Full-compliance mode*
- *Flush-to-zero mode*
- *Default NaN mode.*

11.4.1 Full-compliance mode

In full-compliance mode, the FPU processes all operations according to the IEEE 754 standard in hardware.

11.4.2 Flush-to-zero mode

Setting the FZ bit, FPSCR[24], enables flush-to-zero mode. In this mode, the FPU treats all subnormal input operands of arithmetic CDP operations as zeros in the operation. Exceptions that result from a zero operand are signaled appropriately. VABS, VNEG, and VMOV are not considered arithmetic CDP operations and are not affected by flush-to-zero mode. A result that is *tiny*, as described in the IEEE 754 standard, for the destination-precision is smaller in magnitude than the minimum normal value *before rounding* and is replaced with a zero. The IDC flag, FPSCR[7], indicates when an input flush occurs. The UFC flag, FPSCR[3], indicates when a result flush occurs.

11.4.3 Default NaN mode

Setting the DN bit, FPSCR[25], enables default NaN mode. In this mode, the result of any operation that involves an input NaN, or that generated a NaN result, returns the default NaN. Propagation of the fraction bits is maintained only by VABS, VNEG, and VMOV operations. All other CDP operations ignore any information in the fraction bits of an input NaN.

11.5 Compliance with the IEEE 754 standard

When *Default NaN* (DN) and *Flush-to-Zero* (FZ) modes are disabled, the VFP functionality is compliant with the IEEE 754 standard in hardware. No support code is required to achieve this compliance.

See the *ARM Architecture Reference Manual* for information about VFP architecture compliance with the IEEE 754 standard.

11.5.1 Complete implementation of the IEEE 754 standard

The following operations from the IEEE 754 standard are not supplied by the VFP instruction set:

- remainder
- round floating-point number to integer-valued floating-point number
- binary-to-decimal conversions
- decimal-to-binary conversions
- direct comparison of single-precision and double-precision values.

For complete implementation of the IEEE 754 standard, VFP functionality must be augmented with library functions that implement these operations. See *Application Note 98, VFP Support Code* for information on the available library functions.

11.5.2 IEEE 754 standard implementation choices

Some of the implementation choices permitted by the IEEE 754 standard and used in the VFPv3 architecture are described in the *ARM Architecture Reference Manual*.

NaN handling

All single-precision and double-precision values with the maximum exponent field value and a nonzero fraction field are valid NaNs. A most significant fraction bit of zero indicates a *Signaling NaN* (SNaN). A one indicates a *Quiet NaN* (QNaN). Two NaN values are treated as different NaNs if they differ in any bit. Table 11-9 shows the default NaN values in both single-precision and double-precision.

Table 11-9 Default NaN values

	Single-precision	Double-precision
Sign	0	0
Exponent	0xFF	0x7FF
Fraction	bit [22] = 1, bits [21:0] are all zeros	bit [51] = 1, bits [50:0] are all zeros

Processing of input NaNs for ARM floating-point functionality and libraries is defined as follows:

- In full-compliance mode, NaNs are handled as described in the *ARM Architecture Reference Manual*. The hardware processes the NaNs directly for arithmetic CDP instructions. For data transfer operations, NaNs are transferred without raising the Invalid Operation exception. For the non-arithmetic CDP instructions, VABS, VNEG, and VMOV, NaNs are copied, with a change of sign if specified in the instructions, without causing the Invalid Operation exception.

- In default NaN mode, arithmetic CDP instructions involving NaN operands return the default NaN regardless of the fractions of any NaN operands. SNaNs in an arithmetic CDP operation set the IOC flag, FPSCR[0]. NaN handling by data transfer and non-arithmetic CDP instructions is the same as in full-compliance mode.

Table 11-10 summarizes the effects of NaN operands on instruction execution.

Table 11-10 QNaN and SNaN handling

Instruction type	Default NaN mode	With QNaN operand	With SNaN operand
Arithmetic CDP	Off	The QNaN or one of the QNaN operands, if there is more than one, is returned according to the rules given in the <i>ARM Architecture Reference Manual</i> .	IOC ^a set. The SNaN is quieted and the result NaN is determined by the rules given in the <i>ARM Architecture Reference Manual</i> .
	On	Default NaN returns.	IOC ^a set. Default NaN returns.
Non-arithmetic CDP	Off	NaN passes to destination with sign changed as appropriate.	
	On		
VFCMP	-	Unordered compare.	IOC set. Unordered compare.
VFCMPE	-	IOC set. Unordered compare.	IOC set. Unordered compare.
Load/store	Off	All NaNs transferred.	
	On		

a. IOC is the Invalid Operation exception flag, FPSCR[0].

Comparisons

Comparison results modify the flags in the FPSCR Register. You can use the VMRS APSR_nzcv, FPSCR instruction (formerly FMSTAT) to transfer the current flags from the FPSCR Register to the CPSR Register. See the *ARM Architecture Reference Manual* for mapping of IEEE 754 standard predicates to ARM conditions. The flags used are chosen so that subsequent conditional execution of ARM instructions can test the predicates defined in the IEEE 754 standard.

Underflow

The Cortex-R5F FPU uses the *before rounding* form of *tininess* and the *inexact result* form of *loss of accuracy* as described in the IEEE 754 standard to generate Underflow exceptions.

In flush-to-zero mode, results that are tiny before rounding, as described in the IEEE 754 standard, are flushed to a zero, and the UFC flag, FPSCR[3], is set. See the *ARM Architecture Reference Manual* for information on flush-to-zero mode.

When the FPU is not in flush-to-zero mode, operations are performed on subnormal operands. If the operation does not produce a tiny result, it returns the computed result, and the UFC flag, FPSCR[3], is not set. The IXC flag, FPSCR[4], is set if the operation is inexact. If the operation produces a tiny result, the result is a subnormal or zero value, and the UFC flag, FPSCR[3], is set if the result was also inexact.

11.5.3 Exceptions

The FPU implements the VFPv3 architecture and sets the cumulative exception status flag in the FPSCR register as required for each instruction. The FPU does not support user-mode traps. The exception enable bits in the FPSCR read-as-zero, and cannot be written. The processor also

has six output pins, **FPIXm**, **FPUFm**, **FPOFm**, **FPDZm**, **FPIDm**, and **FPIOm**, that each reflect the status of one of the cumulative exception flags. See *FPU signals* on page A-32 for a description of these outputs. You can mask each of these outputs masked by setting the corresponding bit in the Secondary Auxiliary Control Register.

See *c1, Auxiliary Control Register* on page 4-41 for more information.

Chapter 12

Debug

This chapter describes the processor debug unit. These features assist the development of application software, operating systems, and hardware. This chapter contains the following sections:

- *Debug systems* on page 12-2
- *About the debug unit* on page 12-3
- *Debug register interface* on page 12-5
- *Debug register descriptions* on page 12-10
- *Management registers* on page 12-33
- *Debug events* on page 12-40
- *Debug exception* on page 12-42
- *Debug state* on page 12-45
- *Cache debug* on page 12-50
- *External debug interface* on page 12-51
- *Using the debug functionality* on page 12-54
- *Debugging systems with energy management capabilities* on page 12-70.

12.1 Debug systems

The Cortex-R5 processor is one component of a debug system. Figure 12-1 shows a typical system.




Figure 12-1 Typical debug system

This typical system has three parts, described in the following sections:

- *Debug host*
- *Protocol converter*
- *Debug target.*

12.1.1 Debug host

The debug host is a computer, for example a personal computer, running a software debugger such as RealView™ Debugger. The debug host enables you to issue high-level commands such as setting breakpoint at a certain location, or examining the contents of a memory address.

12.1.2 Protocol converter

The debug host connects to the processor development system using an interface such as Ethernet. The messages broadcast over this connection must be converted to the interface signals of the debug target. A protocol converter performs this function, for example, RealView ICE.

12.1.3 Debug target

The debug target is the lowest level of the system. An example of a debug target is a development system with a Cortex-R5 test chip or a silicon part with a Cortex-R5 processor.

The debug target must implement some system support for the protocol converter to access the processor debug unit using the *Advanced Peripheral Bus* (APB) slave port.

12.2 About the debug unit

The processor debug unit assists in debugging software running on the processor. You can use the processor debug unit, in combination with a software debugger program, to debug:

- application software
- operating systems
- ARM processor-based hardware systems.

The debug unit enables you to:

- stop program execution
- examine and alter processor state
- examine and alter memory and peripheral state
- restart the processor.

You can debug software running on the processor in the following ways:

- *Halting debug-mode debugging*
- *Monitor debug-mode debugging*
- Trace debugging, see *ETM interface* on page 2-11.

The processor debug unit conforms to the ARMv7 debug architecture. For more information see the *ARM Architecture Reference Manual*.

12.2.1 Halting debug-mode debugging

When the processor debug unit is in Halting debug-mode, the processor halts program execution when a debug event, such as a breakpoint, occurs. When the processor is halted, an external debugger can examine and modify the processor state using the APB slave port. This debug mode is invasive to program execution.

12.2.2 Monitor debug-mode debugging

When the processor debug unit is in Monitor debug-mode, the processor takes a debug exception instead of halting. A special piece of software, a monitor target, can then take control to examine or alter the processor state. Monitor debug-mode is essential in real-time systems where the processor cannot be halted to collect information. Examples of these systems are engine controllers and servo mechanisms in hard drive controllers that cannot stop the code without physically damaging the components.

When debugging in Monitor debug-mode, the processor stops execution of the current program and starts execution of a monitor target. The state of the processor is preserved in the same manner as all ARM exceptions. The monitor target communicates with the debugger to access processor and coprocessor state, and to access memory contents and peripherals. Monitor debug-mode requires a debug monitor program to interface between the debug hardware and the software debugger.

12.2.3 Programming the debug unit

The processor debug unit is programmed using the APB slave interface. In a twin-CPU configuration, each CPU has its own APB slave interface and associated registers that operate independently of the other CPU. See Table 12-3 on page 12-6 for a complete list of memory-mapped debug registers accessible using the APB slave interface. Some features of the debug unit that you can access using the memory-mapped registers are:

- instruction address comparators for triggering breakpoints, see *Breakpoint Value Registers* on page 12-23 and *Breakpoint Control Registers* on page 12-24

- data address comparators for triggering watchpoints, see *Watchpoint Value Registers* on page 12-27 and *Watchpoint Control Registers* on page 12-27
- a bidirectional *Debug Communication Channel* (DCC), see *Debug communications channel* on page 12-55
- all other state information associated with the debug unit.

12.3 Debug register interface

You can access the processor debug register map using the APB slave port. The APB slave port conforms to the AMBA3 APBv3 standard as described in the *AMBA 3 APB Protocol Specification*. This is the only way to get full access to the processor debug capability. ARM recommends that if your system requires the processor to access its own debug registers, you choose a system interconnect structure that enables the processor to access the APB slave port by executing load and stores to an appropriate area of physical memory.

This section describes:

- *Coprocessor registers*
- *CP14 access permissions*
- *Coprocessor registers summary*
- *Memory-mapped registers* on page 12-6
- *Memory addresses for breakpoints and watchpoints* on page 12-8
- *Power domains* on page 12-8
- *Effects of resets on debug registers* on page 12-8
- *APB port access permissions* on page 12-8.

12.3.1 Coprocessor registers

Although most of the processor debug registers are accessible through the memory-mapped interface, there are several registers that you can access through a coprocessor interface. This is important for boot-strap access to the register file. It enables software running on the processor to identify the debug architecture version that the device implements.

12.3.2 CP14 access permissions

By default, you can access all CP14 debug registers from a nonprivileged mode. However, you can program the processor to disable user-mode access to all coprocessor registers using bit [12] of the DBGDSCR, see *CP14 c1, Debug Status and Control Register* on page 12-14 for more information. CP14 debug register accesses are always permitted when the processor is in debug state regardless of the processor mode.

Table 12-1 shows access to the CP14 debug registers.

Table 12-1 Access to CP14 debug registers

Debug state	Processor mode	DBGDSCR[12]	CP14 debug access
Yes	X	X	Permitted
No	User	b0	Permitted
No	User	b1	Not permitted ^a
No	Privileged	X	Permitted

a. Instructions attempting to access CP14 registers cause the processor to take an Undefined Instruction exception.

12.3.3 Coprocessor registers summary

Table 12-2 on page 12-6 shows a set of valid CP14 instructions for accessing the debug registers. All CP14 instructions not listed are Undefined.

Note

The CP14 debug instructions are defined as having Opcode_1 set to 0.

Table 12-2 CP14 debug registers summary

Instruction	Mnemonic	Description
MRC p14, 0, <Rd>, c0, c0, 0	DBGDIDR	Debug Identification Register. See <i>CP14 c0, Debug ID Register</i> on page 12-10.
MRC p14, 0, <Rd>, c1, c0, 0	DBGDRAR	Debug ROM Address Register. See <i>CP14 c0, Debug ROM Address Register</i> on page 12-12.
MRC p14, 0, <Rd>, c2, c0, 0	DBGDSAR	Debug Self Address Register. See <i>CP14 c0, Debug Self Address Offset Register</i> on page 12-12.
MRC p14, 0, <Rd>, c0, c5, 0 STC p14, c5, <addressing mode>	DBGDTRRXint	Host to Target Data Transfer Register. See <i>Data Transfer Register</i> on page 12-18.
MCR p14, 0, <Rd>, c0, c5, 0 LDC p14, c5, <addressing mode>	DBGDTRTXint	Target to Host Data Transfer Register. See <i>Data Transfer Register</i> on page 12-18.
MRC p14, 0, <Rd>, c0, c1, 0 MRC p14, 0, APSR_nzcv, c0, c1, 0	DBGDSCRint	Debug Status and Control Register. See <i>CP14 c1, Debug Status and Control Register</i> on page 12-14.

12.3.4 Memory-mapped registers

Table 12-3 shows the complete list of memory-mapped registers accessible at the APB slave interface.

Note

You must ensure that the base address of this 4KB register map is aligned to a 4KB boundary in physical memory.

Table 12-3 Debug memory-mapped registers

Offset (hex)	Register number	Access	Mnemonic	Description
0x000	c0	R	DBGDIDR	<i>CP14 c0, Debug ID Register</i> on page 12-10
0x004-0x014	c1-c5	R	-	RAZ
0x18	c6	RW	DBGWFAR	<i>Watchpoint Fault Address Register</i> on page 12-19
0x01C	c7	RW	DBGVCR	<i>Vector Catch Register</i> on page 12-19
0x020	c8	R	-	RAZ
0x024	c9	RW	DBGECCR	Not implemented in this processor. Reads as zero.
0x028	c10	RW	DBGDSCCR	<i>Debug State Cache Control Register</i> on page 12-21.
0x02C	c11	R	-	RAZ
0x030-0x07C	c12-c31	R	-	RAZ

Table 12-3 Debug memory-mapped registers (continued)

Offset (hex)	Register number	Access	Mnemonic	Description
0x080	c32	RW	DBGDTRRExt	<i>Data Transfer Register</i> on page 12-18.
0x084	c33	W	DBGITR	<i>Instruction Transfer Register</i> on page 12-22.
0x088	c34	RW	DBGDSCRExt	<i>CP14 c1, Debug Status and Control Register</i> on page 12-14.
0x08C	c35	RW	DBGDTRTExt	<i>Data Transfer Register</i> on page 12-18.
0x090	c36	W	DBGDRCR	<i>Debug Run Control Register</i> on page 12-22.
0x094-0x09C	c37-c39	R	-	RAZ.
0x0A0	c40	R	DBGPCSR	Not implemented on this processor. RAZ.
0x0A4	c41	R	DBGCIDSR	Not implemented on this processor. RAZ.
0x0A8-0x0FC	c42-c63	R	-	RAZ.
0x100-0x11C	c64-c71	RW	DBGBVR ^a	<i>Breakpoint Value Registers</i> on page 12-23.
0x120-0x13C	c72-c79	R	-	RAZ.
0x140-0x15C	c80-c87	RW	DBGBCR ^a	<i>Breakpoint Control Registers</i> on page 12-24.
0x160-0x17C	c88-c95	R	-	RAZ.
0x180-0x19C	c96-c103	RW	DBGWVR ^b	<i>Watchpoint Value Registers</i> on page 12-27.
0x1A0-0x1BC	c104-c111	R	-	RAZ.
0x1C0-0x1DC	c112-c119	RW	DBGWCR ^b	<i>Watchpoint Control Registers</i> on page 12-27.
0x1E0-0x1FC	c120-c127	R	-	RAZ.
0x200-0x2FC	c128-c191	R	-	RAZ.
0x300	c192	R	DBGOSLAR	Not implemented in this processor. Reads as zero.
0x304	c193	R	DBGOSLSR	<i>Operating System Lock Status Register</i> on page 12-29.
0x308	c194	R	DBGOSSRR	Not implemented in this processor. Reads as zero.
0x30C	c195	R	-	RAZ.
0x310	c196	RW	DBGPRCR	<i>Device Power-down and Reset Control Register</i> on page 12-31.
0x314	c197	R	DBGPRSR	<i>Device Power-down and Reset Status Register</i> on page 12-32.
0x318-0x7FC	c198-c511	R	-	RAZ.
0x800-0x8FC	c512-575	R	-	RAZ.
0x900-0xCFC	c576-c831	R	-	RAZ.
0xD00-0xDFC	c832-c895	R	-	<i>Processor ID Registers</i> on page 12-33.
0xE00-0xE7C	c896-c927	R	-	RAZ.
0xE80-0xEFC	c928-c959	-	-	Chapter 13 <i>Integration Test Registers</i> .
0xF00-0xFFC	c960-c1023	-	-	<i>Management registers</i> on page 12-33.

a. The actual number of registers depends on the number of breakpoints configured. For non-implemented breakpoints, the corresponding registers are RAZ.

- b. The actual number of registers depends on the number of watchpoints configured. For non-implemented watchpoints, the corresponding registers are RAZ.

12.3.5 Memory addresses for breakpoints and watchpoints

The *Vector Catch Register* (DBGVCR) sets breakpoints on exception vectors as instruction addresses.

The *Watchpoint Fault Address Register* (DBGWFAR) reads an address and a processor state dependent offset, +8 for ARM and +4 for Thumb.

12.3.6 Power domains

Cortex-R5 supports separate debug and core power domains to enable debug over power-down.

The following debug registers are implemented in the debug domain:

- *Debug ID Register* (DBGDIDR)
- *Debug Run Control Register* (DBGDRCR)
- *Device Power-down and Reset Control Register* (DBGPRCR)
- *Device Power-down and Reset Status Register* (DBGPRSR)
- CoreSight management registers.

All other implemented debug registers are in the core domain.

All accesses to core domain debug registers when the CPU is in Dormant or Shutdown modes return an error response on the CPU APB interface.

For more information about these registers and the split between core domain and debug domain registers, see the *ARM Architecture Reference Manual*.

12.3.7 Effects of resets on debug registers

The processor has the following reset signals that affect the processor debug logic:

nSYSPORESET

This signal resets all processor logic including the debug logic.

DBGRESETmn

This signal resets all the core domain debug logic.

PRESETDBGmn

This signal resets all debug domain debug logic.

See *Resets* on page 2-12 for more information on resets and reset requirements.

12.3.8 APB port access permissions

The restrictions for accessing the APB slave port are as follows:

Privilege of memory access

You must configure the system to disable accesses to the memory-mapped registers based on the privilege of the memory access.

Privilege of memory access permission

When non-privileged software attempts to access the APB slave port, the system must ignore the access or generate an error response to the access. You must implement this restriction at the system level because the APB protocol does not have a privileged or user control signal. You can choose to have the system either ignore the access or generate an error response.

You can place additional restrictions on memory transactions that are permitted to access the APB port. However, ARM does not recommend this.

Locks permission

You can lock the APB slave port so that access to some debug registers is restricted. ARM Architecture v7 defines two locks:

Software lock

The external debugger can set this lock to prevent software from modifying the debug registers settings. A debug monitor can also set this lock prior to returning control to the application to reduce the chance of erratic code changing the debug settings. When this lock is set, writes to all debug registers are ignored, except those generated by the external debugger, that override the lock. This is summarized in Table 12-4. For more information, see *Lock Access Register* on page 12-35.

OS Lock The processor does not support OS Lock.

———— Note ————

- These locks are set to their reset values only on reset of the debug logic, provided by **PRESETDBGmn**.
- You must set the **PADDRDBG31m** input signal to 1 for accesses originated from the external debugger for the Software Lock override feature to work.

Table 12-4 External debug interface access permissions

PADDRDBG31m	Lock	Registers			
		DBGDRCR, DBGPRCR, DBGPRSR	Other Debug registers	DBGLAR	Other registers
1	X ^a	OK ^b	OK ^c	OK ^c	OK ^c
0	1 ^c	WI ^d	WI ^c	OK ^c	WI ^c
0	0	OK ^c	OK ^c	OK ^c	OK ^c

- a. X indicates that the outcome does not depend on this condition.
b. OK indicates that the access succeeds.
c. DBGLSR[1] bit is set.
d. WI indicates that writes are ignored.

12.4 Debug register descriptions

Table 12-5 shows definitions of terms used in the register descriptions.

Table 12-5 Terms used in register descriptions

Term	Description
R	Read-only. Written values are ignored.
W	Write-only. This bit cannot be read. Reads return an Unpredictable value.
RW	Read or write.
RAZ	Read-As-Zero. Always zero when read.
RAO	Read-As-One. Always one when read.
SBZP	<i>Should-Be-Zero</i> (SBZ) or <i>Preserved</i> (P). Must be written as 0 or preserved by writing the same value previously read from the same fields on the same processor. These bits are usually reserved for future expansion.
UNP	A read from this bit returns an Unpredictable value.

12.4.1 CP14 c0, Debug ID Register

The DBGDIDR Register characteristics are:

Purpose Identifies the debug architecture version and specifies the number of debug resources that the processor implements.

Usage constraints The DBGDIDR is:

- in CP14 c0
- a 32 bit read-only register
- accessible in User and Privileged modes.

Configurations Available in all processor configurations.

Attributes See Table 12-6 on page 12-11.

Figure 12-2 shows the bit assignments.




Figure 12-2 DBGDIDR Register bit assignments

Table 12-6 shows the bit assignments.

Table 12-6 DBGDIDR Register bit assignments

Bits	Name	Function
[31:28]	WRP	Number of Watchpoint Register Pairs: b0000 = 1 WRP b0001 = 2 WRPs ... b0111 = 8 WRPs.
[27: 24]	BRP	Number of Breakpoint Register Pairs: b0001 = 2 BRPs b0010 = 3 BRPs ... b0111 = 8 BRPs.
[23:20]	Context	Number of Breakpoint Register Pairs (BRP) with context ID comparison capability: b0000 = 1 BRP has context ID comparison capability.
[19:16]	Debug architecture version	Debug architecture version: b0100 denotes ARMv7 Debug.
[15]	DEVID_imp	Indicates whether DBGDEVID is implemented. 0 = not implemented, register 1010 is reserved.
[14]	nSUHD_imp	RAZ.
[13]	PCSR_imp	RAZ.
[12]	SE_imp	RAZ.
[11:8]	-	RAZ.
[7: 4]	Variant	Implementation-defined variant number. This is the major revision number <i>n</i> in the <i>rn</i> part of the <i>rn</i> <i>pn</i> description of the product revision status.
[3: 0]	Revision	Implementation-defined revision number. This is the minor revision number <i>n</i> in the <i>pn</i> part of the <i>rn</i> <i>pn</i> description of the product revision status.

The values of the following fields of the DBGDIDR agree with the values in CP15 c0, Main ID Register:

- DBGDIDR[3:0] is the same as CP15 c0 bits [3:0]
- DBGDIDR[7:4] is the same as CP15 c0 bits [23:20].

See *c0, Main ID Register* on page 4-14 for more information of CP15 c0, Main ID Register.

The reason for duplicating these fields here is that the DBGDIDR is also accessible through the APB slave port. This enables an external debugger to determine the variant and revision numbers without stopping the processor.

To use the DBGDIDR, read CP14 c0 with:

MRC p14, 0, <Rd>, c0, c0, 0 ; Read DBGDIDR

12.4.2 CP14 c0, Debug ROM Address Register

The DBGDSAR Register characteristics are:

Purpose	Returns a 32-bit Debug ROM Address Register value. This is the address that indicates where in memory a debug monitor can locate the debug bus ROM specified by the CoreSight™ multiprocessor trace and debug architecture. Returns a 32-bit Debug ROM Address Register value. This is the address that indicates where in memory a debug monitor can locate the debug bus ROM specified by the CoreSight™ multiprocessor trace and debug architecture.
Usage constraints	The DBGDRAR is: • in CP14 c0, sub-register c1 • a 32 bit read-only register • accessible in User and Privileged modes.
Configurations	Available in all processor configurations.
Attributes	See Table 12-7.

Figure 12-3 shows the bit assignments.




Figure 12-3 DBGDRAR Register bit assignments

Table 12-7 shows the bit assignments.

Table 12-7 DBGDRAR Register bit assignments

Bits	Name	Function
[31:12]	Debug bus ROM address	Indicates bits [31:12] of the debug bus ROM address.
[11: 2]	-	SBZ.
[1:0]	Valid bits	Indicates that the ROM address is valid. Reads b11 if DBGROMADDRV is set to 1, otherwise reads b00. DBGROMADDRV must be set to 1 if DBGROMADDR[31:12] is set to a valid value.

To use the DBGDRAR, read CP14 c0 with:

MRC p14, 0, <Rd>, c1, c0, 0 ; Read DBGDRAR

12.4.3 CP14 c0, Debug Self Address Offset Register

The DBGDSAR Register characteristics are:

Purpose	The DBGDSAR is a read-only register that returns a 32-bit offset value from the Debug ROM Address Register to the address of the CPU debug registers. You can configure the address read in this register during integration using the DBGSELFADDRm[31:12] and DBGSELFADDRVm inputs. DBGSELFADDRVm must be tied off to 1 if DBGSELFADDRm[31:12] is tied off to a valid value.
----------------	---

Usage constraints The DBGDSAR is:

- in CP14 c0, sub-register c2
- a 32 bit read-only register
- accessible in User and Privileged modes.

Configurations Available in all processor configurations.

Attributes See Table 12-8.

Figure 12-4 shows the bit assignments.




Figure 12-4 DBGDSAR Register bit assignments

Table 12-8 shows the bit assignments.

Table 12-8 DBGDSAR Register bit assignments

Bits	Name	Function
[31:12]	Debug bus self address offset value	Indicates bits [31:12] of the two’s complement offset from the debug ROM physical address to the physical address where the debug registers are mapped.
[11: 2]	-	UNP on reads, SBZP on writes.
[1:0]	Valid bits	Reads b11 if DBGSELFADDRV_m is set to 1, otherwise reads b00. DBGSELFADDRV_m must be set to 1 if DBGSELFADDR_m[31:12] is set to a valid value.

To use the DBGDSAR, read CP14 c0 with:

```
MRC p14, 0, <Rd>, c2, c0, 0 ; Read DBGDSAR
```


12.4.4 CP14 c1, Debug Status and Control Register

The DBGDSCR Register characteristics are:

Purpose Contains status and control information about the debug unit.

Usage constraints See *DTR access mode* on page 12-17.

Configurations Available in all processor configurations.

Attributes See Table 12-9.

Figure 12-5 shows the bit assignments.




Figure 12-5 DBGDSCR Register bit assignments

Table 12-9 shows the bit assignments.

Table 12-9 DBGDSCR Register bit assignments

Bits	Name	Function
[31]	-	RAZ on reads, SBZP on writes.
[30]	RXfull	The RXfull flag: 0 = Read-DTR, DBGDTRRX, empty, reset value 1 = Read-DTR, DBGDTRRX, full. When set, this flag indicates to the processor that there is data available to read at the DBGDTRRXint. It is automatically set on writes to the DBGDTRRXext by the debugger, and is cleared when the processor reads the CP14 DTR. If the flag is not set, the DBGDTRRXint returns an Unpredictable value.
[29]	TXfull	The TXfull flag: 0 = Write-DTR, DBGDTRTX, empty, reset value 1 = Write-DTR, DBGDTRTX, full. When clear, this flag indicates to the processor that the DBGDTRTXint is ready to receive data. It is automatically cleared on reads of the DBGDTRTXext by the debugger, and is set when the processor writes to the CP14 DTR. If this bit is set and the processor attempts to write to the DBGDTRTXint, the register contents are overwritten and the TXfull flag remains set.
[28:26]	-	RAZ on reads, SBZP on writes.
[25]	PipeAdv	Sticky pipeline advance read-only bit. This bit enables the debugger to detect whether the processor is idle. In some situations, this might mean that the system bus port is deadlocked. This bit is set to 1 when the processor pipeline retires one instruction. It is cleared by a write to DBGDRCR[3]. 0 = no instruction has completed execution since the last time this bit was cleared 1 = an instruction has completed execution since the last time this bit was cleared.

Table 12-9 DBGDSCR Register bit assignments (continued)

Bits	Name	Function
[24]	InstrCompl_1	Instruction complete read-only bit. This flag determines whether the processor has completed execution of an instruction issued through the APB port. 0 = processor is executing an instruction fetched from the DBGITR Register 1 = processor is not executing an instruction fetched from the DBGITR Register. When the APB port reads the DBGDSCR and this bit is clear, then a subsequent write to the DBGITR Register is ignored unless DBGDSCR[21:20] is not equal to 0. If DBGDSCR[21:20] is not equal to 0, the DBGITR write stalls until the processor completes execution of the current instruction. If the processor is not in debug state, then the value read for this flag is Unpredictable. The flag is set to 1 on entry to debug state.
[23:22]	-	RAZ on reads, SBZP on writes.
[21:20]	ExtDCCmode	DTR access mode. You can use this field to optimize DTR traffic between a debugger and the processor. b00 = Non-blocking mode, this is the reset value b01 = Stall mode b10 = Fast mode b11 = Reserved. Note - This field only affects the behavior of DBGDSCRExt, DBGDTRRXext, DBGDTRTXext, and DBGITR accesses through the APB port, and not through CP14 debug instructions. - Non-blocking mode is the default setting. Improper use of the other modes might result in the debug access bus becoming deadlocked. See <i>DTR access mode</i> on page 12-17 for more information.
[19]	ADAdiscard	The Asynchronous Aborts Discarded bit is set when the processor is in debug state and is cleared on exit from debug state. While this bit is set, the processor does not take asynchronous Data Aborts, instead, the sticky asynchronous Data Abort bit is set to 1. 0 = do not discard asynchronous Data Aborts 1 = discard asynchronous Data Aborts and set ADABORT_I.
[18]	NS	RAZ on reads, SBZP on writes.
[17]	SPNIDdis	This bit is the inverse of bit [6] of the DBGAUTHSTATUS, see <i>Authentication Status Register</i> on page 12-30.
[17]	SPIDdis	This bit is the inverse of bit [4] of the DBGAUTHSTATUS, see <i>Authentication Status Register</i> on page 12-30.
[15]	MDBGGen	The Monitor debug-mode enable bit: 0 = Monitor debug-mode disabled, this is the reset value 1 = Monitor debug-mode enabled. If Halting debug-mode is enabled through bit [14], then the processor is in Halting debug-mode regardless of the value of bit [15]. If the external interface input DBGENm is LOW, this bit reads as 0. The programmed value is masked until DBGENm is HIGH, and at that time the read value reverts to the programmed value.
[14]	HDBGGen	The Halting debug-mode enable bit: 0 = Halting debug-mode disabled, this is the reset value 1 = Halting debug-mode enabled. If the external interface input DBGENm is LOW, this bit reads as 0. The programmed value is masked until DBGENm is HIGH, and at that time the read value reverts to the programmed value.

Table 12-9 DBGDSCR Register bit assignments (continued)

Bits	Name	Function
[13]	ITRen	Execute ARM instruction enable bit: 0 = disabled, this is the reset value 1 = enabled. If this bit is set and an DBGITR write succeeds, the processor fetches an instruction from the DBGITR for execution. If this bit is set to 1 when the processor is not in debug state, the behavior of the processor is Unpredictable.
[12]	UDCCdis	CP14 debug user access disable control bit: 0 = CP14 debug user access enable, this is the reset value 1 = CP14 debug user access disable. If this bit is set and a User mode process attempts to access any CP14 debug registers, an Undefined Instruction exception is taken.
[11]	IntDis	Interrupts disable bit: 0 = interrupts enabled, this is the reset value 1 = interrupts disabled. If this bit is set, the nIRQm and nFIQm input signals are inhibited. The external debugger can optionally use this bit to execute pieces of code in normal state as part of the debugging process and avoid having an interrupt taking control of the program flow.
[10]	DbgAck	Force Debug Acknowledge bit. If this bit is set to 1, the DBGACKm output signal is forced HIGH, regardless of the processor state. The external debugger can optionally use this bit to execute pieces of code in normal state as part of the debugging process for the system to behave as if the processor is in debug state. Some systems rely on DBGACKm to determine whether data accesses are application or debugger generated. This bit is 0 on reset.
[9]	-	RAZ on reads, SBZP on writes.
[8]	UND_I	Sticky Undefined bit: 0 = no Undefined Instruction exception occurred in debug state since the last time this bit was cleared 1 = an Undefined Instruction exception occurred while in debug state since the last time this bit was cleared. This flag detects Undefined Instruction exceptions generated by instructions issued to the processor through the DBGITR. This bit is set to 1 when an Undefined Instruction exception occurs while the processor is in debug state and is cleared by writing a 1 to DBGDRCR[2].
[7]	ADABORT_I	Sticky asynchronous Data Abort bit: 0 = no asynchronous Data Aborts occurred since the last time this bit was cleared 1 = an asynchronous Data Abort occurred since the last time this bit was cleared. This flag detects asynchronous Data Aborts triggered by instructions issued to the processor through the DBGITR. This bit is set to 1 when an asynchronous Data Abort occurs while the processor is in debug state and is cleared by writing a 1 to DBGDRCR[2].
[6]	SDABORT_I	Sticky synchronous Data Abort bit: 0 = no synchronous Data Abort occurred since the last time this bit was cleared 1 = a synchronous Data Abort occurred since the last time this bit was cleared. This flag detects synchronous Data Aborts generated by instructions issued to the processor through the DBGITR. This bit is set to 1 when a synchronous Data Abort occurs while the processor is in debug state and is cleared by writing to the DBGDRCR[2].

Table 12-9 DBGDSCR Register bit assignments (continued)

Bits	Name	Function
[5:2]	MOE	Method of entry bits: b0000 = a DBGDRCR[0] halting debug event occurred b0001 = a breakpoint occurred b0100 = an EDBGRQm halting debug event occurred b0011 = a BKPT instruction occurred b1010 = a synchronous watchpoint occurred others = reserved. These bits are set to indicate any of: - the cause of a debug exception - the cause for entering debug state. A Prefetch Abort or Data Abort handler must check the value of the CP15 Fault Status Register to determine whether a debug exception occurred and then use these bits to determine the specific debug event.
[1] ^a	RESTARTED	CPU restarted bit: 0 = The processor is exiting debug state. 1 = The processor has exited debug state. This is the reset value. The debugger can poll this bit to determine when the processor responds to a request to leave debug state.
[0] ^a	HALTED	CPU halted bit: 0 = The processor is in normal state. This is the reset value. 1 = The processor is in debug state. The debugger can poll this bit to determine when the processor has entered debug state.

- a. These bits always reflect the status of the processor, therefore they only have a reset value if the particular reset event affects the processor. For example, a **PRESETDBGmn** event leaves these bits unchanged and a processor reset event such as **nSYSPORESET** sets DBGDSCR[18] to a 0 and DBGDSCR[1:0] to 10.

To use the DBGDSCR, read or write CP14 c1 with:

```
MRC p14, 0, <Rd>, c0, c1, 0 ; Read DBGDSCR
MCR p14, 0, <Rd>, c0, c1, 0 ; Write DBGDSCR
```

DTR access mode

You can use the ExtDCCmode field to optimize data transfer between a debugger and the processor.

The DCC access mode can be one of the following:

- Nonblocking. This is the default mode.
- Stall.
- Fast.

In Non-blocking mode, reads from DBGDTRTXext and writes to DBGDTRRXext and DBGITR are ignored if the appropriate latched ready flag is not in the ready state. These latched flags are updated on DBGDSCR reads. The following applies:

- writes to DBGDTRRXext are ignored if RXfull_1 is set to b1
- reads from DBGDTRTXext are ignored, and return an Unpredictable value, if TXfull_1 is set to b0

- writes to DBGITR are ignored if InstrCompl_1 is set to b0
- following a successful write to DBGDTRRXext, RXfull and RXfull_1 are set to b1
- following a successful read from DBGDTRTXext, TXfull and TXfull_1 are cleared to b0
- following a successful write to DBGITR, the internal flags InstrCompl and InstrCompl_1 are cleared to b0.

Debuggers accessing these registers must first read DBGDSCRExt. This has the side-effect of copying RXfull and TXfull to RXfull_1 and TXfull_1. The debugger must then:

- write to the DBGDTRRXext if the RXfull flag was b0 (RXfull_1 is b0)
- read from the DBGDTRTXext if the TXfull flag was b1 (TXfull_1 is b1)
- write to the DBGITR if the InstrCompl_1 flag was b1.

However, debuggers can issue both actions together and later determine from the read DBGDSCR value whether the operations were successful.

In Stall mode, the APB accesses to DBGDTRRXext, DBGDTRTXext, and DBGITR stall under the following conditions:

- writes to DBGDTRRXext are stalled until RXfull is cleared
- writes to DBGITR are stalled until InstrCompl is set
- reads from DBGDTRTXext are stalled until TXfull is set.

Fast mode is similar to Stall mode except that in Fast mode, the processor fetches an instruction from the DBGITR when a DBGDTRRXext write or DBGDTRTXext read succeeds. In Stall mode and Nonblocking mode, the processor fetches an instruction from the DBGITR when a DBGITR write succeeds.

12.4.5 Data Transfer Register

The DTR consists of two separate physical registers:

- the DBGDTRRX (Read Data Transfer Register)
- the DBGDTRTX (Write Data Transfer Register).

The register accessed is dependent on the instruction used:

- writes, MCR and LDC instructions, access the DBGDTRTX
- reads, MRC and STC instructions, access the DBGDTRRX.

————— **Note** —————

Read and write are used with respect to the processor.

For information on the use of these registers with the TXfull flag and RXfull flag, see *Debug communications channel* on page 12-55. The Data Transfer Register, bits [31:0] contain the data to be transferred.

Table 12-10 shows how the bit values correspond with the DBGDTRRX and DBGDTRTX functions.

Table 12-10 Data Transfer Register bit assignments

Bits	Name	Function
[31:0]	Data	Reads the Data Transfer Register. This is read-only for the CP14 interface.
Note Reads of the DBGDTRRXint through the coprocessor interface cause the TXfull flag to be cleared. However, reads of the DBGDTRRXext through the APB port do not affect this flag.		
[31:0]	Data	Writes the Data Transfer Register. This is write-only for the CP14 interface.
Note Writes to the DBGDTRTXint through the coprocessor interface cause the RXfull flag to be set. However, writes to the DBGDTRTXext through the APB port do not affect this flag.		

12.4.6 Watchpoint Fault Address Register

The DBGWFAR Register characteristics are:

Purpose Holds the address of the instruction that triggers the watchpoint.

Usage constraints There are no usage constraints.

Configurations Available in all processor configurations.

Attributes See Table 12-11.

Figure 12-6 shows the bit assignments.




Figure 12-6 DBGWFAR Register bit assignments

Table 12-11 shows the bit assignments.

Table 12-11 DBGWFAR Register bit assignments

Bits	Name	Function
[31:1]	Address	This is the address of the watchpointed instruction. When a watchpoint occurs in ARM state, the DBGWFAR contains the address of the instruction causing it plus an offset of 0x8. When a watchpoint occurs in Thumb state, the offset is plus 0x4.
[0]	-	RAZ.

12.4.7 Vector Catch Register

The DBGVCR Register characteristics are:

Purpose Controls efficient exception vector catching.

- Usage constraints**
- If one of the bits in this register is set and the instruction at the corresponding vector is committed for execution, the processor either enters debug state or takes a debug exception.
 - Under this model, any prefetch from an exception vector can trigger a vector catch, not only the ones because of exception entries. An explicit branch to an exception vector might generate a vector catch debug event.
 - If any of the bits are set when the processor is in Monitor debug-mode, then the processor ignores the setting and does not generate a vector catch debug event. This prevents the processor entering an unrecoverable state. The debugger must program these bits to zero when Monitor debug-mode is selected and enabled to ensure forward-compatibility.

Configurations Available in all processor configurations.

Attributes See Table 12-12.

as Figure 12-7 shows.




Figure 12-7 DBGVCR Register bit assignments

Table 12-12 shows the bit assignments.

Table 12-12 DBGVCR Register bit assignments

Bits	Name	Reset value	Normal address	High vectors address	Function	Access
[31:8]	-	0	-	-	Do not modify on writes. On reads, the value returns zero.	RAZ or SBZP
[7]	FIQ	0	0x0000001C	0xFFFF001C	Vector catch enable, FIQ.	RW
[6]	IRQ	-	0x00000018 ^a	0xFFFF0018 ^a	Vector catch enable, IRQ.	-
[5]	-	0	-	-	Do not modify on writes. On reads, the value returns zero.	RAZ or SBZP
[4]	Data Abort	0	0x00000010	0xFFFF0010	Vector catch enable, data abort.	RW
[3]	Prefetch Abort	0	0x0000000C	0xFFFF000C	Vector catch enable, prefetch abort.	RW

Table 12-12 DBGVCR Register bit assignments (continued)

Bits	Name	Reset value	Normal address	High vectors address	Function	Access
[2]	SVC	0	0x00000008	0xFFFF0008	Vector catch enable, SVC.	RW
[1]	Undefined	0	0x00000004	0xFFFF0004	Vector catch enable, Undefined Instruction.	RW
[0]	Reset	0	0x00000000	0xFFFF0000	Vector catch enable, reset.	RW

a. If the VIC interface is enabled, the address is the last IRQ handler address supplied by the VIC, whether or not high vectors are in use.

12.4.8 Debug State Cache Control Register

The DBGDSCCR Register characteristics are:

Purpose Controls the L1 cache behavior when the processor is in debug state.

Usage constraints For information on the usage model of the DBGDSCCR register, see *Cache debug* on page 12-50.

Configurations Available in all processor configurations.

Attributes See Table 12-13.

Figure 12-8 shows the bit assignments.




Figure 12-8 DBGDSCCR Register bit assignments

Table 12-13 shows the bit assignments.

Table 12-13 DBGDSCCR Register bit assignments

Bits	Name	Reset value	Description
[31:3]	-	0	Reserved. Do not modify on writes. On reads, the value returns zero.
[2]	nWT	0	Not write-through: 1 = normal operation of regions marked as write-back in debug state 0 = force write-through behavior for regions marked as write-back in debug state, this is the reset value.
[1]	nIL	0	Instruction cache line-fill: 1 = normal operation of L1 instruction cache in debug state 0 = L1 instruction cache line-fills disabled in debug state, this is the reset value.
[0]	nDL	0	Data cache line-fill: 1 = normal operation of L1 data cache in debug state 0 = L1 data cache line-fills disabled in debug state, this is the reset value.

12.4.9 Instruction Transfer Register

The DBGITR enables the external debugger to feed instructions into the processor for execution while in debug state. The DBGITR is a write-only register. Reads from the DBGITR return an Unpredictable value.

The Instruction Transfer Register, bits [31:0] contain the ARM instruction for the processor to execute while in debug state. The reset value of this register is Unpredictable.

———— Note ————

Writes to the DBGITR when the processor is not in debug state or the DBGDSCR[13] execute instruction enable bit is cleared are Unpredictable. When an instruction is issued to the processor, the debug unit prevents the next instruction from being issued until the DBGDSCR[25] instruction complete bit is set.

12.4.10 Debug Run Control Register

The DBGDSCR Register characteristics are:

- Purpose**
- Requests the processor to enter or leave debug state.
 - Clears the sticky exception bits present in the DBGDSCR.

Usage constraints The DBGDRCR is a write-only register.

Configurations Available in all processor configurations.

Attributes See Table 12-14 on page 12-23.

Figure 12-9 shows the bit assignments.

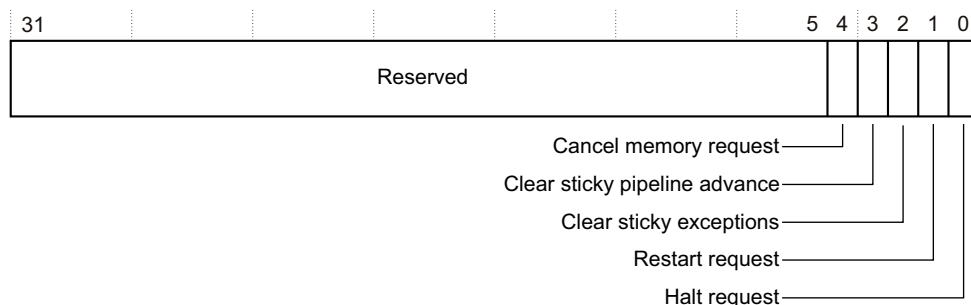


Figure 12-9 DBGDRCR Register bit assignments

Table 12-14 shows the bit assignments.

Table 12-14 DBGDRCR Register functions

Bits	Name	Function
[31:5]	-	RAZ.
[4]	Cancel memory requests	If 1 is written to this bit, the processor abandons any pending memory transactions until it can enter debug state. Debug state entry is the acknowledge event that clears this request. Abandoned transactions have the following behavior: - abandoned stores might write an Unpredictable value to the target address - abandoned loads return an Unpredictable value to the register bank. An abandoned transaction does not cause any exception. Additional instruction fetches or data accesses after the processor entered debug state have an Unpredictable behavior. This bit enables the debugger to progress on a deadlock so the processor can enter debug state. For a debug state entry to occur, a halting debug event must be requested before this bit is set. If you write a 1 to this bit when DBGENm is LOW, the write has no effect.^a
[3]	Clear sticky pipeline advance	Writing a 1 to this bit clears DBGDSCR[25].
[2]	Clear sticky exceptions	Writing a 1 to this bit clears DBGDSCR[8:6].
[1]	Restart request	Writing a 1 to this bit requests that the processor leaves debug state. This request is held until the processor exits debug state. When the debugger makes this request, it polls DBGDSCR[1] until it reads 1. This bit always reads as zero. Writes are ignored when the processor is not in debug state.
[0]	Halt request	Writing a 1 to this bit triggers a halting debug event, that is, a request that the processor enters debug state. This request is held until the debug state entry occurs. When the debugger makes this request, it must poll DBGDSCR[0] until it reads 1. This bit always reads as zero. Writes are ignored when the processor is already in debug state.

a. Entry into debug state is not expected to be recoverable.

12.4.11 Breakpoint Value Registers

Each DBGBVR is associated with a *Breakpoint Control Register* (DBGBCR). DBGBCR_y is the corresponding control register for DBGBVR_y.

A pair of breakpoint registers, DBGBVR_y/DBGBCR_y, is called a *Breakpoint Register Pair* (BRP). DBGBVR0-7 are paired with DBGBCR0-7 to make BRP0-7.

The breakpoint value contained in this register corresponds to either an instruction address or a context ID. Breakpoints can be set on:

- an instruction address
- a context ID value
- an instruction address and context ID pair.

For an instruction address and context ID pair, two BRPs must be linked. A debug event is generated when both the instruction address and the context ID pair match at the same time.

Table 12-15 shows how the bit values correspond with the Breakpoint Value Registers functions.

Table 12-15 Breakpoint Value Register bit assignments

Bits	Reset value	Description
[31:0]	0x0	Breakpoint value

Note

- Only BRP_n supports context ID comparison, where n+1 is the number of breakpoint register pairs implemented in the processor.
- Bits [1:0] of Registers DBGVR0 to DBGVR(n-1) are Do Not Modify on writes and Read-As-Zero because these registers do not support context ID comparisons.
- The contents of the CP15 Context ID Register give the context ID value for a DBGVR to match. For information on the Context ID Register, see Chapter 4 *System Control*.

12.4.12 Breakpoint Control Registers

The DBGBCR Register characteristics are:

Purpose Contains the necessary control bits for setting:

- breakpoints
- linked breakpoints.

Usage constraints There are no usage constraints.

Configurations Available in all processor configurations.

Attributes See Table 12-16 on page 12-25.

Figure 12-10 shows the bit assignments.




Figure 12-10 DBGBCR Registers bit assignments

Table 12-16 shows the bit assignments.

Table 12-16 Breakpoint Control Register bit assignments

Bits	Name	Function
[31:29]	-	Do not modify on writes. On reads, the value returns zero.
[28:24]	Breakpoint address mask	This field sets a breakpoint on a range of addresses by masking lower order address bits out of the breakpoint comparison.^a b00000 = no mask b00001 = Reserved b00010 = Reserved b00011 = 0x00000007 mask for instruction address b00100 = 0x0000000F mask for instruction address b00101 = 0x0000001F mask for instruction address ... b11111 = 0x7FFFFFFF mask for instruction address.
[23]	-	-
[22:20]	M	Meaning of DBGBVR: b000 = instruction address match b001 = linked instruction address match b010 = unlinked context ID b011 = linked context ID b100 = instruction address mismatch b101 = linked instruction address mismatch b11x = Reserved. For more information, see Table 12-17 on page 12-26
[19:16]	Linked BRP number	The binary number encoded here indicates another BRP to link this one with. Note • if a BRP is linked with itself, it is Unpredictable whether a breakpoint debug event is generated • if this BRP is linked to another BRP that is not configured for linked context ID matching, it is Unpredictable whether a breakpoint debug event is generated.
[15:14]	Secure state access control	RAZ or SBZP.
[13:9]	-	Do not modify on writes. On reads, the value returns zero.

Table 12-16 Breakpoint Control Register bit assignments (continued)

Bits	Name	Function
[8:5]	Byte address select	For breakpoints programmed to match an instruction address, the debugger must write a word-aligned address to the DBGBCR. You can then use this field to program the breakpoint so it hits only if certain byte addresses are accessed.^b If the BRP is programmed for instruction address match: b0000 = the breakpoint never hits bxxx1 = the breakpoint hits if the byte at address (DBGBCR & 0xFFFFF0) + 0 is accessed bxx1x = the breakpoint hits if the byte at address (DBGBCR & 0xFFFFF0) + 1 is accessed bx1xx = the breakpoint hits if the byte at address (DBGBCR & 0xFFFFF0) + 2 is accessed b1xxx = the breakpoint hits if the byte at address (DBGBCR & 0xFFFFF0) + 3 is accessed b1111 = the breakpoint hits if any of the four bytes starting at address (DBGBCR & 0xFFFFF0) + 0 is accessed. If the BRP is programmed for instruction address mismatch, the breakpoint hits where the corresponding instruction address breakpoint does not hit, that is, the range of addresses covered by an instruction address mismatch breakpoint is the negative image of the corresponding instruction address breakpoint. If the BRP is programmed for context ID comparison, this field must be set to b1111. Otherwise, breakpoint and watchpoint debug events might not be generated as expected.
[4:3]	-	-
[2:1]	S	Supervisor access control. The breakpoint can be conditioned on the mode of the processor: b00 = User, System, or Supervisor b01 = Privileged b10 = User b11 = any.
[0]	B	Breakpoint enable: 0 = Breakpoint disabled. This is the reset value. 1 = Breakpoint enabled.

- a. If DBGBCR[28:24] is not set to b00000, then DBGBCR[8:5] must be set to b1111. Otherwise the behavior is Unpredictable. In addition, if DBGBCR[28:24] is not set to b00000, then the corresponding DBGBCR bits that are not being included in the comparison Should Be Zero. Otherwise the behavior is Unpredictable. If this BRP is programmed for context ID comparison, this field must be set to b00000. Otherwise the behavior is Unpredictable. There is no encoding for a full 32-bit mask but the same effect of a *break anywhere* breakpoint can be achieved by setting DBGBCR[22] to 1 and DBGBCR[8:5] to b0000.
- b. Writing a value to DBGBCR[8:5] so that DBGBCR[8] is not equal to DBGBCR[7] or DBGBCR[6] is not equal to DBGBCR[5] has Unpredictable results.

Table 12-17 Meaning of DBGBCR bits [22:20]

DBGBCR[22:20]	Meaning
b000	The corresponding DBGBCR[31:2] is compared against the instruction address bus and the state of the processor against this DBGBCR. It generates a breakpoint debug event on a joint instruction address and state match.
b001	The corresponding DBGBCR[31:2] is compared against the instruction address bus and the state of the processor against this DBGBCR. This BRP is linked with the one indicated by DBGBCR[19:16] linked BRP field. They generate a breakpoint debug event on a joint instruction address, context ID, and state match.
b010	The corresponding DBGBCR[31:0] is compared against CP15 Context ID Register, c13 and the state of the processor against this DBGBCR. This BRP is not linked with any other one. It generates a breakpoint debug event on a joint context ID and state match. For this BRP, DBGBCR[8:5] must be set to b1111. Otherwise it is Unpredictable whether a breakpoint debug event is generated.

Table 12-17 Meaning of DBGBVR bits [22:20] (continued)

DBGBVR[22:20]	Meaning
b011	The corresponding DBGBVR[31:0] is compared against CP15 Context ID Register, c13. This BRP links another BRP (of the DBGBCR[21:20]=b01 type), or WRP (with DBGWCR[20]=b1). They generate a breakpoint or watchpoint debug event on a joint instruction address or data address and context ID match. For this BRP, DBGBCR[8:5] must be set to b1111, DBGBCR[15:14] must be set to b00, and DBGBCR[2:1] must be set to b11. Otherwise it is Unpredictable whether a breakpoint debug event is generated.
b100	The corresponding DBGBVR[31:2] and DBGBCR[8:5] are compared against the instruction address bus and the state of the processor against this DBGBCR. It generates a breakpoint debug event on a joint instruction address mismatch and state match.
b101	The corresponding DBGBVR[31:2] and DBGBCR[8:5] are compared against the instruction address bus and the state of the processor against this DBGBCR. This BRP is linked with the one indicated by DBGBCR[19:16] linked BRP field. It generates a breakpoint debug event on a joint instruction address mismatch, state and context ID match.
b11x	Reserved. The behavior is Unpredictable.

12.4.13 Watchpoint Value Registers

Each DBGWVR is associated with a *Watchpoint Control Register* (DBGWCR). DBGWCR_y is the corresponding register for DBGWVR_y.

A pair of watchpoint registers, DBGWVR_y and DBGWCR_y, is called a *Watchpoint Register Pair* (WRP). DBGWVR0-7 are paired with DBGWCR0-7 to make WRP0-7.

The watchpoint value contained in the DBGWVR always corresponds to a data address and can be set either on:

- a data address
- a data address and context ID pair.

For a data address and context ID pair, a WRP and the BRP with context ID comparison capability must be linked. A debug event is generated when both the data address and the context ID pair match simultaneously.

Table 12-18 shows the bit assignments.

Table 12-18 Watchpoint Value Register bit assignments

Bits	Description
[31:2]	Watchpoint address.
[1:0]	Reserved. Do not modify on writes. On reads, the value returns zero.

12.4.14 Watchpoint Control Registers

The DBGWCR Register characteristics are:

Purpose	Contains the necessary control bits for setting: • watchpoints • linked watchpoints.
Usage constraints	There are no usage constraints.
Configurations	Available in all processor configurations.

Attributes See Table 12-19.

Figure 12-11 shows the bit assignments.




Figure 12-11 DBGWCR Register bit assignments

Table 12-19 shows the bit assignments.

Table 12-19 DBGWCR Register bit assignments

Bits	Name	Function
[31:29]	-	Do not modify on writes. On reads, the value returns zero.
[28:24]	Watchpoint address mask	This field watches a range of addresses by masking lower order address bits out of the watchpoint comparison. b00000 = no mask b00001 = Reserved b00010 = Reserved b00011 = 0x00000007 mask for data address b00100 = 0x0000000F mask for data address b00101 = 0x0000001F mask for data address ... b11111 = 0x7FFFFFFF mask for data address. Note - If DBGWCR[28:24] is not set to b00000, then DBGWCR[12:5] must be set to b11111111. Otherwise the behavior is Unpredictable. - If DBGWCR[28:24] is not set to b00000, then the corresponding DBGWVR bits that are not being included in the comparison Should Be Zero. Otherwise the behavior is Unpredictable. - To watch for a write to any byte in an 8-byte aligned object of size 8 bytes, ARM recommends that a debugger sets DBGWCR[28:24] to b00111, and DBGWCR[12:5] to b11111111. This is compatible with both ARMv7 debug compliant implementations that have an 8-bit DBGWCR[12:5] and with those that have a 4-bit DBGWCR[8:5] byte address select field.
[23:21]	-	Do not modify on writes. On reads, the value returns zero.
[20]	E	Enable linking bit: 0 = linking disabled 1 = linking enabled. When this bit is set, this watchpoint is linked with the context ID holding BRP selected by the linked BRP field.
[19:16]	Linked BRP	Linked BRP number. The binary number encoded here indicates a context ID holding BRP to link this WRP with. If this WRP is linked to a BRP that is not configured for linked context ID matching, it is Unpredictable whether a watchpoint debug event is generated.
[15:14]	Secure state access control	RAZ or SBZP.

Table 12-19 DBGWCR Register bit assignments (continued)

Bits	Name	Function
[13]	-	Appear as zero when read. Do not modify on writes.
[12:5]	Byte address select	The DBGWVR is programmed with word-aligned address. You can use this field to program the watchpoint so it only hits if certain byte addresses are accessed: b00000000 The watchpoint never hits. bxxxxxx1 The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +0 is accessed. bxxxxx1x The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +1 is accessed. bxxxx1xx The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +2 is accessed. bxxx1xxx The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +3 is accessed. bxx1xxxx The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +4 is accessed. bx1xxxxx The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +5 is accessed. bx1xxxxx The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +6 is accessed. b1xxxxxx The watchpoint hits if the byte at address (DBGWVR[31:0] & 0xFFFFF8) +7 is accessed.
[4:3]	L/S	Load/store access. The watchpoint can be conditioned to the type of access: b00 = Reserved b01 = load, load exclusive, or swap b10 = store, store exclusive or swap b11 = either. A SWP or SWPB triggers on load, store, or either. A load exclusive instruction triggers on load or either. A store exclusive instruction triggers on store or either, whether it succeeds or not.
[2:1]	S	Privileged access control. The watchpoint can be conditioned to the privilege of the access: b00 = reserved b01 = Privileged, match if the processor does a privileged access to memory b10 = User, match only on non-privileged accesses b11 = either, match all accesses. ———— Note ———— For all cases, the match refers to the privilege of the access, not the mode of the processor.
[0]	W	Watchpoint enable: 0 = Watchpoint disabled. This is the reset value. 1 = Watchpoint enabled.

12.4.15 Operating System Lock Status Register

The DBGOSLSR Register characteristics are:

Purpose	Contains status information about the locked debug registers.
Usage constraints	The DBGOSLSR is a read-only register.
Configurations	Available in all processor configurations.
Attributes	See Table 12-20 on page 12-30.

Figure 12-12 on page 12-30 shows the bit assignments.




Figure 12-12 DBGOSLSR Register bit assignments

Table 12-20 shows the bit assignments.

Table 12-20 DBGOSLSR Register bit assignments

Bits	Name	Function
[31:1]	-	RAZ.
[0]	Lock implemented bit	Indicates whether the OS lock functionality is implemented: 0 = OS lock not implemented.

12.4.16 Authentication Status Register

The DBGAUTHSTATUS Register characteristics are:

Purpose	Reads the current values of the configuration inputs that determine the debug permission level.
Usage constraints	The DBGAUTHSTATUS Register is read-only.
Configurations	Available in all processor configurations.
Attributes	See Table 12-21.

Figure 12-13 shows the bit assignments.

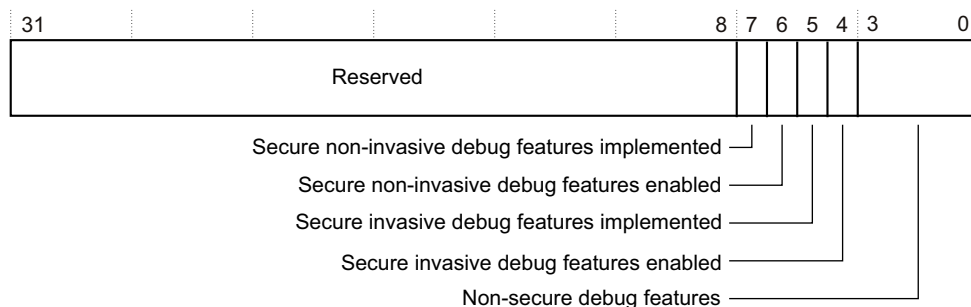


Figure 12-13 DBGAUTHSTATUS Register bit assignments

Table 12-21 shows the bit assignments.

Table 12-21 DBGAUTHSTATUS Register bit assignments

Bits	Name	Value	Function
[31:8]	-	-	RAZ
[7]	Secure non-invasive debug features implemented	0b1	Implemented
[6]	Secure non-invasive debug features enabled	DBGENm NIDENm	Non-invasive debug enable field

Table 12-21 DBGAUTHSTATUS Register bit assignments (continued)

Bits	Name	Value	Function
[5]	Secure invasive debug features implemented	0b1	Implemented
[4]	Secure invasive debug features enabled	DBGENm	Invasive debug enable field
[3:0]	Non-secure debug features ^a	0x0	Not implemented

a. The Cortex-R5 processor does not implement the Security Extensions, so all the debug features are considered secure.

12.4.17 Device Power-down and Reset Control Register

The DBGPRCR Register characteristics are:

Purpose	Controls reset and power-down related functionality.
Usage constraints	The DBGPRCR Register is read-write with more restricted access to some bits.
Configurations	Available in all processor configurations.
Attributes	See Table 12-22.

Figure 12-14 shows the bit assignments.




Figure 12-14 DBGPRCR Register bit assignments

Table 12-22 shows the bit assignments.

Table 12-22 DBGPRCR Register bit assignments

Bits	Name	Function
[31:3]	-	Do not modify on writes. On reads, the value returns zero.
[2]	Hold internal reset	Hold internal reset bit. This bit can be used to prevent the processor from running again before the debugger detects a power-down event and restores the state of the debug registers in the processor. This bit does not have any effect on initial system power-up, because nSYSPORESET clears it. 0 = Do not hold internal reset on power-up or warm reset. This is the reset value. 1 = Hold the processor non-debug logic in reset on warm reset until this flag is cleared.
[1]	Force internal reset	When a 1 is written to this bit, the processor asserts the DBGIRSTREQm output for four cycles. You can connect this output to an external reset controller that, in turn, resets the processor.
[0]	No power-down	When set to 1, the DBGNOPWRDWN output signal is HIGH. This output connects to the system power controller and is interpreted as a request to operate in emulate mode, if the system supports this functionality. In this mode, the processor is not actually powered down when requested by software or hardware handshakes. This mode is useful when debugging applications on top of working operating systems. 0 = DBGNOPWRDWN is LOW. This is the reset value 1 = DBGNOPWRDWN is HIGH.

12.4.18 Device Power-down and Reset Status Register

The DBGPRSR Register characteristics are:

Purpose	Provides information about the reset and power-down state of the processor.
Usage constraints	The DBGPRSR Register is a read-only register, with reads of the register also resetting some register bits.
Configurations	Available in all processor configurations.
Attributes	See Table 12-23.

Figure 12-15 shows the bit assignments.




Figure 12-15 DBGPRSR Register bit assignments

Table 12-23 shows the bit assignments.

Table 12-23 DBGPRSR Register bit assignments

Bits	Name	Function
[31:4]	-	Do not modify on writes. On reads, the value returns zero.
[3]	Sticky reset status	Sticky reset status bit. This bit is cleared on read. 0 = the processor has not been reset since the last time this register was read. This is the reset value. 1 = the processor has been reset since the last time this register was read. This sticky bit is set to 1 when nRESETm is asserted. This bit is reset to 0 by PRESETDBGmn .
[2]	Reset status	Reset status bit: 0 = the processor is not held in reset 1 = the processor is held in reset. This bit reads 1 when nRESETm is asserted.
[1]	Sticky power-down status ^a	Indicates if the core power domain has been powered down since the DBGPRCR was last read. 0 = the CPU has not been powered down since the last read. This is the reset value. 1 = the CPU has been powered down since the last read. If this bit is 1: - The contents of the core domain debug registers have been lost and must be reprogrammed. - Debug-APB transactions that access core domain debug registers receive an error response. This bit is cleared to 0 on a read.
[0]	Power-up status ^a	Indicates the status of the core power domain. 0 = the CPU is powered-down, that is, it is in Dormant or Shutdown mode. Core-domain debug registers cannot be accessed. 1 = the CPU is powered-up, that is, it is in Run or Standby mode. All debug registers can be accessed.

a. If you are implementing a Split/Lock configuration, contact ARM for more information about the functionality of this bit.

12.5 Management registers

The Management Registers define the standardized set of registers that all CoreSight components implement. This section describes these registers.

Table 12-24 shows the contents of the Management Registers for the processor debug unit.

Table 12-24 Management Registers

Offset (hex)	Register number	Access	Mnemonic	Description
0xD00–0xDFC	832-895	R	-	See <i>Processor ID Registers</i> .
0xF00	960	RW	DBGITCTRL	See <i>Integration Mode Control Register (DBGITCTRL)</i> on page 13-8.
0xFA0	1000		DBGCLAIMSET	See <i>Claim Tag Set Register</i> on page 12-34.
0xFA4	1001		DBGCLAIMCLR	See <i>Claim Tag Clear Register</i> on page 12-35.
0xFB0	1004	W	DBGLAR	See <i>Lock Access Register</i> on page 12-35.
0xFB4	1005	R	DBGLSR	See <i>Lock Status Register</i> on page 12-36.
0xFB8	1006	R	DBGAUTHSTATUS	See <i>Authentication Status Register</i> on page 12-30.
0xFB8–0xFC4	1006-1009	R	-	Reserved.
0xFC8	1010	R	DBGDEVID	Device Identifier. Reserved.
0xFCC	1011	R	DBGDEVTYPE	See <i>Device Type Register</i> on page 12-36.
0xFD0–0xFFC	1012-1023	R	-	See <i>Debug Identification Registers</i> on page 12-37.

12.5.1 Processor ID Registers

The Processor ID Registers are read-only registers that return the same values as the corresponding CP15 Main ID Register and Feature ID Registers. See Chapter 4 *System Control* for more information about the information contained in these registers.

Table 12-25 shows the offset value, register number, mnemonic, and description that are associated with each Process ID Register.

Table 12-25 Processor Identifier Registers

Offset (hex)	Register number	Mnemonic	Function
0xD00	832	MIDR	Main ID Register
0xD04	833	CTR	Cache Type Register
0xD08	834	TCMTR	TCM Type Register
0xD0C	835	-	Alias of MIDR
0xD10	836	MPUIR	MPU Type Register
0xD14	837	MPIDR	Multiprocessor Affinity Register
0xD18–0xD1C	838-839	-	Alias of MIDR
0xD20	840	ID_PFR0	Processor Feature Register 0

Table 12-25 Processor Identifier Registers (continued)

Offset (hex)	Register number	Mnemonic	Function
0xD24	841	ID_PFR1	Processor Feature Register 1
0xD28	842	ID_DFR0	Debug Feature Register 0
0xD2C	843	ID_AFR0	Auxiliary Feature Register 0
0xD30	844	ID_MMFR0	Processor Feature Register 0
0xD34	845	ID_MMFR1	Processor Feature Register 1
0xD38	846	ID_MMFR2	Processor Feature Register 2
0xD3C	847	ID_MMFR3	Processor Feature Register 3
0xD40	848	ID_ISAR0	ISA Feature Register 0
0xD44	849	ID_ISAR1	ISA Feature Register 1
0xD48	850	ID_ISAR2	ISA Feature Register 2
0xD4C	851	ID_ISAR3	ISA Feature Register 3
0xD50	852	ID_ISAR4	ISA Feature Register 4
0xD54	853	ID_ISAR5	ISA Feature Register 5
0xD58-0xDFC	854-895	-	Reserved, RAZ/SBZP

12.5.2 Claim Registers

The Claim Tag Set Register and the Claim Tag Clear Register enable an external debugger to claim debug resources.

Claim Tag Set Register

The DBGCLAIMSET Register characteristics are:

- Purpose** Enables an external debugger to claim debug resources.
- Usage constraints** The DBGCLAIMSET Register is a read/write register, in which:
- the CLAIM bits are always RAO
 - writing 0 to a CLAIM bit has no effect.
- Configurations** Available in all processor configurations.
- Attributes** See Table 12-26 on page 12-35.

Figure 12-16 shows the bit assignments.




Figure 12-16 DBGCLAIMSET Register bit assignments

Table 12-26 shows the bit assignments.

Table 12-26 DBGCLAIMSET Register bit assignments

Bits	Name	Function
[31:8]	-	RAZ or SBZP.
[7:0]	Claim tag set	RAO. Sets claim tags on writes.

Writing b1 to a specific claim tag set bit sets that claim tag. Writing b0 to a specific claim tag bit has no effect. This register always reads 0xFF, indicating eight claim tags are implemented.

Claim Tag Clear Register

The DBGCLAIMCLR Register characteristics are:

Purpose	Enables an external debugger to: - read debug resources - clear debug resources.
Usage constraints	The DBGCLAIMCLR Register is a read/write register, in which: - Reading this register returns the current claim tag value - writing 0 to a CLAIM bit has no effect - writing 1 to a specific claim tag clear bit clears that claim tag.
Configurations	Available in all processor configurations.
Attributes	See Table 12-27.

Figure 12-16 on page 12-34 shows the bit assignments.




Figure 12-17 DBGCLAIMCLR Register bit assignments

Table 12-27 shows the bit assignments.

Table 12-27 DBGCLAIMCLR Register bit assignments

Bit	Name	Description
[31:8]	-	RAZ or SBZP.
[7:0]	Claim tag clear	R/W. Reset value is 0x00.

12.5.3 Lock Access Register

The DBGLAR is a write-only register that controls writes to the debug registers. The purpose of the DBGLAR is to reduce the risk of accidental corruption to the contents of the debug registers. It does not prevent all accidental or malicious damage. Because the state of the DBGLAR is in the debug power domain, it is not lost when the processor powers down.

DBGLAR [31:0] contain a key that controls the lock status. To unlock the debug registers, write a 0xC5ACCE55 key to this register. To lock the debug registers, write any other value. Accesses to locked debug registers are ignored. The lock is set on reset.

12.5.4 Lock Status Register

The DBGLSR Register characteristics are:

Purpose	Returns the current lock status of the debug registers.
Usage constraints	The DBGLSR is: • a read-only register • only defined in the memory-mapped interface
Configurations	Available in all processor configurations.
Attributes	See Table 12-28.

Figure 12-18 shows the bit assignments.




Figure 12-18 DBGLSR Register bit assignments

Table 12-28 shows the bit assignments.

Table 12-28 DBGLSR Register bit assignments

Bits	Name	Function
[31:3]	-	Do not modify on writes. On reads, the value returns zero.
[2]	32-bit access	Indicates that a 32-bit access is required to write the key to the DBGLAR. This bit always reads 0.
[1]	Locked bit	Locked bit: 0 = Writes are permitted. 1 = Writes are ignored. This is the reset value.
[0]	Lock implemented bit	Indicates that the OS lock functionality is implemented. This bit always reads 1.

12.5.5 Device Type Register

The DBGDEVTYPE Register characteristics are:

Purpose	Indicates the type of debug component.
Usage constraints	The DBGDEVTYPE Register is a read-only register.
Configurations	Available in all processor configurations.
Attributes	See Table 12-29 on page 12-37.

Figure 12-19 on page 12-37 shows the bit assignments.




Figure 12-19 DBGDEVTYPE Register bit assignments

Table 12-29 shows the bit assignments.

Table 12-29 DBGDEVTYPE Register bit assignments

Bits	Name	Function
[31:8]	-	Do not modify on writes. On reads, the value returns zero.
[7:4]	Subtype	0x1, indicates that the sub-type of the device is <i>processor</i> .
[3:0]	Main class	0x5, indicates that the main class of the device is <i>debug logic</i> .

12.5.6 Debug Identification Registers

The Debug Identification Registers are read-only registers that consist of the Peripheral Identification Registers and the Component Identification Registers. The Peripheral Identification Registers provide standard information that all CoreSight components require. Only bits [7:0] of each register are used. The remaining bits Read-As-Zero.

The Component Identification Registers identify the processor as a CoreSight component. Only bits [7:0] of each register are used, the remaining bits Read-As-Zero. The values in these registers are fixed.

Table 12-30 shows the offset value, register number, and description that are associated with each Peripheral Identification Register.

Table 12-30 Peripheral Identification Registers

Offset (hex)	Register number	Function
0xFD0	1012	Peripheral Identification Register 4
0xFD4	1013	Reserved
0xFD8	1014	Reserved
0xFDC	1015	Reserved
0xFE0	1016	Peripheral Identification Register 0
0xFE4	1017	Peripheral Identification Register 1
0xFE8	1018	Peripheral Identification Register 2
0xFEC	1019	Peripheral Identification Register 3

Table 12-31 shows fields that are in the Peripheral Identification Registers.

Table 12-31 Fields in the Peripheral Identification Registers

Field	Size	Description
4KB Count	4 bits	Indicates the Log_2 of the number of 4KB blocks occupied by the debug device. The processor debug registers occupy a single 4KB block, therefore this field is always 0x0.
JEP106 Identity Code	4+7 bits	Identifies the designer of the processor. This field consists of a 4-bit continuation code and a 7-bit identity code. Because the processor is designed by ARM, the continuation code is 0x4 and the identity code is 0x3B. For more information see <i>JEP106M, Standard Manufacture's Identification Code</i> .
Part number	12 bits	Indicates the part number of the processor. The part number for the processor is 0xC15.
Revision	4 bits	Indicates the major and minor revision of the product. The major revision contains functionality changes and the minor revision contains bug fixes for the product. The revision number starts at 0x0 and increments by 1 at both major and minor revisions: 0x0 = r0p0 0x1 = r1p0 0x2 = r1p1.
RevAnd	4 bits	Indicates the manufacturer revision number. This number starts at 0x0 and increments by the integrated circuit manufacturer on metal fixes. For the Cortex-R5 processor, the initial value is 0x0 but this value can be changed by the manufacturer.
Customer modified	4 bits	Indicates an endorsed modification to the device. On this processor the value is always 0x0.

Table 12-32 shows how the bit values correspond with the Peripheral ID Register 0 functions.

Table 12-32 Peripheral ID Register 0 functions

Bits	Value	Description
[31:8]	-	Reserved
[7:0]	0x15	Indicates bits [7:0] of the Part number for the processor

Table 12-33 shows how the bit values correspond with the Peripheral ID Register 1 functions.

Table 12-33 Peripheral ID Register 1 functions

Bits	Value	Description
[31:8]	-	Reserved
[7:4]	0xB	Indicates bits [3:0] of the JEDEC JEP106 Identity Code
[3:0]	0xC	Indicates bits [11:8] of the Part number for the processor

Table 12-34 shows how the bit values correspond with the Peripheral ID Register 2 functions.

Table 12-34 Peripheral ID Register 2 functions

Bits	Value	Description
[31:8]	-	Reserved.
[7:4]	-	Indicates the revision number for the Cortex-R5 processor. This is the major revision number <i>n</i> in the <i>mn</i> part of the <i>mnpn</i> description of the product revision status.
[3]	0x1	This field is always set to 1. It indicates that the processor uses a JEP 106 identity code.
[2:0]	0x3	Indicates bits [6:4] of the JEDEC JEP106 Identity Code.

Table 12-35 shows how the bit values correspond with the Peripheral ID Register 3 functions.

Table 12-35 Peripheral ID Register 3 functions

Bits	Value	Description
[31:8]	-	Reserved.
[7:4]	0x0	Indicates the manufacturer revision number. This value changes based on the metal fixes made by the manufacturer.
[3:0]	0x0	Customer modified. See Table 12-31 on page 12-38.

Table 12-36 shows how the bit values correspond with the Peripheral ID Register 4 functions.

Table 12-36 Peripheral ID Register 4 functions

Bits	Value	Description
[31:8]	-	Reserved.
[7:4]	0x0	Indicates the number of blocks the debug component occupies. This field is always set to 0.
[3:0]	0x4	Indicates the JEDEC JEP106 continuation code. For the processor, this value is 4.

Table 12-37 shows the offset value, register number, and value that are associated with each Component Identification Register.

Table 12-37 Component Identification Registers

Offset (hex)	Register number	Value	Description
0xFF0	1020	0x0D	Component Identification Register 0
0xFF4	1021	0x90	Component Identification Register 1
0xFF8	1022	0x05	Component Identification Register 2
0xFFC	1023	0xB1	Component Identification Register 3

12.6 Debug events

A processor responds to a debug event in one of the following ways:

- ignores the debug event
- takes a debug exception
- enters debug halt state.

This section describes:

- *Software debug event*
- *Halting debug event* on page 12-41.
- *Behavior of the processor on debug events* on page 12-41
- *Debug event priority* on page 12-41
- *Watchpoint debug events* on page 12-41.

12.6.1 Software debug event

A software debug event is any of the following:

- A watchpoint debug event. This occurs when:
 - The data address for a load or store matches the watchpoint value.
 - All the conditions of the corresponding DBGWCR match.
 - The watchpoint is enabled.
 - The linked context ID-holding BRP, if any, is enabled and its value matches the context ID in CP15 c13. See Chapter 4 *System Control*.
 - The instruction that initiated the memory access is committed for execution.

Watchpoint debug events are only generated if the instruction passes its condition code.
- A breakpoint debug event. This occurs when:
 - An instruction was fetched and the instruction address or the CP15 Context ID register c13 matched the breakpoint value.
 - At the same time the instruction was fetched, all the conditions of the corresponding DBGBCR for unlinked context ID breakpoint generation matched the instruction fetch.
 - The breakpoint is enabled.
 - The instruction is committed for execution. These debug events are generated whether the instruction passes or fails its condition code.
- A BKPT debug event. This occurs when a BKPT instruction is committed for execution. BKPT is an unconditional instruction.
- A vector catch debug event. This occurs when:
 - An instruction was prefetched and the address matched a vector location address. This includes any kind of prefetch, not only the ones because of exception entry.
 - At the same time the instruction was fetched, the corresponding bit of the DBGVCR was set, that is, the vector catch is enabled.
 - The instruction is committed for execution. These debug events are generated whether the instruction passes or fails its condition code.

12.6.2 Halting debug event

The debugger or the system can cause the processor to enter into debug state by triggering any of the following halting debug events:

- assertion of the **EDBGRQm** signal, an External Debug Request
- write to the DBGDRCR[0] Halt Request control bit.

When **EDBGRQm** is asserted while **DBGENm** is HIGH, the device asserting this signal must hold it until the processor enters debug state, that is, until **DBGACKm** is asserted. The state of the processor pipeline determines how long this takes. If the request is not held in this way, the behavior of the processor is Unpredictable. For DBGDRCR[0] halting debug events, the processor records them internally until it is in a state and mode so that they can be taken.

12.6.3 Behavior of the processor on debug events

This section describes how the processor behaves on debug events while not in debug state. See *Debug state* on page 12-45 for information on how the processor behaves while in debug state. When the processor is in Monitor debug-mode, Prefetch Abort and Data Abort vector catch debug events are ignored. All other software debug events generate a debug exception such as Data Abort for watchpoints, and Prefetch Abort for anything else.

When debug is disabled, the BKPT instruction generates a debug exception, Prefetch Abort. All other software debug events are ignored.

When **DBGENm** is LOW, debug is disabled regardless of the value of DBGDSCR[15:14].

Table 12-38 shows the behavior of the processor on debug events.

Table 12-38 Processor behavior on debug events

DBGENm	DBGDSCR[15:14]	Debug mode	Action on software debug event	Action on halting debug event
0	bxx	Debug disabled	Ignore or Prefetch Abort (for BKPT)	Ignore
1	b00	None	Ignore or Prefetch Abort (for BKPT)	Debug state entry
1	bx1	Halting	Debug state entry	Debug state entry
1	b10	Monitor	Debug exception	Debug state entry

12.6.4 Debug event priority

Breakpoint, instruction address or CID match, vector catch, and halting debug events have the same priority. If more than one of these events occurs on the same instruction, it is Unpredictable which event is taken.

Breakpoint, instruction address or CID match and vector catch cancel the instruction that they occur on, therefore a watchpoint cannot be taken on such an instruction.

12.6.5 Watchpoint debug events

A synchronous watchpoint exception has similar behavior to a synchronous data abort exception:

- the processor sets R14_abt to the address of the instruction to return to plus 0x08.
- the processor does not complete the watchpointed instruction.

If the watchpointed access is subject to a synchronous data abort, then the synchronous abort takes priority over the watchpoint because it is a higher priority exception.

12.7 Debug exception

The processor takes a debug exception when a software debug event occurs while in Monitor debug-mode. Prefetch Abort and Data Abort Vector catch debug events are ignored. The debug software must carefully program certain debug events to prevent the processor from entering an unrecoverable state. If the processor takes a debug exception because of a breakpoint, BKPT, or vector catch debug event, the processor performs the following actions:

- sets the DBGDSCR[5:2] method-of-entry bits to indicate that a breakpoint occurred
- sets the CP15 IFSR and IFAR registers as described in *Effect of debug exceptions on CP15 registers and DBGWFAR* on page 12-43
- performs the same sequence of actions as in a Prefetch Abort exception by:
 - updating the SPSR_abt with the saved CPSR
 - changing the CPSR to abort mode and the state indicated by the TE and EE bits with normal interrupts and asynchronous aborts disabled
 - setting R14_abt as for a regular Prefetch Abort exception, that is, this register holds the address of the cancelled instruction plus 0x04
 - setting the PC to the appropriate Prefetch Abort vector.

Note

The Prefetch Abort handler is responsible for checking the IFSR to determine if a debug exception or other kind of Prefetch Abort exception caused the exception entry. If the cause is a debug exception, the Prefetch Abort handler must branch to the debug monitor. The R14_abt register holds the address of the instruction to restart.

If the processor takes a debug exception because of a watchpoint debug event, the processor performs the following actions:

- sets the DBGDSCR[5:2] method-of-entry bits to indicate that a synchronous watchpoint occurred
- sets the CP15 DFSR, DFAR, and DBGWFAR registers as described in *Effect of debug exceptions on CP15 registers and DBGWFAR* on page 12-43
- performs the same sequence of actions as in a Data Abort exception by:
 - updating the SPSR_abt with the saved CPSR
 - changing the CPSR to the state indicated by the TE and EE bits with normal interrupts and asynchronous aborts disabled
 - setting R14_abt as a regular Data Abort exception, that is, this register gets the address of the cancelled instruction plus 0x08
 - setting the PC to the appropriate Data Abort vector.

Note

The Data Abort handler must check the DFSR to determine if the exception entry was caused by a Debug exception or other kind of Data Abort exception. If the cause is a Debug exception, the Data Abort handler must branch to the debug monitor. The R14_abt register holds the address of the instruction to restart.

Table 12-39 shows the values in the link register after exceptions.

Table 12-39 Values in link register after exceptions

Cause of fault	ARM	Thumb	Return address (RA ^a) meaning
Breakpoint	RA+4	RA+4	Breakpointed instruction address
Watchpoint	RA+8	RA+8	Watchpointed instruction address
BKPT instruction	RA+4	RA+4	BKPT instruction address
Vector catch	RA+4	RA+4	Vector address
Prefetch Abort	RA+4	RA+4	Address of the instruction where the execution can resume
Data Abort	RA+8	RA+8	Address of the instruction where the execution can resume

a. This is the address of the instruction that the processor can execute first on debug exception return. The address of the access that hit the watchpoint is in the DBGWFER.

The following sections describe:

- *Effect of debug exceptions on CP15 registers and DBGWFER*
- *Avoiding unrecoverable states* on page 12-44.

12.7.1 Effect of debug exceptions on CP15 registers and DBGWFER

The four CP15 registers that record abort information are:

1. *Data Fault Address Register (DFAR)*
2. *Instruction Fault Address Register (IFAR)*
3. *Instruction Fault Status Register (IFSR)*
4. *Data Fault Status Register (DFSR).*

For more information on these registers, see Chapter 4 *System Control*.

If the processor takes a debug exception because of a watchpoint debug event, the processor performs the following actions on these registers:

- it does not change the IFSR or IFAR
- it updates the DFSR with the debug event encoding
- it writes an Unpredictable value to the DFAR
- it updates the DBGWFER with the address of the instruction that accessed the watchpointed address, plus a processor state dependent offset:
 - + 8 for ARM state
 - + 4 for Thumb state.

If the processor takes a debug exception because of a breakpoint, BKPT, or vector catch debug event, the processor performs the following actions on these registers:

- it updates the IFSR with the debug event encoding
- it writes an Unpredictable value to the IFAR
- it does not change the DFSR, DFAR, or DBGWFER.

12.7.2 Avoiding unrecoverable states

The processor ignores vector catch debug events on the Prefetch or Data Abort vectors while in Monitor debug-mode because these events would otherwise put the processor in an unrecoverable state.

The debuggers must avoid other similar cases by following these rules, that apply only if the processor is in Monitor debug-mode:

- if DBGBCR[22:20] is set to b010, and unlinked context ID breakpoint is selected, then the debugger must program DBGBCR[2:1] for the same breakpoint as stated in this section
- if DBGBCR[22:20] is set to b100 or b101, and instruction address mismatch breakpoint is selected, then the debugger must program DBGBCR[2:1] for the same breakpoint as stated in this section.

The debugger must write DBGBCR[2:1] for the same breakpoint as either b00 or b10, that selects either match in only USR, SYS, or SVC modes or match in only USR mode, respectively. The debugger must not program either b01, that is, match in any Privileged mode, or b11, that is, match in any mode.

You must only request the debugger to write b00 to DBGBCR[2:1] if you know that the abort handler does not switch to one of the USR, SYS, or SVC mode before saving the context that might be corrupted by a later debug event. You must also be careful about requesting the debugger to set a breakpoint or BKPT debug event inside a Prefetch Abort or Data Abort handler, or a watchpoint debug event on a data address that any of these handlers might access.

In general, you must only set breakpoint or BKPT debug events inside an abort handler after it saves the abort context. You can avoid breakpoint debug events in abort handlers by setting DBGBCR[2:1] as previously described.

If the code being debugged is not running in a Privileged mode, you can prevent watchpoint debug events in abort handlers by setting DBGWCR[2:1] to b10 for match only non-privileged accesses.

Failure to follow these guidelines can lead to debug events occurring before the handler is able to save the context of the abort. This causes the corresponding registers to be overwritten, and results in Unpredictable software behavior.

12.8 Debug state

The debug state enables an external agent, usually a debugger, to control the processor following a debug event. While in debug state, the processor behaves as follows:

- The DBGDSCR[0] core halted bit is set.
- The **DBGACKm** signal is asserted, see *DBGACKm* on page 12-51.
- The DBGDSCR[5:2] method of entry bits are set appropriately.
- The processor is halted. The pipeline is flushed and no instructions are fetched.
- The processor does not change the execution mode. The CPSR is not altered.
- Exceptions are treated as described in *Exceptions in debug state* on page 12-48.
- Interrupts are ignored.
- New debug events are ignored.

The following sections describe:

- *Entering debug state*
- *Behavior of the PC and CPSR in debug state* on page 12-46
- *Executing instructions in debug state* on page 12-46
- *Writing to the CPSR in debug state* on page 12-47
- *Privilege* on page 12-47
- *Accessing registers and memory* on page 12-47
- *Coprocessor instructions* on page 12-47
- *Effect of debug state on non-invasive debug* on page 12-48
- *Effects of debug events on processor registers* on page 12-48
- *Exceptions in debug state* on page 12-48
- *Leaving debug state* on page 12-49.

12.8.1 Entering debug state

When a debug event occurs while the processor is in Halting debug-mode, it switches to a special state called *debug state* so the debugger can take control. You can configure Halting debug-mode by setting DBGDSCR[14].

If a halting debug event occurs, the processor enters debug state even when Halting debug-mode is not configured. While the processor is in debug state, the PC does not increment on instruction execution. If the PC is read at any point after the processor has entered debug state, but before an explicit PC write, it returns a value as described in Table 12-40, depending on the previous state and the type of debug event.

Table 12-40 shows the read PC value after debug state entry for different debug events.

Table 12-40 Read PC value after debug state entry

Debug event	ARM	Thumb	Return address (RA) meaning
Breakpoint	RA+8	RA+4	Breakpointed instruction address.
Watchpoint	RA+8	RA+4	Watchpointed instruction address.
BKPT instruction	RA+8	RA+4	BKPT instruction address.
Vector catch	RA+8	RA+4	Vector address.
External debug request signal activation	RA+8	RA+4	Address of the instruction where the execution resumes.

Table 12-40 Read PC value after debug state entry (continued)

Debug event	ARM	Thumb	Return address (RA) meaning
Debug state entry request command	RA+8	RA+4	Address of the instruction where the execution resumes.
OS unlock event	RA+8	RA+4	Address of the instruction where the execution resumes.
CTI debug request signal	RA+8	RA+4	Address of the instruction where the execution resumes.

12.8.2 Behavior of the PC and CPSR in debug state

The behavior of the PC and CPSR registers while the processor is in debug state is as follows:

- The PC is frozen on entry to debug state. That is, it does not increment on the execution of ARM instructions. However, the processor still updates the PC as a response to instructions that explicitly modify the PC.
- If the PC is read after the processor has entered debug state, it returns a value as described in Table 12-40 on page 12-45, depending on the previous state and the type of debug event.
- If the debugger executes a sequence for writing a certain value to the PC and subsequently it forces the processor to restart without any additional write to the PC or CPSR, the execution starts at the address corresponding to the written value.
- If the debugger forces the processor to restart without having performed a write to the PC, the restart address is Unpredictable.
- If the debugger writes to the CPSR, subsequent reads from the PC return an Unpredictable value, and if it forces the processor to restart without having performed a write to the PC, the restart address is Unpredictable. However, CPSR reads after a CPSR write return the written value.
- If the debugger writes to the PC, subsequent reads from the PC return an Unpredictable value.
- If the debugger forces the processor to execute an instruction that writes to the PC and this instruction fails its condition codes, the PC is written with an Unpredictable value. That is, if the debugger forces the processor to restart, the restart address is Unpredictable. Also, if the debugger reads the PC, the read value is Unpredictable.
- While the processor is in debug state, the CPSR does not change unless written to by an instruction. In particular, the CPSR IT execution state bits do not change on instruction execution. The CPSR IT execution state bits do not have any effects on instruction execution.
- If the processor executes a data processing instruction with $Rd=R15$ and $S=0$, then $alu-out[0]$ must equal the current value of the CPSR T bit, otherwise the processor behavior is Unpredictable.

12.8.3 Executing instructions in debug state

In debug state, the processor executes instructions issued through the *Instruction Transfer Register* (DBGITR). Before the debugger can force the processor to execute any instruction, it must enable this feature through DBGDSCR[13].

While the processor is in debug state, it always decodes instructions from the DBGITR as per the ARM instruction set, regardless of the value of the T and J bits of the CPSR.

The following restrictions apply to instructions executed through the DBGITR while in debug state:

- with the exception of branch instructions and instructions that modify the CPSR, the processor executes any ARM instruction in the same manner as if it was not in debug state
- the branch instructions B, BL, BLX(1), and BLX(2) are Unpredictable
- certain instructions that normally update the CPSR are Unpredictable
- instructions that load a value into the PC from memory are Unpredictable.

12.8.4 Writing to the CPSR in debug state

The only instruction that can update the CPSR while in debug state is the MSR instruction. All other ARMv7 instructions that write to the CPSR are Unpredictable, that is, the BX, BXJ, SETEND, CPS, RFE, LDM(exception return), and data processing instructions with Rd==R15 and S==1.

The behavior of the CPSR forms of the MSR and MRS instructions in debug state is different to their behavior in normal state:

- When not in debug state, an MSR instruction that modifies the execution state bits in the CPSR is Unpredictable. However, in debug state an MSR instruction can update the execution state bits in the CPSR. An *Instruction Synchronization Barrier* (ISB) sequence must follow a direct modification of the execution state bits in the CPSR by an MSR instruction.
- When not in debug state, an MRS instruction reads the CPSR execution state bits as zeros. However, in debug state an MRS instruction returns the actual values of the execution state.

The debugger must execute an ISB sequence after it writes to the CPSR execution state bits using an MSR instruction. If the debugger reads the CPSR using an MRS instruction after a write to any of these bits, but before an ISB sequence, the value that MRS returns is Unpredictable. Similarly, if the debugger forces the processor to leave debug state after an MSR writes to the execution state bits, but before any ISB sequence, the behavior of the processor is Unpredictable.

12.8.5 Privilege

When the processor is in debug state, ARM instructions issued through the DBGITR are subject to different rules about whether they can perform privileged actions. The general rule is that all instructions and operations are permitted in debug state.

12.8.6 Accessing registers and memory

The processor always accesses register banks and memory as indicated by the CPSR mode bits, in both normal and debug state. For example, if the CPSR mode bits indicate the processor is in User mode, ARM register reads and returns the User mode banked registers, and memory accesses are presented to the MPU as not privileged.

12.8.7 Coprocessor instructions

CP14 and CP15 instructions can always be executed in debug state regardless of processor mode.

12.8.8 Effect of debug state on non-invasive debug

The processor non-invasive debug features are the ETM and *Performance Monitoring Unit* (PMU). All of these non-invasive debug features are disabled when the processor is in debug state. For more information, see Chapter 4 *System Control* and *ETM interface* on page 2-11.

When the processor is in debug state:

- the ETM ignores all instructions and data transfers
- PMU events are not counted
- events are not visible to the ETM
- the PMU *Cycle Count Register* (CCNT) is stopped.

12.8.9 Effects of debug events on processor registers

On entry to debug state, the processor does not update any general-purpose or program status register. This includes the SPSR_abt and R14_abt registers. In addition, the processor does not update any coprocessor registers, including the CP15 IFSR, DFSR, DFAR, or IFAR registers, except for CP14 DBGDSCR[5:2] method-of-entry bits. These bits indicate the type of debug event that caused the entry into debug state.

———— Note ————

On entry to debug state because of a watchpoint debug event, the processor updates the DBGWFER register with the address of the instruction accessing the watchpointed address plus:

- + 8 in ARM state
- + 4 in Thumb state.

12.8.10 Exceptions in debug state

While in debug state, exceptions are handled as follows:

Reset This exception is taken as in a normal processor state. This means the processor leaves debug state because of the system reset.

Prefetch Abort

This exception cannot occur because the processor does not fetch any instructions while in debug state.

Debug The processor ignores debug events, including BKPT instructions.

SVC The processor ignores SVC exceptions.

Undefined When an Undefined Instruction exception occurs in debug state, the behavior of the processor is as follows:

- PC, CPSR, SPSR_und, and R14_und are unchanged
- the processor remains in debug state
- DBGDSCR[8], sticky Undefined bit, is set.

Synchronous Data abort

When a synchronous Data Abort occurs in debug state, the behavior of the processor is as follows:

- PC, CPSR, SPSR_abt, and R14_abt are unchanged
- the processor remains in debug state
- DBGDSCR[6], sticky synchronous data abort bit, is set

- DFSR and DFAR are set to the same values as if the abort had occurred in normal state.

Asynchronous Data Abort

When an asynchronous Data Abort occurs in debug state, the behavior of the processor is as follows, regardless of the setting of the CPSR A bit:

- PC, CPSR, SPSR_abt, and R14_abt are unchanged
- the processor remains in debug state
- DBGDSCR[7], sticky asynchronous data abort bit, is set
- the DFSR remains unchanged
- the processor does not act on this asynchronous Data Abort on exit from the debug state, that is, the asynchronous abort is discarded.

Asynchronous Data Aborts on entry and exit from debug state

On entering debug state, the processor executes a *Data Synchronization Barrier* (DSB) sequence to ensure that any outstanding asynchronous Data Aborts are detected, before starting debug operations.

If the DSB operation detects an asynchronous Data Abort, the processor records this event and its type as if the CPSR A bit was set. The purpose of latching this event is to ensure that it can be taken on exit from the debug state.

Before forcing the processor to leave debug state, the debugger must execute a DSB sequence to ensure that all debugger-generated asynchronous Data Aborts are detected, and therefore discarded, while still in debug state. After exiting debug state, the processor acts on any previously recorded asynchronous Data Aborts if permitted by the CPSR A bit.

12.8.11 Leaving debug state

The debugger can force the processor to leave debug state:

- by setting the restart request bit, DBGDRCR[1], to 1
- through the *Cross Trigger Interface* (CTI) external restart request mechanism, using the **DBGRESTARTm** and **DBGRESTARTEDm** signals.

When one of those restart requests occurs, the processor:

1. Clears the DBGDSCR[1] core restarted flag.
2. Leaves debug state.
3. Clears the DBGDSCR[0] core halted flag.
4. Drives the **DBGACKm** signal LOW, unless the DBGDSCR[11] DbgAck bit is set to 1.
5. Starts executing instructions from the address last written to the PC in the processor mode and state indicated by the current value of the CPSR. The CPSR IT execution state bit is restarted with the current value applying to the first instruction on restart.
6. Sets the DBGDSCR[1] core restarted flag to 1.

12.9 Cache debug

This section describes cache debug. It consists of:

- *Cache pollution in debug state*
- *Cache coherency in debug state*
- *Cache usage profiling.*

12.9.1 Cache pollution in debug state

If bit [0] of the *Debug State Cache Control Register* (DBGDSCCR) is set to 0 while the processor is in debug state, then the L1 data cache does not perform any line fill.

Note

No special feature is required to prevent L1 instruction cache pollution because instruction side fetches cannot occur while in debug state.

12.9.2 Cache coherency in debug state

The debugger can update memory while in debug state:

- to replace an instruction with a BKPT, or to restore the original instruction
- to download code for the processor to execute on leaving debug state.

The debugger can maintain cache coherency in both these situations with the following features:

- If bit [2] of the DBGDSCCR is set to 0 while the processor is in debug state, then the processor treats any memory access that hits in L1 data cache as write-through, regardless of the memory region attributes. This guarantees that the L1 instruction cache can see the changes to the code region without the debugger executing a sequence of cache clean operations.
- After the code is written to memory, the debugger can execute either a CP15 instruction cache invalidate all operation, or a CP15 instruction cache invalidate line operation.

Note

The processor can normally execute CP15 instruction cache invalidate all operation or CP15 instruction cache invalidate line operation only in Privileged mode. However, in debug state the processor can execute these instructions even when invasive debug is not permitted in Privileged mode. This exception to the rule enables the debugger to maintain coherency.

12.9.3 Cache usage profiling

You can obtain cache usage profiling information using the *Performance Monitoring Unit* (PMU). The processor can count cache accesses and misses over a period of time. See Chapter 6 *Events and Performance Monitor*.

12.10 External debug interface

The system can access memory-mapped debug registers through the processor APB slave ports. This section describes the APB interface and the miscellaneous debug input and output signals:

- *APB signals*
- *Miscellaneous debug signals*
- *Authentication signals* on page 12-52.

12.10.1 APB signals

The APB slave ports are compliant with the *AMBA 3 APB Protocol Specification* and can be connected to the *Debug Access Port* (DAP). This APB slave interface supports 32-bits wide data, stalls, slave-generated aborts, and ten address bits [11:2] mapping 4KB of memory. An extra **PADDRDBG31m** signal indicates to the CPU the source of access.

Table A-21 on page A-26 shows the external debug interface signals.

12.10.2 Miscellaneous debug signals

This section describes the miscellaneous debug signals.

EDBGRQm

This signal generates a halting debug event, that is, it requests the CPU to enter debug state. When this occurs, the DBGDSCR[5:2] method-of-debug entry bits are set to b0100. When **EDBGRQm** is asserted, it must be held until **DBGACKm** is asserted. Failure to do so leads to Unpredictable behavior of the processor.

DBGACKm

The CPU asserts **DBGACKm** to indicate that the system has entered debug state. It serves as a handshake for the **EDBGRQm** signal. The **DBGACKm** signal is also driven HIGH when the debugger sets the DBGDSCR[10] DbgAck bit to 1.

DBGNOPWRDWN

The CPU asserts **DBGNOPWRDWN** when bit [0] of the Device Power down and Reset Control Register is 1 in either CPU. The processor power controller must work in Emulate mode when this signal is HIGH.

DBGROMADDR

The **DBGROMADDR** signal specifies bits [31:12] of the debug ROM physical address. This is a configuration input and must be tied off or only change while the processor is in reset. In a system with multiple debug ROMs, this address must be tied off to point to the top-level ROM address.

DBGROMADDRV is the valid signal for **DBGROMADDR**. If the address cannot be determined, **DBGROMADDR** must be tied off to zero and **DBGROMADDRV** must be tied LOW. The value of these signals can be read from the *Debug ROM Address Register* (DBGDRAR).

DBGSELFADDRm

The **DBGSELFADDRm** signal specifies bits [31:12] of the offset from the debug ROM physical address to the physical address where the CPU APB port is mapped to the base of the 4KB debug register map. This is a configuration input and must be tied off or only change while the CPU is in reset.

DBGSELFADDRV_m is the valid signal for **DBGSELFADDRm**. If the offset cannot be determined, **DBGSELFADDRm** must be tied off to zero and **DBGSELFADDRV_m** must be tied LOW. The value of these signals can be read from the *Debug Self Address Register* (DSAR).

DBGRESTARTm

The **DBGRESTARTm** signal is used to bring the CPU out of debug halt state. The CPU acknowledges **DBGRESTARTm** by asserting **DBGRESTARTEDm**, and then starts fetching instructions when **DBGRESTARTm** is deasserted.

DBGRESTARTEDm

The CPU asserts **DBGRESTARTEDm** in response to a **DBGRESTARTm** request, when it is ready to exit debug halt state and return to normal run state.

DBGTRIGGERm

The CPU asserts **DBGTRIGGERm** to indicate that the system has accepted a debug request and attempts to enter debug state. It is not a handshake for the **EDBGRQm** signal. If **DBGACKm** does not go HIGH following **DBGTRIGGERm**, the memory system has stopped responding and the CPU has not entered debug state.

Table A-22 on page A-26 shows the debug miscellaneous signals.

12.10.3 Authentication signals

Table 12-41 shows a list of the valid authentication signals and the associated debug permissions. Authentication signals are used to configure the CPU so its activity can only be debugged or traced in a certain subset of CPU modes.

Table 12-41 Authentication signal restrictions

DBGENm^a	NIDENm	Non-invasive debug permitted in User and Privileged modes
0	0	No
X	1	Yes
1	0	Yes

- a. When **DBGENm** is LOW, the processor behaves as if DBGDSCR[15:14] equals b00 with the exception that halting debug events are ignored when this signal is LOW.

Changing the authentication signals

The **NIDENm**, and **DBGENm** input signals are either tied off to some fixed value or controlled by some external device.

If software running on the CPU has control over an external device that drives the authentication signals, it must make the change using a safe sequence:

1. Execute an implementation-specific sequence of instructions to change the signal value. For example, this might be a single STR instruction that writes certain value to a control register in a system peripheral.
2. If step 1 involves any memory operation, issue a *Data Synchronization Barrier* (DSB) instruction.
3. Poll the DBGDSCR or DBGAUTHSTATUS to check whether the CPU has already detected the changed value of these signals. This is required because the system might not issue the signal change to the CPU until several cycles after the DSB completes.
4. Issue an *Instruction Synchronization Barrier* (ISB) instruction.

The software cannot perform debug or analysis operations that depend on the new value of the authentication signals until this procedure is complete. The same rules apply when the debugger has control of the CPU through the DBGITR while in debug state.

The values of the **DBGENm** and **NIDENm** signals can be determined by polling DBGDSCR[17:16], DBGDSCR[15:14], or the DBGAUTHSTATUS.

12.11 Using the debug functionality

This section provides some examples of using the processor debug functionality, both from the point of view of a software engineer writing code to run on an ARM processor and of a developer creating debug tools for the processor. In the former case, examples are given in ARM assembly language. In the latter case, the examples are in C pseudo-language, intended to convey the algorithms to be used. These examples are not intended as source code for a debugger.

The debugger examples use a pair of pseudo-functions such as the following:

```
uint32 ReadDebugRegister(int reg_num)
{
    // read the value of the debug register reg_num at address reg_num << 2
}

WriteDebugRegister(int reg_num, uint32 val)
{
    // write the value val to the debug register reg_num at address reg_num >> 2
}
```

A basic function for using the debug state is executing an instruction through the DBGITR. Example 12-1 shows the sequence for executing an ARM instruction through the DBGITR.

Example 12-1 Executing an ARM instruction through the DBGITR

```
ExecuteARMInstruction(uint32 instr)
{
    // Step 1. Poll DBGDSCR until InstrCompl_1 is set.
    repeat
    {
        dbgdscr := ReadDebugRegister(34);
    }
    until (dbgdscr & (1<<24));
    // Step 2. Write the opcode to the DBGITR.
    WriteDebugRegister(33, instr);
    // Step 3. Poll DBGDSCR until InstrCompl is set.
    repeat
    {
        dbgdscr := ReadDebugRegister(34);
    }
    until (dbgdscr & (1<<24));
}
```

This section describes:

- *Debug communications channel* on page 12-55
- *Programming breakpoints and watchpoints* on page 12-57
- *Single-stepping* on page 12-60
- *Debug state entry* on page 12-61
- *Debug state exit* on page 12-62
- *Accessing registers and memory in debug state* on page 12-63.

12.11.1 Debug communications channel

There are two ways that an external debugger can send data to or receive data from the processor:

- The debug communications channel, when the processor is not in debug state. It is defined as the set of resources used for communicating between the external debugger and software running on the processor.
- The mechanism for forcing the processor to execute ARM instructions, when the processor is in debug state. For more information, see *Executing instructions in debug state* on page 12-46.

Rules for accessing the DCC

At the processor side, the debug communications channel resources are:

- CP14 Debug Register c5 (DTR, comprising DBGDTRTXint and DBGDTRRXint)
- CP14 Debug Register c1 (DBGDSCRint).

The ARMv7 debug architecture is implemented on the processor so that:

- If a read of the DBGDSCRint returns 1 for the RXfull flag, a following read of the DBGDTRRXint returns valid data and RXfull is cleared. No ISB is required between these two CP14 instructions.
- If a read of the CP14 DBGDSCRint returns 1 for the TXfull flag, a following write to the DBGDTRTXint is Unpredictable.
- If a read of the CP14 DBGDSCRint returns 0 for the RXfull flag, a following read of the CP14 DTR returns an Unpredictable value.
- If a read of the CP14 DBGDSCRint returns 0 for the TXfull flag, a following write to the CP14 DTR writes the intended 32-bit word, and sets TXfull to 1. No ISB is required between these two CP14 instructions.

When Nonblocking mode is selected for DTR accesses, the following conditions are true for memory-mapped DBGDSCR, DBGDTRRXext, and DBGDTRTXext registers:

- If a read of the DBGDSCRext returns 0 for the TXfull flag, a following read of the memory-mapped DBGDTRTX is ignored. The content of TXfull is unchanged and the read returns an UNKNOWN value.
- If a read of the DBGDSCRext returns 0 for the RXfull flag, a following write of the memory-mapped DBGDTRRX passes valid data to the processor and sets RXfull to 1.
- If a read of the DBGDSCRext returns 1 for the TXfull flag, a following read of the DBGDTRTXext returns valid data and clears TXfull.
- If a read of the DBGDSCRext returns 1 for the RXfull flag, a following write of the memory-mapped DBGDTRRXext is ignored, that is, both RXfull and DBGDTRRX contents are unchanged.

Software access to the DCC

Software running on the processor that sends data to the debugger through the target-to-host channel can use the sequence of instructions that Example 12-2 on page 12-56 shows.

Example 12-2 Target to host data transfer (target end)

```

; r0 -> word to send to the debugger
WriteDCC    MRC      p14, 0, PC, c0, c1, 0
            BEQ      WriteDCC
            MCR      p14, 0, Rd, c0, c5, 0
            BX       lr

```

Example 12-3 shows the sequence of instructions for sending data to the debugger through the host-to-target channel.

Example 12-3 Host to target data transfer (target end)

```

; r0 -> word sent by the debugger
ReadDCC     MRC      p14, 0, PC, c0, c1, 0
            BCC      ReadDCC
            MRC      p14, 0, Rd, c0, c5, 0
            BX       lr

```

Debugger access to the DCC

A debugger can access the DCC through the external interface. The following examples show the pseudo-code operations for these accesses.

Example 12-4 shows the code for target-to-host data transfer.

Example 12-4 Target to host data transfer (host end)

```

uint32      ReadDCC()
{
    // Step 1. Poll DBGDSCR until TXfull is set to 1.
    repeat
    {
        dbgdsr := ReadDebugRegister(34);
    }
    until (dbgdsr & (1<<29));
    // Step 2. Read the value from DBGDTRTX.
    dtr_val := ReadDebugRegister(35);

    return dtr_val;
}

```

Example 12-5 shows the code for host-to-target data transfer.

Example 12-5 Host to target data transfer (host end)

```

WriteDCC(uint32 dtr_val)
{
    // Step 1. Poll DBGDSCR until RXfull is clear.
    repeat
    {
        dbgdsr := ReadDebugRegister(34);
    }
}

```

```

    until (!(dbgdsr & (1<<30)));
    // Step 2. Write the value to DBGDTRRX.
    WriteDebugRegister(32, dtr_val);
}

```

While the processor is running, if the DCC is used as a data channel, it might be appropriate to poll the DCC regularly.

Example 12-6 shows the code for polling the DCC.

Example 12-6 Polling the DCC (host end)

```

PollDCC
{
    dbgdsr := ReadDebugRegister(34);
    if (dbgdsr & (1<<29))
    {
        // DBGDTRTX (target -> host transfer register) full
        dtr := ReadDebugRegister(35)
        ProcessTargetToHostWord(dtr);
    }
    if (!(dbgdsr & (1<<30)))
    {
        // DBGDTRRX (host -> target transfer register) empty
        dtr := GetNextHostToTargetWord()
        WriteDebugRegister(32, dtr);
    }
}

```

12.11.2 Programming breakpoints and watchpoints

This section describes the following operations:

- *Programming simple breakpoints and the byte address select*
- *Setting a simple aligned watchpoint* on page 12-58
- *Setting a simple unaligned watchpoint* on page 12-59.

Programming simple breakpoints and the byte address select

When programming a simple breakpoint, you must set the byte address select bits in the control register appropriately. For a breakpoint in ARM state, this is simple. For Thumb state, you must calculate the value based on the address.

For a simple breakpoint, you can program the settings for the other control bits as Table 12-42 shows:

Table 12-42 Values to write to DBGBCR for a simple breakpoint

Bits	Value to write	Description
[31:29]	0b000	Reserved
[28:24]	0b00000	Breakpoint address mask
[23]	0b0	Reserved
[22:20]	0b000	Meaning of DBGBVR

Table 12-42 Values to write to DBGBCR for a simple breakpoint (continued)

Bits	Value to write	Description
[19:16]	0b0000	Linked BRP number
[15:9]	0b00	Reserved
[8:5]	Derived from address	Byte address select
[4:3]	0b00	Reserved
[2:1]	0b11	Supervisor access control
[0]	0b1	Breakpoint enable

Example 12-7 shows the sequence of instructions for setting a simple breakpoint.

Example 12-7 Setting a simple breakpoint

```

SetSimpleBreakpoint(int break_num, uint32 address, iset_t isa)
{
    // Step 1. Disable the breakpoint being set.
    WriteDebugRegister(80 + break_num, 0x0);
    // Step 2. Write address to the DBGbVR, leaving the bottom 2 bits zero.
    WriteDebugRegister(64 + break_num, address & 0xFFFFFC);
    // Step 3. Determine the byte address select value to use.
    case (isa) of
    {
        // Note: The processor does not support Jazelle or ThumbEE states
        when THUMB:
            byte_address_select := (3 << (address & 2));
        when ARM:
            byte_address_select := 15;
    }
    // Step 4. Write the mask and control register to enable the breakpoint.
    WriteDebugRegister(80 + break_num, 7 | (byte_address_select << 5));
}

```

Setting a simple aligned watchpoint

The simplest and most common type of watchpoint watches for a write to a given address in memory. In practice, a data object spans a range of addresses but is aligned to a boundary corresponding to its size, so you must set the byte address select bits in the same way as for a breakpoint.

For a simple watchpoint, you can program the settings for the other control bits as Table 12-43 shows:

Table 12-43 Values to write to DBGWCR for a simple watchpoint

Bits	Value to write	Description
[31:29]	0b000	Reserved
[28:24]	0b00000	Watchpoint address mask
[23:21]	0b000	Reserved
[20]	0b0	Enable linking

Table 12-43 Values to write to DBGWCR for a simple watchpoint (continued)

Bits	Value to write	Description
[19:16]	0b0000	Linked BRP number
[15:13]	0b00	Reserved
[12:5]	Derived from address	Byte address select
[4:3]	0b10	Load/Store access control
[2:1]	0b11	Privileged access control
[0]	0b1	Watchpoint enable

Example 12-8 shows the code for setting a simple aligned watchpoint.

Example 12-8 Setting a simple aligned watchpoint

```

SetSimpleAlignedWatchpoint(int watch_num, uint32 address, int size)
{
    // Step 1. Disable the watchpoint being set.
    WriteDebugRegister(112 + watch_num, 0);
    // (Step 2. Write address to the DBGWVR, leaving the bottom 3 bits zero.
    WriteDebugRegister(96 + watch_num, address & 0xFFFFF8);
    // Step 3. Determine the byte address select value to use.
    case (size) of
    {
        when 1:
            byte_address_select := (1 << (address & 7));
        when 2:
            byte_address_select := (3 << (address & 6));
        when 4:
            byte_address_select := (15 << (address & 4));
        when 8:
            byte_address_select := 255;
    }
    // Step 4. Write the mask and control register to enable the watchpoint.
    WriteDebugRegister(112 + watch_num, 23 | (byte_address_select << 5));
}

```

Setting a simple unaligned watchpoint

Using the byte address select bits, certain unaligned objects up to a doubleword (64 bits) can be watched in a single watchpoint. However, this cannot cover all cases, and in many cases a second watchpoint might be required.

Table 12-44 shows some examples.

Table 12-44 Example byte address masks for watchpointed objects

Address of object	Object size in bytes	First address value	First byte address mask	Second address value	Second byte address mask
0x00008000	1	0x00008000	0b00000001	Not required	-
0x00008007	1	0x00008000	0b10000000	Not required	-
0x00009000	2	0x00009000	0b00000011	Not required	-

Table 12-44 Example byte address masks for watchpointed objects (continued)

Address of object	Object size in bytes	First address value	First byte address mask	Second address value	Second byte address mask
0x0000900c	2	0x00009000	0b11000000	Not required	-
0x0000900d	2	0x00009000	0b10000000	0x00009008	0b00000001
0x0000A000	4	0x0000A000	0b00001111	Not required	-
0x0000A003	4	0x0000A000	0b01111000	Not required	-
0x0000A005	4	0x0000A000	0b11100000	0x0000A008	0b00000001
0x0000B000	8	0x0000B000	0b11111111	Not required	-
0x0000B001	8	0x0000B000	0b11111110	0x0000B008	0b00000001

Example 12-9 shows the code for setting a simple unaligned watchpoint.

Example 12-9 Setting a simple unaligned watchpoint

```

bool SetSimpleWatchpoint(int watch_num, uint32 address, int size)
{
    // Step 1. Disable the watchpoint being set.
    WriteDebugRegister(112 + watch_num, 0x0);
    // Step 2. Write addresses to the DBGWVRs, leaving the bottom 3 bits zero.
    WriteDebugRegister(96 + watch_num, (address & 0xFFFFF8));
    // Step 3. Determine the byte address select value to use.
    byte_address_select := (1 << size) - 1;
    byte_address_select := (byte_address_select) << (address & 7);
    // Step 4. Write the mask and control register to enable the breakpoint.
    WriteDebugRegister(112 + watch_num, 5'b23 | ((byte_address_select & 0xFF) << 5));
    // Step 5. Set second watchpoint if required. This is the case if the byte
    // address mask is more than 8 bits.
    if (byte_address_select >= 256)
    {
        WriteDebugRegister(112 + watch_num + 1, 0);
        WriteDebugRegister(96 + watch_num + 1, (address & 0xFFFFF8) + 8);
        WriteDebugRegister(112 + watch_num + 1 23 | ((byte_address_select & 0xFF00) >> 3));
    }
    // Step 6. Return flag to caller indicating if second watchpoint was used.
    return (byte_address_select >= 256)
}

```

12.11.3 Single-stepping

You can use the breakpoint mismatch bit to implement single-stepping on the processor. Unlike high-level stepping, single-stepping implements a low-level step that executes a single instruction at a time. With high-level stepping, the instruction is decoded to determine the address of the next instruction and a breakpoint is set at that address.

Example 12-10 on page 12-61 shows the code for single-stepping off an instruction. The processor must be configured for halt-mode debugging.

Example 12-10 Single-stepping off an instruction

```
SingleStepOff(uint32 address)
{
    bkpt := FindUnusedBreakpointWithMismatchCapability();
    SetComplexBreakpoint(bkpt, address, 4 << 20);
}
```

Note

In Example 12-10, the third parameter of SetComplexBreakpoint() indicates the value to set DBGBCR[22:20].

This method of single-stepping steps off the instruction that might not necessarily be the same as stepping to the next instruction executed. In certain circumstances, the next instruction executed might be the same instruction being stepped off.

The simplest example of this is a branch to a self instruction such as (B .). In this case, the wanted behavior is most likely to step off the branch to self because this is often used as a means of waiting for an interrupt.

A more complex example is a return from function that returns to the same point. For example, a simple recursive function might terminate with:

```
BL    ThisFunction
POP   {saved_registers, pc}
```

In this case, the POP instruction loads a link register that is saved at the start of the function, and if that is the link register created by the BL instruction shown, it points back at the POP instruction. Therefore, this single step code unwinds the entire call stack to the point of the original caller, rather than stepping out a level at a time. It is not possible to single step this piece of code using either the high-level or low-level stepping methods.

12.11.4 Debug state entry

On entry to debug state, the debugger can read the processor state, including all registers and the PC, and determine the cause of the exception from the DBGDSCR method-of-entry bits.

Example 12-11 shows the code for entry to debug state.

Example 12-11 Entering debug state

```
OnEntryToDebugState(PROCESSOR_STATE *state)
{
    // Step 1. Read the DBGDSCR to determine the cause of debug entry.
    state->dbgdsr := ReadDebugRegister(34);
    // Step 2. Issue a DataSynchronizationBarrier instruction if required;
    // this is not required by the Cortex-R5 processor but is required for ARMv7
    // debug.
    if ((state->dbgdsr & (1<<19)) == 0)
    {
        ExecuteARMInstruction(0xE57FF040)
        // Step 3. Poll the DBGDSCR for DBGDSCR[19] to be set.
        repeat
        {
            dbgdsr := ReadDebugRegister(34);
        }
    }
}
```

```

    until (dbgdsr & (1<<19));
}
// Step 4. Read the entire processor state. The function ReadAllRegisters
// reads all general-purpose registers for all processor modes, and saves
// the data in "state".
ReadAllRegisters(state);
// Step 5. Based on the CPSR (processor state), determine the actual restart
// address
if (state->cpsr & (1<<5);
{
    // Thumb state
    state->pc := state->pc - 4;
}
else
{
    // ARM state
    state->pc := state->pc - 8;
}
// Step 6. If the method of entry was Watchpoint Occurred, read the DBGWFR
// register
method_of_debug_entry := ((state->dbgdsr >> 2) & 0xF;
if (method_of_debug_entry == 2 || method_of_debug_entry == 10)
{
    state->dbgwfar := ReadDebugRegister(6);
}
}
}

```

12.11.5 Debug state exit

When exiting debug state, the program counter must always be written. If the execution state or CPSR must be changed, this must be done before writing to the PC because writing to the CPSR can affect the PC.

Having restored the program state, the debugger can restart by writing to bit [1] of the DBGDRCR. It must then poll bit [1] of the DBGDSR to determine if the CPU has restarted.

Example 12-12 shows the code for exit from debug state.

Example 12-12 Leaving debug state

```

ExitDebugState(PROCESSOR_STATE *state)
{
    // Step 1. Update the CPSR value
    WriteCPSR(state->cpsr);
    // Step 2. Restore any registers corrupted by debug state. The function
    // WriteAllRegisters restores all general-purpose registers for all
    // processor modes apart from R0.
    WriteAllRegisters(state);
    // Step 3. Write the return address.
    WritePC(state->pc);
    // Step 4. Writing the PC corrupts R0 therefore, restore R0 now.
    WriteRegister(0, state->r0);
    // Step 5. Write the restart request bit in the DBGDRCR.
    WriteDebugRegister(36, 1<<1);
    // Step 6. Poll the RESTARTED flag in the DBGDSR.
    repeat
    {
        dbgdsr := ReadDebugRegister(34);
    }
    until (dbgdsr & (1<<1));
}

```

}

12.11.6 Accessing registers and memory in debug state

This section describes the following:

- *Reading and writing registers through the DCC*
- *Reading the PC in debug state*
- *Writing the CPSR in debug state* on page 12-64
- *Reading memory* on page 12-64
- *Fast register read/write* on page 12-66
- *Fast memory read/write* on page 12-67.

Reading and writing registers through the DCC

To read a single register, the debugger can use the sequence that Example 12-13 shows. This sequence depends on two other sequences, *Executing an ARM instruction through the DBGITR* on page 12-54 and *Target to host data transfer (host end)* on page 12-56.

Example 12-13 Reading an ARM register

```
uint32 ReadARMRegister(int Rd)
{
    // Step 1. Execute instruction MCR p14, 0, Rd, c0, c5, 0 through the DBGITR.
    ExecuteARMInstruction(0xEE00E15 + (Rd<<12));
    // Step 2. Read the register value through DBGDTRTX.
    reg_val := ReadDCC();
    return reg_val;
}
```

Example 12-14 shows a similar sequence for writing an ARM register.

Example 12-14 Writing an ARM register

```
WriteRegister(int Rd, uint32 reg_val)
{
    // Step 1. Write the register value to DBGDTRRX.
    WriteDCC(reg_val);
    // Step 2. Execute instruction MRC p14, 0, Rd, c0, c5, 0 to the DBGITR.
    ExecuteARMInstruction(0xEE10E15 + (Rd<<12));
}
```

Reading the PC in debug state

Example 12-15 shows the code to read the PC.

Example 12-15 Reading the PC

```
ReadPC()
{
    // Step 1. Save R0
    saved_r0 := ReadRegister(0);
}
```

```

// Step 2. Execute the instruction MOV r0, pc through the DBGITR.
ExecuteARMInstruction(0xE1A0000F);
// Step 3. Read the value of R0 that now contains the PC.
pc := ReadRegister(0);
// Step 4. Restore the value of R0.
WriteRegister(0, saved_r0);
return pc;
}

```

Note

You can use a similar sequence to write to the PC to set the return address when leaving debug state or to read the CPSR or coprocessor registers.

Writing the CPSR in debug state

Example 12-16 shows the code for writing the CPSR.

Example 12-16 Writing the CPSR

```

WriteCPSR(uint32 cpsr_val)
{
    // Step 1. Save R0.
    saved_r0 := ReadRegister(0);
    // Step 2. Write the new CPSR value to R0.
    WriteRegister(0, cpsr_val);
    // Step 3. Execute instruction MSR R0, CPSR through the DBGITR.
    ExecuteARMInstruction(0xE12FF000);
    // Step 4. Execute a PrefetchFlush instruction through the DBGITR.
    ExecuteARMInstruction(0xEE070F95);
    // Step 5. Restore the value of R0.
    WriteRegister(0, saved_r0);
}

```

Reading memory

Example 12-17 shows the code for reading a byte of memory.

Example 12-17 Reading a byte of memory

```

uint8 ReadByte(uint32 address, bool &aborted)
{
    // Step 1. Save the values of R0 and R1.
    saved_r0 := ReadRegister(0);
    saved_r1 := ReadRegister(1);
    // Step 2. Write the address to R0.
    WriteRegister(0, address);
    // Step 3. Execute the instruction LDRB R1,[R0] through the DBGITR.
    ExecuteARMInstruction(0xE5D01000);
    // Step 4. Read the value of R1 that contains the data at the address.
    datum := ReadRegister(1);
    // Step 5. Restore the corrupted registers R0 and R1.
    WriteRegister(0, saved_r0);
    WriteRegister(1, saved_r1);
    // Step 6. Check the DBGDSCR for a sticky abort.
    aborted := CheckForAborts();
}

```

```

    return datum;
}

```

Example 12-18 shows the code for checking for aborts after a memory access.

Example 12-18 Checking for an abort after memory access

```

bool CheckForAborts()
{
    // Step 1. Check the DBGDSCR for a sticky abort.
    dbgdsr := ReadDebugRegister(34);
    if (dbgdsr & ((1<<6) + (1<<7)))
    {
        // Step 2. Clear the sticky flag by writing DBGDSCR[2].
        WriteDebugRegister(36, 1<<2);
        return true;
    }
    else
    {
        return false;
    }
}

```

Note

You can use a similar sequence to read a halfword of memory and to write to memory.

To read or write blocks of memory, substitute the data instruction with one that uses post-indexed addressing. For example:

```
LDRB R1, [R0],1
```

This prevents reloading the address value for each sequential word.

Example 12-19 shows the code for reading a block of bytes of memory.

Example 12-19 Reading a block of bytes of memory

```

ReadBytes(uint32 address, bool &aborted, uint8 *data, int nbytes)
{
    // Step 1. Save the value of R0 and R1.
    saved_r0 := ReadRegister(0);
    saved_r1 := ReadRegister(1);
    // Step 2. Write the address to R0
    WriteRegister(0, address);
    while (nbytes > 0)
    {
        // Step 3. Execute instruction LDRB R1,[R0],1 through the DBGITR.
        ExecuteARMInstruction(0xE4D01001);
        // Step 4. Read the value of R1 that contains the data at the
        // address.
        *data++ := ReadRegister(1);
        --nbytes;
    }
    // Step 5. Restore the corrupted registers R0 and R1.
    WriteRegister(0, saved_r0);
    WriteRegister(1, saved_r1);
    // Step 6. Check the DBGDSCR for a sticky abort.
}

```

```

        aborted := CheckForAborts();
        return datum;
    }

```

Example 12-20 shows the sequence for reading a word of memory.

Note

A faster method is available for reading and writing words using the direct memory access function of the DCC. See *Fast memory read/write* on page 12-67.

Example 12-20 Reading a word of memory

```

uint32 ReadWord(uint32 address, bool &aborted)
{
    // Step 1. Save the value of R0.
    saved_r0 := ReadRegister(0);
    // Step 2. Write the address to R0.
    WriteRegister(0, address);
    // Step 3. Execute instruction LDC p14, c5, [R0] through the DBGITR.
    ExecuteARMInstruction(0xED905E00);
    // Step 4. Read the value from the DTR directly.
    datum := ReadDCC();
    // Step 5. Restore the corrupted register R0.
    WriteRegister(0, saved_r0);
    // Step 6. Check the DBGDSCR for a sticky abort.
    aborted := CheckForAborts();
    return datum;
}

```

Fast register read/write

When multiple registers must be read in succession, you can optimize the process by placing the DCC into stall mode and by writing the value 1 to the DCC access mode bits. For more information, see *CP14 c1, Debug Status and Control Register* on page 12-14.

Example 12-21 shows the sequence to change the DTR access mode.

Example 12-21 Changing the DTR access mode

```

SetDTRAccessMode(int mode)
{
    // Step 1. Write the mode value to DBGDSCR[21:20].
    dbgdscr := ReadDebugRegister(34);
    dbgdscr := (dbgdscr & ~(0x3<<20)) | (mode<<20);
    WriteDebugRegister(34, dbgdscr);
}

```

Example 12-22 shows the sequence to read registers in stall mode.

Example 12-22 Reading registers in stall mode

```

ReadRegisterStallMode(int Rd)
{

```

```

// Step 1. Write the opcode for MCR p14, 0, Rd, c5, c0 to the DBGITR.
// Write stalls until the DBGITR is ready.
WriteDebugRegister(33, 0xEE000E15 + (Rd<<12));
// Step 2. Read the register value through the DCC. Read stalls until
// DBGDTRTX is ready
reg_val := ReadDebugRegister(32);
return reg_val;
}

```

Example 12-23 shows the sequence to write registers in stall mode.

Example 12-23 Writing registers in stall mode

```

WriteRegisterInStallMode(int Rd, uint32 value)
{
    // Step 1. Write the value to the DBGDTRRX.
    // Write stalls until the DBGDTRRX is ready.
    WriteDebugRegister(32, value);
    // Step 2. Write the opcode for MRC p14, 0, Rd, c5, c0 to the DBGITR.
    // Write stalls until the DBGITR is ready.
    WriteDebugRegister(33, 0xEE100E15 + (Rd<<12));
}

```

Note

To transfer a register to the CPU when in stall mode, you are not required to poll the DBGDSCR each time an instruction is written to the DBGITR and a value read from or written to the DTR. The CPU stalls using the signal **PREADYDBGm** until the previous instruction has completed or the DTR register is ready for the operation.

Fast memory read/write

This section provides example code to enable faster reads from memory by making use of the DTR access mode.

Example 12-24 shows the sequence for reading a block of words of memory.

Example 12-24 Reading a block of words of memory

```

ReadWords(uint32 address, bool &aborted, uint32 *data, int nwords)
{
    // Step 1. Write the value 0b01 to DBGDSCR[21:20] for stall mode.
    SetDTRAccessMode(1);
    // Step 2. Save the value of R0.
    saved_r0 := ReadRegisterInStallMode(0);
    // Step 3. Write the address to read from to the DBGDTRRX.
    // Write stalls until the DBGDTRRX is ready.
    WriteRegisterInStallMode(0, address);
    // Step 4. Write the opcode for LDC p14, c5, [R0], 4 to the DBGITR.
    // Write stalls until the DBGITR is ready.
    WriteDebugRegister(33, 0xECB05E01);
    // Step 5. Write the value 0b10 to DBGDSCR[21:20] for fast mode.
    SetDCCAccessMode(2);
    // Step 6. Loop reading out the data.
    // Each time a word is read from the DBGDTRTX, the instruction is reissued.
    while (nwords > 1)

```

```

{
    *data++ = ReadDebugRegister(35);
    --nwords;
}
// Step 7. Write the value 0b00 to DBGDSCR[21:20] for non-blocking mode.
SetDTRAccessMode(0);
// Step 8. Must wait for the final instruction to complete. If there
// was an abort, this completes immediately.
do
{
    dbgdsr := ReadDebugRegister(34);
}
until (dbgdsr & (1<<24));
// Step 9: Check for aborts.
aborted := CheckForAborts();
// Step 10: Read the final word from the DCC.
if (!aborted) *data := ReadDCC();
// Step 11. Restore the corrupted register r0.
WriteRegister(0, saved_r0);
}

```

Example 12-25 shows the sequence for writing a block of words to memory.

Example 12-25 Writing a block of words to memory (fast download)

```

WriteWords(uint32 address, bool &aborted, uint32 *data, int nwords)
{
    // Step 1. Save the value of R0.
    saved_r0 := ReadRegister(0);
    // Step 2. Write the value 0b10 to DBGDSCR[21:20] for fast mode.
    SetDTRAccessMode(2);
    // Step 3. Write the opcode for MRC p14, 0, R0, c5, c0 to the DBGITR.
    // Write stalls until the DBGITR is ready but the instruction is not issued.
    WriteDebugRegister(33, 0xEE100E15);
    // Step 4. Write the address to read from to the DBGDTRRX
    // Write stalls until the DBGITR is ready, but the instruction is not reissued.
    WriteDebugRegister(32, address);
    // Step 5. Write the opcode for STC p14, c5, [R0], 4 to the DBGITR.
    // Write stalls until the DBGITR is ready but the instruction is not issued.
    WriteDebugRegister(33, 0xECA05E01);
    // Step 6. Loop writing the data.
    // Each time a word is written to the DBGDTRRX, the instruction is reissued.
    while (nwords > 0)
    {
        WriteDebugRegister(35, *data++);
        --nwords;
    }
    // Step 7. Write the value b00 to DBGDSCR[21:20] for normal mode.
    SetDTRAccessMode(0);
    // Step 8. Restore the corrupted register R0.
    WriteRegister(0, saved_r0);
    // Step 9. Check the DBGDSCR for a sticky abort.
    aborted := CheckForAborts();
}

```

Note

As the amount of data transferred increases, these functions reach an optimum performance of one debug register access per data word transferred.

After writing data to memory, you must execute a data synchronization barrier instruction to ensure that the memory window updates properly

12.12 Debugging systems with energy management capabilities

The processor offers functionality for debugging systems with energy-management capabilities. This section describes scenarios where the OS takes energy-saving measures when in an idle state.

The different measures that the OS can take to save energy during an idle state are divided into two groups:

Standby The OS takes measures that reduce energy consumption but maintain the processor state.

Power down The OS takes measures that reduce energy consumption but do not maintain the processor state, for example, Dormant or Shutdown mode. Recovery involves a reset of the processor after the power level has been restored, and reinstallation of the processor state.

Standby is the least invasive OS energy-saving state because it only implies that the core is unavailable. It does not clear any of the debug settings. For this case, the processor offers the following:

- If the processor is in standby and a halting debug event occurs, the processor:
 - leaves standby
 - retires the *Wait-For-Interrupt* (WFI) or *Wait-For-Event* (WFE) instruction
 - enters debug state.
- If the processor is in standby and detects an APB port access, it temporarily leaves standby state to complete the transaction. While the processor wakes up from standby, the APB access is held by keeping the **PREADYDBGm** signal LOW.

12.12.1 Emulating power down

By writing to bit [0] of the DBGPRCR in either CPU, the debugger causes the processor to assert the **DBGNOPWRDWN** output. The expected usage model of this signal is that it connects to the system power controller and that, when HIGH, it indicates that this controller must work in emulate mode.

On a power-down request from the processor, if the power controller is in emulate mode, it does not remove processor power or ETM power. Otherwise, it behaves exactly the same as in normal mode.

Emulating power down is ideal for debugging applications running on top of operating systems that are free of errors because the debug register settings are not lost on a power-down event. However, you must ensure that:

- **nIRQm** and **nFIQm** interrupts and **EVENTIm** events to the processor are externally masked as part of the emulation to prevent them from retiring the WFI or WFE instruction from the pipeline.
- The reset controller asserts **nRESETm** only on emulated power up, rather than combining it with **DBGRESETmn**. Asserting **DBGRESETmn** clears the debug registers inside the processor.
- The timing effects of power down and voltage stabilization are not factored in the power-down emulation. This is the case for systems with voltage recovery controlled by a closed loop system that monitors the processor supply voltage, rather than a fixed timed for voltage recovery.

- The emulation does not model state lost during power down, making it possible to miss errors in the state storage and recovery routines.
- Attaching the debugger for a postmortem debug session is not possible because setting the **DBGNOPWRDWN** signal to 1 might not cause the processor to power up. The effect of setting **DBGNOPWRDWN** to 1 when the processor is already powered down is implementation-defined, and is up to the system designer.

Chapter 13

Integration Test Registers

This chapter describes how to use the Integration Test Registers in the processor. It contains the following sections:

- *About Integration Test Registers* on page 13-2
- *Summary of the processor registers used for integration testing* on page 13-3
- *Processor integration testing* on page 13-4.

13.1 About Integration Test Registers

The processor contains Integration Test Registers that enable you to verify integration of the design and enable topology detection of the design using debug tools. The *Integration Mode Control Register* (DBGITCTRL), that is also described in this chapter, controls the use of the Integration Test Registers.

The Integration Test Registers are programmed using the debug APB interface. For more information on using the debug APB interface see Chapter 12 *Debug*.

When programming the Integration Test Registers you must enable all the changes at the same time.

For more information about the Integration Test Registers and the Integration Mode Control Register see the *ARM Architecture Reference Manual*.

13.2 Summary of the processor registers used for integration testing

Table 13-1 lists the processor Integration Test Registers and the *Integration Mode Control Register* (DBGITCTRL).

Table 13-1 Integration Test Registers summary

Register name	Base offset	Default value	Type	Description
Integration Test Registers				
DBGITETMIF	0xED8	n/a ^a	WO	See <i>DBGITETMIF Register (ETM interface)</i> on page 13-5
DBGITMISCOUT	0xEF8	n/a	WO	See <i>DBGITMISCOUT Register (Miscellaneous Outputs)</i> on page 13-6
DBGITMISCIN	0xEFC	-a	RO	See <i>DBGITMISCIN Register (Miscellaneous Inputs)</i> on page 13-7
Integration Mode Control Register				
DBGITCTRL	0xF00	0	R/W	See <i>Integration Mode Control Register (DBGITCTRL)</i> on page 13-8

a. See the register description for this value.

13.3 Processor integration testing

This section describes the behavior and use of the Integration Test Registers that are in the processor. It also describes the Integration Mode Control Register that controls the use of the Integration Test Registers. For more information about the DBGITCTRL see the *ARM Architecture Reference Manual*.

If you want to utilise the integration test registers you must first set bit [0] of the Integration Mode Control Register to 1.

- You can use the write-only Integration Test Registers to set the outputs of some of the processor signals. Table 13-2 shows the signals that you can write in this way.
- You can use the read-only Integration Test Registers to read the state of some of the processor inputs. Table 13-3 on page 13-5 shows the signals that you can read in this way.

Various CoreSight components, including ETM-R5, also include Integration Test Registers that you can use in conjunction with processor Integration Test Registers for testing the connectivity between them. For more information see the relevant documentation, for example the *ETM-R5 Technical Reference Manual*

Table 13-2 Output signals that can be controlled by the Integration Test Registers

Signal	Register	Bit	Register description
DBGRESTARTEDm	DBGITMISCOUT	[9]	See <i>DBGITMISCOUT Register (Miscellaneous Outputs)</i> on page 13-6
DBGTRIGGERm	DBGITMISCOUT	[8]	
ETMWFIPENDINGm	DBGITMISCOUT	[5]	
nPMUIRQm	DBGITMISCOUT	[4]	
COMMTXm	DBGITMISCOUT	[2]	
COMMRXm	DBGITMISCOUT	[1]	
DBGACKm	DBGITMISCOUT	[0]	
EVNTBUSm[54, 0]	DBGITETMIF	[13:12]	See <i>DBGITETMIF Register (ETM interface)</i> on page 13-5
ETMCIDm[31, 0]	DBGITETMIF	[11:10]	
ETMDAm[31, 0]	DBGITETMIF	[7:6]	
ETMDCTLm[11, 0]	DBGITETMIF	[5:4]	
ETMDDm[63, 0]	DBGITETMIF	[9:8]	
ETMIAm[31, 1]	DBGITETMIF	[3:2]	
ETMICTLm[13, 0]	DBGITETMIF	[1:0]	

Table 13-3 Input signals that can be read by the Integration Test Registers

Signal	Register	Bit	Register description
DBGRESTARTm	DBGITMISCIN	[11]	See <i>DBGITMISCIN Register (Miscellaneous Inputs)</i> on page 13-7
ETMEXTOUTm[1:0]	DBGITMISCIN	[9:8]	
nETMWFIREADYm	DBGITMISCIN	[5]	
nIRQm	DBGITMISCIN	[2]	
nFIQm	DBGITMISCIN	[1]	
EDBGRQm	DBGITMISCIN	[0]	

This section describes:

- *Using the Integration Test Registers*
- *Performing integration testing*
- *DBGITETMIF Register (ETM interface)*
- *DBGITMISCOUT Register (Miscellaneous Outputs)* on page 13-6
- *DBGITMISCIN Register (Miscellaneous Inputs)* on page 13-7
- *Integration Mode Control Register (DBGITCTRL)* on page 13-8.

13.3.1 Using the Integration Test Registers

When bit [0] of the Integration Mode Control Register (DBGITCTRL) is set to b1:

- Values written to the write-only Integration Test Registers map onto the specified outputs of the macrocell. For example, writing b1 to DBGITMISCOUT[0] causes **DBGACKm** to be asserted HIGH.
- Values read from the read-only Integration Test Registers correspond to the values of the specified inputs of the macrocell. For example, if you read DBGITMISCIN[9:8] you obtain the value of **ETMEXTOUTm[1:0]**.

13.3.2 Performing integration testing

When you perform integration testing or topology detection, ARM strongly recommends that the processor is halted, because toggling input and output pins might have an unwanted effect on the operation of the processor. If you follow this recommendation, you must not set the DBGITCTRL Register until the processor has halted.

After you perform integration testing or topology detection, that is, the Integration Mode Control Register has been set, the system must be reset. This is because the signals that are toggled can have an unwanted effect on connected devices.

13.3.3 DBGITETMIF Register (ETM interface)

The DBGITETMIF Register at offset 0xED8 is write-only. Figure 13-1 on page 13-6 shows the register bit assignments.


**Figure 13-1 DBGITETMIF Register bit assignments**

Table 13-4 shows the fields when writing the DBGITETMIF Register. When this register is written the appropriate output pins take the value written.

Table 13-4 DBGITETMIF Register bit assignments

Bits	Name	Function
[31:14]	-	Reserved. Write as zero.
[13]	EVNTBUSm[54]	Set value of the EVNTBUSm[54] output pin.
[12]	EVNTBUSm[0]	Set value of the EVNTBUSm[0] output pin.
[11]	ETMCIDm[31]	Set value of the ETMCIDm[31] output pin.
[10]	ETMCIDm[0]	Set value of the ETMCIDm[0] output pin.
[9]	ETMDDm[63]	Set value of the ETMDDm[63] output pin.
[8]	ETMDDm[0]	Set value of the ETMDDm[0] output pin.
[7]	ETMDAm[31]	Set value of the ETMDAm[31] output pin.
[6]	ETMDAm[0]	Set value of the ETMDAm[0] output pin.
[5]	ETMDCTLm[11]	Set value of the ETMDCTLm[11] output pin.
[4]	ETMDCTLm[0]	Set value of the ETMDCTLm[0] output pin.
[3]	ETMIAm[31]	Set value of the ETMIAm[31] output pin.
[2]	ETMIAm[1]	Set value of the ETMIAm[1] output pin.
[1]	ETMICTLm[13]	Set value of the ETMICTLm[13] output pin.
[0]	ETMICTLm[0]	Set value of the ETMICTLm[0] output pin.

13.3.4 DBGITMISCOUT Register (Miscellaneous Outputs)

The DBGITMISCOUT Register at offset 0xEF8 is write-only. Figure 13-2 on page 13-7 shows the register bit assignments.




Figure 13-2 DBGITMISCOUT Register bit assignments

Table 13-5 shows the fields when writing the DBGITMISCOUT Register. When this register is written the appropriate output pins take the value written.

Table 13-5 DBGITMISCOUT Register bit assignments

Bits	Name	Function
[31:10]	-	Reserved. Write as zero.
[9]	DBGRESTARTEDm	Set value of the DBGRESTARTEDm output pin.
[8]	DBGTRIGGERm	Set value of the DBGTRIGGERm output pin.
[7:6]	-	Reserved. Write as zero.
[5]	ETMWFIPENDINGm	Set value of the ETMWFIPENDINGm output pin.
[4]	nPMUIRQm	Set value of nPMUIRQm output pin.
[3]	-	Reserved. Write as zero.
[2]	COMMTXm	Set value of COMMTXm output pin.
[1]	COMMRXm	Set value of COMMRXm output pin.
[0]	DBGACKm	Set value of the DBGACKm output pin.

13.3.5 DBGITMISCIN Register (Miscellaneous Inputs)

The DBGITMISCIN Register at offset 0xEFC is read-only. Figure 13-3 on page 13-8 shows the register bit assignments.




Figure 13-3 DBGITMISCIN Register bit assignments

Table 13-6 lists the register bit assignments for the DBGITMISCIN Register.

Table 13-6 DBGITMISCIN Register bit assignments

Bits	Name	Function
[31:12]	-	Reserved. Read Undefined.
[11]	DBGRESTARTm	Read value of the DBGRESTARTm input pin.
[10]	-	Reserved. Read Undefined.
[9:8]	ETMEXTOUTm	Read value of the ETMEXTOUTm[1:0] input pins.
[7:6]	-	Reserved. Read Undefined.
[5]	nETMWFIREADYm	Reads the nETMWFIREADYm input pin. Although this pin is active LOW, the value of this bit matches the physical state of the signal: 0 = input pin is LOW (asserted) 1 = input pin is HIGH (deasserted).
[4:3]	-	Reserved. Read Undefined.
[2]	nFIQm	Read value of nFIQm input pin.
[1]	nIRQm	Read value of nIRQm input pin.
[0]	EDBGRQm	Read value of EDBGRQm input pin.

13.3.6 Integration Mode Control Register (DBGITCTRL)

The DBGITCTRL Register, register 0x3C0 at offset 0xF00, is read/write. Figure 13-4 shows the register bit assignments.




Figure 13-4 DBGITCTRL Register bit assignments

Table 13-7 shows the fields of the DBGITCTRL Register.

Table 13-7 DBGITCTRL Register bit assignments

Bits	Access	Name	Function
[31:1]	RAZ/SBZP	-	Reserved.
[0]	R/W	INTMODE	Controls whether the processor is in normal operating mode or integration mode: b0 = normal operation, this is the reset value b1 = integration mode enabled.

Writing to the DBGITCTRL register controls whether the processor is in its default functional mode, or in integration mode, where the inputs and outputs of the device can be directly controlled for the purpose of integration testing or topology detection. For more information see the *ARM Architecture Reference Manual*.

Appendix A

Signal Descriptions

This appendix describes the processor signals. It contains the following sections:

- *About the processor signal descriptions* on page A-2
- *Global signals* on page A-3
- *Configuration signals* on page A-4
- *Interrupt signals, including VIC interface signals* on page A-8
- *L2 interface signals* on page A-9
- *TCM interface signals* on page A-22
- *Redundant CPU signals* on page A-25
- *Debug interface signals* on page A-26
- *ETM interface signals* on page A-28
- *Test signals* on page A-29
- *MBIST signals* on page A-30
- *Validation signals* on page A-31
- *FPU signals* on page A-32
- *Split/Lock* on page A-33
- *Power modes* on page A-34.

A.1 About the processor signal descriptions

The tables in this appendix list the processor signals, along with their dimensions and direction, input or output, and a high-level description. Unless otherwise specified, all signals are sampled on or driven from the rising edge of the clock, **CLKIN**.

Many of the signal names have an **m** suffix, that appears before the **n** suffix in the case of negative sense signals. This indicates that the processor has two signals, one for each CPU, named with **m** being **0** or **1** for CPU0 and CPU1 respectively.

The Cortex-R5 processor has the same signals regardless of configuration. If a particular feature is not implemented, the signals associated with that feature are not used.

A.2 Global signals

Table A-1 shows the processor global signals.

Table A-1 Global signals

Signal	Direction	Description
CLKIN	Input	Master processor clock.
ACPRESETn	Input	ACP reset. Assert with nRESET0 and nRESET1 to reset the whole processor except the debug registers. This signal can be asserted asynchronously to CLKIN .
ACPIDLE	Output	Indicate when uSCU is empty, for drain-and-power-down.
nRESETm	Input	CPU non-debug logic reset. These signals can be asserted asynchronously to CLKIN .
nSYSPORESET	Input	System power on reset.
nCPUHALTm	Input	Processor halt after reset. These signals can be asserted asynchronously to CLKIN .
DBGNOCLKSTOP	Input	Processor does not stop the clocks when entering standby mode.
nCLKSTOPPEDm	Output	When LOW, this indicates clock has been stopped because processor is in Standby Mode. It is never asserted without one of WFIPIPESTOPPEDm or WFEEPIPESTOPPEDm .
nWFEEPIPESTOPPEDm	Output	When LOW, this indicates that the CPU is in standby mode because of a WFE instruction. The CPU pipeline is inactive..
nWFIPIPESTOPPEDm	Output	When LOW, this indicates the CPU is in standby mode because of a WFI instruction. The CPU pipeline is inactive.
EVENTIm	Input	Event input signal.
EVENTOm	Output	Event output signal.

A.3 Configuration signals

Table A-2 shows the processor configuration signals. These signals must be tied off, or only changed under reset.

Table A-2 Configuration signals

Signal	Direction	Description
VINITHIm	Input	Reset V-bit value. When HIGH indicates HIVECS mode at reset. See <i>c1, System Control Register</i> on page 4-38 for more information.
CFGEE	Input	Reset EE-bit value. When HIGH indicates the implementation uses BE-8 mode for exceptions at reset. See <i>c1, System Control Register</i> on page 4-38 for more information.
CFGIE	Input	Instruction side endianness, reflected in the IE-bit. When HIGH indicates that big endian instruction fetch is used. See <i>c1, System Control Register</i> on page 4-38 for more information.
INITRAMAm	Input	Reset value of ATCM enable bit. When HIGH indicates Tightly-Coupled Memory A, ATCM, enabled at reset. See <i>c9, ATCM Region Register</i> on page 4-64 for more information.
INITRAMBm	Input	Reset value of BTCM bit. When HIGH indicates Tightly-Coupled Memory B, BTCM, enabled at reset. See <i>c9, BTCM Region Register</i> on page 4-63 for more information.
LOCZRAMAm	Input	When HIGH indicates ATCM initial base address is zero and BTCM base address is implementation-defined. When LOW indicates BTCM initial base address is zero and ATCM base address is implementation-defined.
TEINIT	Input	Reset TE-bit value. Determines exception handling state at reset. When set to: 0 = ARM 1 = Thumb. See <i>c1, System Control Register</i> on page 4-38 for more information.
CFGATCMSZm[3:0]	Input	Selects the ATCM size. The encodings for the TCM sizes are: b0000 = 0KB b0011 = 4KB b0100 = 8KB b0101 = 16KB b0110 = 32KB b0111 = 64KB b1000 = 128KB b1001 = 256KB b1010 = 512KB b1011 = 1MB b1100 = 2MB b1101 = 4MB b1110 = 8MB.

Table A-2 Configuration signals (continued)

Signal	Direction	Description
CFGBTCMSZm[3:0]	Input	Selects the BTCM size. The encodings for the TCM sizes are: b0000 = 0KB b0011 = 4KB b0100 = 8KB b0101 = 16KB b0110 = 32KB b0111 = 64KB b1000 = 128KB b1001 = 256KB b1010 = 512KB b1011 = 1MB b1100 = 2MB b1101 = 4MB b1110 = 8MB.
CFGNMFI_m	Input	When HIGH, enable nonmaskable Fast Interrupts. Reflected in the NMFI bit. See <i>c1, System Control Register</i> on page 4-38 for more information.
ENTCM1IF_m	Input	Enable B1TCM interface. Use B0TCM only if this signal not tied HIGH.
PARECCENRAMm[2:0]	Input	TCMs ECC check enable. Tie each bit HIGH to enable ECC checking on the appropriate TCM at reset. The bit allocations are as follows: [2] = B1TCM ^a [1] = B0TCM ^a [0] = ATCM. See <i>c1, Auxiliary Control Register</i> on page 4-41 for more information.
PARITYLEVEL	Input	Selects between odd and even parity for caches and buses. See Chapter 8 <i>Level One Memory System</i> : Tie LOW for even parity Tie HIGH for odd parity.
ERRENRAMm[2:0]	Input	TCMs external error enable. Tie each bit high to enable the external error signals for each TCM at reset. The bit allocations are as follows: [2] = B1TCM [1] = B0TCM [0] = ATCM. See <i>c1, Auxiliary Control Register</i> on page 4-41 for more information.
RMWENRAMm[1:0]^b	Input	RMW enable bits reset values. Tie each bit high to enable read-modify-write for TCM interfaces at reset. ^c The bit allocations are as follows: [1] = BTCM [0] = ATCM. See <i>c1, Auxiliary Control Register</i> on page 4-41 for more information.
SLBTCMSB_m	Input	Use most significant bit of BTCM address to select B1TCM if this signal is HIGH. Use bit [3] of the BTCM address if this signal is LOW.
INITPPX_m	Input	AXI peripheral interface is enabled out-of-reset.
INITPPH_m	Input	AHB peripheral interface is enabled out-of-reset.
GROUPID[3:0]	Input	ID of Cortex-R5 processor group (reflected in MPIDR).

Table A-2 Configuration signals (continued)

Signal	Direction	Description
PPHBASEm[31:12]	Input	Base address of AHB peripheral interface. Must be size-aligned.
PPHSIZEm[4:0]	Input	Size of AHB peripheral interface. See Table A-3 for the size encodings.
PPXBASEm[31:12]	Input	Base address of AXI peripheral interface. Must be size aligned.
PPXSIZEm[4:0]	Input	Size of AXI peripheral interface. See Table A-3 for the size encodings.
PPVBASEm[31:12]	Input	Base address of virtual-AXI peripheral interface. Must be within AXI PP and size-aligned. The virtual AXI peripheral interface region must be the same size or smaller than the AXI peripheral interface.
PPVSIZEm[4:0]	Input	Size of virtual-AXI peripheral interface. See Table A-3 for the size encodings.

a. If the BTCM is configured with ECC, bit[2] and bit[1] must be the same value.

b. Not used if 32-bit ECC is included.

c. Not available in r0px revisions of the processor.

Table A-3 shows the peripheral interface size encodings.

Table A-3 Peripheral interface size encodings

Encoding	Size
b00011	4KB
b00100	8KB
b00101	16KB
b00110	32KB
b00111	64KB
b01000	128KB
b01001	256KB
b01010	512KB
b01011	1MB
b01100	2MB
b01101	4MB
b01110	8MB
b01111	16MB
b10000	32MB
b10001	64MB
b10010	128MB
b10011	256MB
b10100	512MB

Table A-3 Peripheral interface size encodings (continued)

Encoding	Size
b10101	1GB
b10110	2GB
b10111	4GB

A.4 Interrupt signals, including VIC interface signals

Table A-4 shows the interrupt signals including signals used on the VIC interface.

Table A-4 Interrupt signals

Signal	Direction	Description
nFIQm	Input	Fast interrupt ^a . This signal can be asserted asynchronously if INTSYNCEN is HIGH.
nIRQm	Input	Normal interrupt ^a . This signal can be asserted asynchronously if INTSYNCEN is HIGH.
INTSYNCEN	Input	Tie HIGH if the interrupt inputs are asynchronous to CLKIN . Tie LOW if the interrupt inputs are synchronous to CLKIN .
IRQADDRVm	Input	Indicates IRQADDRm is valid. This signal can be asserted asynchronously if IRQADDRVSYNCEN is HIGH.
IRQADDRVSYNCEN	Input	Tie HIGH if the IRQADDRVm input from the VIC is asynchronous to CLKIN . Tie HIGH if the IRQADDRVm input from the VIC is synchronous to CLKIN .
IRQADDRm[31:2]	Input	Address of the IRQ. This signal can be asserted asynchronously but must be stable when IRQADDRVm is asserted.
IRQACKm	Output	Acknowledges interrupt.
nPMUIRQm	Output	Interrupt request by <i>Performance Monitor Unit</i> (PMU).

a. This signal is level-sensitive and must be held LOW until a suitable interrupt response is received from the processor.

A.5 L2 interface signals

This section describes the processor L2 interface AXI signals. For more information on *Advanced Microcontroller Bus Architecture* (AMBA) AXI signals see the *AMBA AXI Protocol Specification*. For more information on the AHB signals, see the *AMBA 3 AHB-Lite Protocol Specification*.

A.5.1 AXI master port

Table A-5 shows the AXI master port signals for the L2 interface. With the exception of the **ACLKENMm**, all signals are only sampled or driven on **CLKIN** edges when **ACLKENMm** is asserted, see *AMBA interface clocking* on page 2-16 for more information.

Table A-5 AXI master port signals for the L2 interface

Signal	Direction	Description
ACLKENMm	Input	Clock enable for the AXI master port.
Write address channel		
AWADDRMm[31:0]	Output	Transfer start address.
AWBURSTMm[1:0]	Output	Write burst type.
AWCACHMm[3:0]	Output	Provides decode information for outer attributes: b0000 = Strongly Ordered. b0001 = Device. b0011 = Normal, Non-cacheable. b0110 = Normal, Cacheable. write-through. b1111 = Normal, Cacheable. write-back, write allocation. b0111 = Normal, Cacheable. write-back, no write allocation.
————— Note ————— The AXI specification describes these encodings using the pre-ARMv6 terms such as <i>cacheable-bufferable</i>. These terms are equivalent to the ARMv6 memory-type descriptions such as <i>Normal</i>, <i>Non-cacheable</i> used here.		
AWIDMm[3:0]	Output	The identification tag for the write address group of signals.
AWLENMm[3:0]	Output	Write transfer burst length.
AWLOCKMm[1:0]	Output	Lock signal.
AWPROTMm[2:0]	Output	Protection type.
AWREADYMm	Input	Address ready. The slave uses this signal to indicate that it can accept the address.
AWSZEMm[2:0]	Output	Indicates the size of the transfer.
AWINNERMm[3:0]	Output	Provides inner attribute information for the write address channel. See Table 9-2 on page 9-6 for information about the encoding of this signal. ^a
AWSHAREMm[0]	Output	Indicates the shareability of the address: 0 = non-shared 1 = shared.
AWVALIDMm	Output	Indicates address and control are valid.

Table A-5 AXI master port signals for the L2 interface (continued)

Signal	Direction	Description
Write data channel		
WDATAMm[63:0]	Output	Write data.
WIDMm[3:0]	Output	The identification tag for the write data group of signals.
WLASTMm	Output	Indicates the last data transfer of a burst.
WREADYMm	Input	Indicates that the slave is ready to accept write data
WSTRBMm[7:0]	Output	Write strobes used to indicate which byte lanes must be updated.
WVALIDMm	Output	Indicates address and control are valid.
Write response channel		
BIDMm[3:0]	Input	The identification tag for the write response signal.
BREADYMm	Output	Indicates that the CPU is ready to accept write response.
BRESPMm[1:0]	Input	Write response.
BVALIDMm	Input	Indicates that a valid write response is available.
Read address channel		
ARADDRMm[31:0]	Output	Instruction fetch burst start address.
ARBURSTMm[1:0]	Output	Burst type.
ARCAHEMm[3:0]	Output	Provides decode information for outer attributes: b0000 = Strongly Ordered. b0001 = Device. b0011 = Normal, Non-cacheable. b0110 = Normal, Cacheable. write-through. b1111 = Normal, Cacheable. write-back, write allocation. b0111 = Normal, Cacheable. write-back, no write allocation.
Note The AXI specification describes these encodings using the pre-ARMv6 terms such as <i>cacheable-bufferable</i>. These terms are equivalent to the ARMv6 memory-type descriptions such as <i>Normal</i>, <i>Non-cacheable</i> used here.		
ARIDMm[3:0]	Output	Identification tag for the read address group of signals
ARLENMm[3:0]	Output	Instruction fetch burst length.
ARLOCKMm[1:0]	Output	Lock signal.
ARPROTMm[2:0]	Output	Protection type.
ARREADYMm	Input	Address ready. The slave uses this signal to indicate that it can accept the address.
ARSIEMm[2:0]	Output	Indicates the size of the transfer.
ARINNERMm[3:0]	Output	Provides inner attribute for the read address channel. See Table 9-2 on page 9-6 for information about the encoding of this signal. ^a

Table A-5 AXI master port signals for the L2 interface (continued)

Signal	Direction	Description
ARSHAREMm	Output	Indicates the shareability of the address: 0 = non-shared 1 = shared.
ARVALIDMm	Output	Indicates address and control are valid.
Read Data Channel		
RDATAMm[63:0]	Input	Read Data.
RIDMm[3:0]	Input	The identification tag for the read data group of signals.
RLASTMm	Input	Indicates the last transfer in a read burst.
RREADYMm	Output	Read ready signal indicating that the bus master can accept read data and response information.
RRESPMm[1:0]	Input	Read response.
RVALIDMm	Input	Indicates that read data is available.

a. This is an AXI extension signal.

A.5.2 AXI master port error detection signals

Table A-6 shows the AXI master port error detection signals. these signals are only generated if the processor is configured to include AXI bus parity. See *Configurable options* on page 1-6 for more information.

Table A-6 AXI master port error detection signals

Signal	Direction	Description
ARADDRPTYMm[3:0]	Output	Parity bits for ARADDRMm ^a
ARCTLPTYMm[3:0]	Output	Parity bits for the rest of the read address channel ^a
ARRPTYMm	Input	Parity bit for ARREADYMm
ARUSERPTYMm	Output	Parity bit for sideband signals ^a
ARVPTYMm	Output	Parity bit for ARVALIDMm
AWADDRPTYMm[3:0]	Output	Parity bits for AWADDRMm . ^a
AWCTLPTYMm[3:0]	Output	Parity bits for the rest of the write address channel ^a
AWRPTYMm	Input	Parity bit for AWREADYMm
AWUSERPTYMm	Output	Parity bit for sideband signals ^a
AWVPTYMm	Output	Parity bit for AWVALIDMm
AXIMCORRm	Output	Correctable error detected on RDATAMm
AXIMFATALm[4:0]	Output	Fatal error detected on AXI master, per channel {R, AR, B, W, AW}
BCTLPTYMm[1:0]	Input	Parity for buffered response channel ^a
BRPTYMm	Output	Parity bit for BREADYMm

Table A-6 AXI master port error detection signals (continued)

Signal	Direction	Description
BVPTYMm	Input	Parity bit for BVALIDMm
MERRADDRm[31:3]^b	Output	Address of correctable error, doubleword
RCTLPTYMm[1:0]	Input	Parity for rest of read data channel ^a
RERRCODEMm[7:0]	Input	ECC code for RDATAMm^a
RRPTYMm	Output	Parity bit for RREADYMm
RVPTYMm	Input	Parity bit for RVALIDMm
WCTLPTYMm[2:0]	Output	Parity bits for the rest of the write data channel ^a
WERRCODEMm[7:0]	Output	ECC code for WDATAMm^a
WRPTYMm	Input	Parity bit for WREADYMm
WVPTYMm	Input	Parity bit for WVALIDMm

a. This is an AXI extension signal.

b. This address bus is also used by other AMBA masters: PPX and PPH.

A.5.3 AXI slave port

Table A-7 shows the AXI slave port signals for the L2 interface. With the exception of the **ACLKENS_m**, all signals are only sampled or driven on **CLKIN** edges when **ACLKENS_m** is asserted, see *AMBA interface clocking* on page 2-16 for more information.

Table A-7 AXI slave port signals for the L2 interface

Signal	Direction	Description
ACLKENS_m	Input	Clock enable for the AXI slave port.
Write Address Channel		
AWADDRSm[31:0]	Input	Transfer start address.
AWBURSTSm[1:0]	Input	Write burst type.
AWCACHESm[3:0]	Input	Write address outer attribute information.
AWCELSm[3:0]	Input	Memory type select data cache, instruction cache, BTCM or ATCM, one hot. ^a
AWIDSm[7:0]	Input	The identification tag for the write address group of signals.
AWLENSm[3:0]	Input	Write transfer burst length.
AWLOCKSm[1:0]	Input	Lock signal.
AWPROTSm[2:0]	Input	Protection information, privileged/normal access.
AWREADYSm	Output	Address ready. The slave uses this signal to indicate that it can accept the address.
AWSIZESm[2:0]	Input	Indicates the size of the transfer.
AWVALIDSm	Input	Indicates address and control are valid.
Write Data Channel		
WDATASm[63:0]	Input	Write data.

Table A-7 AXI slave port signals for the L2 interface (continued)

Signal	Direction	Description
WIDSm[7:0]	Input	The identification tag for the write group of signals.
WLASTSm	Input	Indicates the last data transfer of a burst.
WREADYSm	Output	Indicates that the slave is ready to accept write data.
WSTRBSm[7:0]	Input	Write strobes used to indicate which byte lanes must be updated.
WVALIDSm	Input	Indicates address and control are valid.
Write Response Channel		
BIDSm[7:0]	Output	The identification tag for the write response signal.
BREADYSm	Input	Indicates that the CPU is ready to accept write response.
BRESPSm[1:0]	Output	Write response.
BVALIDSm	Output	Indicates that a valid write response is available.
Read Address Channel		
ARADDRSm[31:0]	Input	Instruction fetch burst start address.
ARBURSTSm[1:0]	Input	Burst type.
ARCACHESm[3:0]	Input	Read address outer attribute information.
ARIDSm[7:0]	Input	Identification tag for the read address group of signals.
ARLENSm[3:0]	Input	Instruction fetch burst length.
ARLOCKSm[1:0]	Input	Lock signal.
ARPROTSm[2:0]	Input	Protection information, privileged/normal access.
ARREADYSm	Output	Address ready. The slave uses this signal to indicate that it can accept the address.
ARSIZEsm[2:0]	Input	Indicates the size of the transfer.
ARCSELSm[3:0]	Input	Memory type select {data cache, instruction cache, BTCM or ATCM}, one hot. ^a
ARVALIDSm	Input	Indicates address and control are valid.
Read Data Channel		
RDATASm[63:0]	Output	Read data.
RIDSm[7:0]	Output	The identification tag for the read data group of signals.
RLASTSm	Output	Indicates the last transfer in a read burst.
RREADYSm	Input	Read ready signal indicating that the bus master can accept read data and response information.
RRESPSm[1:0]	Output	Read response.
RVALIDSm	Output	Indicates address and control are valid.

a. This is an AXI extension signal.

A.5.4 AXI slave port error detection signals

Table A-8 shows the AXI slave port error detection signals. These signals are only generated if the processor is configured to include AXI bus parity. See *Configurable options* on page 1-6 for more information.

Table A-8 AXI slave port error detection signals

Signal	Direction	Description
ARADDRPTYSm[3:0]	Input	Parity bits for ARADDRSm ^a
ARCTLPTYSm[3:0]	Input	Parity bits for the rest of the read address channel ^a
ARRPTYSm	Output	Parity bit for ARREADYSm
ARUSERPTYSm	Input	Parity bit for sideband signals ^a
ARVPTYSm	Input	Parity bit for ARVALIDSm
AWADDRPTYSm[3:0]	Input	Parity bits for AWADDRSm ^a
AWCTLPTYSm[3:0]	Input	Parity bits for the rest of the write address channel ^a
AWRPTYSm	Output	Parity bit for AWREADYSm
AWUSERPTYSm	Input	Parity bit for sideband signals ^a
AWVPTYSm	Input	Parity bit for AWVALIDSm
AXISCORRm	Output	Correctable error, write data channel
AXISFATALm[4:0]	Output	Fatal error, per channel.
BCTLPTYSm[1:0]	Output	Parity for buffered response channel ^a
BRPTYSm	Output	Parity bit for BREADYSm
BVPTYSm	Input	Parity bit for BVALIDSm
RCTLPTYSm[1:0]	Output	Parity for rest of read data channel ^a
RERRCODESm[7:0]	Input	ECC code for RDATASm ^a
RRPTYSm	Output	Parity bit for RREADYSm
RVPTYSm	Output	Parity bit for RVALIDSm
SERRADDRm[22:3]	Output	Address of correctable error, within doubleword
SERRCSELm[3:0]	Output	Chip-select of correctable error.
WCTLPTYSm[2:0]	Input	Parity bits for rest of write data channel ^a
WERRCODESm[7:0]	Input	ECC code for WDATAMm ^a
WRPTYSm	Output	Parity bit for WREADYSm
WVPTYSm	Input	Parity bit for WVALIDSm

a. This is an AXI extension signal.

A.5.5 ACP slave port

Table A-9 shows the ACP slave port signals.

Table A-9 ACP slave port signals

Signal	Direction	Description
ACLKENC	Input	Clock enable, shared between ACP slave and master port.
Write Address Channel		
AWIDCS[1:0]	Input	The identification tag for the write address group of signals.
AWADDRCS[31:0]	Input	Transfer start address.
AWLENC[3:0]	Input	Write transfer burst length.
AWSIZECS[2:0]	Input	Indicates the size of the transfer.
AWBURSTCS[1:0]	Input	Write burst type.
AWLOCKCS[1:0]	Input	Lock signal.
AWCACHECS[3:0]	Input	Provides decode information for outer attributes: b0000 = Strongly Ordered. b0001 = Device. b0011 = Normal, Non-cacheable. b0110 = Normal, Cacheable. write-through. b1111 = Normal, Cacheable. write-back, write allocation. b0111 = Normal, Cacheable. write-back, no write allocation.
Note The AXI specification describes these encodings using the pre-ARMv6 terms such as <i>cacheable-bufferable</i>. These terms are equivalent to the ARMv6 memory-type descriptions such as <i>Normal</i>, <i>Non-cacheable</i> used here.		
AWPROTCS[2:0]	Input	Protection signals provide additional information about a bus access.
AWCOHERENTCS	Input	Require caches to be made coherent with this access. ^a
AWUSERCS[3:0]	Input	For transmission of other sideband information. ^a
AWVALIDCS	Input	Indicates address and control are valid.
AWREADYCS	Output	Address ready. The slave uses this signal to indicate it is ready to accept the address.
Write Response Channel		
BIDCS[1:0]	Output	The identification tag for the write response signal.
BRESPCS[1:0]	Output	Write response.
BVALIDCS	Output	Indicates that a valid write response is available.
BREADYCS	Input	Indicates that the CPU is ready to accept write response.
BMISSCS[1:0]	Output	Access did not hit in either cache, or coherency not required. One bit for each CPU. ^a
BHITDIRTYCS[1:0]	Output	Access hit a dirty line, or Dormant CPU, and was not invalidated. One bit for each CPU. ^a
BUSERCS[3:0]	Output	For transmission of other sideband information. ^a

a. This is an AXI extension signal.

A.5.6 ACP slave port error detection signals

Table A-10 shows the ACP slave port error detection signals. These signals are only generated if the processor is configured to include AXI bus parity. See *Configurable options* on page 1-6 for more information.

Table A-10 ACP slave port error detection signals

Signal	Direction	Description
ACPSFATAL[1:0]	Output	Fatal error, per channel, {B,AW}
AWADDRPTYCS[3:0]	Input	Parity bits for AWADDRCS ^a
AWCTLPTYCS[3:0]	Input	Parity bits for the rest of the write address channel ^a
AWRPTYCS	Output	Parity bit for AWREADYCS
AWUSERPTYCS	Input	Parity bit for sideband signals ^a
AWVPTYCS	Input	Parity bit for AWVALIDCS
BCTLPTYCS[1:0]	Output	Parity for buffered response signal ^a
BRPTYCS	Output	Parity bit for BREADYCS
BVPTYCS	Output	Parity bit for BVALIDCS
BUSERPTYCS	Output	Parity bit for sideband signals ^a

a. This is an AXI extension signal.

A.5.7 ACP master port

Table A-11 shows the ACP master port signals.

Table A-11 ACP master port signals

Signal	Direction	Description
Write Address Channel		
AWIDCM[1:0]	Output	The identification tag for the write address group of signals.
AWADDRCM[31:0]	Output	Transfer start address
AWLENM[3:0]	Output	Write transfer burst length.
AWSIZECM[2:0]	Output	Indicates the size of the transfer.
AWBURSTCM[1:0]	Output	Write burst type.
AWLOCKCM[1:0]	Output	Lock signal.

Table A-11 ACP master port signals (continued)

Signal	Direction	Description
AWCACHECM[3:0]	Output	Provides decode information for outer attributes: b0000 = Strongly Ordered. b0001 = Device. b0011 = Normal, Non-cacheable. b0110 = Normal, Cacheable. write-through. b1111 = Normal, Cacheable. write-back, write allocation. b0111 = Normal, Cacheable. write-back, no write allocation. ——— Note ——— The AXI specification describes these encodings using the pre-ARMv6 terms such as <i>cacheable-bufferable</i> . These terms are equivalent to the ARMv6 memory-type descriptions such as <i>Normal</i> , <i>Non-cacheable</i> used here.
AWPROTCM[2:0]	Output	Protection type.
AWCOHERENTCM	Output	Require caches to be made coherent with this access. ^a
AWUSERCM[3:0]	Output	For transmission of other sideband information. ^a
AWVALIDCM	Output	Indicates address and control are valid.
AWREADYCM	Input	Address ready. The slave uses this signal to indicate it is ready to accept the address.
Write Response Channel		
BIDCM[1:0]	Input	The identification tag for the write response signal.
BRESPCM[1:0]	Input	Write response
BUSERCM[3:0]	Input	For transmission of other sideband information. ^a
BVALIDCM	Input	Indicates that a valid write response is available.
BREADYCM	Output	Indicates that the CPU is ready to accept write response.

a. This is an AXI extension signal.

A.5.8 ACP master port error detection signals

Table A-12 shows the ACP master port error detection signals. These signals are only generated if the processor is configured to include AXI bus parity. See *Configurable options* on page 1-6 for more information.

Table A-12 ACP master port error detection signals

Signal	Direction	Description
AWVPTYCM	Output	Parity bit for AWVALIDCM
AWRPTYCM	Input	Parity bit for AWREADYCM
AWADDRPTYCM[3:0]	Output	Parity bits for AWADDRCM ^a
AWCTLPTYCM[3:0]	Output	Parity bits for the rest of the write address channel ^a
AWUSERPTYCM	Output	Parity bit for sideband signals ^a
BVPTYCM	Input	Parity bit for BVALIDCM

Table A-12 ACP master port error detection signals (continued)

Signal	Direction	Description
BRPTYCM	Output	Parity bit for BREADYCM
BCTLPTYCM[1:0]	Input	Parity for buffered response signal ^a
BUSERPTYCM	Input	Parity bit for sideband signals ^a
ACPMFATAL[1:0]	Output	Fatal error, per channel, {B,AW}

a. This is an AXI extension signal.

A.5.9 AXI peripheral port

Table A-13 shows the AXI peripheral port signals.

Table A-13 AXI peripheral port signals

Signal	Direction	Description
ACLKENPm	Input	Clock enable for the AXI peripheral port.
Write Address Channel		
AWIDPm[3:0]	Output	The identification tag for the write address group of signals.
AWADDRPm[31:0]	Output	Transfer start address.
AWLENPm[3:0]	Output	Write transfer burst length.
AWSIZEPm[2:0]	Output	Indicates the size of the transfer.
AWBURSTPm[1:0]	Output	Write burst type.
AWLOCKPm[1:0]	Output	Lock signal.
AWCACHEPm[3:0]	Output	Provides decode information for outer attributes: b0000 = Strongly Ordered. b0001 = Device. b0011 = Normal, Non-cacheable. b0110 = Normal, Cacheable. write-through. b1111 = Normal, Cacheable. write-back, write allocation. b0111 = Normal, Cacheable. write-back, no write allocation.
Note The AXI specification describes these encodings using the pre-ARMv6 terms such as <i>cacheable-bufferable</i>. These terms are equivalent to the ARMv6 memory-type descriptions such as <i>Normal</i>, <i>Non-cacheable</i> used here.		
AWPROTPm[2:0]	Output	Protection type.
AWVALIDPm	Output	Indicates address and control are valid.
AWREADYPm	Input	Address ready. The slave uses this signal to indicate it is ready to accept the address.
Write Data Channel		
WIDPm[3:0]	Output	The identification tag for the write data group of signals.
WDATAPm[31:0]	Output	Write data.

Table A-13 AXI peripheral port signals (continued)

Signal	Direction	Description
WSTRB_{Pm}[3:0]	Output	Write strobes used to indicate which byte lanes must be updated.
WLAST_{Pm}	Output	Indicates the last data transfer of a burst.
WVALID_{Pm}	Output	Indicates address and control are valid.
WREADY_{Pm}	Input	Indicates that the slave is ready to accept write data.
Write Response Channel		
BID_{Pm}[3:0]	Input	The identification tag for the write response channel.
BRESPP_{Pm}[1:0]	Input	Write response.
BVALID_{Pm}	Input	Indicates that a valid write response is available.
BVPTYP_{Pm}	Input	Parity bit for BVALID_{Pm}
BREADY_{Pm}	Output	Indicates that the CPU is ready to accept a write response.
BRPTYP_{Pm}	Output	Parity bit for BREADY_{Pm}
BCTLPTYP_{Pm}[1:0]	Input	Parity for buffered response channel
Read Address Channel		
ARID_{Pm}[3:0]	Output	Identification tag for the read address group of signals.
ARADDR_{Pm}[31:0]	Output	Instruction fetch burst start address.
ARLEN_{Pm}[3:0]	Output	Instruction fetch burst length.
ARSIZE_{Pm}[2:0]	Output	Indicates the size of the transfer.
ARBURST_{Pm}[1:0]	Output	Burst type.
ARLOCK_{Pm}[1:0]	Output	Lock signal.
ARCACHE_{Pm}[3:0]	Output	Provides decode information for outer attributes: b0000 = Strongly Ordered. b0001 = Device. b0011 = Normal, Non-cacheable. b0110 = Normal, Cacheable. write-through. b1111 = Normal, Cacheable. write-back, write allocation. b0111 = Normal, Cacheable. write-back, no write allocation.
———— Note ———— The AXI specification describes these encodings using the pre-ARMv6 terms such as <i>cacheable-bufferable</i>. These terms are equivalent to the ARMv6 memory-type descriptions such as <i>Normal</i>, <i>Non-cacheable</i> used here.		
ARPROT_{Pm}[2:0]	Output	Protection type.
ARVALID_{Pm}	Output	Indicates address and control are valid.
ARREADY_{Pm}	Input	Address ready. The slave uses this signal to indicate it is ready to accept the address.
Read Data Channel		
RID_{Pm}[3:0]	Input	The identification tag for the read data group of signals.

Table A-13 AXI peripheral port signals (continued)

Signal	Direction	Description
RDATA_{Pm}[31:0]	Input	Read data.
RRESPP_m[1:0]	Input	Read response.
RLAST_{Pm}	Input	Indicates the last transfer in a read burst.
RVALID_{Pm}	Input	Indicates that read data is available.
RREADY_{Pm}	Output	Read ready signal indicating that the bus master can accept read data and response information.

A.5.10 AXI peripheral port error detection signals

Table A-14 shows the AXI peripheral port error detection signals. These signals are only generated if the processor is configured to include AXI bus parity. See *Configurable options* on page 1-6 for more information.

Table A-14 AXI peripheral port error detection signals

Signal	Direction	Description
ARADDRPTYP_m[3:0]	Output	Parity bits for ARADDR_{Pm} ^a
ARCTLPTYP_m[3:0]	Output	Parity bits for the rest of the read address channel ^a
ARRPTYP_m	Input	Parity bit for ARREADY_{Pm}
ARVPTYP_m	Output	Parity bit for ARVALID_{Pm}
AWADDRPTYP_m[3:0]	Output	Parity bits for AWADDR_{Pm} ^a
AWCTLPTYP_m[3:0]	Output	Parity bits for the rest of the write address channel ^a
AWRPTYP_m	Input	Parity bit for AWREADY_{Pm}
AWVPTYP_m	Output	Parity bit for AWVALID_{Pm}
PPXCORR_m	Output	Correctable error on RRESPP_m ^b
PPXFATAL_m[4:0]	Output	Fatal error, one bit for each channel {R,AR,B,W,AW}
RCTLPTYP_m[1:0]	Input	Parity bits for the rest of the read data channel ^a
RERRCODE_{Pm}[6:0]	Input	ECC code for RDATA_{Pm} ^a
RRPTYP_m	Output	Parity bit for RREADY_{Pm}
RVPTYP_m	Input	Parity bit for RVALID_{Pm}
WCTLPTYP_m[2:0]	Output	Parity bits for the rest of the write data channel ^a
WERRCODE_{Pm}[6:0]	Output	ECC code for WDATA_{Pm} ^a
WRPTYP_m	Input	Parity bit for WREADY_{Pm}
WVPTYP_m	Output	Parity bit for WVALID_{Pm}

a. This is an AXI extension signal.

b. Address is reported on **MERRADDR_m**, listed in Table A-6 on page A-11.

A.5.11 AHB peripheral port

Table A-15 shows the AHB peripheral port signals.

Table A-15 AHB peripheral port signals

Signal	Direction	Description
Address Phase		
HCLKENPm	Input	Synchronous enable for AHB transfers.
HADDRPm[31:0]	Output	System address bus
HBURSTPm[2:0]	Output	Burst type
HMASTLOCKPm	Output	Indicates that the current transfer is part of a locked sequence
HPROTPm[3:0]	Output	Protection type
HSIZEPm[2:0]	Output	Indicates the size of the transfer.
HTRANSPm[1:0]	Output	Transfer type
HWDATAPm[31:0]	Output	Write data
HWRITEPm	Output	Indicates the direction of the transfer
Data phase		
HRDATAPm[31:0]	Input	Read data
HREADYPm	Input	Indicates that the previous transfer is finished
HRESPPm	Input	Transfer response

A.5.12 AHB peripheral port error detection signals

Table A-16 shows the AHB peripheral port error detection signals. These signals are only generated if the processor is configured to include AHB bus parity. See *Configurable options* on page 1-6 for more information.

Table A-16 AHB peripheral port error detection signals

Signal	Direction	Description
HWERRCODEPm[6:0]	Output	ECC code for HWDATAPm .
HADDRPTYPm[3:0]	Output	Parity bit for HADDRPm .
HCTLPTYPm[1:0]	Output	Parity bits for the rest of the data channel.
HRERRCODEPm[6:0]	Input	ECC code for HRDATAPm .
HRESPPTYPm	Input	Parity bit for HREADYPm and HRESPPm ^a .
PPHFATALm	Output	Fatal error on: • parity computed for HREADYPm and HRESPPm • two bit error on HRDATAPm.
PPHCORRm	Output	Correctable error on HRDATAPm . Address is reported on MERRADDRm .

a. This is not parity for **HRESP** alone, even though that might be suggested by the name.

A.6 TCM interface signals

Table A-17 shows the ATCM port signals.

Table A-17 ATCM port signals

Signal	Direction	Description
ATCDATAINm[63:0]	Input	Data from ATCM
ATCPARITYINm[13:0]	Input	ECC code from ATCM
ATCERRORm	Input	Error detected by ATCM ^a
ATCWAITm	Input	Wait from ATCM
ATCLATEERRORm	Input	Late error from ATCM ^a
ATCRETRYm	Input	Access to ATCM must be retried ^a
ATCADDRPTYm	Output	Parity formed from ATCM address output ^b
ATCEN0m	Output	Enable for ATCM lower word, bit range [31:0]
ATCEN1m	Output	Enable for ATCM upper word, bit range [64:32]
ATCWEm	Output	Write enable for ATCM
ATCADDRm[22:3]	Output	Address for ATCM data RAM
ATCBYTEWRm[7:0]	Output	Byte strobes for direct write
ATCSEQm	Output	ATCM RAM access is sequential
ATCDATAOUTm[63:0]	Output	Write data for ATCM data RAM
ATCPARITYOUTm[13:0]	Output	Write ECC code for ATCM
ATCACCTYPEm[2:0]	Output	Determines access type: b001 = Load/Store b010 = Fetch b100 = DMA b100 = MBIST ^c .

a. This signal is ignored when bit [0] of the Auxiliary Control Register is set to 0, see *c1, Auxiliary Control Register* on page 4-41.

b. Only generated if the processor is configured to include TCM address bus parity.

c. The MBIST interface has no way of signaling a wait. If it is accessing the TCM, and the TCM signals a wait, the AXI slave pipeline stalls and the data arrives later. However, no signal is sent to the MBIST controller to indicate this.

Table A-18 shows the B0TCM port signals.

Table A-18 B0TCM port signals

Signal	Direction	Description
B0TCDATAINm[63:0]	Input	Data from B0TCM
B0TCPARITYINm[13:0]	Input	ECC code from B0TCM
B0TCERRORm	Input	Error detected by B0TCM ^a
B0TCWAITm	Input	Wait from B0TCM

Table A-18 B0TCM port signals (continued)

Signal	Direction	Description
B0TCLATEERRORm	Input	Late error from B0TCM ^a
B0TCRETRYm	Input	Access to B1TCM must be retried ^a
B0TCADDRPTYm	Output	Parity formed from B0TCM address output ^b
B0TCWE_m	Output	Write enable for B0TCM
B0TCEN0_m	Output	Enable for B0TCM lower word, bit range [31:0]
B0TCEN1_m	Output	Enable for B0TCM upper word, bit range [64:32]
B0TCADDR_m[22:3]	Output	Address for B0TCM data RAM
B0TCBYTEWR_m[7:0]	Output	Byte strobes for direct write
B0TCSEQ_m	Output	B0TCM RAM access is sequential
B0TCDATAOUT_m[63:0]	Output	Write data for B0TCM data RAM
B0TCPARITYOUT_m[13:0]	Output	Write ECC code for B0TCM
B0TCACCTYPE_m[2:0]	Output	Determines access type: b001 = Load/Store b010 = Fetch b100 = DMA b100 = MBIST ^c .

- a. This signal is ignored when bit [1] of the Auxiliary Control Register is set to 0, see *c1, Auxiliary Control Register* on page 4-41.
- b. Only generated if the processor is configured to include TCM address bus parity.
- c. The MBIST interface has no way of signaling a wait. If it is accessing the TCM, and the TCM signals a wait, the AXI slave pipeline stalls and the data arrives later. However, no signal is sent to the MBIST controller to indicate this.

Table A-19 shows the B1TCM port signals.

Table A-19 B1TCM port signals

Signal	Direction	Description
B1TCDATAIN_m[63:0]	Input	Data from B1TCM
B1TCPARITYIN_m[13:0]	Input	ECC code from B1TCM
B1TCERROR_m	Input	Error detected by B1TCM ^a
B1TCRETRY_m	Input	Access to B1TCM must be retried ^a
B1TCLATEERROR_m	Input	Late error from B1TCM ^a
B1TCWAIT_m	Input	Wait from B1TCM
B1TCADDRPTY_m	Output	Parity formed from B1TCM address output ^b
B1TCWE_m	Output	Write enable for B1TCM
B1TCEN0_m	Output	Enable for B1TCM lower word, bit range [31:0]
B1TCEN1_m	Output	Enable for B1TCM upper word, bit range [64:32]

Table A-19 B1TCM port signals (continued)

Signal	Direction	Description
BITCADDRm[22:3]	Output	Address for B1TCM data RAM
BITCBYTEWRm[7:0]	Output	Byte strobes for direct write
BITCSEQm	Output	B1TCM RAM access is sequential
BITCDATAOUTm[63:0]	Output	Write data for B1TCM data RAM
BITCPARITYOUTm[13:0]	Output	Write ECC code for B1TCM
BITCACCTYPEm[2:0]	Output	Determines access type: b001 = Load/Store b010 = Fetch b100 = DMA b100 = MBIST ^c .

- a. This signal is ignored when bit [2] of the Auxiliary Control Register is set to 0, see *c1, Auxiliary Control Register* on page 4-41.
- b. Only generated if the processor is configured to include TCM address bus parity.
- c. The MBIST interface has no way of signaling a wait. If it is accessing the TCM, and the TCM signals a wait, the AXI slave pipeline stalls and the data arrives later. However, no signal is sent to the MBIST controller to indicate this.

A.7 Redundant CPU signals

Table A-20 shows the redundant CPU signals. If you are implementing a redundant CPU configuration, contact ARM for more information about the functionality of these signals.

Table A-20 Redundant CPU signals

Signal	Direction
CLKIN1	Input
DCCMINP[7:0]	Input
DCCMINP2[7:0]	Input
DCCMOUT[7:0]	Output
DCCMOUT2[7:0]	Output

A.8 Debug interface signals

Table A-21 shows the debug interface signals. With the exception of **PCLKENDBG**, **DBGRESETmn**, and **PRESETDBGmn**, all these signals are only sampled or driven on **CLKIN** edges when **PCLKENDBG** is asserted.

Table A-21 Debug interface signals

Signal	Direction	Description
PCLKENDBG	Input	Clock enable for APB buses.
PSELDBGm	Input	Selects the external debug interface.
PADDRDBGm[11:2]	Input	Programming address.
PADDRDBG31m	Input	Programming address.
PRDATADBGm[31:0]	Output	Read data bus.
PWDATADBGm[31:0]	Input	Write data bus.
PENABLEDBGm	Input	Indicates second, and subsequent, cycle of a transfer.
PREADYDBGm	Output	Extends a APB transfer by the inserting wait states.
PSLVERRDBGm	Output	Slave-generated error response.
PWRITEDBGm	Input	Indicates access is a write transfer. Distinguishes between a read, LOW, and a write, HIGH.
PRESETDBGmn	Input	Reset debug domain debug logic. ^a
DBGRESETmn	Input	Reset core domain debug logic. ^a

a. Can be asserted asynchronously.

Table A-22 shows the debug miscellaneous signals.

Table A-22 Debug miscellaneous signals

Signal	Direction	Description
DBGENm	Input	Debug enable ^a
NIDENm	Input	Non-invasive debug enable ^a
EDBGRQm	Input	External debug request ^a
DBGACKm	Output	Debug acknowledge
DBGIRSTREQm	Output	Request for reset from debug logic
DBGTRIGGERm	Output	External debug request taken
COMMRXm	Output	DTRRX full
COMMTXm	Output	DTRTX empty
DBGRESTARTm	Input	External restart request ^a
DBGRESTARTEDm	Output	Handshake for DBGRESTARTm
DBGNOPWRDWN	Output	No power-down request

Table A-22 Debug miscellaneous signals (continued)

Signal	Direction	Description
DBGROMADDR[31:12]	Input	Debug ROM physical address
DBGROMADDRV	Input	Debug ROM physical address valid
DBGSELFADDR_m[31:12]	Input	Debug self-address offset
DBGSELFADDRV_m	Input	Debug self-address offset valid

a. Can be asserted asynchronously.

A.9 ETM interface signals

Table A-23 shows the ETM interface signals.

Table A-23 ETM interface signals

Signal	Direction	Description
ETMICTLm[13:0]	Output	ETM instruction control bus
ETMIAm[31:1]	Output	ETM instruction address
ETMDCTLm[11:0]	Output	ETM data control bus
ETMDAm[31:0]	Output	ETM data address
ETMDDm[63:0]	Output	ETM data-data
ETMCIDm[31:0]	Output	Current value of processor CID register
ETMWFIPENDINGm	Output	Core is attempting to enter standby state because of a WFI or WFE
EVNTBUSm[54:0]	Output	Performance monitor unit output
ETMPWRUPm	Input	Power up ETM interface
nETMWFIREADYm	Input	ETM FIFO is empty, CPU can enter WFI state
ETMEXTOUTm[1:0]	Input	ETM detected events

A.10 Test signals

Table A-24 shows the test signals.

Table A-24 Test signals

Signal	Direction	Description
SEm	Input	Scan Enable
RSTBYPASSm	Input	Bypass pipelined reset

A.11 MBIST signals

Table A-25 shows the MBIST signals.

Table A-25 MBIST signals

Signal	Direction	Description
MBTESTONm	Input	MBIST test is enabled
MBISTDINm[77:0]	Input	MBIST data in
MBISTADDRm[19:0]	Input	MBIST address
MBISTCEm	Input	MBIST chip enable
MBISTSEIm[4:0]	Input	MBIST chip select
MBISTWEIm[7:0]	Input	MBIST write enable
MBISTDOUm[77:0]	Output	MBIST data out

A.12 Validation signals

Table A-26 shows the validation signals.

Table A-26 Validation signals

Signal	Direction	Description
VALEDBGRQm	Output	Debug request
nVALIRQm	Output	Request for an interrupt
nVALFIQm	Output	Request for a Fast Interrupt
nVALRESETm	Output	Request for a reset

A.13 FPU signals

Table A-27 shows the FPU signals. These signals are only driven if the processor is configured to include the floating-point logic.

Table A-27 FPU signals

Signal	Direction	Description
FPIXCm	Output	Masked floating-point inexact exception
FPOFCm	Output	Masked floating-point overflow exception
FPUFCm	Output	Masked floating-point underflow exception
FPIOCm	Output	Masked floating-point invalid operation exception
FPDZCm	Output	Masked floating-point divide-by-zero exception
FPIDCm	Output	Masked floating-point input denormal exception

A.14 Split/Lock

Table A-28 shows the Split/Lock signals. If you are implementing a Split/Lock configuration, contact ARM for more information about the functionality of these signals.

Table A-28 Split/Lock signals

Signal	Direction
SLSPLIT	Input
SLRESETn	Input
SLCLAMP	Input
SLERRACPn	Input
SLERRDBGn	Input

A.15 Power modes

Table A-29 shows the Power mode signal.

Table A-29 Power mode signal

Signal	Direction	Description
RAMCONTROL _m [7:0]	-	Wires only – connected to cortexr5_caches_ramsm module for use controlling physical RAM features of CPU m , where m is 0 or 1, such as retention states.

Appendix B

Cycle Timings and Interlock Behavior

This appendix describes the cycle timings and interlock behavior of instructions on the processor. It contains the following sections:

- *About cycle timings and interlock behavior* on page B-3
- *Register interlock examples* on page B-6
- *Data processing instructions* on page B-7
- *QADD, QDADD, QSUB, and QDSUB instructions* on page B-9
- *Media data-processing* on page B-10
- *Sum of Absolute Differences (SAD)* on page B-11
- *Multiplies* on page B-12
- *Divide* on page B-14
- *Branches* on page B-15
- *Processor state updating instructions* on page B-16
- *Single load and store instructions* on page B-17
- *Load and Store Double instructions* on page B-19
- *Load and Store Multiple instructions* on page B-20
- *RFE and SRS instructions* on page B-23
- *Synchronization instructions* on page B-24
- *Coprocessor instructions* on page B-25
- *SVC, BKPT, Undefined, and Prefetch Aborted instructions* on page B-26
- *Miscellaneous instructions* on page B-27
- *Floating-point register transfer instructions* on page B-28
- *Floating-point load/store instructions* on page B-29
- *Floating-point single-precision data processing instructions* on page B-31

- *Floating-point double-precision data processing instructions* on page B-32
- *Dual issue* on page B-33.

B.1 About cycle timings and interlock behavior

Complex instruction dependencies and memory system interactions make it impossible to describe briefly the exact cycle timing behavior for all instructions in all circumstances. The timings described in this chapter are accurate in most cases. If precise timings are required, you must use a cycle-accurate model of the processor.

Unless stated otherwise, cycle counts and result latencies that this chapter describes are best-case numbers. They assume:

- no outstanding data dependencies between the current instruction and a previous instruction
- the instruction does not encounter any resource conflicts
- all data accesses hit in the data cache, and do not cross protection region boundaries
- all instruction accesses hit in the instruction cache.

This section describes:

- *Instruction execution overview*
- *Conditional instructions* on page B-4
- *Flag-setting instructions* on page B-4
- *Definition of terms* on page B-4.
- *Assembler language syntax* on page B-5.

B.1.1 Instruction execution overview

The instruction execution pipeline has four stages, Iss, Ex1, Ex2, and Wr.

Extensive forwarding to the end of the Iss, Ex1, and Ex2 stages enables many dependent instruction sequences to run without pipeline stalls. General forwarding occurs from the end of the Ex2 and Wr pipeline stages. In addition, the multiplier contains an internal multiply accumulate forwarding path. The address generation unit also contains an internal forwarding path.

Many instructions do not require data from a register until the Ex2 stage. All result latencies are given as the number of cycles until the register is available for a following instruction in the Ex2 stage. Most ALU operations require their source registers at the start of the Ex2 stage, and have a result latency of one. For example, the following sequence takes two cycles:

```
ADD R1,R3,R4          ;Result latency one
ADD R5,R2,R1          ;Register R1 required by ALU
```

The PC is the only register that result latency does not affect. An instruction that alters the PC never causes a pipeline stall because of interlocking with a subsequent instruction that reads the PC.

Most loads have a result latency of two or higher, because they do not forward their results until the Wr stage. For example, the following sequence takes three cycles:

```
LDR R1, [R2]          ;Result latency two
ADD R3, R3, R1         ;Register R1 required by ALU
```

If a subsequent instruction requires the register at the end of the Iss stage then an extra cycle must be added to the result latency of the instruction producing the required register.

Instructions that require a register at the end of these stages are specified by describing that register as an Early Reg. The following sequence, requiring an Early Reg, takes four cycles:

```
LDR R1, [R2]          ;Result latency two
```


ADD R3, R3, R1 LSL#6 ;plus one because Register R1 is Early

The following sequence where R1 is a Late Reg takes two cycles:

LDR R1, [R2] ;Result latency two minus one cycles
STR R1, [R3] ;no penalty because R1 is a Late register

The following sequence where R1 is a Very Early Reg takes four cycles:

ADD R3, R1, R2 ;Result latency one plus two cycles
LDR R4, [R3] ;plus two because register R3 is Very Early

B.1.2 Conditional instructions

Most instructions do not take more or fewer cycles to execute if they fail their condition codes. The exceptions to this are:

- instructions that alter the PC, such as branches
- integer divide instructions, that require only one execute cycle.

The result latency of most instructions that fail their condition codes is one. The exceptions to this are:

- all load and store instructions, that have their result latency unaffected
- integer divide instructions, that have a result latency of three.

B.1.3 Flag-setting instructions

Most instructions do not take more or fewer cycles to execute if they are flag-setting. The exceptions to this are certain multiply instructions.

B.1.4 Definition of terms

Table B-1 gives descriptions of cycle timing terms used in this appendix.

Table B-1 Definition of cycle timing terms

Term	Description
Memory Cycles	This is the number of cycles during which an instruction sends a memory access to the cache.
Cycles	This is the minimum number of cycles required to issue an instruction. Issue cycles that produce memory accesses to the cache are included, so Cycles is always greater than or equal to Memory Cycles.
Result Latency	This is the number of cycles before the result of this instruction is available to a Normal Reg of the following instruction. When the Result Latency of an instruction is greater than Cycles and the following instruction requires the result, the following instruction stalls for a number of cycles equal to Result Latency minus Cycles. If this value is negative, there are zero stall cycles. Note The Result Latency is counted from the first cycle of an instruction.
Normal Reg	The specified registers are required at the start of the Ex2 stage.
Late Reg	The specified registers are not required until the start of the Wr stage. Subtract one cycle from the Result Latency of the instruction producing this register.

Table B-1 Definition of cycle timing terms (continued)

Term	Description
Early Reg	The specified registers are required at the start of the Ex1 stage. Add one cycle to the Result Latency of the instruction producing this register.
Very Early Reg	The specified registers are required at the start of the Iss stage. Add two cycles to the Result Latency of the instruction producing this register, or one cycle if the instruction producing this register is an LDM, LDR, LDRD, LDREX, or LDRT. The lower Result Latency does not apply if this register is the base register of the load instruction producing this register, or if the load instruction is an LDRB, LDRBT, LDRH, LDRSB, or LDRSH.
Interlock	There is a data dependency between two instructions in the pipeline, resulting in the Iss stage being stalled until the processor resolves the dependency.

B.1.5 Assembler language syntax

The syntax used throughout this chapter is unified assembler and the timings apply to ARM and Thumb instructions.

B.2 Register interlock examples

Table B-2 shows register interlock examples using LDR and ADD instructions.

LDR instructions take one cycle, have a result latency of two, and require their base register as a Very Early Reg.

ADD instructions take one cycle and have a result latency of one.

Table B-2 Register interlock examples

Instruction sequence	Behavior
LDR R1, [R2] ADD R6, R5, R4	Takes two cycles because there are no register dependencies.
ADD R1, R2, R3 ADD R9, R6, R1	Takes two cycles because ADD instructions have a result latency of one.
LDR R1, [R2] ADD R6, R5, R1	Takes three cycles because of the result latency of R1.
ADD R2, R5, R6 LDR R1, [R2]	Takes four cycles because of the use of the result of R2 as a Very Early Reg.
LDR R1, [R2] LDR R5, [R1]	Takes four cycles because of the result latency of R1, the use of the result of R1 as a Very Early Reg, and the use of an LDR to generate R1.

B.3 Data processing instructions

This section describes the cycle timing behavior for the ADC, ADD, ADDW, AND, ASR, BIC, CLZ, CMN, CMP, EOR, LSL, LSR, MOV, MOVT, MOVW, MVN, ORN, ORR, ROR, RRX, RSB, RSC, SBC, SUB, SUBW, TEQ, and TST instructions.

This section describes:

- *Cycle counts if destination is not PC*
- *Cycle counts if destination is the PC*
- *Example interlocks on page B-8*

B.3.1 Cycle counts if destination is not PC

Table B-3 shows the cycle timing behavior for data processing instructions if their destination is not the PC. You can substitute ADD with any of the data processing instructions identified in the opening paragraph of this section.

Table B-3 Data Processing Instruction cycle timing behavior if destination is not PC

Example instruction	Cycles	Early Reg	Late Reg	Result latency	Comments
ADD <Rd>, <Rn>, #<immed>	1	-	-	1	Normal cases.
ADD <Rd>, <Rn>, <Rm>	1	-	-	1	
ADD <Rd>, <Rn>, <Rm>, LSL #<immed>	1	<Rm>	-	1	Requires a shifted source register.
ADD <Rd>, <Rn>, <Rm>, LSL <Rs>	1	<Rm>, <Rs>	-	1	Requires a register controlled shifted source register.
MOV <Rd>, <Rm>	1	-	<Rm>	1	Simple MOV case. Must not set the flags or require a shifted source register.

B.3.2 Cycle counts if destination is the PC

Table B-4 shows the cycle timing behavior for data processing instructions if their destination is the PC. You can substitute ADD with any data processing instruction except for a CLZ. A CLZ with the PC as the destination is an Unpredictable instruction.

For condition code failing cycle counts, the cycles for the non-PC destination variants must be used.

Table B-4 Data Processing instruction cycle timing behavior if destination is the PC

Example instruction	Cycles	Early Reg	Late Reg	Result latency	Comments
ADD pc, <Rn>, #<immed>	9	-	-	-	Normal cases to PC
ADD pc, <Rn>, <Rm>	9	-	-	-	
ADD pc, <Rn>, <Rm>, LSL #<immed>	9	<Rm>	-	-	Requires a shifted source register
ADD pc, <Rn>, <Rm>, LSL <Rs>	9	<Rm>, <Rs>	-	-	Requires a register controlled shifted source register

B.3.3 Example interlocks

Most data processing instructions are single-cycle and can be executed back-to-back without interlock cycles, even if there are data dependencies between them. The exceptions to this are when shifts are used.

Shifter

The registers that the shifter requires are Early Regs and require an additional cycle of result availability before use. For example, the following sequence introduces a 1-cycle interlock, and takes three cycles to execute:

```
ADD R1,R2,R3
ADD R4,R5,R1 LSL #1
```

The second source register, that is not shifted, does not incur an extra data dependency check. Therefore, the following sequence takes two cycles to execute:

```
ADD R1,R2,R3
ADD R4,R1,R9 LSL #1
```

Register controlled shifts

The register containing the shift distance is an Early Reg. For example, the following sequence takes three cycles to execute:

```
ADD R1, R2, R3
ADD R4, R2, R4, LSL R1
```

B.4 QADD, QDADD, QSUB, and QDSUB instructions

This section describes the cycle timing behavior for the QADD, QDADD, QSUB, and QDSUB instructions.

These instructions perform saturating arithmetic. They have a result latency of two. The QDADD and QDSUB instructions must double and saturate the register <Rn> before the addition. This register is an Early Reg.

Table B-5 shows the cycle timing behavior for QADD, QDADD, QSUB, and QDSUB instructions.

Table B-5 QADD, QDADD, QSUB, and QDSUB instruction cycle timing behavior

Instructions	Cycles	Early Reg	Result latency
QADD, QSUB	1	-	2
QDADD, QDSUB	1	<Rn>	2

B.5 Media data-processing

Table B-6 shows media data-processing instructions and gives their cycle timing behavior.

All media data-processing instructions are single-cycle issue instructions. These instructions have result latencies of one or two cycles. Some of the instructions require an input register to be shifted, or manipulated in some other way before use and therefore are marked as requiring an Early Reg.

Table B-6 Media data-processing instructions cycle timing behavior

Instructions	Cycles	Early Reg	Result latency
SADD16, SSUB16, SADD8, SSUB8	1	-	1
UADD16, USUB16, UADD8, USUB8	1	-	1
SEL	1	-	1
QADD16, QSUB16, QADD8, QSUB8	1	-	2
SHADD16, SHSUB16, SHADD8, SHSUB8	1	-	1
UQADD16, UQSUB16, UQADD8, UQSUB8	1	-	2
UHADD16, UHSUB16, UHADD8, UHSUB8	1	-	1
SSAT16, USAT16	1	<Rn>	1
SASX, SSAX	1	-	1
UASX, USAX	1	-	1
SXTAB, SXTAB16, SXTAH	1	<Rm>	1
SXTB, SXTB16, SXTBH	1	<Rm> ^a	1
UXTB, UXTB16, UXTBH	1	<Rm> ^a	1
UXTAB, UXTAB16, UXTAH	1	<Rm>	1
REV, REV16, REVSH, RBIT	1	<Rm>	1
PKHBT, PKHTB	1	<Rm>	1
SSAT, USAT	1	<Rm>	1
QASX, QSAX	1	-	2
SHASX, SHSAX	1	-	1
UQASX, UQSAX	1	-	2
UHASX, UHSAX	1	-	1
BFC	1	<Rd>	1
SBFX, UBFX	1	<Rn>	1
BFI	1	<Rd>, <Rn>	1

a. A shift of zero makes <Rm> a Normal Reg for these instructions.

B.6 Sum of Absolute Differences (SAD)

Table B-7 shows SAD instructions and gives their cycle timing behavior.

Table B-7 Sum of absolute differences instruction timing behavior

Instructions	Cycles	Early Reg	Result latency
USAD8	1	<Rn>, <Rm>	2 ^a
USADA8	1	<Rn>, <Rm>	2 ^a

a. Result latency is one fewer if the destination is the accumulate for a subsequent USADA8.

B.6.1 Example interlocks

Table B-8 shows interlock examples using USAD8 and USADA8 instructions.

Table B-8 Example interlocks

Instruction sequence	Behavior
USAD8 R1,R2,R3 ADD R5,R6,R1	Takes three cycles because USAD8 has a Result Latency of two, and the ADD requires the result of the USAD8 instruction.
USAD8 R1,R2,R3 MOV R9,R9 ADD R5,R6,R1	Takes three cycles. The MOV instruction is scheduled during the Result Latency of the USAD8 instruction.
USAD8 R1,R2,R3 USADA8 R1,R4,R5,R1	Takes two cycles. The Result Latency is one less because the result is used as the accumulate for a subsequent USADA8 instruction.

B.7 Multiplies

Most multiply operations cannot forward their result early, except as the accumulate value for a subsequent multiply. For a subsequent multiply accumulate the result is available one cycle earlier than for all other uses of the result.

Certain multiplies require:

- more than one cycle to execute
- more than one pipeline issue to produce a result.

The multiplicand and multiplier are required as Early Regs because they are both required at the end of the Iss stage.

Flag-setting multiplies followed by a conditional instruction interlock the conditional instruction for one cycle, or two cycles if the instruction is a conditional multiply. Flag-setting multiplies followed by a flag-setting instruction interlock the flag-setting instruction for one cycle, unless the instruction is a flag-setting multiply in which case there is no interlock.

Table B-9 shows the cycle timing behavior of example multiply instructions.

Table B-9 Example multiply instruction cycle timing behavior

Example instruction	Cycles	Early Reg	Late Reg	Result latency
MUL(S)	2	<Rn>, <Rm>	-	3
MLA(S), MLS	2	<Rn>, <Rm>	<Ra>	3
SMULL(S)	2	<Rn>, <Rm>	-	3, 3
UMULL(S)	2	<Rn>, <Rm>	-	3, 3
SMLAL(S)	2	<Rn>, <Rm>	<RdLo>, <RdHi>	3, 3
UMLAL(S)	2	<Rn>, <Rm>	<RdLo>, <RdHi>	3, 3
SMULxy	1	<Rn>, <Rm>	-	2
SMLAxy	1	<Rn>, <Rm>	-	2
SMULWy	1	<Rn>, <Rm>	-	2
SMLAWy	1	<Rn>, <Rm>	-	2
SMLALxy	2	<Rn>, <Rm>	<RdLo>, <RdHi>	3, 3
SMUAD, SMUADX	1	<Rn>, <Rm>	-	2
SMLAD, SMLADX	1	<Rn>, <Rm>	-	2
SMUSD, SMUSDx	1	<Rn>, <Rm>	-	2
SMLSD, SMLSDx	1	<Rn>, <Rm>	-	2
SMMUL, SMMULR	2	<Rn>, <Rm>	-	3
SMMLA, SMMLAR	2	<Rn>, <Rm>	<Ra>	3
SMMLS, SMMLSR	2	<Rn>, <Rm>	<Ra>	3

Table B-9 Example multiply instruction cycle timing behavior (continued)

Example instruction	Cycles	Early Reg	Late Reg	Result latency
SMLALD, SMLALDX	1	<Rn>, <Rm>	-	2, 2
SMLSLD, SMLSLDX	1	<Rn>, <Rm>	-	2, 2
UMAAL	2	<Rn>, <Rm>	<RdLo>, <RdHi>	3, 3

Note

Result Latency is one less if the result is used as the accumulate value for a subsequent multiply accumulate. This only applies if the result is the same width as the accumulate value, that is 32 or 64 bits.

B.8 Divide

This section describes the cycle timing behavior of the UDIV and SDIV instructions.

The divider unit is separate from the main execute pipeline so the UDIV and SDIV instructions require one cycle to issue. They execute out-of-order relative to the rest of the pipeline, and require an additional issue cycle at the end of the divide operation to write the result to the destination register. This additional cycle is not required if the divide instruction fails its condition code.

Result Latency for a UDIV instruction A divided by B is given by:

$$\text{Result latency} = 3 + \max\left(\left(\frac{\text{clz}(B) - \text{clz}(A) + 1}{2}\right), 0\right)$$

Result Latency for a SDIV instruction A divided by B is given by:

$$\text{Result latency} = 4 + \max\left(\left(\frac{\text{clz}(B) - \text{clz}(A) + 1}{2}\right), 0\right)$$

Note

- A divide instruction that fails its condition code or attempts to divide by zero has a Result Latency of three.
 - The value of the $(\text{clz}(B) - \text{clz}(A) + 1)/2$ component of these equations must be rounded down.
 - The $\text{clz}(x)$ function counts the number of leading zeros in the 32-bit value x . If x is negative, it is negated before this count occurs.
-

B.9 Branches

This section describes the cycle timing behavior for the B, BL, BLX, BX, BXJ, CBNZ, CBZ, TBB, and TBH instructions. Branches are subject to dynamic and return stack predictions. Table B-10 shows example branch instructions and their cycle timing behavior.

Table B-10 Branch instruction cycle timing behavior

Example instruction	Cycles	Memory cycles	Comments
B<label>, BL<label> ^a , BLX<label> ^a	1	-	Correct dynamic prediction
	8	-	Incorrect dynamic prediction
BX <Rm> ^b	1	-	Correct return stack prediction
	9	-	Incorrect return stack prediction
BX <cond> <Rm> ^b	1	-	Correct condition prediction and correct return stack prediction
	8	-	Incorrect condition prediction
	9	-	Correct condition prediction and incorrect return stack prediction
BXJ <cond> <Rm>	1	-	Condition code fails
	9	-	Condition code passes
BLX <Rm>	9	-	-
BLX <cond> <Rm>	1	-	Condition code fails
	9	-	Condition code passes
CBZ <Rn>, <label>, CBNZ <Rn>, <label>	1	-	Correct condition prediction
	8	-	Incorrectly predicted
TBB [<Rn>, <Rm>] ^c	9	1	Condition code fails
	9	1	Condition code passes
TBH [<Rn>, <Rm>, LSL#1] ^c	9	1	Condition code fails
	9	1	Condition code passes

a. Return stack push.

b. Return stack pop, if condition passes.

c. <Rn> and <Rm> are Very Early Regs.

B.10 Processor state updating instructions

This section describes the cycle timing behavior for the MSR, MRS, CPS, and SETEND instructions. Table B-11 shows processor state updating instructions and their cycle timing behavior.

Table B-11 Processor state updating instructions cycle timing behavior

Instruction	Cycles	Comments
MRS	1	All MRS instructions
MSR SPSR	1	All MSR instructions to the SPSR
MSR	5	All other MSR instructions to the CPSR
CPS<effect> <iflags>	1	Interrupt masks only
CPS<effect> <iflags>, #<mode>	1	Mode changing
SETEND	1	-

B.11 Single load and store instructions

This section describes the cycle timing behavior for LDR, LDRHT, LDRSBT, LDRSHT, LDRT, LDRB, LDRBT, LDRSB, LDRH, LDRSH, STR, STRT, STRB, STRBT, STRH, and PLD instructions.

Table B-12 shows the cycle timing behavior for stores and loads, other than loads to the PC. You can replace LDR with any of these single load or store instructions. The following rules apply:

- They are normally single-cycle issue. Both the base and any offset register are Very Early Regs.
- They are 3-cycle issue if pre-increment addressing with either a negative register offset or a shift other than LSL #1, 2 or 3 is used. Both the base and any offset register are Very Early Regs.
- Accesses to addresses not aligned to the access size that cross a 64-bit aligned boundary generate two memory accesses, and require an additional cycle to issue. This extra cycle is required if the final address is potentially unaligned, even if the final address turns out to be aligned.
- PLD (data preload hint instructions) have cycle timing behavior as for load instructions. Because they have no destination register, the result latency is not-applicable for such instructions.
- For store instructions <Rt> is always a Late Reg.

Table B-12 Cycle timing behavior for stores and loads, other than loads to the PC

Example instruction	Cycles	Memory cycles	Result latency (LDR)	Result latency (base register)	Comments
LDR <Rt>, <addr_md_1cycle> ^a	1	1	2	1	Aligned access
LDR <Rt>, <addr_md_3cycle> ^a	3	1	4	3	Aligned access
LDR <Rt>, <addr_md_1cycle> ^a	2	2	3	2	Potentially unaligned access
LDR <Rt>, <addr_md_3cycle> ^a	4	2	5	4	Potentially unaligned access

a. See Table B-14 on page B-18 for an explanation of <addr_md_1cycle> and <addr_md_3cycle>.

Table B-13 shows the cycle timing behavior for loads to the PC.

Table B-13 Cycle timing behavior for loads to the PC

Example instruction	Cycles	Memory cycles	Result latency	Comments
LDR pc, [sp, #<imm>] (!)	1	1	-	Correctly return stack predicted, or conditional predicted correctly
LDR pc, [sp], #<imm>	1	1	-	
LDR pc, [sp, #<imm>] (!)	9	1	-	Return stack mispredicted, conditional predicted correctly
LDR pc, [sp], #<imm>	9	1	-	
LDR <cond> pc, [sp, #<imm>] (!)	8	1	-	Conditional predicted incorrectly, but return stack predicted correctly
LDR <cond> pc, [sp], #cns	8	1	-	
LDR pc, <addr_md_1cycle> ^a	9	1	-	-
LDR pc, <addr_md_3cycle> ^a	11	1	-	-

- a. See Table B-14 for an explanation of <addr_md_1cycle> and <addr_md_3cycle>. For condition code failing cycle counts, you must use the cycles for the non-PC destination variants.

Only cycle times for aligned accesses are given because Unaligned accesses to the PC are not supported.

The processor includes a 4-entry return stack that can predict procedure returns. Any LDR instruction to the PC with an immediate post-indexed offset of plus four, and the stack pointer R13 as the base register is considered a procedure return.

Table B-14 shows the explanation of <addr_md_1cycle> and <addr_md_3cycle> used in Table B-12 on page B-17 and Table B-13 on page B-17.

Table B-14 <addr_md_1cycle> and <addr_md_3cycle> LDR example instruction explanation

Example instruction	Very Early Reg	Comments
<addr_md_1cycle>		
LDR <Rt>, [<Rn>, #<imm>] (!)	<Rn>	If post-increment addressing or pre-increment addressing with an immediate offset, or a positive register offset with no shift or shift LSL #1, 2 or 3, then 1-issue cycle
LDR <Rt>, [<Rn>, <Rm>] (!)	<Rn>, <Rm>	
LDR <Rt>, [<Rn>, <Rm>, LSL #1, 2 or 3] (!)	<Rn>, <Rm>	
LDR <Rt>, [<Rn>], #<imm>	<Rn>	
LDR <Rt>, [<Rn>], +/-<Rm>	<Rn>, <Rm>	
LDR <Rt>, [<Rn>], +/-<Rm> <shift> <cns>	<Rn>, <Rm>	
<addr_md_3cycle>		
LDR <Rt>, [<Rn>, -<Rm>] (!)	<Rn>, <Rm>	If pre-increment addressing with a negative register offset or shift other than LSL #1, 2 or 3, then 3-issue cycles
LDR <Rt>, <Rn>, +/-<Rm> <shift> <cns>] (!)	<Rn>, <Rm>	

B.11.1 Base register update

The base register update for load or store instructions occurs in the ALU pipeline. To prevent an interlock for back-to-back load or store instructions reusing the same base register, there is a local forwarding path to recycle the updated base register around the address generator. This only applies when the load or store instruction with base write-back uses pre-increment addressing, and is a single load or store instruction that is not a load or store double instruction or load or store multiple instruction.

For example, with R2 aligned the following instruction sequence take three cycles to execute:

```
LDR R5, [R2, #4]!
LDR R6, [R2, #0x10]!
LDR R7, [R2, #0x20]!
```

B.12 Load and Store Double instructions

This section describes the cycle timing behavior for the LDRD and STRD instructions.

The LDRD and STRD instructions:

- Are normally single-cycle issue. Both the base and any offset register are Very Early Regs.
- Are 3-cycle issue if offset or pre-increment addressing with a negative register offset is used. Both the base and any offset register are Very Early Regs.
- Take only one memory cycle if the address is doubleword aligned.
- Take two memory cycles if the address is not doubleword aligned.

Table B-15 shows the cycle timing behavior for LDRD and STRD instructions.

Table B-15 Load and Store Double instructions cycle timing behavior

Example instruction	Cycles	Cycles with base writeback	Memory cycles	Result latency (LDRD)	Result latency (base register)
Address is doubleword aligned					
LDRD R0, R1, <addr_md_1cycle> ^a	1	2	1	2, 2	2
LDRD R0, R1, <addr_md_3cycle> ^a	3	4	1	4, 4	4
Address not doubleword aligned					
LDRD R0, R1, <addr_md_1cycle> ^a	2	2	2	2, 3	2
LDRD R0, R1, <addr_md_3cycle> ^a	4	4	2	4, 5	4

a. See Table B-16 for an explanation of <addr_md_1cycle> and <addr_md_3cycle>.

Table B-16 shows the explanation of <addr_md_1cycle> and <addr_md_3cycle> used in Table B-15.

Table B-16 <addr_md_1cycle> and <addr_md_3cycle> LDRD example instruction explanation

Example instruction	Very Early Reg	Comments
<addr_md_1cycle>		
LDRD <Rt>, <Rt2>, [<Rn>, #<imm>] (!)	<Rn>	If post-increment addressing, pre-increment addressing with an immediate offset or a positive register offset, then 1-issue cycle
LDRD <Rt>, <Rt2>, [<Rn>, <Rm>] (!)	<Rn>, <Rm>	
LDRD <Rt>, <Rt2>, [<Rn>], #<imm>	<Rn>	
LDRD <Rt>, <Rt2>, [<Rn>], +/-<Rm>	<Rn>, <Rm>	
<addr_md_3cycle>		
LDRD <Rt>, <Rt2>, [<Rn>, -<Rm>] (!)	<Rn>, <Rm>	If pre-increment addressing with a negative register offset, then 3-issue cycles

B.13 Load and Store Multiple instructions

This section describes the cycle timing behavior for the LDM, STM, PUSH, and POP instructions. These instructions take multiple cycles to issue, and then use multiple memory cycles to load and store all the registers. Because the memory datapath is 64-bits wide, two registers can be loaded or stored on each cycle.

This section describes:

- *Load and Store Multiples, other than load multiples including the PC*
- *Load Multiples, where the PC is in the register list on page B-21*
- *Example Interlocks on page B-21*

B.13.1 Load and Store Multiples, other than load multiples including the PC

In all cases the base register, <Rn>, is a Very Early Reg.

Table B-17 shows the cycle timing behavior of load and store multiples including the PC.

Table B-17 Cycle timing behavior of Load and Store Multiples, other than load multiples including the PC

Example instruction	Cycles	Cycles with base register write-back	Memory cycles	Result latency (LDM)	Result latency (base register)
First address 64-bit aligned					
LDMIA <Rn>, {R1}	1	1	1	2	1
LDMIA <Rn>, {R1,R2}	1	2	1	2,2	2
LDMIA <Rn>, {R1,R2,R3}	2	2	2	2,2,3	2
LDMIA <Rn>, {R1,R2,R3,R4}	2	3	2	2,2,3,3	3
LDMIA <Rn>, {R1,R2,R3,R4,R5}	3	3	3	2,2,3,3,4	3
LDMIA <Rn>, {R1,R2,R3,R4,R5,R6}	3	4	3	2,2,3,3,4,4	4
LDMIA <Rn>, {R1,R2,R3,R4,R5,R6,R7}	4	4	4	2,2,3,3,4,4,5	4
First address not 64-bit aligned					
LDMIA <Rn>, {R1}	1	2	1	2	2
LDMIA <Rn>, {R1,R2}	2	2	2	2,3	2
LDMIA <Rn>, {R1,R2,R3}	2	3	2	2,3,3	3
LDMIA <Rn>, {R1,R2,R3,R4}	3	3	3	2,3,3,4	3
LDMIA <Rn>, {R1,R2,R3,R4,R5}	3	4	3	2,3,3,4,4	4
LDMIA <Rn>, {R1,R2,R3,R4,R5,R6}	4	4	4	2,3,3,4,4,5	4
LDMIA <Rn>, {R1,R2,R3,R4,R5,R6,R7}	4	5	4	2,3,3,4,4,5,5	5

Note

The Cycle timing behavior that Table B-17 shows also covers PUSH and POP instructions that behave like store and load multiple instructions with base register write-back.

B.13.2 Load Multiples, where the PC is in the register list

The processor includes a 4-entry return stack that can predict procedure returns. Any LDM to the PC that does not restore the SPSR to the CPSR, is predicted as a procedure return.

In all cases the base register, <Rn>, is a Very Early Reg.

Table B-18 shows the cycle timing behavior of Load Multiples, where the PC is in the register list.

Table B-18 Cycle timing behavior of Load Multiples, with PC in the register list (64-bit aligned)

Example instruction	Cycles	Memory cycles	Result latency	Comments
LDMIA <Rn>, {..., pc}	m ^a	n ^b	2,...	Correct return stack prediction
LDMIA <Rn>, {..., pc}	m ^a + 8	n ^b	2,...	Incorrect return stack prediction
LDMIA <cond> <Rn>, {..., pc}	m ^a	n ^b	2,...	Correct condition prediction and correct return stack prediction
LDMIA <cond> <Rn>, {..., pc}	m ^a + 7	n ^b	2,...	Incorrect condition prediction
LDMIA <cond> <Rn>, {..., pc}	m ^a + 8	n ^b	2,...	Correct condition prediction and incorrect return stack prediction

a. Where m is the number of cycles for this instruction if the PC were treated as a normal register.

b. Where n is the number of memory cycles for this instruction if the PC were treated as a normal register.

———— Note ————

The Cycle timing behavior that Table B-18 shows also covers PUSH and POP instructions that behave like store and load multiple instructions with base register writeback.

B.13.3 Example Interlocks

The following sequence that has an LDM instruction takes six cycles to execute, because R7 has a result latency of five cycles:

```
LDMIA R0, {R1-R7}
ADD R10, R10, R7
```

The following sequence that has an STM instruction takes five cycles to execute:

```
STMIA R0, {R1-R7}
ADD R7, R10, R11
```

The following sequence has a result latency hidden by issue cycles. It takes five cycles to execute.

```
LDMIA R0, {R1-R7}
ADD R10, R10, R3
```

The following sequence that has a POP instruction takes seven cycles to execute, because R9 has a result latency of six cycles:

```
POP {R1-R9}
ADD R10, R10, R9
```

The following sequence that has a PUSH instruction takes five cycles to execute:

```
PUSH {R1-R7}
ADD R10, R10, R7
```

Note

In the examples, *R0* and *sp* are 64-bit aligned addresses. The instructions *PUSH* and *POP* always use the *sp* register for the base address.

B.14 RFE and SRS instructions

This section describes the cycle timing for the RFE and SRS instructions.

These instructions:

- return from an exception and save exception return state respectively
- take one or two memory cycles depending on doubleword alignment first address location.

In all cases the base register is a Very Early Reg.

Table B-19 shows the cycle timing behavior for RFE and SRS instructions.

Table B-19 RFE and SRS instructions cycle timing behavior

Example instruction	Cycles	Memory cycles
Address doubleword aligned		
RFEIA <Rn>	10	1
SRSIA #<mode>	1	1
Address not doubleword aligned		
RFEIA <Rn>	11	2
SRSIA #<mode>	2	2

B.15 Synchronization instructions

This section describes the cycle timing behavior for the CLREX, DMB, DSB, ISB, LDREX, LDREXB, LDREXD, LDREXH, STREX, STREXB, STREXD, STREXH, SWP, and SWPB instructions

In all cases the base register, Rn, is a Very Early Reg. Table B-20 shows the synchronization instructions cycle timing behavior.

Table B-20 Synchronization instructions cycle timing behavior

Instruction	Cycles	Memory cycles	Result latency
CLREX	1	-	-
LDREX <Rt>, <Rn>	1	1	2
LDREXB <Rt>, <Rn>	1	1	2
LDREXH <Rt>, <Rn>	1	1	2
LDREXD <Rt>, <Rn> ^a	1	1	2
STREX <Rd>, <Rt>, <Rn>	1	1	2
STREXB <Rd>, <Rt>, <Rn>	1	1	2
STREXH <Rd>, <Rt>, <Rn>	1	1	2
STREXD <Rd>, <Rt>, <Rt2>, <Rn> ^a	1	1	2
SWP <Rt>, <Rt2>, <Rn>	2	2	3
SWPB <Rt>, <Rt2>, <Rn>	2	2	3

a. Address must be 64-bit aligned.

The synchronization instructions DMB, DSB, and ISB stall the pipeline for a variable number of cycles, depending on the current state of the memory system.

B.16 Coprocessor instructions

This section describes the cycle timing behavior for the MCR and MRC instructions to CP14, the debug coprocessor or CP15, the system control coprocessor.

The precise timing of coprocessor instructions is tightly linked with the behavior of the relevant coprocessor. Table B-21 shows the coprocessor instructions cycle timing behavior. Table B-21 shows the best case numbers.

Table B-21 Coprocessor instructions cycle timing behavior

Instruction	Cycles	Result latency	Comments
MCR	6	-	-
MCR<cond>	6	-	Condition code passes
	4	-	Condition code fails
MRC	6	6	-
MRC<cond>	6	6	Condition code passes
	4	4	Condition code fails

Note

Some instructions such as cache operations take more cycles.

B.17 SVC, BKPT, Undefined, and Prefetch Aborted instructions

This section describes the cycle timing behavior for SVC, Undefined Instruction, BKPT and Prefetch Abort.

In all cases the exception is taken in the Wr stage of the pipeline. SVC and most Undefined Instructions that fail their condition codes take one cycle. A small number of Undefined Instructions that fail their condition codes take two cycles. Table B-22 shows the SVC, BKPT, Undefined, prefetch aborted instructions cycle timing behavior.

Table B-22 SVC, BKPT, Undefined, prefetch aborted instructions cycle timing behavior

Instruction	Cycles
SVC (formerly SWI)	9
BKPT	9
Prefetch Abort	9
Undefined Instruction	9

B.18 Miscellaneous instructions

Table B-23 shows the cycle timing behavior for *If-Then* (IT) and *No Operation* (NOP) instructions.

Table B-23 IT and NOP instructions cycle timing behavior

Example instructions	Cycles	Early Reg	Late Reg	Result latency	Comments
IT{<v>{<w>{<z>}}} <cond>	1	-	-	-	-
NOP	1	-	-	-	-

The DBG, PLI, and YIELD instructions are all treated the same as NOP, and so have the same cycle timing behavior.

The WFI and WFE instructions stall the pipeline for a variable number of cycles, depending on the current state of the memory system.

B.19 Floating-point register transfer instructions

This section describes the cycle timing behavior for the various VFP instruction that transfer data between the VFP register file and the integer register file, including the system registers.

All source operands are Normal Regs, and the result latency for non-system register transfers is always 1 cycle.

Instructions that write data from the integer register file to the VFP system registers (VMSR) are blocking, that is, no subsequent instruction can start execution before the VMSR has completed execution. Consequently, the VMSR instructions take six cycles to execute.

All transfers to and from the VFP system registers are also serializing. This means that if there are any outstanding out-of-order-completion VFP instructions, the system register transfer instruction stalls in the iss-stage until these instructions are complete.

VFP instructions that complete out-of-order are VMLA.F32, VMLS.F32, VNMLS.F32, VNMLA.F32, VDIV.F32, VSQRT.F32, VCVT.F64.F32, and double-precision arithmetic and conversion instructions.

Table B-24 shows the floating-point register transfer instructions cycle timing behavior.

Table B-24 Floating-point register transfer instructions cycle timing behavior

Example instruction	Cycles	Result latency	Comments
VMOV <Sn>, <Rt>	1	1	-
VMOV <Rt>, <Sn>	1	2	-
VMOV <Dn[x]>, <Rt>	1	1	-
VMOV.<32> <Rt>, <Dn[x]>	1	2	-
VMOV <Sm>, <Sm1>, <Rt>, <Rt2>	1	1	-
VMOV <Rt>, <Rt2>, <Sm>, <Sm1>	1	2	-
VMOV <Dm>, <Rt>, <Rt2>	1	1	-
VMOV <Rt>, <Rt2>, <Dm>	1	2	-
VMSR <spec_reg>, <Rt>	6	-	Blocking and serializing
VMRS <Rt>, <spec_reg>	1	2	Serializing
VMRS APSR_nzcv, FPSCR	1	-	Serializing

B.20 Floating-point load/store instructions

This section describes the cycle timing behavior for all load and store instructions that operate on the VFP register file:

- The base address register, and any offset register are Very Early Regs for both loads and stores.
- For store instructions, the data register (Sd or Dd), or registers are always Late Regs.
- The cycle timing of load and store instructions is affected by the starting address for the transfer.

Note

The starting address is not always the same as the base address.

- The cycle timing of load and store multiple instructions is also affected by whether or not the base address register is updated by the instruction, that is, base register writeback.

Table B-25 shows the number of cycles and result latencies for single load and store instructions and load multiple instructions. Values are shown for each instruction with and without base register writeback, and with different starting address alignments. Cycle counts and base register result latencies for store multiple instructions are the same as for the equivalent load multiple instruction.

Table B-25 Floating-point load/store instructions cycle timing behavior

Example instruction	Cycles/ memory cycles	Cycles with writeback (!)	Result latency (load)	Result latency (base register, <Rn>)	Comments
VLDR.32 <Sd>, [<Rn>{, #+/-<imm>}]	1	-	1	-	-
VLDR.64 <Dd>, [<Rn>{, #+/-<imm>}]	1	-	1	-	64-bit aligned address
VLDR.64 <Dd>, [<Rn>{, #+/-<imm>}]	2	-	2	-	Not aligned
VSTR.32 <Sd>, [<Rn>{, #+/-<imm>}]	1	-	-	-	-
VSTR.64 <Dd>, [<Rn>{, #+/-<imm>}]	1	-	-	-	64-bit aligned address
VSTR.64 <Dd>, [<Rn>{, #+/-<imm>}]	2	-	-	-	Not aligned
First address 64-bit aligned					
VLDM{mode}.32 <Rn>{!}, {s1}	1	1	1	1	-
VLDM{mode}.32 <Rn>{!}, {s1,s2}	1	2	1,1	2	-
VLDM{mode}.32 <Rn>{!}, {s1-s3}	2	2	1,1,2	2	-
VLDM{mode}.32 <Rn>{!}, {s1-s4}	2	3	1,1,2,2	3	-
VLDM{mode}.64 <Rn>{!}, {d1}	1	2	1	2	-
VLDM{mode}.64 <Rn>{!}, {d1,d2}	2	3	1,2	3	-
VLDM{mode}.64 <Rn>{!}, {d1-d3}	3	4	1,2,3	4	-
VLDM{mode}.64 <Rn>{!}, {d1-d4}	4	5	1,2,3,4	5	-
First address not 64-bit aligned					
VLDM{mode}.32 <Rn>{!}, {s1}	1	1	1	1	-

Table B-25 Floating-point load/store instructions cycle timing behavior (continued)

Example instruction	Cycles/ memory cycles	Cycles with writeback (!)	Result latency (load)	Result latency (base register, <Rn>)	Comments
VLDM{mode}.32 <Rn>{!}, {s1,s2}	2	2	1,2	2	-
VLDM{mode}.32 <Rn>{!}, {s1-s3}	2	3	1,2,2	3	-
VLDM{mode}.32 <Rn>{!}, {s1-s4}	3	3	1,2,2,3	3	-
VLDM{mode}.64 <Rn>{!}, {d1}	2	2	2	2	-
VLDM{mode}.64 <Rn>{!}, {d1,d2}	3	3	2,3	3	-
VLDM{mode}.64 <Rn>{!}, {d1-d3}	4	4	2,3,4	4	-
VLDM{mode}.64 <Rn>{!}, {d1-d4}	5	5	2,3,4,5	5	-

B.21 Floating-point single-precision data processing instructions

This section describes the cycle timing behavior for all single-precision VFP CDP instructions. This includes arithmetic instructions such as VMUL.F32, data and immediate moving instructions such as "VMOV.F32 <Sd>, #<imm>", VABS.F32, VNEG.F32, and "VMOV <Sd>, <Sm>", and comparison instructions and conversion instructions.

Table B-26 shows the floating-point single-precision data processing instructions cycle timing behavior.

Table B-26 Floating-point single-precision data processing instructions cycle timing behavior

Example instruction	Cycles	Early Reg	Result latency
VMLA.F32 <Sd>, <Sn>, <Sm> ^a	1 ^b	<Sn>, <Sm>	5 ^c
VADD.F32 <Sd>, <Sn>, <Sm> ^d	1	<Sn>, <Sm>	2
VDIV.F32 <Sd>, <Sn>, <Sm>	2	<Sn>, <Sm>	16
VSQRT.F32 <Sd>, <Sm>	2	<Sm>	16
VMOV.F32 <Sd>, #<imm>	1	-	1
VMOV.F32 <Sd>, <Sm> ^e	1	-	1
VCMP.F32 <Sd>, <Sm> ^f	1	<Sd>, <Sm>	-
VCMP.F32 <Sd>, #0.0 ^f	1	<Sd>	-
VCVT.F32.U32 <Sd>, <Sm> ^g	1	<Sm>	2
VCVT.F32.U32 <Sd>, <Sd>, #<fbits> ^h	1	<Sd>	2
VCVTR.U32.F32 <Sd>, <Sm> ⁱ	1	<Sm>	2
VCVT.U32.F32 <Sd>, <Sd>, #<fbits> ^j	1	<Sd>	2
VCVT.F64.F32 <Dd>, <Sn>	3	<Sm>	5

a. Also VMLS.F32, VNMLS.F32, and VNMLA.F32.

b. VMLA.F32 completes out-of-order, and can take an extra cycle (two in total) if an add instruction (VADD) or certain dual-issued instruction pairs are in the iss-stage when the instruction completes.

c. Except when the instruction dependent on the result <Sd> is another VMLA.F32 instruction, and the dependent operand is the accumulate operand, <Sd>. In this case, the result latency is reduced to 3 cycles.

d. Also VSUB.F32, VMUL.F32, and VNMUL.F32.

e. Also VABS.F32 and VNEG.F32.

f. Also VCMPE.F32.

g. Also VCVT.F32.S32.

h. Also VCVT.F32.U16, VCVT.F32.S32, and VCVT.F32.S16.

i. Also VCVT.U32.F32, VCVTR.S32.F32, and VCVT.S32.F32.

j. Also VCVT.U16.F32, VCVT.S32.F32, and VCVT.S16.F32.

B.22 Floating-point double-precision data processing instructions

This section describes the cycle timing behavior for all double-precision VFP CDP instructions. This includes arithmetic instructions such as `VMUL.F64`, data and immediate moving instructions such as `VMOV.F64 <Dd>, #<imm>`, `VABS.F64`, `VNEG.F64`, and `VMOV <Dd>, <Dm>`, and comparison instructions and conversion instructions.

Table B-27 shows the floating-point double-precision data processing instructions cycle timing behavior

Table B-27 Floating-point double-precision data processing instructions cycle timing behavior

Example instruction	Cycles	Early Reg	Result latency
<code>VMLA.F64 <Dd>, <Dn>, <Dm>^a</code>	13	<code><Dn>, <Dm></code>	19
<code>VADD.F64 <Dd>, <Dn>, <Dm>^b</code>	3	<code><Dn>, <Dm></code>	9
<code>VDIV.F64 <Dd>, <Dn>, <Dm></code>	3	<code><Dn>, <Dm></code>	96
<code>VSQRT.F64 <Dd>, <Dm></code>	3	<code><Dm></code>	96
<code>VMOV.F64 <Dd>, #<imm></code>	1	-	1
<code>VMOV.F64 <Dd>, <Dm>^c</code>	1	-	1
<code>VCMP.F64 <Dd>, <Dm>^d</code>	2	<code><Dd>, <Dm></code>	-
<code>VCMP.F64 <Dd>, #0.0^d</code>	2	<code><Dm></code>	-
<code>VCVT.F64.U32 <Dd>, <Sm>^e</code>	3	<code><Dm></code>	7
<code>VCVT.F64.U32 <Dd>, <Dd>, #<fbits>^f</code>	3	<code><Dd></code>	7
<code>VCVTR.U32.F64 <Sd>, <Dm>^g</code>	3	<code><Dm></code>	7
<code>VCVT.U32.F64 <Dd>, <Dd>, #<fbits>^h</code>	3	<code><Dd></code>	7
<code>VCVT.F32.F64 <Sd>, <Dn></code>	3	<code><Dm></code>	7

a. Also `VMLS.F64`, `VNMLS.F64`, and `VNMLA.F64`.

b. Also `VSUB.F64`, `VMUL.F64`, and `VNMUL.F64`.

c. Also `VABS.F64` and `VNEG.F64`.

d. Also `VCMP.E.F64`.

e. Also `VCVT.F64.S32`.

f. Also `VCVT.F64.U16`, `VCVT.F64.S32`, and `VCVT.F64.S16`.

g. Also `VCVT.U32.F64`, `VCVTR.S32.F64`, and `VCVT.S32.F64`.

h. Also `VCVT.U16.F64`, `VCVT.S32.F64`, and `VCVT.S16.F64`.

B.23 Dual issue

To increase instruction throughput, the processor can issue certain pairs of instructions simultaneously. This is called dual issue. When this happens, the instruction with the smaller cycle count is assumed to execute in zero cycles. If a pair of instructions can be dual-issued, they are always dual-issued unless dual-issuing is disabled, see *c1, Auxiliary Control Register* on page 4-41. If one instruction of the pair is interlocked, both are interlocked.

This section describes:

- *Dual issue rules*
- *Permitted combinations* on page B-34

B.23.1 Dual issue rules

The following rules apply to dual-issue instructions:

- Both instructions must be available to the issue stage at the same time. This is unlikely if there are many branches.
- The second instruction must not use the PC as a source register unless it is B #immed.
- The first instruction must not use the PC as a destination register.
- Both instructions must belong to the same instruction set, ARM or Thumb.
- There must be no data dependency between the two instructions. That is, the second instruction must not have any source registers that are destination registers of the first instruction.

B.23.2 Permitted combinations

Table B-28 lists the permitted instruction combinations. Any instruction can be conditional or flag-setting unless otherwise stated. Only the exact instruction combinations listed in Table B-28 can be dual issued, provided you ensure the instruction combinations obey the rules specified in *Dual issue rules* on page B-33.

Table B-28 Permitted instruction combinations

Dual issue case	First instruction	Second instruction
Case A	Any instruction other than load/store multiple/double, flag-setting multiply, non-VFP coprocessor operations, miscellaneous processor control instructions ^a , or floating point instructions if floating point logic is not included in the processor	B #immed IT NOP
Case A-F ^b	Any floating point instructions, excluding load/store multiple, double-precision CDP instructions, VCVT.F64.F32, and VMRS and VMSR.	
Case B1	LDR <Rt>, [<Rn>, #<imm>] ^c LDR <Rt>, [<Rn>, <Rm>] ^c LDR <Rt>, [<Rn>, <Rm>, LSL #1, 2 or 3] ^c	Any data processing instruction that does not require a shift by a register value. ^d Any bitfield, saturate or bit-packing instruction. ^e Any signed or unsigned extend instruction. ^f Any SIMD add or subtract instruction. ^g Other miscellaneous instructions. ^h
Case B1-F ^b		Any single-precision CDP ⁱ , excluding "VMOV.F32 <Sd>, #<imm>", VNEG.F32, VABS.F32, VCVT.F64.F32, VDIV.F32, and VSQRT.F32. 32-bit transfers to and from the floating-point register file ^l .
Case B2	STR <Rt>, [<Rn>, #<imm>] ^c	As for Case B1.
Case B2-F ^b		As for Case B1-F
Case C	MOV <Rd>, #immedi ^k MOVW <Rd>, #immedi ^j MOV <Rd>, <Rm> ^j	Any data processing instruction. ^d Any bitfield, saturate or bit-packing instruction. ^e Any signed or unsigned extend instruction. ^f Any SIMD add or subtract instruction. ^g Other miscellaneous instructions. ^h
Case C-F ^b		32-bit transfers to and from the floating-point register file ^l .
Case F1 ^{b,m}	Any single-precision CDP ⁱ , excluding "VMOV.S32 <Sd>, #<imm>", VCVT.F64.F32, VABS.F32, and VNEG.F32.	As for case C or C-F.
Case F2_ld ^b	VLDR.F32 ⁿ	As for Case B1 or Case B1-F
Case F2_st ^b	VSTR.F32 ⁿ	As for Case B1. Any single-precision CDP ⁱ , excluding multiply-accumulate instructions ^o . 32-bit transfers to and from the floating-point register file ^l .
Case F2D ^b	VLDR.F64 ⁿ	As for Case B1.

Table B-28 Permitted instruction combinations (continued)

Dual issue case	First instruction	Second instruction
Case F3 ^b	32-bit transfers to and from the floating-point register file ^l "VMOV.F32, <Sd>, <Sm>", VABS.F32, and VNEG.F32.	As for Case F2_st.
Case F4 ^b	Any instruction that does not set flags, other than load/store multiple/double, non-VFP coprocessor operations, multi-cycle multiply instructions ^p , double-precision floating point CDP instructions, VCVT.F64.F32, or a miscellaneous processor control instruction ^a	Any single-precision CDP ⁱ , excluding "VMOV.F32 <Sd>, #<imm>", VNEG.F32, VABS.F32, VCVT.F64.F32, VDIV.F32, and VSQRT.F32. 32-bit transfers to and from the floating-point register file ^l .
Case F6 ^b	VMRS r15, FPSCR	As for Case A.

- a. These are processor state updating instructions, synchronization instructions, SVC, BKPT, prefetch abort and Undefined Instructions.
- b. This case can only occur if the optional floating-point functionality has been configured for the Cortex-R5F processor, see *Configurable options* on page 1-6.
- c. You can substitute LDR with LDRB, LDRH, LDRSB, or LDRSH. You can also substitute STR with STRB or STRH.
- d. Data processing instructions are ADC, ADD, ADDW, AND, ASR, BIC, CLZ, CMN, CMP, EOR, LSL, LSR, MOV, MOVT, MOVW, MVN, ORN, ORR, ROR, RRX, RSB, SBC, SUB, SUBW, TEQ, and TST.
- e. Bitfield, saturate, and bit-packing instructions are BFC, BFI, PKHBT, PKHTB, QADD, QDADD, QDSUB, QSUB, SBFX, SSAT, SSAT16, UBFX, USAT, and USAT16.
- f. Signed or unsigned extend instructions are SXTAB, SXTAB16, SXTAH, SXTB, SXTB16, SXTH, UXTAB, UXTAB16, UXTAH, UXTB, UXTB16, and UXTH.
- g. SIMD add and subtract instructions are QADD16, QADD8, QASX, SQUB16, QSUB8, QSAX, SADD16, SADD8, SASX, SHADD16, SHADD8, SHASX, SHSUB16, SHSUB8, SHSAX, SSUB16, SSUB8, SSAX, UADD16, UADD8, UASX, UHADD16, UHADD8, UHASX, UHSUB16, UHSUB8, UHSAX, UQADD16, UQADD8, UQASX, UQSUB16, UQSUB8, UQSAX, USUB16, USUB8, and USAX.
- h. Other miscellaneous instructions are RBIT, REV, REV16, REVSH, and SEL.
- i. Single-precision CDPs are VABS.F32, VNEG.F32, "VMOV.F32 <Sd>, #<imm>", VMLA.F32, VMLS.F32, VNMLS.F32, VNMLA.F32, VMUL.F32, VNMUL.F32, VADD.F32, VSUB.F32, VDIV.F32, VSQRT.F32, VCMPE.F32, VCMPE.F32, VCVT.F64.F32, VCVT.F32.U32, VCVT.F32.S32, VCVT.F32.U16, VCVT.F32.S16, VCVTR.U32.F32, VCVT.U32.F32, VCVTR.S32.F32, VCVT.S32.F32, VCVT.U16.F32, and VCVT.S16.F32.
- j. Must not be flag-setting.
- k. Immediate value must not require a shift.
- l. 32-bit transfers to or from the floating point register file include single or half-double floating point register transfers, including "VMOV <Sn>, <Rt>", "VMOV.F32 <Dn[x]>, <Rt>", "VMOV.F32 <Rt>, <Dn[x]>", and "VMOV <Rt>, <Sn>", but excluding VMRS and VMSR.
- m. When the first instruction is a floating point multiply-accumulate, and the second instruction is a 32-bit transfer to the floating-point register file, case F1 can only occur if the two instructions have different destination registers.
- n. Any addressing modes.
- o. Single-precision floating-point multiply-accumulate instructions are VMLA.F32, VMLS.F32, VNMLS.F32, and VNMLA.F32.
- p. Multi-cycle multiply instructions are SMMUL, SMMULA, SMMLS, MUL, MLA, MLS, SMULL, SMLAL, UMAAL, UMULL, and UMLAL.

Appendix C

ECC Schemes

This appendix describes some of the advantages and disadvantages of the different *Error Checking and Correction* (ECC) schemes for the TCMs. It contains the following section:

- *ECC scheme selection guidelines* on page C-2.

C.1 ECC scheme selection guidelines

When deciding to implement a Cortex-R5 processor with an ECC scheme on one or both of the TCM interfaces, give careful consideration between using 32-bit or 64-bit ECC. To calculate or check the ECC code for data, the processor must know the value of all bytes in the data chunk protected by the scheme. Therefore, when using these schemes, the processor must perform additional read accesses to calculate and check the ECC code stored with the data.

For example, if the ATCM is implemented with 32-bit ECC and a program performs an aligned STR to the memory, the processor can calculate the error correction code using only the data stored by the program.

If the same memory was implemented with 64-bit ECC, the processor cannot calculate the ECC code for the doubleword memory chunk being written using only the data stored by the program. To calculate the ECC code and store the data, the processor must first perform a read of the other word in that memory chunk. This increases the number of memory accesses required to execute the program. This increases power consumption, and can also lead to a decrease in performance.

Use the following guidelines to decide which scheme to use. If you are in any doubt, benchmark your system running typical software to find the best balance between area, power, and performance for your application.

- For a TCM interface that contains mainly instructions, use 64-bit ECC. The vast majority of reads requested by the prefetch unit are doubleword.
- Use 64-bit ECC when a TCM contains data that is accessed using:
 - LDRD or STRD instructions where the start address is doubleword aligned
 - LDM or STM instructions where the start address is doubleword aligned and there are an even number of registers in the register list.

64-bit ECC requires less RAM area, and does not provide any performance loss or increased power consumption over 32-bit ECC in these cases.
- When LDM and STM instructions are used to access many registers, the majority of TCM accesses do not require additional reads with 64-bit ECC.
- 32-bit ECC provides better power consumption and generally better performance compared to 64-bit ECC when:
 - a program performs many unaligned accesses to data in a TCM
 - a program performs many byte, halfword, and word accesses to data in a TCM.

You might be able to obtain optimal results by using a different error detection scheme on each TCM interface, and allocating instructions and data to each interface based on the guidelines given in this section.

Appendix D

Memory Ordering

This appendix describes the processor memory ordering. It contains the following sections:

- *Memory ordering* on page D-2
- *Virtual AXI peripheral interface* on page D-3.

D.1 Memory ordering

The ARM architecture requires that transactions to locations in Device-type memory be ordered. The Cortex-R5 processor has an in-order pipeline, so any non-cached read blocks, preventing any subsequent read or write from starting until the current read is complete. On an AXI bus, as used by the Cortex-R5 processor, a series of writes issued in order, is kept in order by using the same ID for all the transactions.

To maintain ordering between a write and a subsequent read, the Cortex-R5 processor waits for the write transaction to complete before starting the read. The writes that the Cortex-R5 processor must wait for are any Device-type writes in its write buffer or bus interface and writes for which the address and data have been accepted by the bus but for which no response has been received, that is *AXI outstanding writes*. The latency of the Device read depends on how many writes must complete before it starts.

The architectural ordering requirements apply only to individual peripherals so, for example, an outstanding write to a UART does not have to be completed before a read from an interrupt controller can be started. However, the Cortex-R5 processor views the memory attached the each interface as flat, so ordering is preserved for all accesses to a given interface. Accesses to different Cortex-R5 interfaces are not ordered, so selecting which interface is used can improve the latency of critical Device read accesses.

For example, if a CPU has a number of write transactions outstanding on the AXI master interface, a read from an interrupt controller attached to the AXI master interface must wait for those writes to complete and the latency incurred might impact the interrupt handling performance. Alternatively, if the interrupt controller were attached to the AXI peripheral interface, the read could start without waiting for the outstanding writes on the AXI master interface. However, the read would have to wait for any outstanding writes on the AXI peripheral interface or its buffers.

Note

- The transaction ordering provided by Device memory is useful in situations where the access has side effects. For example, if the processor writes to a memory-mapped FIFO, and then reads a different memory-mapped register that indicates whether the FIFO is full, the value read must reflect the state of the FIFO after the write otherwise a further write could be performed that causes an overflow.
 - If a write to a peripheral on one interface causes a side effect on a peripheral on a different interface, there is no implicit ordering to ensure the side effect is observed by a subsequent access to the second peripheral, even if both are in Device-type memory. In this situation, you must perform a read from the first peripheral to ensure that the write has completed, followed by a DMB to ensure ordering before performing the second access. On the Cortex-R5 processor, a DMB alone is sufficient to force this ordering, but this is not architectural and cannot be relied on in the general case.
-

Writes to Device-type memory always drain from the Cortex-R5 buffers as quickly as possible. If the memory system attached to a port is perfect, that is the write response is returned in the cycle after the address and data have been received, outstanding accesses cannot accumulate. Selecting different interfaces for different peripherals does not improve read latencies in such a system.

D.2 Virtual AXI peripheral interface

Each Cortex-R5 CPU can perform memory transactions using the AXI master interface, the AXI peripheral interface or, if included, the AHB peripheral interface. Each of these interfaces is treated independently from an ordering point of view. The virtual AXI peripheral interface provides an additional interface that, although it shares the same physical port as the AXI peripheral interface, is treated independently from an ordering point of view.

The two AXI peripheral interfaces use different AXI IDs to enable the memory system to return responses out of order. They also have different limits on the number of outstanding writes permitted so, by selecting a particular interface for a peripheral, you can have some control over the maximum latency of accesses to that peripheral. If your AXI peripheral port memory system accepts outstanding write transactions, ARM recommends that you configure the peripheral interfaces so that the most latency critical peripheral, possibly an interrupt controller, is on the virtual AXI peripheral interface and all others elsewhere.

Note

- The AXI peripheral interface and virtual AXI peripheral interface share write buffer logic, and write data is drained in order from this buffer. The interfaces use different IDs, so write responses can be received out-of-order. If the buffer contains writes to both interfaces, and the AXI peripheral interface writes are older, a virtual AXI peripheral interface read cannot start until the virtual AXI peripheral interface writes have all completed, and this in turn requires that the AXI peripheral interface writes have posted address and data to the bus though not necessarily completed.
 - Similarly, if the memory system on the AXI peripheral port returns all write responses in order, regardless of ID, this can force reads on one interface to wait for writes on a different interface. The same effect is possible if two CPU ports connect to a common memory bus that forces ordering.
-

Appendix E

Revisions

This appendix describes the technical changes between released issues of this book.

Table E-1 Issue A

Change	Location	Affects
First release	-	-

Table E-2 Differences between issue A and issue B

Change	Location	Affects
Add ID values for r1p0	Table 1-3 on page 1-16	r1p0
Updated AMBA interface clock gating	<i>Clock gating</i> on page 2-16	r1p0
System control register enables SWP and SWPB to be Undefined	Table 4-24 on page 4-39	r1p0
Single-precision only option for Cortex-R5F	<i>Features</i> on page 1-4 Table 1-1 on page 1-6 Figure 4-45 on page 4-69 Table 4-58 on page 4-80 <i>About the FPU programmers model</i> on page 11-2 <i>Media and VFP Feature Registers, MVFR0 and MVFR1</i> on page 11-9	r1p0
Accessibility of Slave Port Control Register	page 4-65	All revisions

Table E-2 Differences between issue A and issue B (continued)

Change	Location	Affects
Additional description for c15, Build Options 1 Register	<i>c15, Build Options 1 Register</i> on page 4-79	All revisions
Update AXI slave address decode information	<i>AXI slave interface for cache RAMs</i> on page 9-18	All revisions
Update AXI slave characteristics	<i>AXI slave characteristics</i> on page 9-20	All revisions
Changed RAM access using AXI slave interface	<i>Accessing RAMs using the AXI slave interface</i> on page 9-21	r1p0
Register name corrections	<i>STRH</i> on page 9-38 <i>DTR access mode</i> on page 12-17	All revisions
MVFR1.LS change of usage	Table 11-8 on page 11-10	r1p0

Table E-3 Differences between issue B and issue C

Change	Location	Affects
Update revision information	Table 1-3 on page 1-16	r1p1
	Table 4-7 on page 4-18	
	Table 4-15 on page 4-28	
	Table 4-17 on page 4-31	
	<i>AXI master interface transfers</i> on page 9-7	
Correct RVPTYSm signal name	Table A-8 on page A-14	All revisions
Add BVPTYCS signal description	Table A-10 on page A-16	All revisions
Add ARCTLPTYS[3:0] signal description	Table A-8 on page A-14	All revisions
Update RAM-Access space reference	<i>Cache RAM access</i> on page 9-23	All revisions
Update validation register short names	<i>Validation Registers</i> on page 4-68	All revisions
Update descriptions of product revisions	Table 4-3 on page 4-15	All revisions
	Table 12-6 on page 12-11	
	Table 12-31 on page 12-38	
	Table 11-4 on page 11-6	
Update register descriptions	Throughout manual	All revisions

Glossary

This glossary describes some of the terms used in technical documents from ARM.

- Abort** An exception caused by an illegal memory access. Aborts can be caused by the external memory system or by the memory-management hardware, that might include a *Memory Management Unit* (MMU) or a *Memory Protection Unit* (MPU).
- See also* Data abort, External abort and Prefetch abort.
- Abort model** Describes what happens to the processor state when a Data abort exception occurs. Different abort models behave differently with regard to load and store instructions that specify base register write-back.
- Addressing mode** A method for generating the memory address used by a load or store instruction.
- Advanced eXtensible Interface (AXI)** A bus protocol that supports separate phases for address or control and data, unaligned data transfers using byte strobes, burst-based transactions with only start address issued, separate read and write data channels, issuing multiple outstanding addresses, out-of-order transaction completion, and easy addition of register stages to provide timing closure.
- The AXI protocol includes optional extensions for signaling for low-power operation.
- Advanced High-performance Bus (AHB)** A bus protocol with a fixed pipeline between the address or control and data phases. It supports a subset of the functionality of the AMBA AXI protocol. The full AMBA AHB protocol specification includes a number of features that are not commonly required for master and slave implementations and ARM recommends using the AMBA AHB-Lite subset of the protocol.
- See also* Advanced Microcontroller Bus Architecture (AMBA) and AHB-Lite.

Advanced Microcontroller Bus Architecture (AMBA)

The AMBA family of protocol specifications is the ARM open standard for on-chip buses. AMBA provides a strategy for the interconnection and management of the functional blocks that make up a *System-on-Chip* (SoC). Applications include the development of embedded systems with one or more processors or signal processors and multiple peripherals. AMBA defines a common backbone for SoC modules, and therefore complements a reusable design methodology.

Advanced Peripheral Bus (APB)

A bus protocol that is designed for use with ancillary or general-purpose peripherals such as timers, interrupt controllers, UARTs, and I/O ports. It connects to the main system bus through a system-to-peripheral bus bridge that helps reduce system power consumption.

Advanced SIMD

An extension to the ARM architecture that provides *Single Instruction Multiple Data* (SIMD) operations on a bank of extension registers. If a floating-point extension is also implemented, the two extensions share a common extension register bank. The Advanced SIMD extension implements NEON technology, and is often called NEON.

AHB

See Advanced High-performance Bus (AHB).

AHB-Lite

A subset of the full AMBA AHB protocol specification. It provides all of the basic functions required by the majority of AMBA AHB slave and master designs, particularly when used with a multi-layer AMBA interconnect. In most cases, the extra facilities provided by a full AMBA AHB interface are implemented more efficiently using an AMBA AXI protocol interface.

Aligned

A data item stored at an address that is divisible by the number of bytes that defines its data size is said to be aligned. Aligned doublewords, words, and halfwords have addresses that are divisible by eight, four, and two respectively. The terms doubleword-aligned, word-aligned, and halfword-aligned therefore stipulate addresses that are divisible by eight, four, and two respectively. An aligned access is one where the address of the access is aligned to the size of an element of the access.

AMBA

See Advanced Microcontroller Bus Architecture (AMBA).

APB

See Advanced Peripheral Bus (APB).

ARM instruction

A word that specifies an operation for a processor in ARM state to perform. ARM instructions must be word-aligned.

ARM state

In ARM state the processor executes the ARM instruction set.

AXI

See Advanced eXtensible Interface (AXI).

AXI channels, channel order and interfaces

The block diagram shows:

- the order in which AXI channel signals are described
- the master and slave interface conventions for AXI components.

AXI signal names have a one or two letter prefix that denotes the AXI channel as follows:

AW	Write address channel.
W	Write data channel.
B	Write response channel.
AR	Read address channel.
R	Read data channel.

General descriptions of AXI signals use x to represent this prefix, for example, **xVALID** and **xREADY**.

AXI terminology

The following general AXI terms apply to both masters and slaves:

Active read transaction

A transaction for which the read address transfer has been completed, but the last read data transfer has not been completed.

Active transfer

A transfer for which the transmitting interface has asserted the **xVALID** handshake signal, but the receiving interface has not asserted the **xREADY** handshake signal.

Active write transaction

A transaction for which the write address or leading write data transfer has been completed, but the write response has not been completed.

Completed transfer

A transfer for which the handshake using **xVALID** and **xREADY** is complete.

Payload The non-handshake signals in a transfer.

Transaction An entire burst of transfers, comprising an address transfer, one or more data transfers and, for write transactions only, a response transfer.

Transmitting interface

An initiator driving the payload and asserting the relevant **xVALID** signal.

Transfer A single exchange of information. That is, a transfer with a single handshake using **xVALID** and **xREADY**.

The following AXI terms are master interface attributes. To permit system performance optimization, they must be specified for every component with an AXI master interface:

Combined issuing capability

The maximum number of active transactions that the interface can generate. It is specified for master interfaces that use combined storage for active write and read transactions. If not specified you can assume it is equal to the sum of the write and read issuing capabilities.

Read ID capability

The maximum number of different **ARID** values that the interface can generate for all active read transactions at any one time.

Read ID width

The number of bits in the **ARID** bus.

Read issuing capability

The maximum number of active read transactions that the interface can generate. Must be specified if the combined issuing capability is not specified.

Write ID capability

The maximum number of different **AWID** values that the interface can generate for all active write transactions at any one time.

Write ID width

The number of bits in the **AWID** and **WID** buses.

Write interleave capability

The number of active write transactions for which the interface can transmit data. This is counted from the earliest transaction.

Write issuing capability

The maximum number of active write transactions that a master interface can generate. Must be specified if the combined issuing capability is not specified.

The following AXI terms are slave interface attributes. To permit performance optimization, they must be specified for every component with an AXI slave interface:

Combined acceptance capability

The maximum number of active transactions that the interface can accept. It is specified for slave interfaces that use combined storage for active write and read transactions. If not specified then you can assume it is equal to the sum of the write and read acceptance capabilities.

Read acceptance capability

The maximum number of active read transactions that the interface can accept. Must be specified if the combined acceptance capability is not specified.

Read data reordering depth

The number of active read transactions for which the interface can transmit data. This is counted from the earliest transaction.

Write acceptance capability

The maximum number of active write transactions that the interface can accept. Must be specified if the combined acceptance capability is not specified.

Write interleave depth

The number of active write transactions for which the interface can receive data. This is counted from the earliest transaction.

Banked registers

A register that has multiple instances, with the instance that is in use dependent on a property of the state of the device, for example the processor mode or security state.

Base register

A register specified by a load or store instruction that is used as the base value for the address calculation for the instruction. Depending on the instruction and its addressing mode, an offset can be added to or subtracted from the base register value to form the virtual address that is sent to memory.

Base register write-back

Writing back a modified value to the base register used in an address calculation.

Beat

Alternative word for an individual transfer within a burst. For example, an INCR4 burst comprises four beats.

See also Burst.

Boundary scan chain

A boundary scan chain is made up of serially-connected devices that implement boundary scan technology using a standard JTAG TAP interface. Each device contains at least one TAP controller containing shift registers that form the chain, connected between **TDI** and **TDO**, through which test data is shifted. A processor can contain several shift registers, enabling you to access selected parts of the device.

Branch prediction	A technique where a processor chooses a future execution path to prefetch along. For example, after a branch instruction, the processor can choose to prefetch either the instruction following the branch or the instruction at the branch target. <i>See also</i> Prefetching.
Breakpoint	A breakpoint is a debug event triggered by the execution of a particular instruction. It is specified in terms of one or both of the address of the instruction and the state of the processor when the instruction is executed. <i>See also</i> Watchpoint.
Burst	A group of transfers to consecutive addresses. Because the addresses are consecutive, the device transmitting the data does not have to supply an address for any transfer after the first one. This increases the speed at which the burst occurs. If using an AMBA interface, the transmitting device controls the burst using signals that indicate the length of the burst and how the addresses are incremented. <i>See also</i> Beat.
Byte lane strobe	A signal that determines which byte lanes are active, or valid, in a data transfer. Each bit of this signal corresponds to eight bits of the data bus.
Byte-invariant	In a byte-invariant system, the address of each byte of memory remains unchanged when switching between little-endian and big-endian operation. When a data item larger than a byte is loaded from or stored to memory, the bytes making up that data item are arranged into the correct order depending on the endianness of the memory access. The ARM architecture supports byte-invariant systems in ARMv6 and later versions. When byte-invariant support is selected, unaligned halfword and word memory accesses are also supported. The architecture requires multi-word accesses to be word-aligned. <i>See also</i> Word-invariant.
Cache hit	A memory access that can be processed at high speed because the instruction or data that it addresses is already held in the cache.
Cache line	The basic unit of storage in a cache. Its size in words is always a power of two, usually four or eight words. A cache line must be aligned to a suitable memory boundary. <i>See also</i> Cache terminology diagram.
Cache miss	A memory access that cannot be processed at high speed because the instruction or data it addresses is not in the cache.
Cache sets	Areas of a cache, divided up to simplify and speed up the process of determining whether a cache hit occurs. In the ARM architecture, the number of cache sets is always a power of two. <i>See also</i> Cache terminology diagram.
Cache terminology	<i>See</i> the Cache terminology diagram and the entries for terms used in that diagram.
Cache terminology diagram	The diagram illustrates the following cache terminology: • block address • cache line • cache set • cache way • index • tag.

Cache way	A cache way consists of one cache line from each cache set. The cache ways are indexed from 0 to ASSOCIATIVITY-1. Each cache lines in a cache way has the same index as the cache way. For example cache way n consists of the cache line with index n from each cache set. <i>See also</i> Cache terminology diagram.
CDP instruction	Coprocessor data processing instruction. For the VFP coprocessor, CDP instructions are arithmetic instructions and FCPY, FABS, and FNEG.
Clean	A cache line that has not been modified while it is in the cache is said to be clean. To clean a cache is to write dirty cache entries into main memory. If a cache line is clean, it is not written on a cache miss because the next level of memory contains the same data as the cache. <i>See also</i> Dirty.
Clock gating	Gating a clock signal for a macrocell or functional block with a control signal and using the modified clock that results to control the operating state of the macrocell or block.
Clocks Per Instruction (CPI)	<i>See</i> Cycles Per Instruction (CPI).
Coherency	<i>See</i> Memory coherency.
Cold reset	Also known as power-on reset. Starting the processor by turning power on. Turning power off and then back on again clears main memory and many internal settings. Some program failures can lock up the processor and require a cold reset to restart the system. In other cases, only a warm reset is required. <i>See also</i> Warm reset.
Communications channel	The hardware used for communicating between the software running on the processor, and an external host, using the debug interface. When this communication is for debug purposes, it is called the <i>Debug Communications Channel</i> (DCC). From ARMv6, the DCC includes the Data Transfer Register, some bits in the Data Status and Control Register, and the external debug interface controller, such as the DBGTap controller in the case of a JTAG interface.
Condition field	A four-bit field in an ARM instruction that specifies a condition under which the instruction executes. <i>See also</i> Conditional execution.
Conditional execution	For ARM instructions, if the condition field indicates that the corresponding condition is true when the instruction starts executing, it executes normally. Otherwise, the instruction does nothing. In the Thumb instruction set, the IT instruction makes up to four of the following instructions conditional.
Context switch	The saving and restoring of computational state when switching between different threads or processes. In ARM documentation, the term context switch describes any situation where the context is switched by an operating system and might or might not include changes to the address space.
Coprocessor	A processor that supplements the main processor to carry out additional functions that the main processor cannot perform. The ARM architecture defines an interface to up to 16 coprocessors, CP0-CP15 for use by ARM: • CP15 instructions access the System Control processor

	• CP14 instructions access control registers for debug, trace, and execution environment features • CP10 and CP11 instruction space is for floating-point and Advanced SIMD instructions if supported.
Core register	One of the 32-bit general-purpose integer registers, R0 to R15. R15 is an alias for PC, the Program Counter. R14 is an alias for LR, the Link Register, and R13 is an alias for SP, the Stack Pointer. See the appropriate ARM Architectural Reference Manual for the constraints on the use of PC, LR, and SP.
CoreSight	ARM on-chip debug and trace components, that provide the infrastructure for monitoring, tracing, and debugging a complete system on chip.
CPI	See Cycles Per Instruction (CPI).
CPSR	See Current Program Status Register (CPSR).
Cross Trigger Interface (CTI)	Part of an <i>Embedded Cross Trigger</i> (ECT) device. In an ECT, the CTI provides the interface between a processor or ETM and the CTM.
Cross Trigger Matrix (CTM)	In an ECT device, the CTM combines the trigger requests generated by CTIs and broadcasts them to all CTIs as channel triggers.
CTI	See Cross Trigger Interface (CTI).
CTM	See Cross Trigger Matrix (CTM).
Current Program Status Register (CPSR)	The register that holds the current operating processor status. See also Program Status Register and Saved Program Status Register.
Cycles Per Instruction (CPI)	A measure of the number of computer instructions that can be performed in one clock cycle, also called clocks per instruction. This value can be used to compare the performance of different processors that implement the same instruction set. The lower the value, the better the performance.
DAP	See Debug Access Port.
Data abort	An indication from a memory system to the processor of an attempt to access an illegal data memory location. An exception must be taken if the processor attempts to use the data that caused the abort. See also Abort, External abort, and Prefetch abort.
DBGTAP	See Debug Test Access Port.
Debug Access Port (DAP)	A block that acts as a master on a system bus and provides access to the bus from an external debugger.
Debug Test Access Port (DBGTAP)	A debug control and data interface based on the IEEE 1149.1 JTAG <i>Test Access Port</i> (TAP). The interface has four or five signals.

Debugger	A debugging system that includes a program, used to detect, locate, and correct software faults, together with custom hardware that supports software debugging.				
Default NaN mode	In floating-point operation, a mode in which all operations that result in a NaN return the default NaN, regardless of the cause of the NaN result. This mode is compliant with the IEEE 754 standard but implies that all information contained in any input NaNs to an operation is lost. <i>See also</i> NaN.				
Digital Signal Processing (DSP)	A variety of algorithms to process signals that have been sampled and converted to digital form. Saturated arithmetic is often used in such algorithms.				
Dirty	A dirty cache line in a write-back cache is a line that has been modified while it is in the cache. Typically, a cache line is marked as dirty by setting the dirty bit to 1. <i>See also</i> Clean.				
DNM	<i>See</i> Do Not Modify.				
Do Not Modify (DNM)	A value that must not be altered by software. DNM fields read as unknown values, and must only be written with the value read from the same field on the same processor.				
Double-precision value	In floating-point operation, consists of two 32-bit words that must appear consecutively in memory and are both word-aligned. The value is interpreted as a basic double-precision floating-point number according to the IEEE 754-1985 standard.				
Doubleword	A 64-bit data item. Doublewords are normally at least word-aligned in ARM systems.				
Doubleword-aligned	A data item having a memory address that is divisible by eight.				
DSP	<i>See</i> Digital Signal Processing.				
Embedded Trace Macrocell (ETM)	A hardware macrocell that, when connected to a processor, outputs trace information on a trace port. The ETM provides processor driven trace through a trace port compliant to the ATB protocol. An ETM always supports instruction trace, and might support data trace.				
EmbeddedICE logic	An on-chip logic block that provides TAP-based debug support for an ARM processor. It is accessed through the DAP on the ARM processor.				
EmbeddedICE-RT	Hardware provided by an ARM processor to aid debugging in real-time.				
Endianness	The scheme that determines the order of successive bytes of a data word when it is stored in memory.				
ETM	<i>See</i> Embedded Trace Macrocell.				
Event	In an ARM trace macrocell, event has a particular meaning and these events can be simple or complex: <table data-bbox="483 1650 1500 1843"> <tr> <td>Simple</td><td>An observable condition that a trace macrocell can use to control aspects of a trace.</td></tr> <tr> <td>Complex</td><td>A boolean combination of simple events that a trace macrocell can use to control aspects of a trace.</td></tr> </table>	Simple	An observable condition that a trace macrocell can use to control aspects of a trace.	Complex	A boolean combination of simple events that a trace macrocell can use to control aspects of a trace.
Simple	An observable condition that a trace macrocell can use to control aspects of a trace.				
Complex	A boolean combination of simple events that a trace macrocell can use to control aspects of a trace.				
Exception	A mechanism to handle a fault or error event. For example, exceptions handle external interrupts and undefined instructions.				

Exception vector	A fixed address that contains the address of the first instruction of the corresponding exception handler.
External abort	An abort generated by the external memory system. <i>See also</i> Abort, Data abort and Prefetch abort.
Fast Context Switch Extension (FCSE)	An extension to the ARM architecture that modifies the behavior of the memory system. It enables multiple programs running on the processor to use identical address ranges, while ensuring that the addresses they present to the rest of the memory system differ. From ARMv6, use of the FCSE is deprecated. The FCSE is optional in ARMv7, and obsolete from the ARMv7 Multiprocessing Extensions.
Fault	An abort generated by the memory system, for example by the <i>Memory Management Unit</i> (MMU).
FCSE	<i>See</i> Fast Context Switch Extension.
Flat address mapping	A system of organizing memory where the physical address for every access is equal to its virtual address.
Flush-to-zero mode	In floating-point operation, a special processing mode that optimizes the performance of some floating-point algorithms by replacing the denormalized operands and intermediate results with zeros, without significantly affecting the accuracy of their final results.
General-purpose register	<i>See</i> Core register.
Halfword	A 16-bit data item. Halfwords are normally halfword-aligned in ARM systems.
Halfword-aligned	A data item having a memory address that is divisible by 2.
Halting debug-mode	One of two mutually exclusive debug modes. In Halting debug-mode all processor execution halts when a breakpoint or watchpoint is encountered. You can examine and alter all processor state, coprocessor state, memory, input and output locations using the debug interface. <i>See also</i> Monitor debug-mode.
High registers	<i>See</i> Core register.
High vectors	Alternative locations for exception vectors. The high vector address range is near the top of the address space, rather than at the bottom.
Hint instruction	A hint instruction provides information that the hardware can take advantage of. A processor implementation can choose whether to implement hint instructions or not. If they are not implemented, they execute as NOP.
Host	A computer that provides data and other services to another computer. Especially, a computer providing debugging services to a target being debugged.
Immediate values	Values that are encoded directly in the instruction and used as numeric data when the instruction is executed. Many ARM and Thumb instructions can be used with an immediate argument.
Implementation-defined	Behavior that is not defined by the architecture, but is defined and documented by the implementation.

Implementation-specific	<i>See</i> Implementation-defined
Index	<i>See</i> Cache index.
Instruction cycle count	The number of cycles for which an instruction occupies the Execute stage of the pipeline.
Internal scan chain	A series of registers connected together to form a path through a device, used during production testing to import test patterns into internal nodes of the device and export the resulting values.
Interrupt handler	<i>See</i> Exception handler.
Invalidate	Marking a cache line as being not valid, by clearing the valid bit to 0. This must be done whenever the line does not contain a valid cache entry. For example, after a cache flush all lines are invalid.
Jazelle state	In Jazelle state the processor executes Java bytecodes as part of a <i>Java Virtual Machine</i> (JVM). <i>See also</i> ARM state, Thumb state, and ThumbEE state.
JTAG Access Port (JTAG-AP)	An optional component of the DAP that provides debugger access to on-chip scan chains.
Load/store architecture	A processor architecture where data-processing operations only operate on register contents, not directly on memory contents.
Macrocell	A complex logic block with a defined interface and behavior. A typical VLSI system comprises several macrocells, such as a processor, an ETM, and a memory block integrated with application-specific logic.
Memory coherency	A memory is coherent if the value read by a data read or instruction fetch is the value that was most recently written to that location. Memory coherency is made difficult when the memory system includes multiple possible physical locations, such as main memory, a write buffer and one or more caches.
Memory Management Unit (MMU)	A hardware unit that provides detailed control of a memory system. Most of the control is provided by translation tables held in memory.
Memory Protection Unit (MPU)	A hardware unit that provides simple control of a limited number of protection regions in memory.
Miss	<i>See</i> Cache miss.
MMU	<i>See</i> Memory Management Unit.
Modified Virtual Address (MVA)	The address produced by the FCSE that is sent to the rest of the memory system to be used in place of the normal virtual address. When the FCSE is absent or disabled, the MVA and the <i>Virtual Address</i> (VA) have the same value. <i>See also</i> Fast Context Switch Extension.

Monitor debug-mode

One of two mutually exclusive debug modes. In Monitor debug-mode the processor enables a software abort handler provided by the debug monitor or operating system debug task. When a breakpoint or watchpoint is encountered, system interrupts continue to be serviced while normal program execution is suspended.

See also Halting debug-mode.

MPU

See Memory Protection Unit.

MVA

See Modified Virtual Address.

NaN

Not a number. In floating-point operation, NaNs are special floating-point values that can be used when neither a numeric value nor an infinity is appropriate. NaNs can be quiet NaNs that propagate through most floating-point operations, or signaling NaNs that cause Invalid Operation floating-point exceptions when used.

PA

See Physical Address.

Penalty

The number of cycles in which no useful Execute stage pipeline activity can occur because an instruction flow is different from that assumed or predicted.

Physical Address (PA)

The address that identifies a main memory location.

Power-on reset

See Cold reset.

Prefetch abort

An indication from a memory system to the processor that an instruction has been fetched from an illegal memory location. An exception must be taken if the processor attempts to execute the instruction. A Prefetch abort can be caused by the external or internal memory system as a result of attempting to access invalid instruction memory.

See also Data abort, External abort and Abort.

Prefetching

The process of fetching instructions from memory before the instructions that precede them have finished executing. Prefetching an instruction does not mean that the instruction must be executed.

Privileged mode

Any processor mode other than User mode. Memory systems typically check memory accesses from privileged modes against supervisor access permissions rather than the more restrictive user access permissions. The use of some instructions is also restricted in privileged modes.

Programming Language Interface (PLI)

For Verilog simulators, an interface by which foreign code can be included in a simulation. Foreign code is code written in a different language.

Protection region

A memory region whose position, size, and other properties are defined by the Memory Protection Unit registers.

Read

Memory operations that have the semantics of a load. *See the ARM Architecture Reference Manual* for more information.

RealView ICE

ARM JTAG interface unit for debugging embedded processor cores that uses a DBGTap or Serial Wire interface.

Remapping

Changing the address of physical memory or devices after the application has started executing. This might be done to permit RAM to replace ROM when the initialization has completed.

Reserved

Registers and instructions that are reserved are Unpredictable unless otherwise stated. Bit positions described as Reserved are UNK/SBZP.

Round to Nearest (RN) mode

In floating-point operation, the rounded result is the nearest representable number to the unrounded result. The tie case is rounded up if it would clear the least significant bit of the significand, making it even.

See also Rounding mode, Rounding error.

Round towards Minus infinity (RM) mode

In floating-point operation, the rounded result is the nearest representable number that is less than or equal to the exact result. This rounding mode is used in interval arithmetic.

See also Rounding mode, Rounding error.

Round towards Plus infinity (RP) mode

In floating-point operation, the rounded result is the nearest representable number that is greater than or equal to the exact result. This rounding mode is used in interval arithmetic.

See also Rounding mode, Rounding error.

Round towards Zero (RZ) mode

In floating-point operation, results are rounded to the nearest representable number that is no greater in magnitude than the unrounded result. This rounding mode chops any bits to the right of the significand, always rounding down, and is used by the C, C++, and Java languages in integer conversions.

See also Rounding mode, Rounding error.

Rounding error

Is defined to be the value of the rounded result of an arithmetic operation minus the exact result of the operation.

See also Rounding mode.

Rounding mode

In floating-point operation, specifies how the exact result of a floating-point operation is rounded to a value that is representable in the destination format.

See also Round to Nearest (RN) mode, Round towards Minus Infinity (RM) mode, Round towards Plus infinity (RP) mode, and Round towards Zero (RZ) mode.

Saved Program Status Register (SPSR)

The register that holds the CPSR of the task immediately before the exception occurred that caused the switch to the current mode. Each exception mode has its own SPSR.

SBO

See Should Be One.

SBZ

See Should Be Zero.

SBZP

See Should Be Zero or Preserved.

Set

See Cache set.

Short vector operation

A floating-point coprocessor operation involving more than one destination register and perhaps more than one source register in the generation of the result for each destination.

Should Be One (SBO)

Software must write as 1, or all 1s for bit fields. Writing any other value produces Unpredictable results.

Should Be Zero (SBZ)

Software must write as 0, or all 0s for bit fields. Writing any other value produces Unpredictable results.

Should Be Zero or Preserved (SBZP)

Software must write as 0, or all 0s for a bit field, if the value is being written without having previously been read, or if the register has not been initialized. If the register has previously been read, software must preserve the field value by writing back the value that was read from the same field on the same processor.

Signaling NaN

In floating-point operation, the floating-point coprocessor causes an Invalid Operation exception whenever any floating-point operation receives a signaling NaN as an operand. You can use signaling NaNs in debugging, to track down some uses of uninitialized variables.

SIMD

Single-Instruction, Multiple-Data operation.

SPSR

See Saved Program Status Register.

Subnormal value

In floating-point operation, a value in the range $(-2^{E_{min}} < x < 2^{E_{min}})$, except for plus or minus 0. In the IEEE 754 standard format for single-precision and double-precision operands, a subnormal value has a zero exponent and a nonzero fraction field. The IEEE 754 standard requires that the generation and manipulation of subnormal operands be performed with the same precision as normal operands.

Support code

In a floating-point implementation, system software that complements the hardware VFP implementation to provide compatibility with the IEEE 754 standard. The support code has a library of routines that perform supported functions, such as divide with unsupported inputs or inputs that might generate an exception, in addition to operations beyond the scope of the hardware. The support code has a set of exception handlers to process exceptional conditions in compliance with the IEEE 754 standard.

SVC

See Supervisor Call.

SWI

See Supervisor Call.

Synchronization primitive

An instruction that is used to ensure memory synchronization, for example **LDREX** or **STREX**. See the *ARM Architecture Reference Manual* for more information.

Tag bits

In a cache implementation, bits $[31:(L+S)]$ of a virtual address, where $L = \log_2$ (cache line length) and $S = \log_2$ (number of cache sets). A cache hit occurs if the tag bits of the virtual address supplied by the processor match the tag bits associated with a valid line in the selected cache set.

See also Cache terminology diagram on the last page of this glossary.

TCM

See Tightly Coupled Memory.

Thumb instruction

One or two halfwords that specify an operation for a processor in Thumb state to perform. Thumb instructions must be halfword-aligned.

See also Thumb state, ThumbEE state.

Thumb state

In Thumb state the processor executes the Thumb instruction set.

ThumbEE state

In ThumbEE state the processor executes a variation of the Thumb instruction set specifically targeted for use with dynamic compilation techniques associated with an execution environment.

See also ARM state, Jazelle state, Thumb state.

Tightly Coupled Memory (TCM)

An area of low latency memory that provides predictable instruction execution or data load timing in cases where deterministic performance is required. TCMs are suited to holding:

- critical routines such as for interrupt handling
- scratchpad data

	• data types whose locality is not suited to caching • critical data structures, such as interrupt stacks.
Tiny	In a floating-point operation, a nonzero result or value that is between the positive and negative minimum normal values for the destination precision.
TLB	See Translation Lookaside Buffer.
Trace hardware	A term for a device that contains an ARM trace macrocell.
Translation Lookaside Buffer (TLB)	A memory structure containing the results of translation table walks. TLBs help to reduce the average cost of memory accesses. Usually, there is a TLB for each memory interface of the processor implementation.
Trigger instruction	A floating-point instruction that causes a bounce when it is issued. A potentially exceptional instruction causes the floating-point coprocessor to enter the exceptional state. A subsequent instruction, unless it is an FMXR or FMRX instruction accessing the FPEXC, FPINST, or FPSID register, causes a bounce, starting exception processing. The trigger instruction might not be exceptional, and is not processed. It is retried at the return from the exception processing of the potentially exceptional instruction. See also Bounce, Potentially exceptional instruction, and Exceptional state.
Unaligned	An unaligned access is an access where the address of the access is not aligned to the size of an element of the access.
Undefined	Indicates an instruction that generates an Undefined Instruction exception. See the <i>ARM Architecture Reference Manual</i> for more information.
Unknown	An Unknown value does not contain valid data, and can vary from moment to moment, instruction to instruction, and implementation to implementation. An Unknown value must not be a security hole.
UNP	See Unpredictable.
Unpredictable	For a processor means the behavior cannot be relied on. Unpredictable behavior must not represent security holes. Unpredictable behavior must not halt or hang the processor, or any parts of the system.
Unpredictable	For an ARM trace macrocell, means that the behavior of the macrocell cannot be relied on. Such conditions have not been validated. When applied to the programming of an event resource, only the output of that event resource is Unpredictable. Unpredictable behavior can affect the behavior of the entire system, because the trace macrocell can cause the processor to enter debug state, and external outputs can be used for other purposes.
VA	See Virtual Address.
VFP	A coprocessor extension to the ARM architecture that provides floating-point arithmetic. For ARMv7, more accurately described as the Floating-Point Extension.
Victim	A cache line, selected to be discarded to make room for a replacement cache line that is required because of a cache miss. The way that the victim is selected for eviction is processor-specific. A victim is also known as a cast out.
Virtual Address (VA)	An address generated by an ARM processor. For a <i>Protected Memory System Architecture</i> (PMSA) implementation, the virtual address is identical to the physical address.
WA	See Write-Allocate cache.

Warm reset	Also known as a core reset. Initializes the majority of the processor excluding the debug controller and debug logic. This type of reset is useful if you are using the debugging features of a processor.
Watchpoint	A debug event triggered by an access to memory, specified in terms of the address of the location in memory being accessed. <i>See also</i> Breakpoint.
Way	<i>See</i> Cache way.
WB	<i>See</i> Write-Back cache.
Word	A 32-bit data item. Words are normally word-aligned in ARM systems.
Word-aligned	A data item having a memory address that is divisible by four.
Word-invariant	In a word-invariant system, the address of each byte of memory changes when switching between little-endian and big-endian operation, in such a way that the byte with address A in one endianness has address A EOR 3 in the other endianness. As a result, each aligned word of memory always consists of the same four bytes of memory in the same order, regardless of endianness. The change of endianness occurs because of the change to the byte addresses, not because the bytes are rearranged. The ARM architecture supports word-invariant systems in ARMv3 and later versions. When word-invariant support is selected, the behavior of load or store instructions with unaligned addresses is instruction-specific, and is in general not the expected behavior for an unaligned access. ARM strongly recommends that word-invariant systems use the endianness that produces the required byte addresses at all times, apart possibly from very early in their reset handlers before they have set up the endianness, and that this early part of the reset handler uses only aligned word memory accesses. <i>See also</i> Byte-invariant.
Write	Operations that have the semantics of a store. See the <i>ARM Architecture Reference Manual</i> for more information.
Write buffer	A block of high-speed memory implemented to optimize stores to main memory.
Write interleave capability	The number of data-active write transactions for which the interface can transmit data. This is counted from the earliest transaction.
Write interleave depth	The number of data-active write transactions for which the interface can receive data.
Write-Allocate cache	A cache where a cache miss on storing data causes a cache line to be allocated and main memory contents to be read into it, followed by writing the stored data into the cache line.
Write-Back cache	A cache where, when a cache hit occurs on a store access, the data is only written to the cache. Data in the cache can therefore be more up-to-date than data in main memory. Any such data is written back to main memory when the cache line is cleaned or re-allocated. Also called <i>copy-back cache</i>.
Write-Through cache	A cache in which, when a cache hit occurs on a store access, the data is written both to the cache and to main memory. This is normally done using a write buffer, to avoid slowing down the processor.
WT	<i>See</i> Write-Through cache.