

Cortex™-A9

Revision: r2p2

Technical Reference Manual

ARM®

Cortex-A9

Technical Reference Manual

Copyright © 2008-2010 ARM. All rights reserved.

Release Information

The following changes have been made to this book.

Change history

Date	Issue	Confidentiality	Change
31 March 2008	A	Non-Confidential	First release for r0p0
08 July 2008	B	Non-Confidential Restricted Access	First release for r0p1
17 December 2008	C	Non-Confidential Restricted Access	First release for r1p0
30 September 2009	D	Non-Confidential Restricted Access	First release for r2p0
27 November 2009	E	Non-Confidential	Second release for r2p0
30 April 2010	F	Non-Confidential	First release for r2p2

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-A9 Technical Reference Manual

Preface

About this manual	xii
Conventions	xiv
Additional reading	xv
Feedback	xvii

Chapter 1

Introduction

1.1 About the Cortex-A9 processor	1-2
1.2 Cortex-A9 variants	1-4
1.3 Compliance	1-5
1.4 Features	1-6
1.5 Interfaces	1-7
1.6 Configurable options	1-8
1.7 Test features	1-9
1.8 Product documentation, design flow, and architecture	1-10
1.9 Product revisions	1-13

Chapter 2

Functional Description

2.1 About the functions	2-2
2.2 Interfaces	2-4
2.3 Clocking and resets	2-6
2.4 Power management	2-10
2.5 Constraints and limitations of use	2-15

Chapter 3

Programmers Model

3.1 About the programmers model	3-2
3.2 ThumbEE architecture	3-3
3.3 Advanced SIMD architecture	3-4
3.4 Security Extensions architecture	3-5

	3.5	Multiprocessing Extensions	3-6
	3.6	The Jazelle Extension	3-7
	3.7	Memory model	3-8
	3.8	Addresses in the Cortex-A9 processor	3-9
Chapter 4	System Control		
	4.1	About system control	4-2
	4.2	Register summary	4-3
	4.3	Register descriptions	4-8
Chapter 5	Jazelle DBX registers		
	5.1	About coprocessor CP14	5-2
	5.2	CP14 Jazelle register summary	5-3
	5.3	CP14 Jazelle register descriptions	5-4
Chapter 6	Memory Management Unit		
	6.1	About the MMU	6-2
	6.2	TLB Organization	6-4
	6.3	Memory Access Sequence	6-6
	6.4	MMU enabling or disabling	6-7
	6.5	External aborts	6-8
Chapter 7	Level 1 Memory System		
	7.1	About the L1 memory system	7-2
	7.2	Security Extensions support	7-4
	7.3	About the L1 instruction side memory system	7-5
	7.4	About the L1 data side memory system	7-8
	7.5	About DSB	7-9
	7.6	Data prefetching	7-10
	7.7	Parity error support	7-11
Chapter 8	Level 2 Memory Interface		
	8.1	Cortex-A9 L2 interface	8-2
	8.2	Optimized accesses to the L2 memory interface	8-7
	8.3	STRT instructions	8-9
Chapter 9	Preload Engine		
	9.1	About the Preload Engine	9-2
	9.2	PLE control register descriptions	9-3
	9.3	PLE operations	9-4
Chapter 10	Debug		
	10.1	About the debug interface	10-2
	10.2	About the Cortex-A9 debug interface	10-4
	10.3	Debug register descriptions	10-7
	10.4	Debug management registers	10-13
	10.5	External debug interface	10-16
Chapter 11	Performance Monitoring Unit		
	11.1	About the Performance Monitoring Unit	11-2
	11.2	PMU management registers	11-3
	11.3	Performance monitoring events	11-7
Appendix A	Signal Descriptions		
	A.1	Clock and clock control signals	A-2
	A.2	Resets and reset control	A-3
	A.3	Interrupts	A-4
	A.4	Configuration signals	A-5

A.5	Standby and Wait For Event signals	A-6
A.6	Power management signals	A-7
A.7	AXI interfaces	A-8
A.8	Performance monitoring signals	A-14
A.9	Exception flags signal	A-17
A.10	Parity signal	A-18
A.11	MBIST interface	A-19
A.12	Scan test signal	A-20
A.13	External Debug interface	A-21
A.14	PTM interface signals	A-24

Appendix B Instruction Cycle Timings

B.1	About instruction cycle timing	B-2
B.2	Data-processing instructions	B-3
B.3	Load and store instructions	B-4
B.4	Multiplication instructions	B-7
B.5	Branch instructions	B-8
B.6	Serializing instructions	B-9

Appendix C Revisions

Glossary

List of Tables

Cortex-A9 Technical Reference Manual

	Change history	ii
Table 1-1	Configurable options for the Cortex-A9 processor	1-8
Table 2-1	Reset modes	2-7
Table 2-2	Cortex-A9 processor power modes	2-10
Table 3-1	Address types in the processor system	3-9
Table 4-1	Summary of CP15 system control coprocessor registers	4-3
Table 4-2	c0 system control registers	4-8
Table 4-3	TLBTR bit assignments	4-9
Table 4-4	MPIDR bit assignments	4-11
Table 4-5	CCSIDR bit assignments	4-12
Table 4-6	CLIDR bit assignments	4-13
Table 4-7	CSSELR bit assignments	4-14
Table 4-8	c1 system control registers	4-15
Table 4-9	SCTLR bit assignments	4-16
Table 4-10	ACTLR bit assignments	4-19
Table 4-11	CPACR bit assignments	4-21
Table 4-12	SDER bit assignments	4-22
Table 4-13	NSACR bit assignments	4-23
Table 4-14	VCR bit assignments	4-25
Table 4-15	c2 system control registers	4-25
Table 4-16	c3 system control register	4-26
Table 4-17	c5 system control registers	4-26
Table 4-18	c6 system control registers	4-26
Table 4-19	c7 system control registers	4-27
Table 4-20	c8 system control registers	4-28
Table 4-21	c9 system control registers	4-28
Table 4-22	c10 system control registers	4-29
Table 4-23	TLB Lockdown Register bit assignments	4-29
Table 4-24	c11 system control registers	4-30
Table 4-25	PLEIDR bit assignments	4-31

Table 4-26	PLEASR bit assignments	4-31
Table 4-27	PLESFR bit assignments	4-32
Table 4-28	PLEUAR bit assignments	4-33
Table 4-29	PLEPCR bit assignments	4-34
Table 4-30	c12 system control registers	4-34
Table 4-31	Virtualization Interrupt Register bit assignments	4-35
Table 4-32	c13 system control registers	4-35
Table 4-33	c15 system control registers	4-36
Table 4-34	Power Control Register bit assignments	4-37
Table 4-35	Neon busy Register bit assignments	4-38
Table 4-36	TLB lockdown operations	4-39
Table 4-37	TLB VA Register bit assignments	4-40
Table 4-38	TLB PA Register bit assignments	4-40
Table 4-39	TLB Attributes Register bit assignments	4-41
Table 5-1	CP14 Jazelle registers summary	5-3
Table 5-2	JIDR bit assignments	5-4
Table 5-3	JOSCR bit assignments	5-6
Table 5-4	JMCR bit assignments	5-7
Table 5-5	Jazelle Parameters Register bit assignments	5-8
Table 5-6	Jazelle Configurable Opcode Translation Table Register bit assignments	5-9
Table 8-1	AXI master 0 interface attributes	8-2
Table 8-2	AXI master 1 interface attributes	8-2
Table 8-3	ARUSERM0[6:0] encodings	8-4
Table 8-4	ARUSERM1[6:0] encodings	8-5
Table 8-5	ARUSERM0[8:0] encodings	8-5
Table 8-6	Cortex-A9 mode and AxPROT values	8-9
Table 9-1	PLE program new channel operation bit assignments	9-5
Table 10-1	CP14 interface registers	10-5
Table 10-2	BVRs and corresponding BCRs	10-7
Table 10-3	Breakpoint Value Registers bit functions	10-7
Table 10-4	Breakpoint Control Registers bit assignments	10-8
Table 10-5	Meaning of BVR bits [22:20]	10-10
Table 10-6	WVRs and corresponding WCRs	10-10
Table 10-7	Watchpoint Value Registers bit functions	10-11
Table 10-8	Watchpoint Control Registers bit assignments	10-11
Table 10-9	Debug management registers	10-13
Table 10-10	Processor Identifier Registers	10-13
Table 10-11	Peripheral Identification Registers	10-14
Table 10-12	Component Identification Registers	10-15
Table 10-13	Authentication signal restrictions	10-16
Table 10-14	PMU register names and Debug APB interface addresses	10-18
Table 11-1	Performance monitoring instructions and Debug APB mapping	11-2
Table 11-2	PMU Management registers	11-3
Table 11-3	Processor Identifier Registers	11-3
Table 11-4	Peripheral Identification Registers	11-4
Table 11-5	Component Identification Registers	11-5
Table 11-6	PMU register names and APB addresses	11-5
Table 11-7	Cortex-A9 specific events	11-7
Table A-1	Clock and clock control signals for Cortex-A9	A-2
Table A-2	Cortex-A9 processor reset signals	A-3
Table A-3	Interrupt line signals	A-4
Table A-4	Configuration signals	A-5
Table A-5	CP15SDISABLE signal	A-5
Table A-6	Standby and wait for event signals	A-6
Table A-7	Power management signals	A-7
Table A-8	AXI-AW signals for AXI Master0	A-8
Table A-9	AXI-W signals for AXI Master0	A-9
Table A-10	AXI-B signals for AXI Master0	A-10
Table A-11	AXI-AR signals for AXI Master0	A-10
Table A-12	AXI-R signals for AXI Master0	A-11

Table A-13	AXI Master0 clock enable signal	A-11
Table A-14	AXI-AR signals for AXI Master1	A-12
Table A-15	AXI-R signals for AXI Master1	A-13
Table A-16	AXI Master1 clock enable signal	A-13
Table A-17	Performance monitoring signals	A-14
Table A-18	Event signals and event numbers	A-14
Table A-19	DEFLAGS signal	A-17
Table A-20	Parity signal	A-18
Table A-21	MBIST interface signals	A-19
Table A-22	MBIST signals with parity support implemented	A-19
Table A-23	MBIST signals without parity support implemented	A-19
Table A-24	Scan test signal	A-20
Table A-25	Authentication interface signals	A-21
Table A-26	APB interface signals	A-22
Table A-27	CTI signals	A-22
Table A-28	Miscellaneous debug signals	A-23
Table A-29	PTM interface signals	A-24
Table B-1	Data-processing instructions cycle timings	B-3
Table B-2	Single load and store operation cycle timings	B-4
Table B-3	Load multiple operations cycle timings	B-5
Table B-4	Store multiple operations cycle timings	B-6
Table B-5	Multiplication instruction cycle timings	B-7
Table C-1	Issue A	C-1
Table C-2	Differences between issue A and issue B	C-1
Table C-3	Differences between issue B and issue C	C-3
Table C-4	Differences between issue C and issue D	C-4
Table C-5	Differences between issue D and issue F	C-6

List of Figures

Cortex-A9 Technical Reference Manual

	Key to timing diagram conventions	xiv
Figure 1-1	Cortex-A9 uniprocessor system.	1-2
Figure 2-1	Cortex-A9 processor top-level diagram	2-2
Figure 2-2	PTM interface signals	2-4
Figure 2-3	ACLKENM0 used with a 3:1 clock ratio	2-6
Figure 2-4	Power domains for Cortex-A9 r2 designs	2-14
Figure 4-1	TLBTR bit assignments	4-9
Figure 4-2	MPIDR bit assignments	4-10
Figure 4-3	CCSIDR bit assignments	4-12
Figure 4-4	CLIDR bit assignments	4-13
Figure 4-5	CSSELR bit assignments	4-14
Figure 4-6	SCTLR bit assignments	4-16
Figure 4-7	ACTLR bit assignments	4-19
Figure 4-8	CPACR bit assignments	4-20
Figure 4-9	SDER bit assignments	4-22
Figure 4-10	NSACR bit assignments	4-23
Figure 4-11	VCR bit assignments	4-25
Figure 4-12	TLB Lockdown Register bit assignments	4-29
Figure 4-13	PLEIDR bit assignments	4-30
Figure 4-14	PLEASR bit assignments	4-31
Figure 4-15	PLESFR bit assignments	4-32
Figure 4-16	PLEUAR bit assignments	4-33
Figure 4-17	PLEPCR bit assignments	4-33
Figure 4-18	VIR bit assignments	4-35
Figure 4-19	Power Control Register bit assignments	4-37
Figure 4-20	NEON busy register bit assignments	4-37
Figure 4-21	Configuration Base Address Register bit assignments	4-38
Figure 4-22	Lockdown TLB index bit assignments	4-39
Figure 4-23	TLB VA Register bit assignments	4-39
Figure 4-24	Memory space identifier format	4-40

Figure 4-25	TLB PA Register bit assignments	4-40
Figure 4-26	Main TLB Attributes Register bit assignments	4-41
Figure 5-1	JIDR bit assignment	5-4
Figure 5-2	JOSCR bit assignments	5-5
Figure 5-3	JMCR bit assignments	5-6
Figure 5-4	Jazelle Parameters Register bit assignments	5-8
Figure 5-5	Jazelle Configurable Opcode Translation Table Register bit assignments	5-9
Figure 7-1	Branch prediction and instruction cache	7-5
Figure 7-2	Parity support	7-11
Figure 9-1	Program new channel operation bit assignments	9-5
Figure 10-1	Debug registers interface and CoreSight infrastructure	10-4
Figure 10-2	Breakpoint Control Registers bit assignments	10-8
Figure 10-3	Watchpoint Control Registers bit assignments	10-11
Figure 10-4	External debug interface signals	10-16
Figure 10-5	Debug request restart-specific connections	10-19

Preface

This preface introduces the *Cortex-A9 Technical Reference Manual (TRM)*. It contains the following sections:

- *About this manual* on page xii
- *Feedback* on page xvii

About this manual

This book is for the Cortex-A9 processor.

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 hardware and software engineers implementing Cortex-A9 system designs. It provides information that enables designers to integrate the processor into a target system.

Note

- The Cortex-A9 processor is a single core processor.
 - The multiprocessor variant, the Cortex-A9 MPCore™ processor, consists of between one and four Cortex-A9 processors and a *Snoop Control Unit* (SCU). See the *Cortex-A9 MPCore Technical Reference Manual* for a description.
-

Using this manual

This book is organized into the following chapters:

Chapter 1 *Introduction*

Read this for an introduction to the Cortex-A9 processor and descriptions of the major functional blocks.

Chapter 2 *Functional Description*

Read this for a description of the functionality of the Cortex-A9.

Chapter 3 *Programmers Model*

Read this for a description of the Cortex-A9 registers and programming details.

Chapter 4 *System Control*

Read this for a description of the Cortex-A9 system registers and programming details.

Chapter 5 *Jazelle DBX registers*

Read this for a description of the CP14 coprocessor and its non-debug use for Jazelle DBX.

Chapter 6 *Memory Management Unit*

Read this for a description of the Cortex-A9 *Memory Management Unit* (MMU) and the address translation process.

Chapter 7 *Level 1 Memory System*

Read this for a description of the Cortex-A9 level one memory system, including caches, *Translation Lookaside Buffers* (TLB), and store buffer.

Chapter 8 *Level 2 Memory Interface*

Read this for a description of the Cortex-A9 level two memory interface, the AXI interface attributes, and information about STRT instructions.

Chapter 9 *Preload Engine*

Read this for a description of the *Preload Engine* (PLE) and PLE operations.

Chapter 10 *Debug*

Read this for a description of the Cortex-A9 support for debug.

Chapter 11 *Performance Monitoring Unit*

Read this for a description of the Cortex-A9 *Performance Monitoring Unit* (PMU) and associated events.

Appendix A *Signal Descriptions*

Read this for a summary of the Cortex-A9 signals.

Appendix B *Instruction Cycle Timings*

Read this for a description of the Cortex-A9 instruction cycle timing.

Appendix C *Revisions*

Read this for a description of 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*
- *Signals* on page xv.

Typographical


The typographical conventions are:

<i>italic</i>	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.
monospace	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: <ul style="list-style-type: none"> • MRC p15, 0 <Rd>, <CRn>, <CRm>, <opc2>

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

Additional reading

This section lists publications by ARM and by third parties.

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

ARM publications

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

- *ARM Architecture Reference Manual, ARMv7-A and ARMv7-R edition* (ARM DDI 0406)
- *Cortex™-A9 MPCore Technical Reference Manual* (ARM DDI 0407)
- *Cortex-A9 Floating-Point Unit (FPU) Technical Reference Manual* (ARM DDI 0408)
- *Cortex-A9 NEON® Media Processing Engine Technical Reference Manual* (ARM DDI 0409)
- *Cortex-A9 Configuration and Sign-Off Guide* (ARM DII 00146)
- *Cortex-A9 MBIST Controller Technical Reference Manual* (ARM DDI 0414)
- *CoreSight™ PTM™-A9 Technical Reference Manual* (ARM DDI 0401)
- *CoreSight PTM-A9 Integration Manual* (ARM DII 0162)
- *CoreSight Program Flow Trace™ Architecture Specification, v1.0* (ARM IHI 0035)
- *AMBA® Level 2 Cache Controller (L2C-310) Technical Reference Manual* (ARM DDI 0246)
- *AMBA AXI Protocol v10 Specification* (ARM IHI 0022)
- *ARM Generic Interrupt Controller Architecture Specification* (ARM IHI 0048)
- *PrimeCell® Generic Interrupt Controller (PL390) Technical Reference Manual* (ARM DDI 0416)
- *RealView ICE User Guide* (ARM DUI 0155)
- *CoreSight Architecture Specification* (ARM IHI 0029)
- *CoreSight Technology System Design Guide* (ARM DGI 0012)
- *ARM Debug Interface v5 Architecture Specification* (ARM IHI 0031)
- *The ARM Cortex-A9 Processors White paper.*

Other publications

This section lists relevant documents published by third parties:

- *ANSI/IEEE Std 754-1985, IEEE Standard for Binary Floating-Point Arithmetic*
- *IEEE Std 1500-2005, IEEE Standard Testability Method for Embedded Core-based Integrated Circuits.*

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 if appropriate.

Feedback on this book

If you have any comments on this book, send e-mail to errata@arm.com. Give:

- the title
- the number, ARM DDI 0388F
- the relevant page number(s) 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 Cortex-A9 processor and its features. It contains the following sections:

- *About the Cortex-A9 processor* on page 1-2
- *Cortex-A9 variants* on page 1-4
- *Compliance* on page 1-5
- *Features* on page 1-6
- *Interfaces* on page 1-7
- *Configurable options* on page 1-8
- *Test features* on page 1-9
- *Product documentation, design flow, and architecture* on page 1-10
- *Product revisions* on page 1-13.

1.1 About the Cortex-A9 processor

The Cortex-A9 processor is a high-performance, low-power, ARM macrocell with an L1 cache subsystem that provides full virtual memory capabilities. The Cortex-A9 processor implements the ARMv7-A architecture and runs 32-bit ARM instructions, 16-bit and 32-bit Thumb instructions, and 8-bit Java™ bytecodes in Jazelle state.

Figure 1-1 shows a Cortex-A9 uniprocessor in a design with a PL390 Interrupt Controller and an L2C-310 L2 Cache Controller,

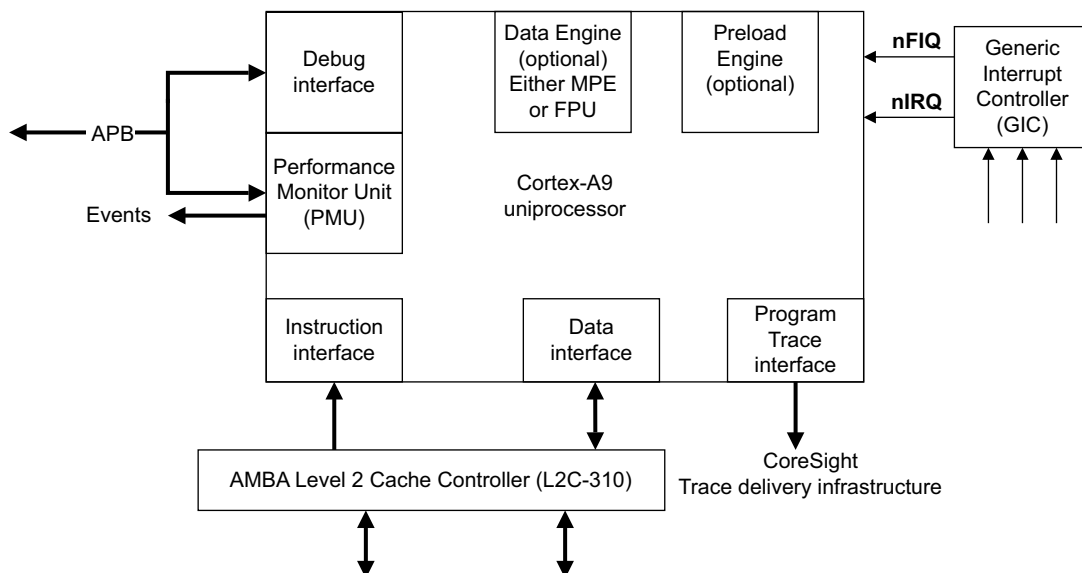


Figure 1-1 Cortex-A9 uniprocessor system.

1.1.1 Data Engine

The design can include a Data Engine. The following sections describe the Data Engine options:

- *Media Processing Engine*
- *Floating-Point Unit.*

Media Processing Engine

The optional *NEON Media Processing Engine* (MPE) is the ARM *Advanced Single Instruction Multiple Data* (SIMD) media processing engine extension to the ARMv7-A architecture. It provides support for integer and floating-point vector operations. NEON MPE can accelerate the performance of multimedia applications such as 3-D graphics and image processing.

When implemented, the NEON MPE option extends the processor functionality to provide support for the ARMv7 Advanced SIMD and VFPv3 D-32 instruction sets.

See the *Cortex-A9 NEON Media Processing Engine Technical Reference Manual*.

Floating-Point Unit

When the design does not include the optional MPE, you can include the optional ARMv7 VFPv3-D16 FPU, without the Advanced SIMD extensions. It provides trapless execution and is optimized for scalar operation. The Cortex-A9 FPU hardware does not support the deprecated VFP short vector feature. Attempts to execute VFP data-processing instructions when the

FPSCR.LEN field is non-zero result in the FPSCR.DEX bit being set and a synchronous Undefined instruction exception being taken. You can use software to emulate the short vector feature, if required.

See the *Cortex-A9 Floating-Point Unit Technical Reference Manual*.

1.1.2 System design components

This section describes the PrimeCell components in Figure 1-1 on page 1-2 in the following sections:

- *PrimeCell Generic Interrupt Controller*
- *AMBA Level 2 Cache Controller (L2C-310)*.

PrimeCell Generic Interrupt Controller

A generic interrupt controller such as the *PrimeCell Generic Interrupt Controller (PL390)* can be attached to the Cortex-A9 uniprocessor. The Cortex-A9 MPCore contains an integrated interrupt controller that shares the same programmers model as the PL390 although there are implementation-specific differences.

See the *Cortex-A9 MPCore Technical Reference Manual* for a description of the Cortex-A9 MPCore Interrupt Controller.

AMBA Level 2 Cache Controller (L2C-310)

The addition of an on-chip secondary cache, also referred to as a Level 2 or L2 cache, is a recognized method of improving the performance of ARM-based systems when significant memory traffic is generated by the processor. The AMBA Level 2 Cache Controller reduces the number of external memory accesses and has been optimized for use with Cortex-A9 processors and Cortex-A9 MPCore processors.

1.2 Cortex-A9 variants

Cortex-A9 processors can be used in both a uniprocessor configuration and multiprocessor configurations.

In the multiprocessor configuration, up to four Cortex-A9 processors are available in a cache-coherent cluster, under the control of a *Snoop Control Unit* (SCU), that maintains L1 data cache coherency.

The Cortex-A9 MPCore multiprocessor has:

- up to four Cortex-A9 processors
- an SCU responsible for:
 - maintaining coherency among L1 data caches
 - *Accelerator Coherency Port* (ACP) coherency operations
 - routing transactions on Cortex-A9 MPCore AXI master interfaces
 - Cortex-A9 uniprocessor accesses to private memory regions.
- an *Interrupt Controller* (IC) with support for legacy ARM interrupts
- a private timer and a private watchdog per processor
- a global timer
- AXI high-speed *Advanced Microprocessor Bus Architecture* (AMBA3) L2 interfaces.
- an *Accelerator Coherency Port* (ACP), that is, an optional AXI 64-bit slave port that can be connected to a DMA engine or a noncached peripheral.

See the *Cortex-A9 MPCore Technical Reference Manual* for more information.

The following system registers have Cortex-A9 MPCore uses:

- *Multiprocessor Affinity Register* on page 4-10
- *Auxiliary Control Register* on page 4-18
- *Configuration Base Address Register* on page 4-38.

Some PMU event signals have Cortex-A9 MPCore uses. See *Performance monitoring signals* on page A-14.

1.3 Compliance

The Cortex-A9 processor implements the ARMv7-A architecture that includes the following features:

- ARM Thumb[®]-2 32-bit instruction set architecture for overall code density comparable with Thumb and performance comparable with ARM instructions. See the *ARM Architecture Reference Manual* for information on both the ARM and Thumb instruction sets.
- *Thumb Execution Environment* (ThumbEE) architecture for execution environment acceleration. See the *ARM Architecture Reference Manual* for information on the ThumbEE instruction set.
- Security Extensions for enhanced security. See the *ARM Architecture Reference Manual* for information on Security Extensions..
- Advanced SIMD architecture extension to accelerate the performance of multimedia applications such as 3-D graphics and image processing. See the *ARM Architecture Reference Manual* for information on the Advanced SIMD architecture extension . See the *Cortex-A9 NEON Media Processing Engine Technical Reference Manual* for implementation-specific information.
- *Vector Floating-Point version 3* (VFPv3) architecture extension for floating-point computation that is fully compliant with the IEEE 754 standard. See the *ARM Architecture Reference Manual* for information on the VFPv3 extension. See the *Cortex-A9 Floating-Point Unit Technical Reference Manual* for implementation-specific information.
- ARMv7 Debug architecture that includes support for Security Extensions and CoreSight. See the *ARM Architecture Reference Manual* for information on the ARMv7 Debug architecture.

1.4 Features

The Cortex-A9 processor includes the following features :

- superscalar, variable length, out-of-order pipeline with dynamic branch prediction
- full implementation of the ARM architecture v7-A instruction set
- Security Extensions
- Harvard level 1 memory system with *Memory Management Unit* (MMU).
- two 64-bit AXI master interfaces with Master 0 for the data side bus and Master 1 for the instruction side bus
- ARMv7 Debug architecture
- support for trace with the *Program Trace Macrocell* (PTM) interface
- support for advanced power management with up to 3 power domains
- optional Preload Engine
- optional Jazelle hardware acceleration
- optional Data Engine with MPE and VFPv3.

1.5 Interfaces

The processor has the following external interfaces:

- AMBA AXI interfaces
- v7 compliant debug interface, including a debug APBv3 external debug interface
- DFT.

See the *AMBA AXI Protocol Specification*, the *CoreSight Architecture Specification*, and the *Cortex-A9 MBIST Controller Technical Reference Manual* for more information on these interfaces.

1.6 Configurable options

Table 1-1 shows the configurable options for the Cortex-A9 processor.

Table 1-1 Configurable options for the Cortex-A9 processor

Feature	Range of options	Default value
Instruction cache size	16KB, 32KB, or 64KB	32KB
Data cache size	16KB, 32KB, or 64KB	32KB
TLB entries	64 entries or 128 entries	128 entries
Jazelle Architecture Extension	Full or trivial	Full
Media Processing Engine with NEON technology	Included or not ^a	Not included
FPU	Included or not ^a	
PTM interface	Included or not	
Wrappers for power off and dormant modes	Included or not	
Support for parity error detection	-	Inclusion of this feature is a configuration and design decision.
Preload Engine	Included or not	
Preload Engine FIFO size ^b	16, 8, or 4 entries	16 entries
ARM_BIST	Included or not	Included
USE DESIGNWARE	Use or not	Use

- a. The MPE and FPU RTL options are mutually exclusive. If you choose the MPE option, the MPE is included along with its VFPv3-D32 FPU, and the FPU RTL option is not available in this case. When the MPE RTL option is not implemented, you can implement the VFPv3-D16 FPU by choosing the FPU RTL option.
- b. Only when the design includes the Preload Engine.

The MBIST solution must be configured to match the chosen Cortex-A9 cache sizes. In addition, the form of the MBIST solution for the RAM blocks in the Cortex-A9 design must be determined when the processor is implemented.

For details, see the *Cortex-A9 MBIST Controller Technical Reference Manual*.

1.7 Test features

There are no test features for the Cortex-A9 processor.

1.8 Product documentation, design flow, and architecture

This section describes the Cortex-A9 family books, how they relate to the design flow, and the relevant architectural standards and protocols.

See *Additional reading* on page xv for more information about the books described in this section.

1.8.1 Documentation

The Cortex-A9 family 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-A9 family. It is required at all stages of the design flow. Some behavior described in the TRM might not be relevant because of the way that the Cortex-A9 processor is implemented and integrated.

- the *Cortex-A9 TRM* describes the uniprocessor variant.
- the *Cortex-A9 MPCore TRM* describes the multiprocessor variant of the Cortex-A9 processor.
- the *Cortex-A9 Floating-Point Unit (FPU) TRM* describes the implementation-specific FPU parts of the Data Engine.
- the *Cortex-A9 NEON Media Processing Engine TRM* describes the Advanced SIMD Cortex-A9 implementation-specific parts of the Data Engine.

If you are programming the Cortex-A9 processor then contact:

- the implementer to determine the build configuration of the implementation
- the integrator to determine the pin configuration of the SoC 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) description 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 documentation from your EDA tools vendor complements the CSG.

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

1.8.2 Design flow

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

1. Implementation. The implementer configures and synthesizes the RTL to produce a hard macrocell. If appropriate, this includes integrating the RAMs into the design.

2. Integration. The integrator connects the implemented design into a SoC, This includes connecting it to a memory system and peripherals.
3. Programming. The system programmer develops the software required to configure and initialize the processor, and tests the required application software.

Each stage of the process:

- can be performed by a different party
- can include options that affect the behavior and features at the next stage:

Build configuration

The implementer chooses the options that affect how the RTL source files are pre-processed. They usually include or exclude logic that can affect the area or maximum frequency of the resulting macrocell.

Configuration inputs

The integrator configures some features of the 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.

Software configuration

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

Note

This manual refers to implementation-defined features that are applicable to build configuration options. References to a feature that is included mean that the appropriate build and pin configuration options have been selected, while references to an enabled feature mean one that has also been configured by software.

1.8.3 Architecture and protocol information

The Cortex-A9 processor complies with, or implements, the specifications described in:

- *ARM architecture*
- *Advanced Microcontroller Bus Architecture.*
- *Trace macrocell*
- *Debug Architecture* on page 1-12.

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

ARM architecture

The Cortex-A9 processor implements the ARMv7-A architecture profile. See the *ARM Architecture Reference Manual*.

Advanced Microcontroller Bus Architecture

This Cortex-A9 processor complies with the AMBA 3 protocol. See the *AMBA AXI Protocol Specification* and the *AMBA 3 APB Protocol Specification*.

Trace macrocell

The v1.0 PFT architecture. See the *CoreSight Program Flow Trace Architecture Specification*.

Debug Architecture

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

1.9 Product revisions

This section summarizes the differences in functionality between the different releases of this processor:

- *Differences in functionality between r0p0 and r0p1*
- *Differences in functionality between r0p1 and r1p0*
- *Differences in functionality between r1p0 and r2p0* on page 1-14.

1.9.1 Differences in functionality between r0p0 and r0p1

There is no change in the described functionality between r0p0 and r0p1.

The only differences between the two revisions are:

- r0p1 includes fixes for all known engineering errata relating to r0p0
- r0p1 includes an upgrade of the micro TLB entries from 8 to 32 entries, on both the Instruction and Data side.

Neither of these changes affect the functionality described in this document.

1.9.2 Differences in functionality between r0p1 and r1p0

The differences between the two revisions are:

- r1p0 includes fixes for all known engineering errata relating to r0p1.
- In r1p0 **CPUCLKOFF** and **DECLKOFF** enable control of Cortex-A9 processors during reset sequences. See *Configuration signals* on page A-5
 - In a multiprocessor implementation of the design there are as many **CPUCLKOFF** pins as there are Cortex-A9 processors.
 - **DECLKOFF** controls the Data Engine clock during reset sequences.
- r1p0 includes dynamic high level clock gating of the Cortex-A9 processor. See *Dynamic high level clock gating* on page 2-8
 - **MAXCLKLATENCY[2:0]** bus added. See *Configuration signals* on page A-5
 - Addition of CP15 power control register. See *Power Control Register* on page 4-36
- Extension of the Performance Monitoring Event bus. In r1p0, **PMUEVENT** is 52 bits wide:
 - Addition of Cortex-A9 specific events. See Table 2-2 on page 2-5.
 - Event descriptions extended. See Table 2-2 on page 2-5.
- Addition of **PMUSECURE** and **PMUPRIV**. See *Performance monitoring signals* on page A-14.
- Main TLB options for 128 entries or 64 entries. See *TLB Type Register* on page 4-9.
- **DEFLAGS[6:0]** added. See *DEFLAGS[6:0]* on page 4-37
- The power management signal **BISTSCLAMP** is removed.
- The scan test signal **SCANTEST** is removed.
- Addition of a second replacement strategy. Selection done by SCTLRRR bit. See *System Control Register* on page 4-15.

- Addition of PL310 cache controller optimization description. See *Optimized accesses to the L2 memory interface* on page 8-7.
- Change to the serializing behavior of DMB. See *Serializing instructions* on page B-9.
- ID Register values changed to reflect correct revision.

1.9.3 Differences in functionality between r1p0 and r2p0

The differences between the revisions are:

- Addition of optional Preload Engine hardware feature and support.
 - PLE bit added to NSACR. See *Non-secure Access Control Register* on page 4-23.
 - Preload Engine registers added. See *CP15 c11 register summary* on page 4-30.
 - Preload operations added and MCRR instruction added. See Chapter 9 *Preload Engine*.
 - Addition of Preload Engine events.
See *Performance monitoring* on page 2-3, Table 11-7 on page 11-7, and Table A-18 on page A-14.
- Change to voltage domains. See Figure 2-4 on page 2-14
- NEON busy register. See *NEON busy Register* on page 4-37
- ID Register values changed to reflect correct revision.

1.9.4 Differences in functionality between r2p0 and r2p1

- None.

1.9.5 Differences in functionality between r2p1 and r2p2

- None. Documentation updates and corrections only. See *Differences between issue D and issue F* on page C-6.

Chapter 2

Functional Description

This chapter describes the functionality of the product. It contains the following sections:

- *About the functions* on page 2-2
- *Interfaces* on page 2-4
- *Clocking and resets* on page 2-6
- *Power management* on page 2-10
- *Constraints and limitations of use* on page 2-15.

2.1 About the functions

The Cortex-A9 processor is a high-performance, low-power, ARM macrocell with an L1 cache subsystem that provides full virtual memory capabilities.

Figure 2-1 shows a top-level diagram of the Cortex-A9 processor.

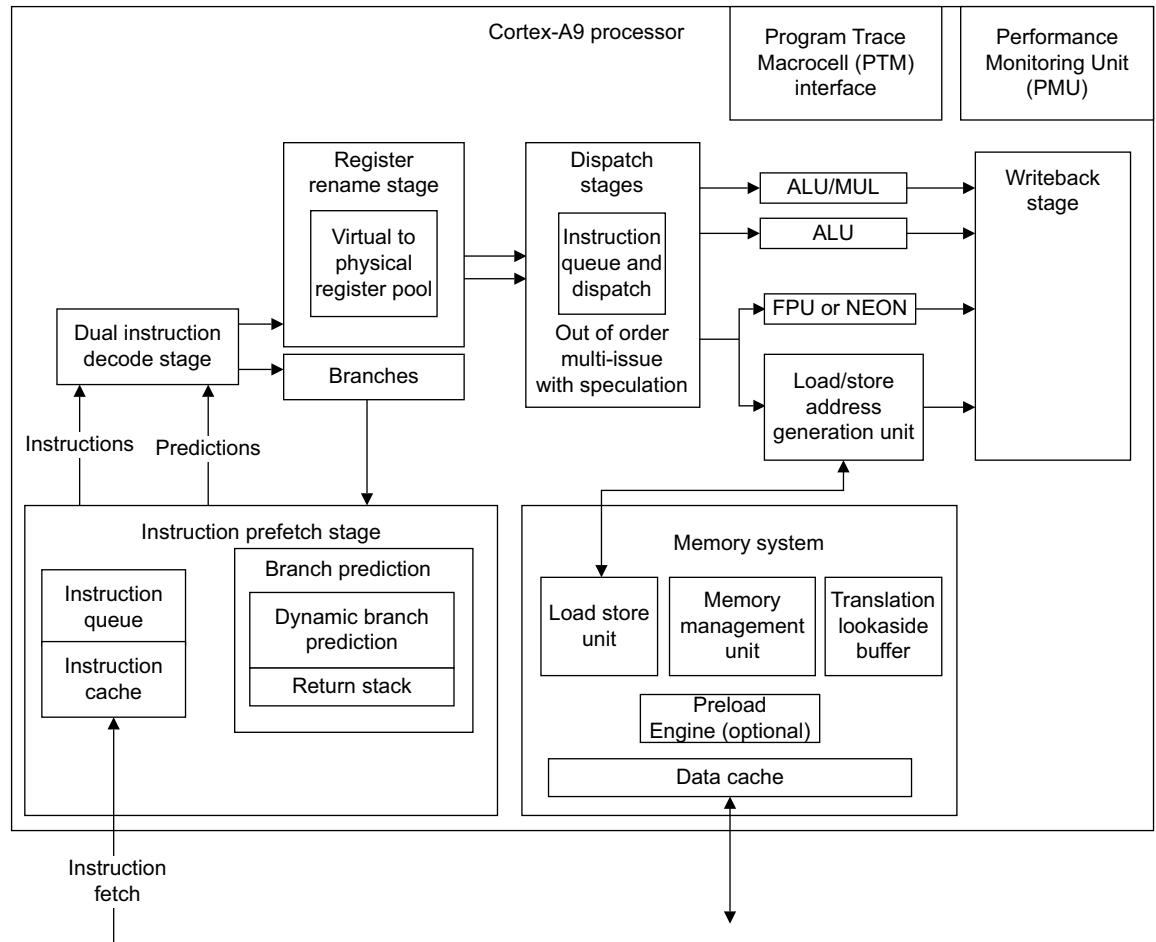


Figure 2-1 Cortex-A9 processor top-level diagram

2.1.1 Register renaming

The register renaming scheme facilitates out-of-order execution in *Write-after-Write* (WAW) and *Write-after-Read* (WAR) situations for the general purpose registers and the flag bits of the Current Program Status Register (CPSR).

The scheme maps the 32 ARM architectural registers to a pool of 56 physical 32-bit registers, and renames the flags (N, Z, C, V, Q, and GE) of the CPSR using a dedicated pool of eight physical 9-bit registers.

2.1.2 Instruction queue

In the instruction queue small loop mode provides low power operation while executing small instruction loops. See *Energy efficiency features* on page 2-10.

2.1.3 Dynamic branch prediction

The Prefetch Unit implements two-level dynamic branch prediction with a *Global History Buffer* (GHB), a *Branch Target Address Cache* (BTAC) and a return stack. See *About the L1 instruction side memory system* on page 7-5.

2.1.4 PTM interface

The Cortex-A9 processor optionally implements a *Program Trace Macrocell* (PTM) interface, which is compliant with the *Program Flow Trace* (PFT) instruction-only architecture protocol. Waypoints, changes in the program flow or events such as changes in context ID, are output to enable the trace to be correlated with the code image. See *Program Flow Trace and the Program Trace Macrocell interface* on page 2-4.

2.1.5 Performance monitoring

The Cortex-A9 processor provides program counters and event monitors that can be configured to gather statistics on the operation of the processor and the memory system.

You can access performance monitoring counters and their associated control registers from the CP15 coprocessor interface and from the APB Debug Interface. See Chapter 11 *Performance Monitoring Unit*

2.1.6 Virtualization of interrupts

With virtualized interrupts a guest *Operating System* (OS) can use a modified version of the exception behavior model to speed up handling of interrupts

See *Virtualization Control Register* on page 4-24.

The behavior of the Virtualization Control Register depends on whether the processor is in Secure or Non-Secure state.

If the exception occurs when the processor is in Secure state the AMO, IMO and IFO bits in the Virtualization Control Register are ignored. Whether the exception is taken or not depends solely on the setting of the CPSR A, I, and F bits.

If the exception occurs when the processor is in Non-secure state if the SCR EA bit, FIQ bit, or IRQ bit is not set, whether the corresponding exception is taken or not depends solely on the setting of the CPSR A, I, and F bits.

See *Non-secure Access Control Register* on page 4-23.

If the SCR.EAbit, FIQ bit or IRQ bit is set, then the corresponding exception is trapped to Monitor mode. In this case, the corresponding exception is taken or not depending on the CPSR.A bit, I bit, or F bits masked by the AMO, IMO, or IFO bits in the Virtualization Control Register.

2.2 Interfaces

The processor has the following external interfaces:

- AMBA AXI interfaces
- APB CoreSight interface
- DFT interface.

See the *AMBA AXI Protocol Specification*, the *CoreSight Architecture Specification*, the *CoreSight PFT Architecture Specification*, and the *Cortex-A9 MBIST Controller Technical Reference Manual* for more information on these interfaces.

2.2.1 Program Flow Trace and the Program Trace Macrocell interface

In addition, the Cortex-A9 processor implements the *Program Flow Trace* (PFT) architecture protocol. The following sections describe the Cortex-A9 *Program Trace Macrocell* (PTM) interface:

- *Program Flow Trace*
- *Program Trace Macrocell signals*.

Program Flow Trace

PFT is an instruction-only trace protocol that uses waypoints to correlate the trace to the code image. Waypoints are changes in the program flow or events such as branches or changes in context ID that must be output to enable the trace.

See the *CoreSight Program Flow Trace Architecture Specification* and the *CoreSight PTM-A9 Technical Reference Manual* for more information about tracing with waypoints.

Program Trace Macrocell signals

Figure 2-2 shows the PTM interface signals.

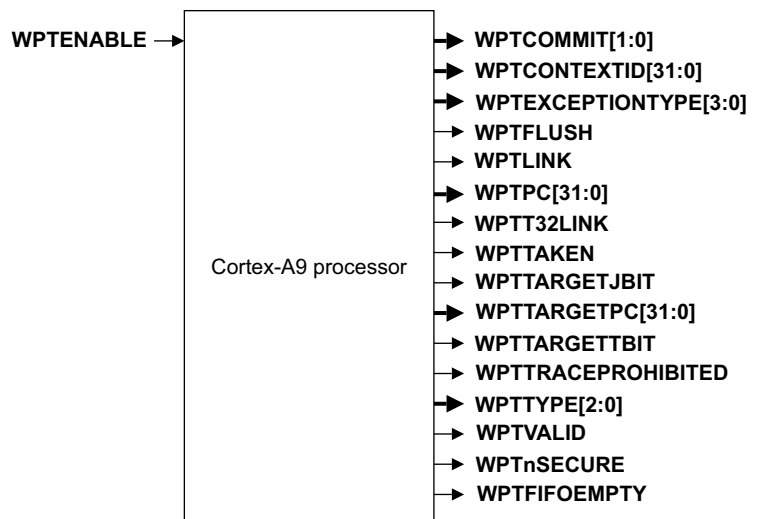


Figure 2-2 PTM interface signals

See Appendix A *Signal Descriptions* and the *CoreSight PTM-A9 Technical Reference Manual* for more information.

Prohibited regions

Trace must be disabled in some regions. The prohibited regions are described in the *ARM Architecture Reference Manual*. The Cortex-A9 processor must determine prohibited regions for non-invasive debug in regions, including trace, performance monitoring, and PC sampling. No waypoints are generated for instructions that are within a prohibited region.

Only entry to and exit from Jazelle state are traced. A waypoint to enter Jazelle state is followed by a waypoint to exit Jazelle state.

2.3 Clocking and resets

This section describes the clocks and resets of the processor in:

- *Synchronous clocking*
- *Reset*
- *Dynamic high level clock gating* on page 2-8.

2.3.1 Synchronous clocking

The Cortex-A9 uniprocessor has one functional clock input, **CLK**.

The Cortex-A9 uniprocessor does not have any asynchronous interfaces. All the bus interfaces and the interrupt signals must be synchronous with reference to **CLK**.

The AXI bus clock domain can be run at n:1 (AXI: processor ratio to **CLK**) using the **ACLKEN** signal.

Figure 2-3 shows a timing example with **ACKLENM0** used with a 3:1 clock ratio between **CLK** and **ACLK** in a Cortex-A9 uniprocessor .




Figure 2-3 ACLKENM0 used with a 3:1 clock ratio

The master port, Master0, changes the AXI outputs only on the **CLK** rising edge when **ACLKENM0** is HIGH.

2.3.2 Reset

The Cortex-A9 processor has the following reset inputs:

- nCPURESET** The **nCPURESET** signal is the main Cortex-A9 processor reset. It initializes the Cortex-A9 processor logic and the FPU logic including the FPU register file when the MPE or FPU option is present.
- nNEONRESET** The **nNEONRESET** signal is the reset that controls the NEON SIMD independently of the main Cortex-A9 processor reset.
- nDBGRESET** The **nDBGRESET** signal is the reset that initializes the debug logic. See Chapter 10 *Debug*.

All of these are active-LOW signals.

Reset modes

The reset signals present in the Cortex-A9 design enable you to reset different parts of the processor 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

Mode	nCPURESET	nNEONRESET	nDBGRESET
Power-on reset, cold reset	0	0	0
Processor reset, soft or warm reset	0	0	1
SIMD MPE power-on reset	1	0	1
Debug logic reset	1	1	0
No reset, normal run mode	1	1	1

Power-on reset

You must apply power-on or *cold* reset to the Cortex-A9 uniprocessor when power is first applied to the system. In the case of power-on reset, the leading edge, that is the falling edge, of the reset signals do not have to be synchronous to **CLK**, but the rising edge must be.

You must assert the reset signals for at least nine **CLK** cycles to ensure correct reset behavior.

ARM recommends the following reset sequence:

1. Apply **nCPURESET** and **nDBGRESET**, plus **nNEONRESET** if the SIMD MPE is present.
2. Wait for at least nine **CLK** cycles, plus at least one cycle in each other clock domain, or more if the documentation for other components requires it. There is no harm in applying more clock cycles than this, and maximum redundancy can be achieved by applying 15 cycles on every clock domain.
3. Stop the **CLK** clock input to the Cortex-A9 uniprocessor. If there is a Data Engine present, use **NEONCLKOFF**. See *Configuration signals* on page A-5.
4. Wait for the equivalent of approximately 10 cycles, depending on your implementation. This compensates for clock and reset tree latencies.
5. Release all resets.
6. Wait for the equivalent of another approximately 10 cycles, again to compensate for clock and reset tree latencies.
7. Restart the clock.

Software reset

A processor or warm reset initializes the majority of the Cortex-A9 processor, apart from its debug logic. Breakpoints and watchpoints are retained during a processor reset. Processor reset is typically used for resetting a system that has been operating for some time.

Use the same reset sequence described in *Power-on reset* with the only difference that **nDBGRESET** must remain HIGH during the sequence, to ensure that all values in the debug registers are maintained.

Processor reset

A processor or *warm* reset initializes the majority of the Cortex-A9 processor, apart from its debug logic. Breakpoints and watchpoints are retained during a processor reset. Processor reset is typically used for resetting a system that has been operating for some time. Use **nCPURESET** and **nNEONRESET** for a warm reset.

MPE SIMD logic reset

This reset initializes all the SIMD logic of the MPE. It is expected to be applied when the SIMD part of the MPE exits from powerdown state.

This reset only applies to configurations where the SIMD MPE logic is implemented in its own dedicated power domain, separated from the rest of the processor logic.

ARM recommends the following reset sequence for an MPE SIMD reset:

1. Apply **nNEONRESET**.
2. Wait for at least nine CLK cycles. . There is no harm in applying more clock cycles than this, and maximum redundancy can be achieved by for example applying 15 cycles on every clock domain.
3. Assert **NEONCLKOFF** with a value of 1'b1.
4. Wait for the equivalent of approximately 10 cycles, depending on your implementation. This compensates for clock and reset tree latencies.
5. Release **nNEONRESET**.
6. Wait for the equivalent of another approximately 10 cycles, again to compensate for clock and reset tree latencies.
7. Deassert **NEONCLKOFF**. This ensures that all registers in the SIMD MPE part of the processor see the same CLK edge on exit from the reset sequence.

Use **nNEONRESET** to control the SIMD part of the MPE logic independently of the Cortex-A9 processor reset. Use this reset to hold the SIMD part of the MPE in a reset state so that the power to the SIMD part of the MPE can be safely switched on or off. See Table 2-2 on page 2-10.

Debug reset

This reset initializes the debug logic in the Cortex-A9 uniprocessor, including breakpoints and watchpoints values.

To perform a debug reset, you must assert the **nDBGRESET** signal LOW during a few CLK cycles.

2.3.3 Dynamic high level clock gating

The following sections describe dynamic high level clock gating:

- *Gated blocks* on page 2-9
- *Power Control Register* on page 2-9
- *Effects of max_clk latency bits* on page 2-9
- *Dynamic high level clock gating activity* on page 2-9.

Gated blocks

The Cortex-A9 processor or each processor in a CortexA9MP Core design supports dynamic high level clock gating of:

- the integer core
- the system control block.
- the Data Engine, if implemented.

Power Control Register

The Power Control Register controls dynamic high level clock gating. This register contains fields that are common to these blocks:

- the enable bit for clock gating
- the max_clk latency bits.

See *Power Control Register* on page 4-36.

Effects of max_clk latency bits

The max_clk latency bits determine the length of the delay between when one of these blocks has its clock cut and the time when it can receive new active signals.

If the value determined by max_clk latency is lower than the real delay, the block that had its clock cut can receive active signals even though it does not have a clock. This can cause the device to malfunction.

If the value determined by max_clk latency is higher than the real delay, the master block waits extra cycles before sending its signals to the block that had its clock cut. This can have some performance impact.

When the value is correctly set, the block that had its clock cut receives active signals on the first clock edge of the wake-up. This gives optimum performance.

Dynamic high level clock gating activity

When dynamic high level clock gating is enabled the clock of the integer core is cut in the following cases:

- the integer core is empty and there is an instruction miss causing a linefill
- the integer core is empty and there is an instruction TLB miss
- the integer core is full and there is a data miss causing a linefill
- the integer core is full and data stores are stalled because the linefill buffers are busy.

When dynamic clock gating is enabled, the clock of the system control block is cut in the following cases:

- there are no system control coprocessor instructions being executed
- there are no system control coprocessor instructions present in the pipeline
- performance events are not enabled
- debug is not enabled.

When dynamic clock gating is enabled, the clock of the Data Engine is cut when there is no Data Engine instruction in the Data Engine and no Data Engine instruction in the pipeline.

2.4 Power management

The processor provides mechanisms to control both dynamic and static power dissipation. Static power control is implementation-specific. The following sections describe:

- Energy efficiency features
- Cortex-A9 processor power control.

2.4.1 Energy efficiency features

The features of the Cortex-A9 processor that improve energy efficiency include:

- accurate branch and return prediction, reducing the number of incorrect instruction fetch and decode operations
- the use of physically addressed caches, reducing the number of cache flushes and refills, saving energy in the system
- the use of micro TLBs reduces the power consumed in translation and protection look-ups for each cycle
- caches that use sequential access information to reduce the number of accesses to the tag RAMs and to unnecessary accesses to data RAMs
- instruction loops that are smaller than 64 bytes often complete without additional instruction cache accesses, so lowering power consumption.

2.4.2 Cortex-A9 processor power control

Place holders for level-shifters and clamps are inserted around the Cortex-A9 processor to ease the implementation of different power domains.

The Cortex-A9 processor can have the following power domains:

- a power domain for Cortex-A9 processor logic
- a power domain for Cortex-A9 processor MPE.
- a power domain for Cortex-A9 processor RAMs.

Table 2-2 shows the power modes.

Table 2-2 Cortex-A9 processor power modes

Mode	Cortex-A9 processor RAM arrays	Cortex-A9 processor logic	Cortex-A9 Data Engine	Comments
Full Run Mode	Powered-up	Powered-up	Powered-up	-
		Clocked	Clocked	
Run Mode with MPE disabled	Powered-up	Powered-up	Powered-up	See <i>Coprocessor Access Control Register</i> on page 4-20 for information about disabling the MPE.
		Clocked	No clock	
Run Mode with MPE powered off	Powered-up	Powered-up	Powered off	The MPE can be implemented in a separate power domain and be powered off separately
		Clocked		

Table 2-2 Cortex-A9 processor power modes (continued)

Mode	Cortex-A9 processor RAM arrays	Cortex-A9 processor logic	Cortex-A9 Data Engine	Comments
Standby	Powered-up	Powered-up	Powered Up	Standby modes, see <i>Standby modes</i> .
		Only wake-up logic is clocked.	Clock is disabled, or powered off	
Dormant	Retention state/voltage	Powered-off	Powered-off	External wake-up event required to wake up.
Shutdown	Powered-off	Powered-off	Powered-off	External wake-up event required to wake up.

Entry to Dormant or Shutdown mode must be controlled through an external power controller.

Run mode

Run mode is the normal mode of operation, where all of the functionality of the Cortex-A9 processor is available.

Standby modes

WFI and WFE Standby modes disable most of the clocks of a processor, while keeping its logic powered up. This reduces the power drawn to the static leakage current, leaving a tiny clock power overhead requirement to enable the device to wake up.

Entry into WFI Standby mode is performed by executing the WFI instruction.

The transition from the WFI Standby mode to the Run mode is caused by:

- An interrupt, masked or unmasked.
- An asynchronous data abort, regardless of the value of the CPSR.A bit. A pending wake-up event prevents the processor from entering low power mode.
- A debug request, regardless of whether debug is enabled.
- A reset.

Entry into WFE Standby mode is performed by executing the WFE instruction.

The transition from the WFE Standby mode to the Run mode is caused by:

- An interrupt, unless masked.
- A debug request, regardless of whether debug is enabled.
- A previous exception return on the same processor.
- A reset.
- The assertion of the **EVENTI** input signal.

The debug request can be generated by an externally generated debug request, using the **EDBGRQ** pin on the Cortex-A9 processor, or from a Debug Halt instruction issued to the Cortex-A9 processor through the APB debug port.

The debug channel remains active throughout a WFI.

Dormant mode

Dormant mode enables the Cortex-A9 processor to be powered down, while leaving the caches powered up and maintaining their state.

The RAM blocks that must remain powered up during Dormant mode are:

- all data RAMs associated with the cache
- all tag RAMs associated with the cache
- Outer RAMs.

The RAM blocks that are to remain powered up must be implemented on a separate power domain..

Before entering Dormant mode, the state of the Cortex-A9 processor, excluding the contents of the RAMs that remain powered up in dormant mode, must be saved to external memory. These state saving operations must ensure that the following occur:

- All ARM registers, including CPSR and SPSR registers are saved.
- All system registers are saved.
- All debug-related state must be saved.
- A Data Synchronization Barrier instruction is executed to ensure that all state saving has completed.
- The Cortex-A9 processor then communicates with the power controller, using the **STANDBYWFI**, to indicate that it is ready to enter dormant mode by performing a WFI instruction. See *Communication to the power management controller* on page 2-13 for more information.
- Before removing the power, the Reset signals to the Cortex-A9 processor must be asserted by the external power control mechanism.

The external power controller triggers the transition from Dormant state to Run state. The external power controller must assert reset to the Cortex-A9 processor until the power is restored. After power is restored, the Cortex-A9 processor leaves reset and can determine that the saved state must be restored.

Shutdown mode

Shutdown mode powers down the entire device, and all state, including cache, must be saved externally by software. This state saving is performed with interrupts disabled, and finishes with a Data Synchronization Barrier operation. The Cortex-A9 processor then communicates with a power controller that the device is ready to be powered down in the same manner as when entering Dormant Mode. The processor is returned to the run state by asserting reset.

Note

You must power up the processor before performing a reset.

Communication to the power management controller

Communication between the Cortex-A9 processor and the external power management controller can be performed using the Standby signals, Cortex-A9 input clamp signals, and **DBGNOPWRDWN**.

Standby signals

These signals control the external power management controller.

The **STANDBYWFI** signal indicates that the Cortex-A9 processor is ready to enter Power Down mode. See *Standby and Wait For Event signals* on page A-6.

Cortex-A9 input signals

The external power management controller uses **NEONCLAMP** and **CPURAMCLAMP** to isolate Cortex-A9 power domains from one another before they are turned off. These signals are only meaningful if the Cortex-A9 processor implements power domain clamps. See *Power management signals* on page A-7.

DBGNOPWRDWN

DBGNOPWRDWN is connected to the system power controller and is interpreted as a request to operate in emulate mode. In this mode, the Cortex-A9 processor and PTM are not actually powered down when requested by software or hardware handshakes. See *Miscellaneous debug interface signals* on page A-23.

2.4.3 Power domains

The Cortex-A9 uniprocessor contains optional placeholders between the Cortex-A9 logic and RAM arrays, or between the Cortex-A9 logic and the NEON SIMD logic, when NEON is present, so that these parts can be implemented in different voltage domains

2.4.4 Cortex-A9 voltage domains

The Cortex-A9 processor can have the following power domains:

- a power domain for Cortex-A9 processor logic cells
- a power domain for Cortex-A9 processor data engines
- a power domain for Cortex-A9 processor RAMs.

Figure 2-4 on page 2-14 shows the power domains.

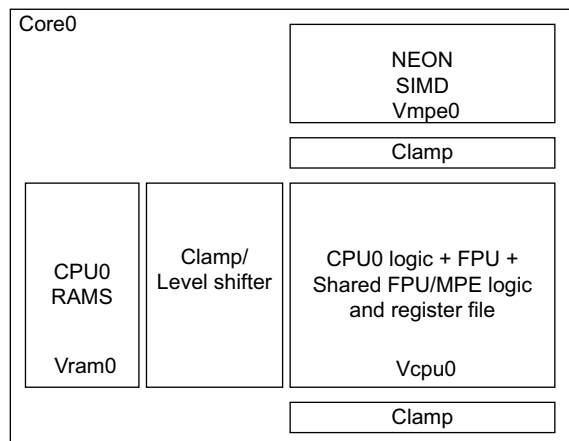


Figure 2-4 Power domains for Cortex-A9 r2 designs

The FPU is part of the CPU power domain. The FPU clock is based on the CPU clock. There is static and dynamic high-level clock-gating.

NEON SIMD data paths and logic are in a separate power domain, with dedicated clock and reset signals. There is static and dynamic high-level clock-gating.

When NEON is present, you can run FPU (non-SIMD) code without powering the SIMD part or clocking the SIMD part.

2.5 Constraints and limitations of use

This section describes memory consistency.

Memory coherency in a Cortex-A9 processor is maintained following a weakly ordered memory consistency model.

———— **Note** —————

When the Shareable attribute is applied to a memory region that is not Write-Back, Normal memory, data held in this region is treated as Non-cacheable.

Chapter 3

Programmers Model

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

- *About the programmers model* on page 3-2
- *ThumbEE architecture* on page 3-3
- *Advanced SIMD architecture* on page 3-4
- *Security Extensions architecture* on page 3-5
- *Multiprocessing Extensions* on page 3-6
- *The Jazelle Extension* on page 3-7
- *Memory model* on page 3-8
- *Addresses in the Cortex-A9 processor* on page 3-9.

3.1 About the programmers model

The Cortex-A9 processor implements the ARMv7-A architecture.

See the *ARM Architecture Reference Manual* for more information about the ARMv7-A architecture.

3.2 ThumbEE architecture

The *Thumb Execution Environment* (ThumbEE) extension is a variant of the Thumb instruction set that is designed as a target for dynamically generated code.

See the *ARM Architecture Reference Manual* for information on the ThumbEE instruction set.

3.3 Advanced SIMD architecture

The Advanced SIMD architecture extension is a media and signal processing architecture that adds instructions targeted primarily at audio, video, 3-D graphics, image, and speech processing.

———— **Note** —————

The Advanced SIMD architecture extension, its associated implementations, and supporting software, are commonly referred to as NEON MPE.

NEON MPE includes both Advanced SIMD instructions and the ARM VFPv3 instructions. All Advanced SIMD instructions and VFP instructions are available in both ARM and Thumb states.

See the *ARM Architecture Reference Manual* for more information.

See the *Cortex-A9 NEON Media Processing Engine Technical Reference Manual* for implementation-specific information.

3.4 Security Extensions architecture

The purpose of the security extensions is to enable the construction of a secure software environment. This section describes the following:

- *System boot sequence.*

See the *ARM Architecture Reference Manual* for more information.

3.4.1 System boot sequence

Caution

The Security Extensions enable the construction of an isolated software environment for more secure execution, depending on a suitable system design around the processor. The technology does not protect the processor from hardware attacks, and you must make sure that the hardware containing the reset handling code is appropriately secure.

The processor always boots in the privileged Supervisor mode in the Secure state, with the NS bit set to 0. This means that code that does not attempt to use the Security Extensions always runs in the Secure state. If the software uses both Secure and Non-secure states, the less trusted software, such as a complex operating system, executes in Non-secure state, and the more trusted software executes in the Secure state.

The following sequence is expected to be typical use of the security extensions:

1. Exit from reset in Secure state.
2. Configure the security state of memory and peripherals. Some memory and peripherals are accessible only to the software running in Secure state.
3. Initialize the secure operating system. The required operations depend on the operating system, and typically include initialization of caches, MMU, exception vectors, and stacks.
4. Initialize Secure Monitor software to handle exceptions that switch execution between the Secure and Non-Secure operating systems.
5. Optionally lock aspects of the secure state environment against further configuration.
6. Pass control through the Secure Monitor software to the non-secure OS with an SMCs instruction to enable the Non-secure operating system to initialize. The required operations depend on the operating system, and typically include initialization of caches, MMU, exception vectors, and stacks.

The overall security of the secure software depends on the system design, and on the secure software itself.

3.5 Multiprocessing Extensions

The Multiprocessing Extensions are a set of features that enhance multiprocessing functionality. See the *ARM Architecture Reference Manual* for more information.

3.6 The Jazelle Extension

The Cortex-A9 processor provides hardware support for the Jazelle Extension. The processor accelerates the execution of most bytecodes. Some bytecodes are executed by software routines.

See the *ARM Architecture Reference Manual* for more information.

See also Chapter 5 *Jazelle DBX registers*.

3.7 Memory model

The Cortex-A9 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 store words in memory in either big-endian format or little-endian format.

Instructions are always treated as little-endian.

———— **Note** —————

ARMv7 does not support the BE-32 memory model.

3.8 Addresses in the Cortex-A9 processor

In the Cortex-A9 the VA and MVA are identical.

When the Cortex-A9 processor is executing in Non-secure state, the processor performs translation table look-ups using the Non-secure versions of the Translation Table Base Registers. In this situation, any VA can only translate into a Non-secure PA. When it is in Secure state, the Cortex-A9 processor performs translation table look-ups using the Secure versions of the Translation Table Base Registers. In this situation, the security state of any VA is determined by the NS bit of the translation table descriptors for that address.

Table 3-1 shows the address types in the processor system.

Table 3-1 Address types in the processor system

Processor	Caches	Translation Lookaside Buffers	AXI bus
Data VA	Data cache is <i>Physically Indexed Physically Tagged</i> (PIPT)	Translates Virtual Address to Physical Address	Physical Address
Instruction VA	Instruction cache is <i>Virtually Indexed Physically Tagged</i> (VIPT)		

This is an example of the address manipulation that occurs when the Cortex-A9 processor requests an instruction.

1. The Cortex-A9 processor issues the VA of the instruction as Secure or Non-secure VA according to the state the processor is in.
2. The instruction cache is indexed by the lower bits of the VA. The TLB performs the translation in parallel with the cache look-up. The translation uses Secure descriptors if the processor is in the Secure state. Otherwise it uses the Non-secure descriptors.
3. If the protection check carried out by the TLB on the VA does not abort and the PA tag is in the instruction cache, the instruction data is returned to the processor.
4. If there is a cache miss, the PA is passed to the AXI bus interface to perform an external access. The external access is always Non-secure when the core is in the Non-secure state. In the Secure state, the external access is Secure or Non-secure according to the NS attribute value in the selected descriptor. In Secure state, both L1 and L2 table walks accesses are marked as Secure, even if the first level descriptor is marked as NS.

———— **Note** ————

Secure L2 look-ups are secure even if the L1 entry is marked Non-secure.

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 4-3
- *Register descriptions* on page 4-8.

4.1 About system control

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
- MMU configuration and management
- cache configuration and management
- system performance monitoring.

4.1.1 Deprecated registers

In ARMv7-A the following have instruction set equivalents:

- Instruction Synchronization Barrier
- Data Synchronization Barrier
- Data Memory Barrier
- Wait for Interrupt.

The use of the registers is optional and deprecated.

In addition, the Fast Context Switch Extensions are deprecated in ARM Architecture v7, and are not implemented in the Cortex-A9 processor.

4.2 Register summary

Table 4-1 shows all the CP15 system control registers, ordered by the parameters used to access the register:

- the primary CP15 coprocessor register, CRn
- the opcode_1 value
- the secondary CP15 coprocessor register, CRm
- the opcode_2 value.

The table includes references to the summaries of the attributes of each CRn group of registers.

For registers described in the ARM Architecture Reference Manual see *ARM Architecture Reference Manual, ARMv7-A and ARMv7-M editions*, <http://silver.arm.com/browse/AR570>.

Table 4-1 Summary of CP15 system control coprocessor registers

CRn	Name	Register description
c0	-	<i>CP15 c0 register summary</i> on page 4-8
	MIDR	Main ID Register, see the <i>ARM Architecture Reference Manual</i>
	CTR	Cache Type Register, see the <i>ARM Architecture Reference Manual</i>
	TCMTR	TCM Type Register, see the <i>ARM Architecture Reference Manual</i>
	TLBTR	<i>TLB Type Register</i> on page 4-9
	MPIDR	<i>Multiprocessor Affinity Register</i> on page 4-10
	ID_PFR0	Processor Feature Register 0, see the <i>ARM Architecture Reference Manual</i>
	ID_PFR1	Processor Feature Register 1, see the <i>ARM Architecture Reference Manual</i>
	ID_DFR0	Debug Feature Register , see the <i>ARM Architecture Reference Manual</i>
	ID_MMFR0	Memory Model Feature Register 0, see the <i>ARM Architecture Reference Manual</i>
	ID_MMFR1	Memory Model Feature Register 1, see the <i>ARM Architecture Reference Manual</i>
	ID_MMFR2	Memory Model Feature Register 2, see the <i>ARM Architecture Reference Manual</i>
	ID_MMFR3	Memory Model Feature Register 3, see the <i>ARM Architecture Reference Manual</i>
	ID_ISAR0	Instruction Set Attributes Register 0, see the <i>ARM Architecture Reference Manual</i>
c0	ID_ISAR1	Instruction Set Attributes Register 1, see the <i>ARM Architecture Reference Manual</i>
	ID_ISAR2	Instruction Set Attributes Register 2, see the <i>ARM Architecture Reference Manual</i>
	ID_ISAR3	Instruction Set Attributes Register 3, see the <i>ARM Architecture Reference Manual</i>
	ID_ISAR4	Instruction Set Attributes Register 4, see the <i>ARM Architecture Reference Manual</i>
	CCSIDR	<i>Cache Size Identification Register</i> on page 4-11
	CLIDR	<i>Cache Level ID Register</i> on page 4-13
	AIDR	<i>Auxiliary ID Register</i> on page 4-13
	CSSELR	<i>Cache Size Selection Register</i> on page 4-14

Table 4-1 Summary of CP15 system control coprocessor registers (continued)

CRn	Name	Register description
c1	-	<i>CP15 c1 register summary on page 4-15</i>
	SCTLR	<i>System Control Register on page 4-15</i>
	ACTLR	<i>Auxiliary Control Register on page 4-18</i>
	CPACR	<i>Coprocessor Access Control Register on page 4-20</i>
	SCR	Secure Configuration Register, see the <i>ARM Architecture Reference Manual</i>
	SDER	<i>Secure Debug Enable Register on page 4-22</i>
	NSACR	<i>Non-secure Access Control Register on page 4-23</i>
	VCR	<i>Virtualization Control Register on page 4-24</i>
c2	-	<i>CP15 c2 register summary on page 4-25</i>
	TTBR0	Translation Table Base Register 0, see the <i>ARM Architecture Reference Manual</i>
	TTBR1	Translation Table Base Register 1, see the <i>ARM Architecture Reference Manual</i>
	TTBCR	Translation Table Base Control Register, see the <i>ARM Architecture Reference Manual</i>
c3	-	<i>CP15 c3 register summary on page 4-26</i>
	DACR	Domain Access Control Register, see the <i>ARM Architecture Reference Manual</i>
c4	-	<i>CP15 c4, not used on page 4-26</i>
c5	-	<i>CP15 c5 register summary on page 4-26</i>
	DFSR	Data Fault Status Register, see the <i>ARM Architecture Reference Manual</i>
	IFSR	Instruction Fault Status Register, see the <i>ARM Architecture Reference Manual</i>
	ADFSR	Auxiliary Data Fault Status Register, see the <i>ARM Architecture Reference Manual</i>
	AIFSR	Auxiliary Instruction Fault Status Register, see the <i>ARM Architecture Reference Manual</i>
c6	-	<i>CP15 c6 register summary on page 4-26</i>
	DFAR	Data Fault Address Register, see the <i>ARM Architecture Reference Manual</i>
	IFAR	Instruction Fault Address Register, see the <i>ARM Architecture Reference Manual</i>

Table 4-1 Summary of CP15 system control coprocessor registers (continued)

CRn	Name	Register description
c7	-	<i>CP15 c7 register summary on page 4-27</i>
	ICIALLUIS	See the <i>ARM Architecture Reference Manual</i>
	BPIALLIS	
	PAR	
	ICIALLU	
	ICIMVAU	
	BPIALL	
	DCIMVAC	
	DCISW	
	V2PCWPR	
	DCCVAC	
	DCCSW	
	DCCVAU	
	DCCIMVAC	
	DCCISW	
c8	-	<i>CP15 c8 register summary on page 4-28</i>
	TLBIALLIS	See the <i>ARM Architecture Reference Manual</i>
	TLBIMVAIS	
	TLBIASIDIS	
	TLBIMVAAIS	
	TLBIALL	
	TLBIMVA	
	TLBIASID	
	TLBIMVAA	

Table 4-1 Summary of CP15 system control coprocessor registers (continued)

CRn	Name	Register description
c9	-	<i>CP15 c9 register summary on page 4-28</i>
	PMCR	Performance Monitor Control Register, see the <i>ARM Architecture Reference Manual</i>
	PMCNTENSET	Count Enable Set Register, see the <i>ARM Architecture Reference Manual</i>
	PMCNTENCLR	Count Enable Clear Register, see the <i>ARM Architecture Reference Manual</i>
	PMOVSr	Overflow Flag Status Register, see the <i>ARM Architecture Reference Manual</i>
	PMSWINC	Software Increment Register, see the <i>ARM Architecture Reference Manual</i>
	PMSELR	Event Counter Selection Register, see the <i>ARM Architecture Reference Manual</i>
	PMCCNTR	Cycle Count Register, see the <i>ARM Architecture Reference Manual</i>
	PMXEVTYPER	Event Selection Register, see the <i>ARM Architecture Reference Manual</i>
	PMXEVCNTR	Performance Monitor Count Registers, see the <i>ARM Architecture Reference Manual</i>
	PMUSERENR	User Enable Register, see the <i>ARM Architecture Reference Manual</i>
	PMINTENSET	Interrupt Enable Set Register, see the <i>ARM Architecture Reference Manual</i>
	PMINTENCLR	Interrupt Enable Clear Register, see the <i>ARM Architecture Reference Manual</i>
	c10	-
-		<i>TLB Lockdown Register on page 4-29</i>
PRRR		Primary Region Remap Register, see the <i>ARM Architecture Reference Manual</i>
NRRR		Normal Memory Remap Register, see the <i>ARM Architecture Reference Manual</i>
c11	-	<i>CP15 c11 register summary on page 4-30</i>
	PLEIDR	<i>PLE ID Register on page 4-30</i>
	PLEASR	<i>PLE Activity Status Register on page 4-31</i>
	PLEFSR	<i>PLE FIFO Status Register on page 4-32</i>
	PLEUAR	<i>Preload Engine User Accessibility Register on page 4-32</i>
	PLEPCR	<i>Preload Engine Parameters Control Register on page 4-33</i>
c12	-	<i>CP15 c12 register summary on page 4-34</i>
	VBAR	Vector Base Address Register, see the <i>ARM Architecture Reference Manual</i>
	MVBAR	Monitor Vector Base Address Register, see the <i>ARM Architecture Reference Manual</i>
	ISR	Interrupt Status Register, see the <i>ARM Architecture Reference Manual</i>
	-	<i>Virtualization Interrupt Register on page 4-34</i>
c13	-	<i>CP15 c13 register summary on page 4-35</i>
	FCSEIDR	FCSE PID Register, see the <i>ARM Architecture Reference Manual</i>
	CONTEXTIDR	Context ID Register, see the <i>ARM Architecture Reference Manual</i>
	TPIDRURW	User Read/Write Software Thread Register, see the <i>ARM Architecture Reference Manual</i>

Table 4-1 Summary of CP15 system control coprocessor registers (continued)

CRn	Name	Register description
	TPIDRURO	User Read Only Software Thread Register, see the <i>ARM Architecture Reference Manual</i>
	TPIDRPRW.	Privileged Only Software Thread Register, see the <i>ARM Architecture Reference Manual</i>
c14	-	<i>CP15 c14, not used on page 4-36</i>
c15	-	<i>CP15 c15 register summary on page 4-36</i>
	-	<i>Power Control Register on page 4-36</i>
	-	<i>NEON busy Register on page 4-37</i>
	-	<i>Configuration Base Address Register on page 4-38</i>
	-	<i>TLB lockdown operations on page 4-39</i>

4.3 Register descriptions

This section shows summary tables of the register allocation and reset values of the system control coprocessor, grouped by the primary CP15 register number, CRn, used to access the register.

In the summaries of the registers in each CRn group:

- 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.
- Type is:
 - Read-only (RO)
 - Write-only (WO)
 - Read/write (RW).
- Reset is the reset value of the register.

All system control coprocessor registers are 32 bits wide, except for the Program New Channel operation described in *PLE Program New Channel operation* on page 9-5. Reserved registers are RAZ/WI.

This section does not reproduce information about registers already described in the *ARM Architecture Reference Manual*. This chapter describes the implementation-defined control coprocessor registers.

4.3.1 CP15 c0 register summary

Table 4-2 shows the system control registers you can access when CRn is c0.

Table 4-2 c0 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	MIDR	RO	0x412FC090	Main ID Register for r2p0
			MIDR	RO	0x412FC091	Main ID Register for r2p1
			MIDR	RO	0x412FC092	Main ID Register for r2p2
		1	CTR	RO	0x83338003	Cache Type Register
		2	TCMTR	RO	0x00000000	TCM Type Register
		3	TLBTR ^a	RO	-	<i>TLB Type Register</i> on page 4-9
5	MPIDR	RO	-	<i>Multiprocessor Affinity Register</i> on page 4-10		
c1	c1	0	ID_PFR0	RO	0x00001231	Processor Feature Register 0
		1	ID_PFR1	RO	0x00000011	Processor Feature Register 1
		2	ID_DFR0	RO	0x00010444	Debug Feature Register 2
		4	ID_MMFR0	RO	0x00100103	Memory Model Feature Register 0
		5	ID_MMFR1	RO	0x20000000	Memory Model Feature Register 1
		6	ID_MMFR2	RO	0x01230000	Memory Model Feature Register 2
		7	ID_MMFR3	RO	0x00102111	Memory Model Feature Register 3

Table 4-2 c0 system control registers (continued)

Op1	CRm	Op2	Name	Type	Reset	Description
	c2	0	ID_ISAR0	RO	0x00101111	Instruction Set Attributes Register 0
		1	ID_ISAR1	RO	0x13112111	Instruction Set Attributes Register 1
		2	ID_ISAR2	RO	0x21232041	Instruction Set Attributes Register 2
		3	ID_ISAR3	RO	0x11112131	Instruction Set Attributes Register 3
		4	ID_ISAR4	RO	0x00011142	Instruction Set Attributes Register 4
1	c0	0	CCSIDR	RO	-	Cache Size Identification Register on page 4-11
		1	CLIDR	RO	0x09000003	Cache Level ID Register on page 4-13
		7	AIDR	RO	0x00000000	Auxiliary ID Register on page 4-13
2	c0	0	CSSELR	RW	-	Cache Size Selection Register on page 4-14

a. Depends on TLBSIZE. See *TLB Type Register*.

4.3.2 TLB Type Register

The TLBTR characteristics are:

Purpose Returns the number of lockable entries for the TLB

Usage constraints The TLBTR is:

- common to the Secure and Non-secure states.
- only accessible in privileged mode.

Configurations Available in all configurations.

Attributes See the register summary in Table 4-2 on page 4-8.

Figure 4-1 shows the TLBTR bit assignments.




Figure 4-1 TLBTR bit assignments

Table 4-3 shows the TLBTR bit assignments.

Table 4-3 TLBTR bit assignments

Bits	Name	Function
[31:24]	SBZ	-
[23:16]	ILsize	Specifies the number of instruction TLB lockable entries. For the Cortex-A9 processor this is 0.
[15:8]	DLsize	Specifies the number of unified or data TLB lockable entries. For the Cortex-A9 processor this is 4.

Table 4-3 TLBTR bit assignments (continued)

Bits	Name	Function
[7:2]	SBZ or UNP	-
[1]	TLB_size	0 = TLB has 64 entries 1 = TLB has 128 entries.
[0]	nU	Specifies if the TLB is unified, 0, or if there are separate instruction and data TLBs. 0 = The Cortex-A9 processor has a unified TLB.

To access the TLBTR, use:

MRC p15,0,<Rd>,c0,c0,3; returns TLB details

4.3.3 Multiprocessor Affinity Register

The MPIDR characteristics are:

Purpose	To identify: <ul style="list-style-type: none"> whether the processor is part of a Cortex-A9 MPCore implementation. Cortex-A9 processor accesses within a Cortex-A9 MPCore processor the target Cortex-A9 processor in a multi-processor cluster system.
Usage constraints	The MPIDR is: <ul style="list-style-type: none"> only accessible in privileged mode. common to the Secure and Non-secure states.
Configurations	Available in all configurations. The value of the U bit, bit [30], indicates if the configuration is a multiprocessor configuration or a uniprocessor configuration.
Attributes	See the register summary in Table 4-2 on page 4-8.

Figure 4-2 shows the MPIDR bit assignments.




Figure 4-2 MPIDR bit assignments

Table 4-4 shows the MPIDR bit assignments.

Table 4-4 MPIDR bit assignments

Bits	Name	Function
[31]	-	Indicates the register uses the new multiprocessor format. This is always 1.
[30]	U bit	Multiprocessing Extensions 0 = Processor is part of an MPCore cluster. 1 = Processor is a uniprocessor .
[29:12]	-	SBZ.
[11:8]	Cluster ID	Value read in CLUSTERID configuration pins ^a . It identifies a Cortex-A9 MPCore processor in a system with more than one Cortex-A9 MPCore processor present. SBZ for a uniprocessor configuration.
[7:2]	-	SBZ.
[1:0]	CPU ID	Indicates the CPU number in the Cortex-A9 MPCore configuration: <ul style="list-style-type: none"> • 0x0 processor is CPU0. • 0x1 processor is CPU1. • 0x2 processor is CPU2. • 0x3 processor is CPU3. In the uniprocessor version this value is fixed at 0x0.

a. A uniprocessor implementation does not include any **CLUSTERID** pins.

To access the MPIDR, use:

```
MRC p15,0,<Rd>,c0,c0,5; read Multiprocessor ID register
```

4.3.4 Cache Size Identification Register

The CCSIDR characteristics are:

Purpose Provides information about the architecture of the caches selected by CSSELR.

Usage constraints The CCSIDR is:

- only accessible in privileged modes.
- common to the Secure and Non-secure states.

Configurations Available in all configurations.

Attributes See the register summary in Table 4-2 on page 4-8.

Figure 4-3 on page 4-12 shows the CCSIDR bit assignments.




Figure 4-3 CCSIDR bit assignments

Table 4-5 shows how the CSSIDR bit assignments.

Table 4-5 CCSIDR bit assignments

Bits	Name	Function
[31]	WT	Indicates support available for Write-Through: 0 = Write-Through support not available 1 = Write-Through support available.
[30]	WB	Indicates support available for Write-Back: 0 = Write-Back support not available 1 = Write-Back support available.
[29]	RA	Indicates support available for read allocation: 0 = Read allocation support not available 1 = Read allocation support available.
[28]	WA	Indicates support available for write allocation: 0 = Write allocation support not available 1 = Write allocation support available.
[27:13]	NumSets	Indicates number of sets. 0x7F = 16KB cache size 0xFF = 32KB cache size 0x1FF = 64KB cache size.
[12:3]	Associativity	Indicates number of ways. b000000011. Four ways.
[2:0]	LineSize	Indicates number of words. b001 = Eight words per line.

To access the CCSIDR, use:

MRC p15, 1, <Rd>, c0, c0, 0; Read current Cache Size Identification Register

If the CSSELR reads the instruction cache values, then bits[31:28] are b0010.

If the CSSELR reads the data cache values, then bits[31:28] are b0111. See *Cache Size Selection Register* on page 4-14.

4.3.5 Cache Level ID Register

The CLIDR characteristics are:

Purpose	Indicates the cache levels that are implemented in the processor and under the control of the System Control Coprocessor.
Usage constraints	The CLIDR is: <ul style="list-style-type: none"> only accessible in privileged modes. common to the Secure and Non-secure states.
Configurations	Available in all configurations.
Attributes	See the register summary in Table 4-2 on page 4-8.

Figure 4-4 shows the CLIDR bit assignments.




Figure 4-4 CLIDR bit assignments

Table 4-6 shows the CLIDR bit assignments.

Table 4-6 CLIDR bit assignments

Bits	Name	Function
[31:30]	-	UNP or SBZ
[29:27]	LoU	b001 = Level of unification
[26:24]	LoC	b001 = Level of coherency
[23:21]	LoUIS	b001 = Level of Unification Inner Shareable
[20:18]	CL 7	b000 = No cache at CL 7
[17:15]	CL 6	b000 = No cache at CL 6
[14:12]	CL 5	b000 = No cache at CL 5
[11:9]	CL 4	b000 = No cache at CL 4
[8:6]	CL 3	b000 = No cache at CL 3
[5:3]	CL 2	b000 = No unified cache at CL 2
[2:0]	CL 1	b011 = Separate instruction and data caches at CL 1

To access the CLIDR, use:

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

4.3.6 Auxiliary ID Register

The AIDR characteristics are:

Purpose	Provides implementation-specific information.
----------------	---

Usage constraints The AIDR is:

- only accessible in privileged modes
- common to the Secure and Non-secure states.

Configurations Available in all configurations.

Attributes See the register summary in Table 4-2 on page 4-8.

To access the Auxiliary Level ID Register, use:

MRC p15,1,<Rd>,c0,c0,7; Read Auxiliary ID Register

———— **Note** —————

The AIDR is unused in this implementation.

4.3.7 Cache Size Selection Register

The CSSELR characteristics are:

Purpose Selects the current CCSIDR.

Usage constraints The CSSELR is:

- only accessible in privileged modes
- banked for Secure and Non-secure states

Configurations Available in all configurations.

Attributes See the register summary in Table 4-2 on page 4-8.

Figure 4-5 shows the CSSELR bit assignments.




Figure 4-5 CSSELR bit assignments

Table 4-7 shows the CSSELR bit assignments.

Table 4-7 CSSELR bit assignments

Bits	Name	Function
[31:4]	-	UNP or SBZ.
[3:1]	Level	Cache level selected, RAZ/WI. There is only one level of cache in the Cortex-A9 processor so the value for this field is b000.
[0]	InD	1 = Instruction cache 0 = Data cache.

To access the CSSELR, use:

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

4.3.8 CP15 c1 register summary

Table 4-8 shows the CP15 registers you can access when CRn is c1.

Table 4-8 c1 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	SCTLR	RW	- ^a	<i>System Control Register</i>
		1	ACTLR ^b	RW	0x00000000	<i>Auxiliary Control Register</i> on page 4-18
		2	CPACR	RW	c	<i>Coprocessor Access Control Register</i> on page 4-20
c1	c1	0	SCR ^d	RW	0x00000000	Secure Configuration Register
		1	SDER ^c	RW	0x00000000	<i>Secure Debug Enable Register</i> on page 4-22
		2	NSACR	RW ^e	f	<i>Non-secure Access Control Register</i> on page 4-23
		3	VCR ^c	RW	0x00000000	<i>Virtualization Control Register</i> on page 4-24

a. Depends on input signals. See *System Control Register*.

b. RO in Non-secure state if NSACR[18]=0 and RW if NSACR[18]=1.

c. 0x00000000 if NEON present and 0xC0000000 if NEON not present or powered down.

d. No access in Non-secure state.

e. This is a read and write register in Secure state and a read-only register in the Non-secure state.

f. 0x00000000 if NEON present and 0x0000C000 if NEON not present..

4.3.9 System Control Register

The SCTLR characteristics are:

Purpose	Provides control and configuration of: <ul style="list-style-type: none"> memory alignment and endianness, memory protection and fault behavior MMU and cache enables interrupts and behavior of interrupt latency location for exception vectors program flow prediction.
Usage constraints	The SCTLR is: <ul style="list-style-type: none"> Only accessible in privileged modes. Partially banked. Table 4-9 on page 4-16 shows banked and secure modify only bits.
Configurations	Available in all configurations.
Attributes	See the register summary in Table 4-8.

Figure 4-6 on page 4-16 shows the SCTLR bit assignments.




Figure 4-6 SCTLR bit assignments

Table 4-9 shows the SCTLR bit assignments.

Table 4-9 SCTLR bit assignments

Bits	Name	Access	Function
[31]	-	-	SBZ.
[30]	TE	Banked	Thumb exception Enable: 0 = exceptions including reset are handled in ARM state. 1 = exceptions including reset are handled in Thumb state. The TEINIT signal defines the reset value.
[29]	AFE	Banked	Access Flag Enable bit: 0 = Full access permissions behavior. This is the reset value. The software maintains binary compatibility with ARMv6K behavior. 1 = Simplified access permissions behavior. The Cortex-A9 processor redefines the AP[0] bit as an access flag. The TLB must be invalidated after changing the AFE bit.
[28]	TRE	Banked	This bit controls the TEX remap functionality in the MMU: 0 = TEX remap disabled. This is the reset value. 1 = TEX remap enabled.
[27]	NMFI	Read-only	Nonmaskable FIQ support. The bit cannot be configured by software. The CFGNMFI signal defines the reset value.
[26]	-	-	RAZ/SBZP.
[25]	EE bit	Banked	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. This value also indicates the endianness of the translation table data for translation table look-ups. 0 = little-endian 1 = big-endian. The CFGEND signal defines the reset value.
[24]	-	-	RAZ/WI.
[23:22]	-	-	RAO/SBOP.
[21]	-	-	RAZ/WI.
[20:19]	-	-	RAZ/SBZP.
[18]	-	-	RAO/SBOP.

Table 4-9 SCTLR bit assignments (continued)

Bits	Name	Access	Function
[17]	HA	-	RAZ/WI.
[16]	-	-	RAO/SBOP.
[15]	-	-	RAZ/SBZP.
[14]	RR	Secure modify only	Replacement strategy for caches, BTAC, and micro TLBs. This bit is read/write in Secure state and read-only in Non-secure state: 0 = Random replacement. This is the reset value. 1 = Round-robin replacement.
[13]	V	Banked	Vectors bit. This bit selects the base address of the exception vectors: 0 = Normal exception vectors, base address 0x00000000. The Security Extensions are implemented, so this base address can be re-mapped. 1 = High exception vectors, Hivecs, base address 0xFFFF0000. This base address is never remapped. At reset the value for the secure version if this bit is taken from VINITHI .
[12]	I bit	Banked	Determines if instructions can be cached at any available cache level: 0 = Instruction caching disabled at all levels. This is the reset value. 1 = Instruction caching enabled.
[11]	Z bit	Banked	Enables program flow prediction: 0 = Program flow prediction disabled. This is the reset value. 1 = Program flow prediction enabled.
[10]	SW bit	Banked	SWP/SWPB enable bit: 0 = SWP and SWPB are Undefined. This is the reset value. 1 = SWP and SWPB perform normally.
[9:7]	-	-	RAZ/SBZP.
[6:3]	-	-	RAO/SBOP.
[2]	C bit	Banked	Determines if data can be cached at any available cache level: 0 = Data caching disabled at all levels. This is the reset value. 1 = Data caching enabled.
[1]	A bit	Banked	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 bit	Banked	Enables the MMU: 0 = MMU disabled. This is the reset value. 1 = MMU enabled.

Attempts to read or write the SCTLR from secure or Non-secure User modes result in an Undefined Instruction exception.

Attempts to write to this register in secure privileged mode when **CP15SDISABLE** is HIGH result in an Undefined Instruction exception.

Attempts to write secure modify only bits in Non-secure privileged modes are ignored.

Attempts to read secure modify only bits return the secure bit value.

Attempts to modify read-only bits are ignored.

To access the SCTRL, use:

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

4.3.10 Auxiliary Control Register

The ACTLR characteristics are:

Purpose	<p>Controls:</p> <ul style="list-style-type: none"> • parity checking, if implemented • allocation in one way • exclusive caching with the L2 cache • coherency mode, <i>Symmetric Multiprocessing (SMP)</i> or <i>Asymmetric Multiprocessing (AMP)</i> • speculative accesses on AXI. • broadcast of cache, branch predictor, and TLB maintenance operations. • L2C-310 cache allocation: <ul style="list-style-type: none"> — write full line of zeros mode.
Usage constraints	<p>The ACTLR is:</p> <ul style="list-style-type: none"> • Only accessible in privileged modes. • Common to the Secure and Non-secure states. • RW in Secure state. • RO in Non-secure state if NSACR.NS_SMP = 0. • RW in Non-secure state if NSACR.NS_SMP = 1. In this case all bits are Write Ignore except for the SMP bit.
Configurations	<p>Available in all configurations.</p> <ul style="list-style-type: none"> • In all configurations when the SMP bit = 0, Inner Cacheable Shareable attributes are treated as Non-cacheable. • In multiprocessor configurations when the SMP bit is set: <ul style="list-style-type: none"> — broadcasting cache and TLB maintenance operations is permitted if the FW bit is set. — receiving cache and TLB maintenance operations broadcast by other Cortex-A9 processors in the same coherent cluster is permitted if the FW bit is set — the Cortex-A9 processor can send and receive coherent requests for Shared Inner Write-back Write-Allocate accesses from the other Cortex-A9 processors in the same coherent cluster.
Attributes	<p>See the register summary in Table 4-8 on page 4-15.</p>

Figure 4-7 on page 4-19 shows the ACTLR bit assignments.

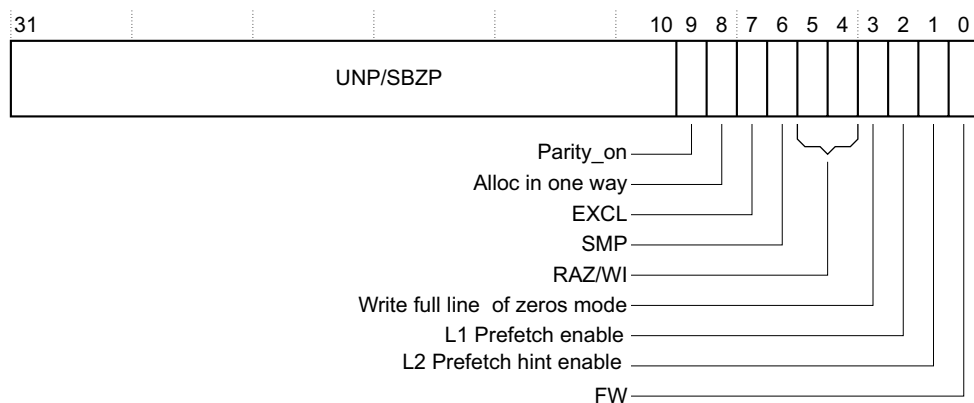


Figure 4-7 ACTLR bit assignments

Table 4-10 shows the ACTLR bit assignments.

Table 4-10 ACTLR bit assignments

Bits	Name	Function
[31:10]	-	UNP or SBZP.
[9]	Parity on	Support for parity checking, if implemented: 0 = Disabled. This is the reset value. 1 = Enabled. If parity checking is not implemented this bit reads as zero and writes are ignored.
[8]	Alloc in one way	Enable allocation in one cache way only. For use with memory copy operations to reduce cache pollution. The reset value is zero.
[7]	EXCL	Exclusive cache bit. The exclusive cache configuration does not permit data to reside in L1 and L2 at the same time. The exclusive cache configuration provides support for only caching data on an eviction from L1 when the inner cache attributes are Write-Back, Cacheable and allocated in L1. Ensure that your cache controller is also configured for exclusive caching. 0 = Disabled. This is the reset value. 1 = Enabled.
[6]	SMP	Signals if the Cortex-A9 processor is taking part in coherency or not. In uniprocessor configurations, if this bit is set, then Inner Cacheable Shared is treated as Cacheable. The reset value is zero.
[5:4]	-	RAZ/WI.
[3]	Write full line of zeros mode	Enable write full line of zeros mode ^a . The reset value is zero.

Table 4-10 ACTLR bit assignments (continued)

Bits	Name	Function
[2]	L1 prefetch enable	Dside prefetch. 0 = Disabled. This is the reset value. 1 = Enabled.
[1]	L2 prefetch enable	Prefetch hint enable ^a . The reset value is zero
[0]	FW	Cache and TLB maintenance broadcast: 0 = Disabled. This is the reset value. 1 = Enabled. RAZ/WI if only one Cortex-A9 processor present.

- a. This feature must be enabled only when the slaves connected on the Cortex-A9 AXI master port support it. The L2-310 Cache Controller supports this feature. See *Optimized accesses to the L2 memory interface* on page 8-7.

To access the ACTLR you must use a read modify write technique. To access the ACTLR, use:

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

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

Attempts to write to this register in secure privileged mode when **CP15SDISABLE** is HIGH result in an Undefined Instruction exception.

4.3.11 Coprocessor Access Control Register

The CPACR characteristics are:

- Purpose**
- sets access rights for the coprocessors CP11 and CP10.
 - enables software to determine if any particular coprocessor exists in the system.

———— **Note** —————

This register has no effect on access to CP14 or CP15.

- Usage constraints** The CPACR is:
- only accessible in privileged modes
 - Common to Secure and Non-secure states.

Configurations Available in all configurations.

Attributes See the register summary in Table 4-8 on page 4-15.

Figure 4-8 shows the CPACR bit assignments.




Figure 4-8 CPACR bit assignments

Table 4-11 shows the CPACR bit assignments.

Table 4-11 CPACR bit assignments

Bits	Name	Function
[31]	ASEDIS	<p>Disable Advanced SIMD Extension functionality:</p> <p>0 = this bit does not cause any instructions to be undefined.</p> <p>1 = all instruction encodings identified in the <i>ARM Architecture Reference Manual</i> as being part of the Advanced SIMD Extensions but that are not VFPv3 instructions are undefined.</p> <p>See the <i>Cortex-A9 Floating-Point Unit Technical Reference Manual</i> and <i>Cortex-A9 NEON Media Processing Engine Technical Reference Manual</i> for more information.</p> <p>If implemented with VFP only, RAO/WI.</p> <p>If implemented without both VFP and NEON, UNK/SBZP.</p>
[30]	D32DIS	<p>Disable use of D16-D31 of the VFP register file:</p> <p>0 = this bit does not cause any instructions to be undefined.</p> <p>1 = all instruction encodings identified in the <i>ARM Architecture Reference Manual</i> as being VFPv3 instructions are undefined if they access any of registers D16-D31.</p> <p>See the <i>Cortex-A9 Floating-Point Unit Technical Reference Manual</i> and <i>Cortex-A9 NEON Media Processing Engine Technical Reference Manual</i> for more information.</p> <p>If implemented with VFP only RAO/WI. If implemented without both VFP and NEON, UNK/SBZP.</p>
[29:24]	-	RAZ/WI.
[23:22]	cp11	<p>Defines access permissions for the coprocessor. Access denied is the reset condition and is the behavior for nonexistent coprocessors.</p> <p>b00 = Access denied. This is the reset value. Attempted access generates an Undefined Instruction exception.</p> <p>b01 = Privileged mode access only.</p> <p>b10 = Reserved.</p> <p>b11 = Privileged and User mode access.</p>
[21:20]	cp10	<p>Defines access permissions for the coprocessor. Access denied is the reset condition and is the behavior for nonexistent coprocessors.</p> <p>b00 = Access denied. This is the reset value. Attempted access generates an Undefined Instruction exception.</p> <p>b01 = Privileged mode access only.</p> <p>b10 = Reserved.</p> <p>b11 = Privileged and User mode access.</p>
[19:0]	-	RAZ/WI.

Access to coprocessors in the Non-secure state depends on the permissions set in the *Non-secure Access Control Register* on page 4-23.

Attempts to read or write the CPACR access bits depend on the corresponding bit for each coprocessor in *Non-secure Access Control Register* on page 4-23.

To access the CPACR, use:

MRC p15, 0, <Rd>, c1, c0, 2; Read Coprocessor Access Control Register
MCR p15, 0, <Rd>, c1, c0, 2; Write Coprocessor Access Control Register

You must execute an ISB immediately after an update of the CPACR. See Memory Barriers in the *ARM Architecture Reference Manual*. You must not attempt to execute any instructions that are affected by the change of access rights between the ISB and the register update.

To determine if any particular coprocessor exists in the system, write the access bits for the coprocessor of interest with b11. If the coprocessor does not exist in the system the access rights remain set to b00.

———— **Note** ————

You must enable both coprocessor 10 and coprocessor 11 before accessing any NEON or VFP system registers.

4.3.12 Secure Debug Enable Register

The SDER characteristics are:

Purpose	Controls Cortex-A9 debug.
Usage constraints	The SDER is: <ul style="list-style-type: none"> • only accessible in privileged modes. • only accessible in Secure state, accesses in Non-secure state cause an Undefined Instruction exception.
Configurations	Available in all configurations.
Attributes	See the register summary in Table 4-8 on page 4-15.

Figure 4-9 shows the SDER bit assignments.




Figure 4-9 SDER bit assignments

Table 4-12 shows the SDER bit assignments.

Table 4-12 SDER bit assignments

Bits	Name	Function
[31:2]	-	Reserved.
[1]	Secure User Non-Invasive Debug Enable	0 = Non-invasive debug not permitted in Secure User mode. This is the reset value. 1 = Non-invasive debug permitted in Secure User mode.
[0]	Secure User Invasive Debug Enable	0 = Invasive debug not permitted in Secure User mode. This is the reset value. 1 = Invasive debug permitted in Secure User mode.

To access the SDER, use:

MRC p15,0,<Rd>,c1,c1,1; Read Secure debug enable Register
MCR p15,0,<Rd>,c1,c1,1; Write Secure debug enable Register

4.3.13 Non-secure Access Control Register

The NSACR characteristics are:

Purpose Sets the Non-secure access permission for coprocessors.

Usage constraints The NSACR is:

- only accessible in privileged modes
- a read and write register in Secure state
- a read-only register in Non-secure state.

Note

This register has no effect on Non-secure access permissions for the debug control coprocessor, or the system control coprocessor.

Configurations Available in all configurations.

Attributes See the register summary in Table 4-8 on page 4-15.

Figure 4-10 shows the NSACR bit assignments.

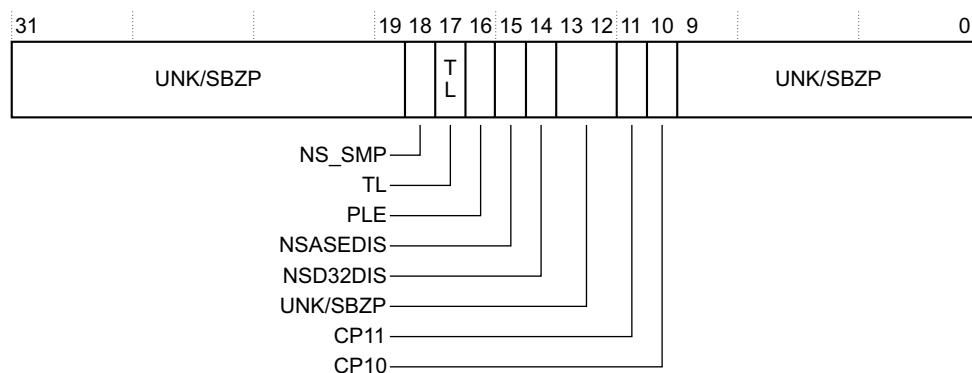


Figure 4-10 NSACR bit assignments

Table 4-13 shows the NSACR bit assignments.

Table 4-13 NSACR bit assignments

Bits	Name	Function
[31:19]	-	UNK/SBZP.
[18]	NS_SMP	Determines if the SMP bit of the Auxiliary Control Register is writable in Non-secure state: 0 = A write to Auxiliary Control Register in Non-secure state takes an undefined exception and the SMP bit is write ignored. This is the reset value. 1 = A write to Auxiliary Control Register in Non-secure state can modify the value of the SMP bit. Other bits are write ignored.
[17]	TL	Determines if lockable TLB entries can be allocated in Non-secure state: 0 = Lockable TLB entries cannot be allocated. This is the reset value. 1 = Lockable TLB entries can be allocated.

Table 4-13 NSACR bit assignments (continued)

Bits	Name	Function
[16]	PLE	Controls NS accesses to the Preload Engine resources: 0 = Only Secure accesses to CP15 c11 are permitted. All Non-secure accesses to CP15 c11 are trapped to UNDEF. This is the default value. 1 = Non-secure accesses to the CP15 c11 domain are permitted. That is, PLE resources are available in the Non-secure state. If the Preload Engine is not implemented this bit is RAZ/WI. See Chapter 9 <i>Preload Engine</i> .
[15]	NSASEDIS	Disable Non-secure Advanced SIMD Extension functionality: 0 = This bit has no effect on the ability to write CPACR.ASEDIS. This is the reset value. 1 = The CPACR.ASEDIS bit when executing in Non-secure state has a fixed value of 1 and writes to it are ignored. See the <i>Cortex-A9 Floating-Point Unit Technical Reference Manual</i> and <i>Cortex-A9 NEON Media Processing Engine Technical Reference Manual</i> for more information.
[14]	NSD32DIS	Disable the Non-secure use of D16-D31 of the VFP register file: 0 = This bit has no effect on the ability to write CPACR.D32DIS. This is the reset value. 1 = The CPACR.D32DIS bit when executing in Non-secure state has a fixed value of 1 and writes to it are ignored. See the <i>Cortex-A9 Floating-Point Unit Technical Reference Manual</i> and <i>Cortex-A9 NEON Media Processing Engine Technical Reference Manual</i> for more information.
[13:12]	-	UNK/SBZP.
[11]	CP11	Determines permission to access coprocessor 11 in the Non-secure state: 0 = Secure access only. This is the reset value. 1 = Secure or Non-secure access.
[10]	CP10	Determines permission to access coprocessor 10 in the Non-secure state: 0 = Secure access only. This is the reset value. 1 = Secure or Non-secure access.
[9:0]	-	UNK/SBZP.

To access the NSACR, use:

MCR p15, 0, <Rd>, c1, c1, 2; Read NSACR data
MCR p15, 0, <Rd>, c1, c1, 2; Write NSACR data

See the *Cortex-A9 Floating-Point Unit Technical Reference Manual* and *Cortex-A9 NEON Media Processing Engine Technical Reference Manual* for more information.

4.3.14 Virtualization Control Register

The VCR characteristics are:

Purpose	Forces an exception regardless of the value of the A, I, or F bits in the <i>Current Program Status Register (CPSR)</i> .
Usage constraints	The VCR is: <ul style="list-style-type: none"> only accessible in privileged modes only accessible in Secure state.
Configurations	Available in all configurations.
Attributes	See the register summary in Table 4-8 on page 4-15.

Figure 4-11 shows the VCR bit assignments.




Figure 4-11 VCR bit assignments

Table 4-14 shows the VCR bit assignments.

Table 4-14 VCR bit assignments

Bits	Name	Function
[31:9]	-	UNK/SBZP
[8]	AMO	Abort Mask Override When the processor is in Non-secure state and the SCR.EA bit is set, if the AMO bit is set, this enables an asynchronous Data Abort exception to be taken regardless of the value of the CPSR.A bit. When the processor is in Secure state, or when the SCR.EA bit is not set, the AMO bit is ignored.
[7]	IMO	IRQ Mask Override When the processor is in Non-secure state and the SCR.IRQ bit is set, if the IMO bit is set, this enables an IRQ exception to be taken regardless of the value of the CPSR.I bit. When the processor is in Secure state, or when the SCR.IRQ bit is not set, the IMO bit is ignored.
[6]	IFO	FIQ Mask Override When the processor is in Non-secure state and the SCR.FIQ bit is set, if the IFO bit is set, this enables an FIQ exception to be taken regardless of the value of the CPSR.F bit. When the processor is in Secure state, or when the SCR.FIQ bit is not set, the IFO bit is ignored.
[5:0]	-	UNK/SBZP

To access the VCR, use:

MRC p15, 0, <Rd>, c1, c1, 3; Read VCR data

MCR p15, 0, <Rd>, c1, c1, 3; Write VCR data

4.3.15 CP15 c2 register summary

Table 4-15 shows the system control registers you can access when CRn is c2.

Table 4-15 c2 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	TTBR0	RW	-	
		1	TTBR1	RW	-	Translation Table Base Register 1
		2	TTBCR	RW	0x00000000 ^a	Translation Table Base Control Register

a. In Secure state only. You must program the Non-secure version with the required value.

4.3.16 CP15 c3 register summary

Table 4-16 shows the system control register you can access when CRn is c3.

Table 4-16 c3 system control register

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	DACR	RW	-	Domain Access Control Register

4.3.17 CP15 c4, not used

No CP15 registers are accessed with CRn set to c4.

4.3.18 CP15 c5 register summary

Table 4-17 shows the system control registers you can access when CRn is c5.

Table 4-17 c5 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	DFSR	RW	-	Data Fault Status Register
		1	IFSR	RW	-	Instruction Fault Status Register
	c1	0	ADFSR	-	-	Auxiliary Data Fault Status Register
		1	AIFSR	-	-	Auxiliary Instruction Fault Status Register

4.3.19 CP15 c6 register summary

Table 4-18 shows the system control registers you can access when CRn is c6.

Table 4-18 c6 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	DFAR	RW	-	Data Fault Address Register
		2	IFAR	RW	-	Instruction Fault Address Register

4.3.20 CP15 c7 register summary

Table 4-19 shows the system control registers you can access when CRn is c7.

Table 4-19 c7 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description		
0	c0	0-3	Reserved	WO	-	-		
		4	NOP ^a	WO	-	-		
c1	0	0	ICIALUIS	WO	-	Cache operations registers		
		6	BPIALLIS	WO	-			
		7	Reserved	WO	-			
c4	0	PAR	RW	-	-			
c5	0	0	ICIALLU	WO	-	Cache operations registers		
		1	ICIMVAU	WO	-			
		2-3	Reserved	WO	-			
		4	ISB	WO	User		<i>Deprecated registers on page 4-2</i>	
		6	BPIALL	WO	-		Cache operations registers	
c6	1	1	DCIMVAC	WO	-			
		2	DCISW	WO	-			
0	c8	0-7	V2PCWPR	WO	-	VA to PA operations		
		c10	1	DCCVAC	WO	-	Cache operations registers	
			2	DCCSW	WO	-		
			4	DSB	WO	User		<i>Deprecated registers on page 4-2</i>
			5	DMB	WO	User		
		c11	1	DCCVAU	WO	-	Cache operations registers	
		c14	1	1	DCCIMVAC	WO	-	
2	DCCISW			WO	-			

a. This operation is performed by the WFI instruction. See also *Deprecated registers on page 4-2*.

4.3.21 CP15 c8 register summary

Table 4-20 shows the system control registers you can access when CRn is c8.

Table 4-20 c8 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c3	0	TLBIALLIS ^a	WO	-	-
		1	TLBIMVAIS ^b	WO	-	-
		2	TLBIASIDIS ^b	WO	-	-
		3	TLBIMVAAIS ^a	WO	-	-
	c5, c6, or c7	0	TLBIALL ^a	WO	-	-
		1	TLBIMVA ^b	WO	-	-
		2	TLBIASID ^b	WO	-	-
		3	TLBIMVAA ^a	WO	-	-

a. Has no effect on entries that are locked down.

b. Invalidates the locked entry when it matches.

See also *Invalidate TLB Entries on ASID Match* on page 4-41.

4.3.22 CP15 c9 register summary

Table 4-21 shows the system control registers you can access when CRn is c9.

Table 4-21 c9 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c12	0	PMCR	RW	0x41093000	Performance Monitor Control Register
		1	PMCNTENSET	RW	0x00000000	Count Enable Set Register
		2	PMCNTENCLR	RW	0x00000000	Count Enable Clear Register
		3	PMOVSr	RW	-	Overflow Flag Status Register
		4	PMSWINC	WO	-	Software Increment Register
		5	PMSELR	RW	0x00000000	Event Counter Selection Register
	c13	0	PMCCNTR	RW	-	Cycle Count Register
		1	PMXEVTYPER	RW	-	Event Selection Register
		2	PMXVCNTR	RW	-	Performance Monitor Count Registers
	c14	0	PMUSERENR	RW ^a	0x00000000	User Enable Register
		1	PMINTENSET	RW	0x00000000	Interrupt Enable Set Register
		2	PMINTENCLR	RW	0x00000000	Interrupt Enable Clear Register

a. RO in user mode

See Chapter 11 *Performance Monitoring Unit*.

4.3.23 CP15 c10 register summary

Table 4-22 shows the system control registers you can access when CRn is c10.

Table 4-22 c10 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	TLB Lockdown Register ^a	RW	0x00000000	TLB Lockdown Register
	c2	0	PRRR	RW	0x00098AA4	Primary Region Remap Register
		1	NRRR	RW	0x44E048E0	Normal Memory Remap Register

a. No access in Non-secure state if NSCAR.TL=0 and RW if NSACR.TL=1.

4.3.24 TLB Lockdown Register

The TLB Lockdown Register characteristics are:

Purpose	Controls where hardware translation table walks place the TLB entry. The TLB entry can be in either: <ul style="list-style-type: none"> the set-associative region of the TLB the lockdown region of the TLB, and if in the lockdown region, the entry to write The lockdown region of the TLB contains four entries.
Usage constraints	The TLB Lockdown Register is: <ul style="list-style-type: none"> only accessible in privileged modes. common to Secure and Non-secure states. not accessible if NSACR.TL is 0.
Configurations	Available in all configurations.
Attributes	See the register summary in Table 4-22.

Figure 4-12 shows the TLB Lockdown Register bit assignments.




Figure 4-12 TLB Lockdown Register bit assignments

Table 4-23 shows the TLB Lockdown Register bit assignments

Table 4-23 TLB Lockdown Register bit assignments

Bits	Name	Function
[31:30]	-	UNK/SBZP.
[29:28]	Victim	Lockdown region.
[27:1]	-	UNK/SBZP.
[0]	P	Preserve bit. 0 is the reset value.

To access the TLB Lockdown Register use:

MRC p15, 0, <Rd>, c10, c0, 0; Read TLB Lockdown victim
MCR p15, 0, <Rd>, c10, c0, 0; Write TLB Lockdown victim

Writing the TLB Lockdown Register with the preserve bit (P bit) set to:

- 1** Means subsequent hardware translation table walks place the TLB entry in the lockdown region at the entry specified by the victim, in the range 0 to 3.
- 0** Means subsequent hardware translation table walks place the TLB entry in the set-associative region of the TLB.

See *Invalidate TLB Entries on ASID Match* on page 4-41.

4.3.25 CP15 c11 register summary

Table 4-24 shows the system control registers you can access where CRn is c11.

Table 4-24 c11 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	PLEIDR	RO ^a	-	<i>PLE ID Register</i>
		2	PLEASR	RO ^a	-	<i>PLE Activity Status Register</i> on page 4-31
		4	PLEFSR	RO ^a	-	<i>PLE FIFO Status Register</i> on page 4-32
c1	c1	0	PLEUAR	Privileged R/W User RO	-	<i>Preload Engine User Accessibility Register</i> on page 4-32
		1	PLEPCR	Privileged R/W User RO	-	<i>Preload Engine Parameters Control Register</i> on page 4-33

a. RAZ if the PLE is not present.

4.3.26 PLE ID Register

The PLEIDR characteristics are:

Purpose Indicates whether the PLE is present or not and the size of its FIFO.

Usage constraints The PLEIDR is:

- Common to Secure and Non-secure states
- Accessible in User and Privileged modes, regardless of any configuration bit.

Configurations Available in all Cortex-A9 configurations regardless of whether a PLE is present or not.

Attributes See Table 4-24.

Figure 4-13 shows the PLEIDR bit assignments.




Figure 4-13 PLEIDR bit assignments

Table 4-25 shows the PLEIDR bit assignments.

Table 4-25 PLEIDR bit assignments

Bits	Name	Function
[31:21]	-	-
[20:16]	PLE FIFO size	Permitted values are: <ul style="list-style-type: none"> 5'b00000 indicates the PLE is not present 5'b00100 indicates a PLE is present with a FIFO size of 4 entries 5'b01000 indicates a PLE is present with a FIFO size of 8 entries 5'b10000 indicates a PLE is present with a FIFO size of 16 entries.
[15:1]	-	RAZ.
[0]	-	1 indicates that the Preload Engine is present in the given configuration.

To access the PLEIDR, use:

MRC p15, 0, <Rt>, c11, c0, 0; Read PLEIDR

4.3.27 PLE Activity Status Register

The PLEASR characteristics are:

Purpose	Indicates whether the PLE engine is currently active.
Usage constraints	The PLEASR is: <ul style="list-style-type: none"> Common to Secure and Non-secure states Accessible in User and Privileged modes, regardless of any configuration bit.
Configurations	Available in all Cortex-A9 configurations regardless of whether a PLE is present or not.
Attributes	See Table 4-24 on page 4-30.

Figure 4-14 shows the PLEASR bit assignments.




Figure 4-14 PLEASR bit assignments

Table 4-26 shows the PLEASR bit assignments.

Table 4-26 PLEASR bit assignments

Bits	Name	Function
[31:1]	-	-
[0]	R	PLE Channel running 1 = The Preload Engine is currently handling a PLE request.

To access the PLEASR, use:

MRC p15, 0, <Rt>, c11, c0, 2; Read PLEASR

4.3.28 PLE FIFO Status Register

The PLEFSR characteristics are:

Purpose	Indicates how many entries remain available in the PLE FIFO.
Usage constraints	The PLEFSR is: <ul style="list-style-type: none"> • Common to Secure and Non-secure states • Accessible in User and Privileged modes, regardless of any configuration bit. NSAC.PLE controls Non-secure accesses.
Configurations	Available in all Cortex-A9 configurations regardless of whether a PLE is present or not.
Attributes	See Table 4-24 on page 4-30.

Figure 4-15 shows the PLEFSR bit assignments.




Figure 4-15 PLESFR bit assignments

Table 4-27 shows the PLEFSR bit assignments.

Table 4-27 PLESFR bit assignments

Bits	Name	Function
[31:5]	-	-
[4:0]	Available entries	Number of available entries in the PLE FIFO This is the difference between the total number of entries in the FIFO, which is configuration-specific, and the number of entries already programmed.

Use the PLESFR to check that an entry is available before programming a new PLE channel.

To access the PLESFR, use:

MRC p15, 0, <Rt>, c11, c0, 4; Read the PLESFR

4.3.29 Preload Engine User Accessibility Register

The PLEUAR characteristics are:

Purpose	Controls whether PLE operations are available in User mode.
Usage constraints	The PLEUAR is: <ul style="list-style-type: none"> • Common to Secure and Non-secure states • Accessible in User and Privileged modes, regardless of any configuration bit.
Configurations	Only available in configurations where the Preload Engine is present, otherwise an Undefined Instruction exception is taken.

Attributes See Table 4-24 on page 4-30.

Figure 4-16 shows the PLEUAR bit assignments.




Figure 4-16 PLEUAR bit assignments

Table 4-28 shows the PLEUAR bit assignments.

Table 4-28 PLEUAR bit assignments

Bits	Name	Function
[31:1]	-	RAZ.
[0]	U	User accessibility 1 = User modes can access PLE registers and execute PLE operations.

To access the PLEUAR, use:

MCR p15, 0, <Rt>, c11, c1, 0; Read PLEAUR
MRC p15, 0, <Rt>, c11, c1, 0; Write PLEAUR

4.3.30 Preload Engine Parameters Control Register

The PLEPCR characteristics are:

Purpose Contains PLE control parameters, available only in Privilege modes, to limit the issuing rate and transfer size of the PLE.

Usage constraints The PLEPCR is:

- Read/Write register
- only accessible in Privileged mode.
- Common to Secure and Non-secure states
- NSACR.PLE controls Non-secure accesses.

Configurations Only available in configurations where the Preload Engine is present, otherwise an Undefined Instruction exception is taken.

Attributes See Table 4-24 on page 4-30.

Figure 4-17 shows the PLEPCR bit assignments.




Figure 4-17 PLEPCR bit assignments

Table 4-29 shows the PLEPCR bit assignments.

Table 4-29 PLEPCR bit assignments

Bits	Name	Function
[31:30]	-	RAZ
[29:16]	Block size mask	Permits Privilege modes to limit the maximum block size for PLE transfers. The transferred block size is: (Block size) & (Block size mask). For example, a block size mask of 14'b1111111111111111 authorizes the transfer of block sizes with the maximum value of 16k * 4 bytes. A block size mask of 14'b0000000000000000 limits block sizes to 1 * 4 bytes.
[15:8]	Block number mask	Permits Privilege modes to limit the maximum number of blocks for a single PLE transfer. The transferred block number is: (Block number) & (Block number mask). For example, a block number mask of 8'b11111111 authorizes the transfer of a maximum possible number of 256 blocks. A block number mask of 8'b00000000 limits the transfer to only one block of data.
[7:0]	PLE wait states	Permit Privilege modes to limit the issuing rate of PLD requests performed by the PLE engine to prevent saturation of the external memory bandwidth. PLE wait states specifies the number of cycles inserted between two PLD requests performed by the PLE engine. When PLE wait states = 8'b11111111, the PLE engine can issue one PLD request, a cache line, every 256 cycles. When PLE wait states = 8'b00000000, the PLE engine can issue one PLD request every cycle.

To access the PLEPCR, use:

MCR p15, 0, <Rt>, c11, c1, 1; Read PLEPCR
MRC p15, 0, <Rt>, c11, c1, 1; Write PLEPCR

4.3.31 CP15 c12 register summary

Table 4-30 shows the system control registers you can access when CRn is c12.

Table 4-30 c12 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	VBAR	RW	0x00000000 ^a	Vector Base Address Register
		1	MVBAR	RW	-	Monitor Vector Base Address Register
c1		0	ISR	RO	0x00000000	Interrupt Status Register
		1	Virtualization Interrupt Register	RW	0x00000000	<i>Virtualization Interrupt Register</i>

a. Only the secure version is reset to 0. The Non-secure version must be programmed by software

4.3.32 Virtualization Interrupt Register

The VIR characteristics are:

Purpose Indicates that there is a virtual interrupt pending.

Usage constraints The VIR is:

- only accessible in privileged modes-
- only accessible in Secure state.

Configurations Available in all configurations.

Attributes See the register summary in Table 4-30 on page 4-34.

The virtual interrupt is delivered as soon as the processor is in NS state. Figure 4-18 shows the VIR bit assignments.




Figure 4-18 VIR bit assignments

Table 4-31 shows the Virtualization Interrupt Register bit assignments.

Table 4-31 Virtualization Interrupt Register bit assignments

Bits	Name	Function
[31:9]	-	UNK/SBZP.
[8]	VA	Virtual Abort bit. When set the corresponding Abort is sent to software in the same way as a normal Abort. The virtual abort happens only when the processor is in Non-secure state.
[7]	VI	Virtual IRQ bit. When set the corresponding IRQ is sent to software in the same way as a normal IRQ. The virtual IRQ happens only when the processor is in Non-secure state.
[6]	VF	Virtual FIQ bit. When set the corresponding FIQ is sent to software in the same way as a normal FIQ. The FIQ happens only when the processor is in Non-secure state.
[5:0]	-	UNK/SBZP.

To access the VIR, use:

MRC p15, 0, <Rd>, c12, c1, 1 ; Read Virtualization Interrupt Register
MCR p15, 0, <Rd>, c12, c1, 1 ; Write Virtualization Interrupt Register

4.3.33 CP15 c13 register summary

Table 4-32 shows the system control registers you can access when CRn is c13.

Table 4-32 c13 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	FCSEIDR	RW	0x00000000	Deprecated registers on page 4-2
		1	CONTEXTIDR	RW	-	Context ID Register
		2	TPIDRURW	RW ^a	-	Software Thread ID registers
		3	TPIDRURO	RO ^b	-	
		4	TPIDRPRW	RW	-	

- a. RW in User mode
- b. RO in User mode

4.3.34 CP15 c14, not used

No CP15 registers are accessed with CRn set to c4.

4.3.35 CP15 c15 register summary

Table 4-33 shows the system control registers you can access when CRn is c15.

Table 4-33 c15 system control registers

Op1	CRm	Op2	Name	Type	Reset	Description
0	c0	0	Power Control Register	RW ^a	b	<i>Power Control Register</i>
	c1	0	NEON busy Register	RO	0x00000000	<i>NEON busy Register</i> on page 4-37
4	c0	0	Configuration Base Address	RO ^c	d	<i>Configuration Base Address Register</i> on page 4-38
5	c4	2	Select Lockdown TLB Entry for read	WO ^e	-	<i>TLB lockdown operations</i> on page 4-39
		4	Select Lockdown TLB Entry for write	WO ^e	-	
	c5	2	Main TLB VA register	RW ^e	-	
	c6	2	Main TLB PA register	RW ^e	-	
	c7	2	Main TLB Attribute register	RW	-	

a. RW in Secure state. Read only in Non-secure state.

b. Reset value depends on the **MAXCLKLATENCY[2:0]** value. See *Configuration signals* on page A-5.

c. RW in Secure privileged mode and RO in Non-secure state and user secure state.

d. In Cortex-A9 uniprocessor implementations the configuration base address is set to zero.

In Cortex-A9 MPCore implementations the configuration base address is reset to **PERIPHBASE[31:13]** so that software can determine the location of the private memory region.

e. No access in Non-secure state.

4.3.36 Power Control Register

The Power Control Register characteristics are:

Purpose	Enables you to set: <ul style="list-style-type: none"> • the clock latency for your implementation of the Cortex-A9 processor. • dynamic clock gating.
Usage constraints	<ul style="list-style-type: none"> • a read and write register in Secure state • a read-only register in Non-secure state
Configurations	Available in all configurations.
Attributes	See the register summary in Table 4-33.

Figure 4-19 on page 4-37 shows the Power Control Register bit assignments.




Figure 4-19 Power Control Register bit assignments

Table 4-34 shows the Power Control Register bit assignments.

Table 4-34 Power Control Register bit assignments

Bits	Name	Function
[31:11]	-	Reserved.
[10:8]	Max clock latency	Samples the value present on the MAXCLKLATENCY pins on exit from reset. This value reflects an implementation specific parameter, and ARM recommends that the software does not modify it.
[7:1]	-	Reserved.
[0]	Enable dynamic clock gating	Disabled at reset.

To access the Power Control Register, use:

MRC p15,0,<Rd>,c15,c0,0; Read Power Control Register

MCR p15,0,<Rd>,c15,c0,0; Write Power Control Register

4.3.37 NEON busy Register

The NEON busy Register characteristics are:

Purpose Enables software to determine if a NEON instruction is executing.

Usage constraints

- a read-only register in Secure state
- a read-only register in Non-secure state

Configurations Available in all configurations.

Attributes See the register summary in Table 4-33 on page 4-36.

Figure 4-20 shows the NEON busy register bit assignments




Figure 4-20 NEON busy register bit assignments

Table 4-35 shows the NEON busy Register bit assignments.

Table 4-35 Neon busy Register bit assignments

Bits	Name	Function
[31:1]	-	Reserved.
[0]	NEON busy	Software can use this to determine if a NEON instruction is executing. This bit is set to 1 if there is a NEON instruction in NEON pipeline, or in the core pipeline

To access the NEON busy Register, use:

```
MRC p15,0,<Rd>,c15,c1,0; Read NEON busy Register
```

4.3.38 Configuration Base Address Register

The Configuration Base Address Register characteristics are:

Purpose Takes the physical base address value at reset

Usage constraints The Configuration Base Address Register is:

- read and write in Secure privileged modes.
- read only in Non-secure state.
- read only in user mode.

Configurations In Cortex-A9 uniprocessor implementations the base address is set to zero. In Cortex-A9 MPCore implementations it is reset to **PERIPHBASE[31:13]** so that software can determine the location of the private memory region.

Attributes See the register summary in Table 4-33 on page 4-36.

Figure 4-21 shows the Configuration Base Address Register bit assignments.




Figure 4-21 Configuration Base Address Register bit assignments

To access the Configuration Base Address Register, use:

```
MRC p15,4,<Rd>,c15,c0,0; Read Configuration Base Address Register
MCR p15,4,<Rd>,c15,c0,0; Write Configuration Base Address Register
```

4.3.39 TLB lockdown operations

TLB lockdown operations enable saving or restoring lockdown entries in the TLB. Table 4-36 shows the defined TLB lockdown operations.

Table 4-36 TLB lockdown operations

Description	Data	Instruction
Select Lockdown TLB Entry for Read	Main TLB Index	MCR p15, 5, <Rd>, c15, c4, 2
Select Lockdown TLB Entry for Write	Main TLB Index	MCR p15, 5, <Rd>, c15, c4, 4
Read Lockdown TLB VA Register	Data	MRC p15, 5, <Rd>, c15, c5, 2
Write Lockdown TLB VA Register	Data	MCR p15, 5, <Rd>, c15, c5, 2
Read Lockdown TLB PA Register	Data	MRC p15, 5, <Rd>, c15, c6, 2
Write Lockdown TLB PA Register	Data	MCR p15, 5, <Rd>, c15, c6, 2
Read Lockdown TLB attributes Register	Data	MRC p15, 5, <Rd>, c15, c7, 2
Write Lockdown TLB attributes Register	Data	MCR p15, 5, <Rd>, c15, c7, 2

The Select Lockdown TLB entry for a read operation is used to select the entry that the data read by a read Lockdown TLB VA/PA/attributes operations are coming from. The Select Lockdown TLB entry for a write operation is used to select the entry that the data write Lockdown TLB VA/PA/attributes data are written to. The TLB PA register must be the last written/read register when accessing TLB lockdown registers. Figure 4-22 shows the bit assignment of the index register used to access the lockdown TLB entries.




Figure 4-22 Lockdown TLB index bit assignments

Figure 4-23 shows the bit arrangement of the TLB VA Register format.




Figure 4-23 TLB VA Register bit assignments

Table 4-37 shows the TLB VA Register bit assignments.

Table 4-37 TLB VA Register bit assignments

Bits	Name	Function
[31:12]	VPN	Virtual page number. Bits of the virtual page number that are not translated as part of the page table translation because the size of the tables is Unpredictable when read and SBZ when written.
[11]	-	UNK/SBZP.
[10]	NS	NS bit.
[9:0]	Process	Memory space identifier.

Figure 4-24 shows the bit arrangement of the memory space identifier.




Figure 4-24 Memory space identifier format

Figure 4-25 shows the TLB PA Register bit assignment.



Figure 4-25 TLB PA Register bit assignments

Table 4-38 describes the functions of the TLB PA Register bits.

Table 4-38 TLB PA Register bit assignments

Bits	Name	Function
[31:12]	PPN	Physical Page Number. Bits of the physical page number that are not translated as part of the page table translation are unpredictable when read and SBZP when written.
[11:8]	-	UNK/SBZP.
[7:6]	SZ	Region Size. b00 = 16MB Supersection. b01 = 4KB page. b10 = 64KB page. b11 = 1MB section. All other values are reserved.

Table 4-38 TLB PA Register bit assignments (continued)

Bits	Name	Function
[5:4]	-	UNK/SBZP.
[3:1]	AP	Access permission: b000 = All accesses generate a permission fault. b001 = Supervisor access only, User access generates a fault. b010 = Supervisor read and write access, User write access generates a fault. b011 = Full access, no fault generated. b100 = Reserved. b101 = Supervisor read only. b110 = Supervisor/User read only. b111 = Supervisor/User read only.
[0]	V	Value bit. Indicates that this entry is locked and valid.

Figure 4-26 shows the bit assignments of the TLB Attributes Register.




Figure 4-26 Main TLB Attributes Register bit assignments

Table 4-39 shows the TLB Attributes Register bit assignments. The Cortex-A9 processor does not support subpages.

Table 4-39 TLB Attributes Register bit assignments

Bits	Name	Function
[31:12]	-	UNK/SBZP.
[11]	NS	Non-secure description.
[10:7]	Domain	Domain number of the TLB entry.
[6]	XN	Execute Never attribute.
[5:3]	TEX	Region type encoding. See the <i>ARM Architecture Reference Manual</i> .
[2:1]	CB	
[0]	S	Shared attribute.

Invalidate TLB Entries on ASID Match

This is a single interruptible operation that invalidates all TLB entries that match the provided *Address Space Identifier* (ASID) value. This function invalidates locked entries. Entries marked as global are not invalidated by this function.

In the Cortex-A9 processor, this operation takes several cycles to complete and the instruction is interruptible. When interrupted the r14 state is set to indicate that the MCR instruction has not executed. Therefore, r14 points to the address of the MCR + 4. The interrupt routine then

automatically restarts at the MCR instruction. If this operation is interrupted and later restarted, any entries fetched into the TLB by the interrupt that uses the provided ASID are invalidated by the restarted invalidation.

Chapter 5

Jazelle DBX registers

This chapter introduces the CP14 coprocessor and describes the non-debug use of CP14. It contains the following sections:

- *About coprocessor CP14* on page 5-2
- *CP14 Jazelle register summary* on page 5-3
- *CP14 Jazelle register descriptions* on page 5-4.

5.1 About coprocessor CP14

The purpose of non-debug use of coprocessor CP14 is to provide support for the hardware acceleration of Java bytecodes.

5.2 CP14 Jazelle register summary

In the Cortex-A9 implementation of the Jazelle Extension:

- Jazelle state is supported.
- The BXJ instruction enters Jazelle state.

Table 5-1 shows the CP14 Jazelle registers. For all Jazelle register accesses, CRm and Op2 are zero. All Jazelle registers are 32 bits wide.

Table 5-1 CP14 Jazelle registers summary

Op1	CRn	Name	Type	Reset	Page
7	0	Jazelle ID Register (JIDR)	RW ^a	0xF4100168	page 5-4
7	1	Jazelle OS Control Register (JOSCR)	RW	-	page 5-5
7	2	Jazelle Main Configuration Register (JMCR)	RW	-	page 5-6
7	3	Jazelle Parameters Register	RW	-	page 5-7
7	4	Jazelle Configurable Opcode Translation Table Register	WO	-	page 5-8

a. See *Write operation of the JIDR* on page 5-5 for the effect of a write operation.

See the *ARM Architecture Reference Manual* for details of the Jazelle Extension.

5.3 CP14 Jazelle register descriptions

The following sections describe the CP14 Jazelle DBX registers arranged in numerical order, as shown in Table 5-1 on page 5-3:

- *Jazelle ID Register*
- *Jazelle Operating System Control Register* on page 5-5
- *Jazelle Main Configuration Register* on page 5-6
- *Jazelle Parameters Register* on page 5-7
- *Jazelle Configurable Opcode Translation Table Register* on page 5-8.

5.3.1 Jazelle ID Register

The JIDR characteristics are:

Purpose	Enables software to determine the implementation of the Jazelle Extension provided by the processor.
Usage constraints	The JIDR is: <ul style="list-style-type: none"> • accessible in privileged modes. • also accessible in user mode if the CD bit is clear. See <i>Jazelle Operating System Control Register</i> on page 5-5.
Configurations	Available in all configurations.
Attributes	See the register summary in Table 5-1 on page 5-3.

Figure 5-1 shows the JIDR bit assignments.




Figure 5-1 JIDR bit assignment

Table 5-2 shows the JIDR bit assignments.

Table 5-2 JIDR bit assignments

Bits	Name	Function
[31:28]	Arch	This uses the same architecture code that appears in the Main ID register.
[27:20]	Design	Contains the implementor code of the designer of the subarchitecture.
[19:12]	SArchMajor	The subarchitecture code.
[11:8]	SArchMinor	The subarchitecture minor code.
[7]	-	RAZ
[6]	TrTbleFrm	Indicates the format of the Jazelle Configurable Opcode Translation Table Register.
[5:0]	TrTbleSz	Indicates the size of the Jazelle Configurable Opcode Translation Table Register.

To access the JIDR, use:

MRC p14, 7, <Rd>, c0, c0, 0; Read Jazelle Identity Register

Write operation of the JIDR

A write to the JIDR clears the translation table. This has the effect of making all configurable opcodes executed in software only. See *Jazelle Configurable Opcode Translation Table Register* on page 5-8.

5.3.2 Jazelle Operating System Control Register

The JOSCR characteristics are:

- Purpose** Enables operating systems to control access to Jazelle Extension hardware.
- Usage constraints** The JOSCR is:
 - only accessible in privileged modes.
 - set to zero after a reset and must be written in privileged modes.
- Configurations** Available in all configurations.
- Attributes** See the register summary in Table 5-1 on page 5-3.

Figure 5-2 shows the JOSCR bit assignments.




Figure 5-2 JOSCR bit assignments

Table 5-3 shows the JOSCR bit assignments.

Table 5-3 JOSCR bit assignments

Bits	Name	Function
[31:2]	-	Reserved, RAZ.
[1]	CV	<p>Configuration Valid bit.</p> <p>0 = The Jazelle configuration is invalid. Any attempt to enter Jazelle state when the Jazelle hardware is enabled:</p> <ul style="list-style-type: none"> generates a configuration invalid Jazelle exception sets this bit, marking the Jazelle configuration as valid. <p>1 = The Jazelle configuration is valid. Entering Jazelle state succeeds when the Jazelle hardware is enabled.</p> <p>The CV bit is automatically cleared on an exception.</p>
[0]	CD	<p>Configuration Disabled bit.</p> <p>0 = Jazelle configuration in User mode is enabled:</p> <ul style="list-style-type: none"> reading the JIDR succeeds reading any other Jazelle configuration register generates an Undefined Instruction exception writing the JOSCR generates an Undefined Instruction exception writing any other Jazelle configuration register succeeds. <p>1 = Jazelle configuration from User mode is disabled:</p> <ul style="list-style-type: none"> reading any Jazelle configuration register generates an Undefined Instruction exception writing any Jazelle configuration register generates an Undefined Instruction exception.

To access the JOSCR, use:

```
MRC p14, 7, <Rd>, c1, c0, 0; Read JOSCR
MCR p14, 7, <Rd>, c1, c0, 0; Write JOSCR
```

5.3.3 Jazelle Main Configuration Register

The JMCR characteristics are:

- Purpose** Describes the Jazelle hardware configuration and its behavior.
- Usage constraints** Only accessible in privileged modes.
- Configurations** Available in all configurations.
- Attributes** See the register summary in Table 5-1 on page 5-3.

Figure 5-3 shows the JMCR bit assignments.




Figure 5-3 JMCR bit assignments

Table 5-4 shows the JMCR bit assignments.

Table 5-4 JMCR bit assignments

Bits	Name	Function
[31]	nAR	<i>not Array Operations</i> (nAR) bit. 0 = Execute array operations in hardware, if implemented. Otherwise, call the appropriate handlers in the VM Implementation Table. 1 = Execute all array operations by calling the appropriate handlers in the VM Implementation Table.
[30]	FP	The FP bit controls how the Jazelle hardware executes JVM floating-point opcodes: 0 = Execute all JVM floating-point opcodes by calling the appropriate handlers in the VM Implementation Table. 1 = Execute JVM floating-point opcodes by issuing VFP instructions, where possible. Otherwise, call the appropriate handlers in the VM Implementation Table. In this implementation FP is set to zero and is read only.
[29]	AP	The <i>Array Pointer</i> (AP) bit controls how the Jazelle hardware treats array references on the operand stack: 0 = Array references are treated as handles. 1 = Array references are treated as pointers.
[28]	OP	The <i>Object Pointer</i> (OP) bit controls how the Jazelle hardware treats object references on the operand stack: 0 = Object references are treated as handles. 1 = Object references are treated as pointers.
[27]	IS	The <i>Index Size</i> (IS) bit specifies the size of the index associated with quick object field accesses: 0 = Quick object field indices are 8 bits. 1 = Quick object field indices are 16 bits.
[26]	SP	The <i>Static Pointer</i> (SP) bit controls how the Jazelle hardware treats static references: 0 = Static references are treated as handles. 1 = Static references are treated as pointers.
[25:1]	-	UNK/SBZP.
[0]	JE	The <i>Jazelle Enable</i> (JE) bit controls whether the Jazelle hardware is enabled, or is disabled: 0 = The Jazelle hardware is disabled: <ul style="list-style-type: none"> • BXJ instructions behave like BX instructions • setting the J bit in the CPSR generates a Jazelle-Disabled Jazelle exception. 1 = The Jazelle hardware is enabled: <ul style="list-style-type: none"> • BXJ instructions enter Jazelle state • setting the J bit in the CPSR enters Jazelle state.

To access the JMCR, use:

```
MRC p14, 7, <Rd>, c2, c0, 0; Read JMCR
MCR p14, 7, <Rd>, c2, c0, 0; Write JMCR
```

5.3.4 Jazelle Parameters Register

The Jazelle Parameters Register characteristics are:

Purpose Describes the parameters that configure how the Jazelle hardware behaves.

Usage constraints Only accessible in privileged modes.

Configurations Available in all configurations.

Attributes See the register summary in Table 5-1 on page 5-3.

Figure 5-4 shows the Jazelle Parameters Register bit assignments.

31	22 21	17 16	12 11	8 7	4 3	0
UNK/SBZP	BSH	sADO	ARO	STO	ODO	

Figure 5-4 Jazelle Parameters Register bit assignments

Table 5-5 shows the Jazelle Parameters Register bit assignments.

Table 5-5 Jazelle Parameters Register bit assignments

Bits	Name	Function
[31:22]	-	UNK/SBZP.
[21:17]	BSH	The <i>Bounds SHift</i> (BSH) bits contain the offset, in bits, of the array bounds (number of items in the array) within the array descriptor word.
[16:12]	sADO	The <i>signed Array Descriptor Offset</i> (sADO) bits contain the offset, in words, of the array descriptor word from an array reference. The offset is a sign-magnitude signed quantity: <ul style="list-style-type: none"> • Bit [16] gives the sign of the offset. The offset is positive if the bit is clear, and negative if the bit is set. • Bits [15:12] give the absolute magnitude of the offset.
[11:8]	ARO	The <i>Array Reference Offset</i> (ARO) bits contain the offset, in words, of the array data or the array data pointer from an array reference.
[7:4]	STO	The <i>Static Offset</i> (STO) bits contain the offset, in words, of the static or static pointer from a static reference.
[3:0]	ODO	The <i>Object Descriptor Offset</i> (ODO) bits contain the offset, in words, of the field from the base of an object data block.

To access the Jazelle Parameters Register, use:

MRC p14, 7, <Rd>, c3, c0, 0; Read Jazelle Parameters Register
MCR p14, 7, <Rd>, c3, c0, 0; Write Jazelle Parameters Register

5.3.5 Jazelle Configurable Opcode Translation Table Register

The Jazelle Configurable Opcode Translation Table Register characteristics are:

Purpose Provides translations between the configurable opcodes in the range 0xCB-0xFD and the operations that are provided by the Jazelle hardware.

Usage constraints Only accessible in privileged modes.

Configurations Available in all configurations.

Attributes See the register summary in Table 5-1 on page 5-3.

Figure 5-5 on page 5-9 shows the Jazelle Configurable Opcode Translation Table Register bit assignments.




Figure 5-5 Jazelle Configurable Opcode Translation Table Register bit assignments

Table 5-6 shows the Jazelle Configurable Opcode Translation Table Register bit assignments.

Table 5-6 Jazelle Configurable Opcode Translation Table Register bit assignments

Bits	Name	Function
[31:16]	-	UNK/SBZP.
[15:10]	Opcode	Contains the bottom bits of the configurable opcode.
[9:4]	-	UNK/SBZP.
[3:0]	Operation	Contains the code for the operation 0x0-0x9.

To access this register, use:

MRC p14, 7, <Rd>, c4, c0, 0; Read Jazelle Configurable Opcode Translation Table Register

MCR p14, 7, <Rd>, c4, c0, 0; Write Jazelle Configurable Opcode Translation Table Register

Chapter 6

Memory Management Unit

This chapter describes the MMU. It contains the following sections:

- *About the MMU* on page 6-2
- *TLB Organization* on page 6-4
- *Memory Access Sequence* on page 6-6
- *MMU enabling or disabling* on page 6-7
- *External aborts* on page 6-8.

6.1 About the MMU

The MMU works with the L1 and L2 memory system to translate virtual addresses to physical addresses. It also controls accesses to and from external memory.

The *Virtual Memory System Architecture version 7* (VMSAv7) features include the following:

- page table entries that support 4KB, 64KB, 1MB, and 16MB
- 16 domains
- global and address space identifiers to remove the requirement for context switch TLB flushes
- extended permissions check capability.

See the *ARM Architecture Reference Manual* for a full architectural description of the VMSAv7.

The processor implements the ARMv7-A MMU enhanced with security extensions and multiprocessor extensions to provide address translation and access permission checks. The MMU controls table walk hardware that accesses translation tables in main memory. The MMU enables fine-grained memory system control through a set of virtual-to-physical address mappings and memory attributes.

———— Note —————

In VMSAv7 first level descriptor formats page table base address bit 9 is implementation defined. In Cortex-A9 processor designs this bit is unused.

The MMU features include the following:

- Instruction side micro TLB
 - 32 fully associative entries
- Data side micro TLB
 - 32 fully associative entries
- Unified main TLB
 - unified, 2-way associative, 2x32 entry TLB for the 64-entry TLB and 2x64 entry TLB for the 128-entry TLB.
 - 4 lockable entries using the lock-by-entry model
 -
 - supports hardware page table walks to perform look-ups in the L1 data cache.

6.1.1 Memory Management Unit

The MMU performs the following operations:

- checking of Virtual Address and ASID
- checking of domain access permissions
- checking of memory attributes
- virtual-to-physical address translation
- support for four page (region) sizes
- mapping of accesses to cache, or external memory
- TLB loading for hardware and software.

Domains

The Cortex-A9 processor supports sixteen access domains.

TLB

The Cortex-A9 processor implements a two-level TLB structure. Four entries in the main TLB are lockable.

ASIDs

Main TLB entries can be global, or can be associated with particular processes or applications using Address Space Identifiers (ASIDs). ASIDs enable TLB entries to remain resident during context switches, avoiding the requirement of reloading them subsequently. See *Invalidate TLB Entries on ASID Match* on page 4-41.

System control coprocessor

TLB maintenance and configuration operations are controlled through a dedicated coprocessor, CP15, integrated within the core. This coprocessor provides a standard mechanism for configuring the level one memory system.

6.2 TLB Organization

The following sections describe the organization of the TLB:

- *Micro TLB*
- *Main TLB.*

6.2.1 Micro TLB

The first level of caching for the page table information is a micro TLB of 32 entries that is implemented on each of the instruction and data sides. These blocks provide a fully associative look-up of the virtual addresses in a single **CLK** cycle.

The micro TLB returns the physical address to the cache for the address comparison, and also checks the protection attributes to signal either a Prefetch Abort or a Data Abort.

All main TLB related operations affect both the instruction and data micro TLBs, causing them to be flushed. In the same way, any change of the Context ID Register causes the micro TLBs to be flushed.

6.2.2 Main TLB

The main TLB catches the misses from the Micro TLBs. It also provides a centralized source for lockable translation entries.

Accesses to the main TLB take a variable number of cycles, according to competing requests from each of the micro TLBs and other implementation-dependent factors. Entries in the lockable region of the main TLB are lockable at the granularity of a single entry. As long as the lockable region does not contain any locked entries, it can be allocated with non-locked entries to increase overall main TLB storage size.

The main TLB is implemented as a combination of:

- a fully-associative, lockable array of four elements
- a two-way associative structure of 2x32 or 2x64 entries.

TLB match process

Each TLB entry contains a virtual address, a page size, a physical address, and a set of memory properties. Each is marked as being associated with a particular application space, or as global for all application spaces. CONTEXIDR determines the currently selected application space. A TLB entry matches if bits [31:N] of the modified virtual address match, where N is \log_2 of the page size for the TLB entry. It is either marked as global, or the ASID matches the current ASID.

A TLB entry matches when these conditions are true:

- its virtual address matches that of the requested address
- its Non-secure TLB ID (NSTID) matches the Secure or Non-secure state of the MMU request
- its ASID matches the current ASID or is global.

The operating system must ensure that, at most, one TLB entry matches at any time.

Supersections, sections, and large pages are supported to permit mapping of a large region of memory while using only a single entry in a TLB. If no mapping for an address is found in the TLB, then the translation table is automatically read by hardware and a mapping is placed in the TLB.

TLB lockdown

The TLB supports the TLB lock-by-entry model as described in the *ARM Architecture Reference Manual*. See *TLB lockdown operations* on page 4-39 for more information.

6.3 Memory Access Sequence

When the processor generates a memory access, the MMU:

1. Performs a look-up for the requested virtual address and current ASID and security state in the relevant instruction or data micro TLB.
2. If there is a miss in the micro TLB, performs a look-up for the requested virtual address and current ASID and security state in the main TLB.
3. If there is a miss in main TLB, performs a hardware translation table walk.

You can configure the MMU to perform hardware translation table walks in cacheable regions by setting the IRGN bits in the Translation Table Base Registers. If the encoding of the IRGN bits is write-back, then an L1 data cache look-up is performed and data is read from the data cache. If the encoding of the IRGN bits is write-through or non-cacheable then an access to external memory is performed.

The MMU might not find a global mapping, or a mapping for the currently selected ASID, with a matching *Non-secure TLB ID* (NSTID) for the virtual address in the TLB. In this case, the hardware does a translation table walk if the translation table walk is enabled by the PD0 or PD1 bit in the TTB Control Register. If translation table walks are disabled, the processor returns a Section Translation fault.

If the MMU finds a matching TLB entry, it uses the information in the entry as follows:

1. The access permission bits and the domain determine if the access is enabled. If the matching entry does not pass the permission checks, the MMU signals a memory abort. See the *ARM Architecture Reference Manual* for a description of access permission bits, abort types and priorities, and for a description of the IFSR and *Data Fault Status Register* (DFSR).
2. The memory region attributes specified in both the TLB entry and the CP15 c10 remap registers control the cache and write buffer, and determine if the access is
 - Secure or Non-secure
 - Shared or not
 - Normal memory, Device, or Strongly-ordered.
3. The MMU translates the virtual address to a physical address for the memory access.

If the MMU does not find a matching entry, a hardware table walk occurs.

6.4 MMU enabling or disabling

You can enable or disable the MMU as described in the *ARM Architecture Reference Manual*.

6.5 External aborts

External memory errors are defined as those that occur in the memory system rather than those that are detected by the MMU. External memory errors are expected to be extremely rare. External aborts are caused by errors flagged by the AXI interfaces when the request goes external to the processor. External aborts can be configured to trap to Monitor mode by setting the EA bit in the Secure Configuration Register.

6.5.1 External aborts on data read or write

Externally generated errors during a data read or write can be asynchronous. This means that the `r14_abt` on entry into the abort handler on such an abort might not hold the address of the instruction that caused the exception.

The DFAR is Unpredictable when an asynchronous abort occurs.

In the case of a load multiple or store multiple operation, the address captured in the DFAR is that of the address that generated the synchronous external abort.

6.5.2 Synchronous and asynchronous aborts

To determine a fault type, read the DFSR for a data abort or the IFSR for an instruction abort.

The processor supports an Auxiliary Fault Status Register for software compatibility reasons only. The processor does not modify this register because of any generated abort.

Chapter 7

Level 1 Memory System

This chapter describes the L1 Memory System. It contains the following sections:

- *About the L1 memory system* on page 7-2
- *Security Extensions support* on page 7-4
- *About the L1 instruction side memory system* on page 7-5
- *About the L1 data side memory system* on page 7-8
- *About DSB* on page 7-9
- *Data prefetching* on page 7-10
- *Parity error support* on page 7-11.

7.1 About the L1 memory system

The L1 memory system has:

- separate instruction and data caches each with a fixed line length of 32 bytes
- 64-bit data paths throughout the memory system
- support for four sizes of memory page
- export of memory attributes for external memory systems
- support for Security Extensions.

The data side of the L1 memory system has:

- two 32-byte linefill buffers and one 32-byte eviction buffer
- a 4-entry, 64-bit merging store buffer.

———— Note —————

You must invalidate the instruction cache, the data cache, TLB, and BTAC before using them. You are not required to invalidate the main TLB, even though it is recommended for safety reasons. This ensures compatibility with future revisions of the processor.

7.1.1 Memory system

This section describes:

- *Cache features*
- *Store buffer* on page 7-3.

Cache features

The Cortex-A9 processor has separate instruction and data caches. The caches have the following features:

- Each cache can be disabled independently. See *System Control Register* on page 4-15.
- Cache replacement policy is either pseudo round-robin or pseudo random.
- Both caches are 4-way set-associative.
- The cache line length is eight words.
- On a cache miss, critical word first filling of the cache is performed.
- You can configure the instruction and data caches independently during implementation to sizes of 16KB, 32KB, or 64KB.
- To reduce power consumption, the number of full cache reads is reduced by taking advantage of the sequential nature of many cache operations. If a cache read is sequential to the previous cache read, and the read is within the same cache line, only the data RAM set that was previously read is accessed.

Instruction cache features

The instruction cache is virtually indexed and physically tagged.

Data cache features

The data cache is physically indexed and physically tagged.

Both data cache read misses and write misses are non-blocking with up to four outstanding data cache read misses and up to four outstanding data cache write misses being supported.

Store buffer

The Cortex-A9 CPU has a store buffer with four 64-bit slots with data merging capability.

7.2 Security Extensions support

The Cortex-A9 processor supports the Security Extensions, and exports the Secure or Non-secure status of its memory requests to the memory system.

7.3 About the L1 instruction side memory system

The L1 instruction side memory system is responsible for providing an instruction stream to the Cortex-A9 processor. To increase overall performance and to reduce power consumption, it contains the following functionality:

- dynamic branch prediction
- instruction caching.

Figure 7-1 shows this.

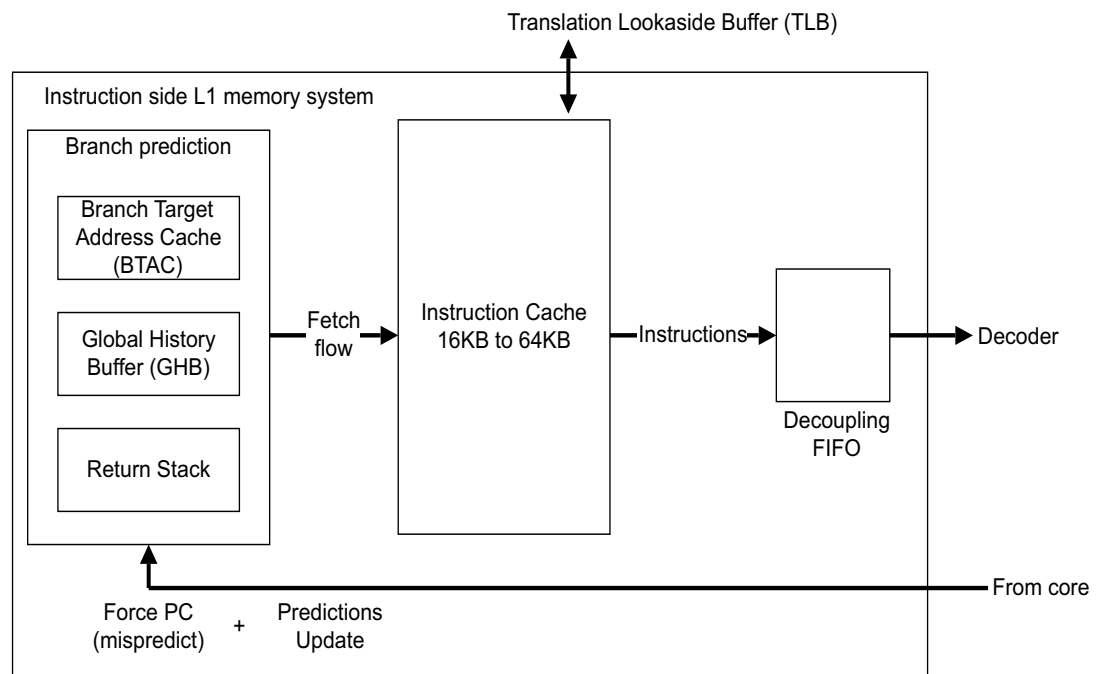


Figure 7-1 Branch prediction and instruction cache

The ISide comprises the following:

The Prefetch Unit (PFU)

The Prefetch Unit implements a two-level prediction mechanism, comprising:

- a two-way BTAC of 512 entries organized as two-way x 256 entries implemented in RAMs.
- a *Global History Buffer* (GHB) containing 4096 2-bit predictors implemented in RAMs
- a return stack with eight 32-bit entries.

The prediction scheme is available in ARM state, Thumb state, ThumbEE state, and Jazelle state. It is also capable of predicting state changes from ARM to Thumb, and from Thumb to ARM. It does not predict any other state changes. Nor does it predict any instruction that changes the mode of the core. See *Program flow prediction* on page 7-6.

Instruction Cache Controller

The instruction cache controller fetches the instructions from memory depending on the program flow predicted by the prefetch unit.

The instruction cache is 4-way set associative. It comprises the following features:

- configurable sizes of 16KB, 32KB, or 64KB
- *Virtually Indexed Physically Tagged (VIPT)*
- 64-bit native accesses so as to provide up to four instructions per cycle to the prefetch unit
- security extensions support
- no lockdown support.

7.3.1 Enabling program flow prediction

You can enable program flow prediction by setting the Z bit in the CP15 c1 Control Register to 1. See *System Control Register* on page 4-15. Before switching program flow prediction on, you must perform a BTAC flush operation.

This has the additional effect of setting the GHB into a known state.

7.3.2 Program flow prediction

The following sections describe program flow prediction:

- *Predicted and non-predicted instructions*
- *Thumb state conditional branches*
- *Return stack predictions* on page 7-7.

Predicted and non-predicted instructions

This section shows the instructions that the processor predicts. Unless otherwise specified, the list applies to ARM, Thumb, ThumbEE, and Jazelle instructions. As a general rule, the flow prediction hardware predicts all branch instructions regardless of the addressing mode, including:

- conditional branches
- unconditional branches
- indirect branches
- PC-destination data-processing operations
- branches that switch between ARM and Thumb states.

However, some branch instructions are nonpredicted:

- branches that switch between states (except ARM to Thumb transitions, and Thumb to ARM transitions)
- Instructions with the S suffix are not predicted as they are typically used to return from exceptions and have side effects that can change privilege mode and security state.
- All mode changing instructions.

Thumb state conditional branches

In Thumb state, a branch that is normally encoded as unconditional can be made conditional by inclusion in an *If-Then-Else* (ITE) block. Then it is treated as a normal conditional branch.

Return stack predictions

The return stack stores the address and the instruction execution state of the instruction after a function-call type branch instruction. This address is equal to the link register value stored in r14. The following instructions cause a return stack push if predicted:

- BL immediate
- BLX(1) immediate
- BLX(2) register
- HBL (ThumbEE state)
- HBLP (ThumbEE state).

The following instructions cause a return stack pop if predicted:

- BX r14
- MOV pc, r14
- LDM r13, {...pc}
- LDR pc, [r13].

The LDR instruction can use any of the addressing modes, as long as r13 is the base register. Additionally, in ThumbEE state you can also use r9 as a stack pointer so the LDR and LDM instructions with pc as a destination and r9 as a base register are also treated as a return stack pop.

Because return-from-exception instructions can change processor privilege mode and security state, they are not predicted. This includes the LDM(3) instruction, and the MOVS pc, r14 instruction.

7.4 About the L1 data side memory system

The L1 data cache is organized as a physically indexed and physically tagged cache. The micro TLB produces the physical address from the virtual address before performing the cache access.

7.4.1 Local Monitor

The Cortex-A9 processor L1 memory system has a Local 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 to construct semaphores, ensuring synchronization between different processes running on the CPU, and also between different processors that are using the same coherent memory locations for the semaphore.

———— **Note** —————

A store exclusive can generate an MMU fault or cause the processor to take a data watchpoint exception regardless of the state of the local monitor. See Table 10-8 on page 10-11

See the *ARM Architecture Reference Manual* for more information about these instructions.

Treatment of intervening STR operations

In cases where there is an intervening STR operation in an LDREX/STREX code sequence, the intermediate STR does not produce any effect on the internal exclusive monitor. The local monitor is in the Exclusive Access state after the LDREX, remains in the Exclusive Access state after the STR, and returns to the Open Access state only after the STREX.

LDREX/STREX operations using different sizes

In cases where the LDREX and STREX operations are of different sizes a check is performed to ensure that the tagged address bytes match or are within the size range of the store operation.

The granularity of the tagged address for an LDREX instruction is eight words, aligned on an eight-word boundary. This size is implementation defined, and as such, software must not rely on this granularity remaining constant on other ARM cores.

7.4.2 External aborts handling

The L1 data cache handles two types of external abort depending on the attributes of the memory region of the access:

- All Strongly-ordered accesses use the synchronous abort mechanism.
- All Cacheable, Device, and Normal Non-cacheable memory requests use the asynchronous abort mechanism. For example, an abort returned on a read miss, issuing a linefill, is flagged as asynchronous.

7.5 About DSB

The Cortex-A9 processor only implements the SY option of the DSB instruction. All other DSB options execute as a full system DSB operation, but software must not rely on this operation.

7.6 Data prefetching

This section describes:

- *The PLD instruction*
- *Data prefetching and monitoring.*

7.6.1 The PLD instruction

All PLD instructions are handled in a dedicated unit in the Cortex-A9 processor with dedicated resources. This avoids using resources in the integer core or the Load Store Unit

7.6.2 Data prefetching and monitoring

The Cortex-A9 data cache implements an automatic prefetcher that monitors cache misses done by the processor. This unit can monitor and prefetch two independent data streams. It can be activated in software using a CP15 Auxiliary Control Register bit. See *Auxiliary Control Register* on page 4-18.

When the software issues a PLD instruction the PLD prefetch unit always takes precedence over requests from the data prefetch mechanism. Prefetched lines in the speculative prefetcher can be dropped before they are allocated. PLD instructions are always executed and never dropped.

7.7 Parity error support

If your configuration implements parity error support, the features are as follows:

- the parity scheme is even parity. For byte 0000000 parity is 0.
- each RAM in the design generates parity information. As a general rule each RAM byte generates one parity bit. Where RAM bit width is not a multiple of eight, the remaining bits produce one parity bit.

There is also support for parity bit-writable data.

- RAM arrays in a design with parity support store parity information alongside the data in the RAM banks. As a result RAM arrays are wider when your design implements parity support.
- The Cortex-A9 logic includes the additional parity generation logic and the parity checking logic.

Figure 7-2 shows the parity support design features and stages. In stages 1 and 2 RAM writes and parity generation take place in parallel. RAM reads and parity checking take place in parallel in stages 3 and 4.

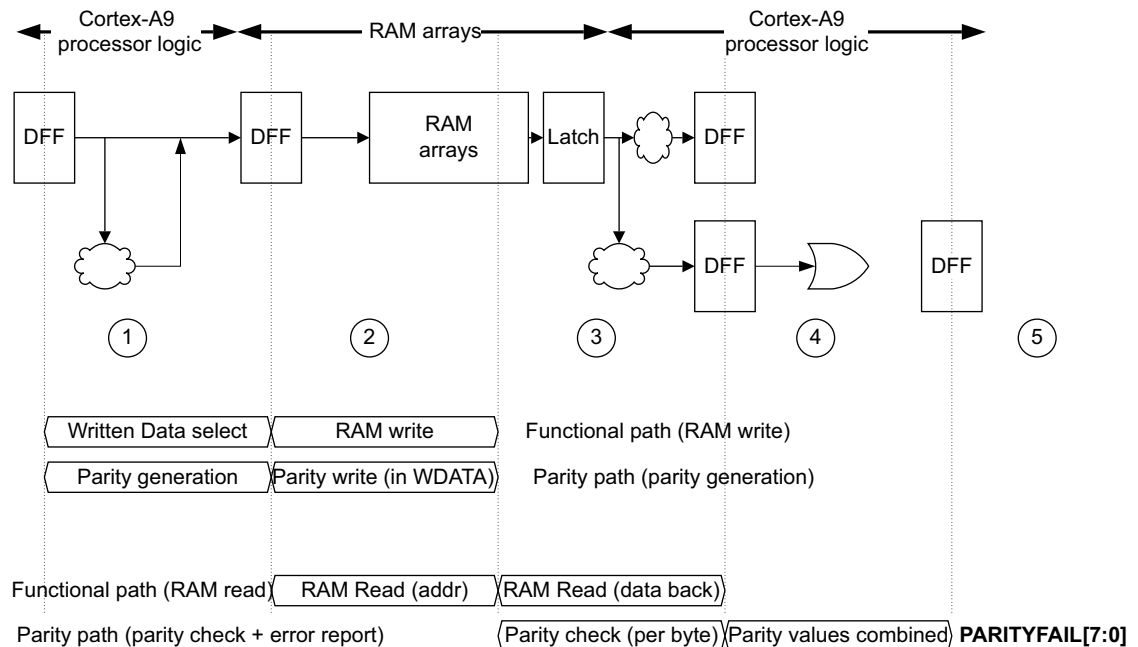


Figure 7-2 Parity support

The output signals **PARITYFAIL[7:0]** report parity errors. Typically, **PARITYFAIL[7:0]** reports parity errors 3 clock cycles after the corresponding RAM read.

Note

This is not a precise error detection scheme. Designers can implement a precise error detection scheme by adding address register pipelines for RAMs. It is the responsibility of the designer to correctly implement this logic.

7.7.1 GHB and BTAC data corruption

The scheme provides parity error support for GHB RAMs and BTAC RAMs but this support has limited diagnostic value. Corruption in GHB data or BTAC data does not generate functional errors in the Cortex-A9 processor. Corruption in GHB data or BTAC data results in a branch misprediction, that is detected and corrected.

Chapter 8

Level 2 Memory Interface

This chapter describes the L2 memory interface. It contains the following sections:

- *Cortex-A9 L2 interface* on page 8-2
- *Optimized accesses to the L2 memory interface* on page 8-7
- *STRT instructions* on page 8-9.

8.1 Cortex-A9 L2 interface

This section describes the Cortex-A9 Level 2 interface in:

- *About the Cortex-A9 L2 interface*
- *Supported AXI transfers* on page 8-3
- *AXI transaction IDs* on page 8-3
- *STRT instructions* on page 8-9.

8.1.1 About the Cortex-A9 L2 interface

The Cortex-A9 L2 interface consists of two 64-bit wide AXI bus masters:

- M0 is the data side bus
- M1 is the instruction side bus and has no write channels.

Table 8-1 shows the AXI master 0 interface attributes.

Table 8-1 AXI master 0 interface attributes

Attribute	Format
Write issuing capability	12, including: <ul style="list-style-type: none"> • eight noncacheable writes • four evictions
Read issuing capability	10, including: <ul style="list-style-type: none"> • six linefill reads. • four noncacheable read
Combined issuing capability	22
Write ID capability	2
Write interleave capability	1
Write ID width	2
Read ID capability	3
Read ID width	2

Table 8-2 shows the AXI master 1 interface attributes.

Table 8-2 AXI master 1 interface attributes

Attribute	Format
Write issuing capability	None
Read issuing capability	4 instruction reads
Combined issuing capability	4
Write ID capability	None
Write interleave capability	None

Table 8-2 AXI master 1 interface attributes (continued)

Attribute	Format
Write ID width	None
Read ID capability	4
Read ID width	2

Note

The numbers in Table 8-1 on page 8-2 and Table 8-2 on page 8-2 are the theoretical maximums for the Cortex-A9 MP processor. A typical system is unlikely to reach these numbers. ARM recommends that you perform profiling to tailor your system resources appropriately for optimum performance.

The AXI protocol and meaning of each AXI signal are not described in this document. For more information see *AMBA AXI Protocol v1.0 Specification*.

Supported AXI transfers

Cortex-A9 master ports generate only a subset of all possible AXI transactions.

For cacheable transactions:

- WRAP4 64-bit for read transfers (linefills)
- INCR4 64-bit for write transfers (evictions)

For noncacheable transactions:

- INCR N (N:1- 9) 64-bit read transfers
- INCR 1 for 64-bit write transfers
- INCR N (N:1-16) 32-bit read transfers
- INCR N (N:1-2) for 32-bit write transfers
- INCR 1 for 8-bit and 16-bit read/write transfers
- INCR 1 for 8-bit, 16-bit, 32-bit, 64-bit exclusive read/write transfers
- INCR 1 for 8-bit and 32-bit read/write (locked) for swap
-
-

The following points apply to AXI transactions:

- WRAP bursts are only read transfers, 64-bit, 4 transfers
- INCR 1 can be any size for read or write
- INCR burst (more than one transfer) are only 32-bit or 64-bit
- No transaction is marked as FIXED
- Write transfers with all byte strobes low can occur.

8.1.2 AXI transaction IDs

The AXI ID signal is encoded as follows:

- For the data side read bus, **ARIDM0**, is encoded as follows:
 - 2'b00 for noncacheable accesses
 - 2'b01 is unused
 - 2'b10 for linefill 0 accesses

- 2'b11 for linefill 1 accesses.
- For the instruction side read bus, **ARIDM1**, is encoded as follows:
 - 2'b00 for outstanding transactions
 - 2'b01 for outstanding transactions
 - 2'b10 for outstanding transactions
 - 2'b11 for outstanding transactions.
- For the data side write bus, **AWIDM0**, is encoded as follows:
 - 2'b00 for noncacheable accesses
 - 2'b01 is unused
 - 2'b10 for linefill 0 evictions
 - 2'b11 for linefill 1 evictions.

8.1.3 AXI USER bits

The AXI USER bits encodings are as follows:

Data side read bus, **ARUSERM0[6:0]**

Table 8-3 shows the bit encodings for **ARUSERM0[6:0]**

Table 8-3 ARUSERM0[6:0] encodings

Bits	Name	Description
[6]	Reserved	b0
[5]	L2 Prefetch hint	Indicates that the read access is a prefetch hint to the L2, and does not expect any data back.
[4:1]	Inner attributes	b0000 = Strongly Ordered b0001 = Device b0011 = Normal Memory Non-Cacheable b0110 = Write-Through b0111 = Write Back no Write Allocate b1111 = Write Back Write Allocate
[0]	Shared bit	b0 = Non-shared b1 = Shared

Instruction side read bus, ARUSERM1[6:0]

Table 8-4 shows the bit encodings for ARUSERM1[6:0].

Table 8-4 ARUSERM1[6:0] encodings

Bits	Name	Description
[6]	Reserved	b0
[5]	Reserved	b0
[4:1]	Inner attributes	b0000 = Strongly Ordered b0001 = Device b0011 = Normal Memory Non-Cacheable b0110 = Write-Through b0111 = Write Back no Write Allocate b1111 = Write Back Write Allocate.
[0]	Shared bit	b0 = Non-shared b1 = Shared

Data side write bus, AWUSERM0[8:0]

Table 8-5 shows the bit encodings for AWUSERM0[8:0].

Table 8-5 ARUSERM0[8:0] encodings

Bits	Name	Description
[8]	Early BRESP Enable bit	Indicates that the L2 slave can send an early BRESP answer to the write request. See <i>Early BRESP</i> on page 8-7.
[7]	Full line of write zeros bit	Indicates that the access is an entire cache line write full of zeros. See <i>Write full line of zeros</i> on page 8-8.
[6]	Clean eviction	Indicates that the write access is the eviction of a clean cache line.
[5]	L1 eviction	Indicates that the write access is a cache line eviction from the L1.
[4:1]	Inner attributes	b0000 = Strongly Ordered b0001 = Device b0011 = Normal Memory Non-Cacheable b0110 = Write-Through b0111 = Write Back no Write Allocate b1111 = Write Back Write Allocate.
[0]	Shared bit	b0 = Non-shared b1 = Shared

8.1.4 Exclusive L2 cache

The Cortex-A9 processor can be connected to an L2 cache that supports an exclusive cache mode. This mode must be activated both in the Cortex-A9 processor and in the L2 cache controller.

In this mode, the data cache of the Cortex-A9 processor and the L2 cache are exclusive. At any time, a given address is cached in either L1 data caches or in the L2 cache, but not in both. This has the effect of greatly increasing the usable space and efficiency of an L2 cache connected to the Cortex-A9 processor. When exclusive cache configuration is selected:

- Data cache line replacement policy is modified so that the victim line always gets evicted to L2 memory, even if it is clean.
- If a line is dirty in the L2 cache controller, a read request to this address from the processor causes writeback to external memory and a linefill to the processor.

8.2 Optimized accesses to the L2 memory interface

This section describes optimized accesses to the L2 memory interface. These optimized accesses can generate non-AXI compliant requests on the Cortex-A9 AXI master ports. These non-AXI compliant requests must be generated only when the slaves connected on the Cortex-A9 AXI master ports can support them. The L2 cache controller supports these kinds of requests. The following subsections describe the requests:

- *Prefetch hint to the L2 memory interface*
- *Early BRESP*
- *Write full line of zeros* on page 8-8.
- *Speculative coherent requests* on page 8-8.

8.2.1 Prefetch hint to the L2 memory interface

The Cortex-A9 processor can generate prefetch hint requests to the L2 memory controller. The prefetch hint requests are non-compliant AXI read requests generated by the Cortex-A9 processor which do not expect any data return.

You can generate prefetch hint requests to the L2 by:

- Enabling the L2 Prefetch Hint feature, bit [1] in the ACTLR. When enabled, this feature enables the Cortex-A9 processor to automatically issue L2 prefetch hint requests when it detects regular fetch patterns on a coherent memory. This feature is only triggered in a Cortex-A9 MPCore processor, and not in a uniprocessor.
- Programming PLE operations, when this feature is available in the Cortex-A9 processor. In this case, the PLE engine issues a series of L2 prefetch hint requests at the programmed addresses. See Chapter 9 *Preload Engine*.

L2 prefetch hint requests are identified by having their ARUSER[5] bit set.

———— **Note** —————

No additional programming of the L2C-310 is required.

8.2.2 Early BRESP

According to the AXI specification, **BRESP** answers on response channels must be returned to the master only once the last data has been sent by the master. Cortex-A9 processors can also deal with **BRESP** answers returned as soon as address has been accepted by the slave, regardless of whether data is sent or not. This enables the Cortex-A9 processor to provide a higher bandwidth for writes if the slave can support the Early BRESP feature. Cortex-A9 processors set the **AWUSER[8]** bit to indicate to the slave that it can accept an early **BRESP** answer for this access. This feature can optimize the performance of the processor, but the Early BRESP feature generates non-AXI compliant requests. When a slave receives a write request with **AWUSER[8]** set, it can either give the BRESP answer after the last data is received, AXI compliant, or in advance, non-AXI compliant. The L2C-310 cache controller supports this non-AXI compliant feature.

The Cortex-A9 does not require any programming to enable this feature, which is always on by default.

Note

You must program the L2 cache controller to benefit from this optimization. See the *AMBA Level 2 Cache Controller (L2C-310) Technical Reference Manual*.

8.2.3 Write full line of zeros

When this feature is enabled, the Cortex-A9 processor can write entire non-coherent cache lines full of zero to the L2C-310 cache controller with a single request. This provides a performance improvement and some power savings. This feature can optimize the performance of the processor, but it requires a slave that is optimized for this special access. The requests are marked as full line of write zeros by having the associated **AWUSERM0[7]** bit set.

Setting bit[3] of the ACTLR enables this feature. See *Auxiliary Control Register* on page 4-18.

You must program the L2C-310 Cache Controller first, prior to enabling the feature in the Cortex-A9 processor, to support this feature. See the *AMBA Level 2 Cache Controller (L2C-310) Technical Reference Manual*.

8.2.4 Speculative coherent requests

This optimization is available for Cortex-A9 MPCore processors only. See the *Cortex-A9 MPCore TRM*.

8.3 STRT instructions

Take particular care with noncacheable write accesses when using the STRT instruction. To put the correct information on the external bus ensure one of the following:

- The access is to Strongly-ordered memory.
This ensures that the STRT instruction does not merge in the store buffer.
- The access is to Device memory.
This ensures that the STRT instruction does not merge in the store buffer.
- A DSB instruction is issued before the STRT and after the STRT.
This prevents an STRT from merging into an existing slot at the same 64-bit address, or merging with another write at the same 64-bit address.

Table 8-6 shows Cortex-A9 modes and corresponding **AxPROT** values.

Table 8-6 Cortex-A9 mode and AxPROT values

Processor mode	Type of access	Value of AxPROT
User	Cacheable read access	User
Privileged		Privileged
User	Noncacheable read access	User
Privileged		Privileged
-	Cacheable write access	Always marked as Privileged
User	Noncacheable write access	User
Privileged	Noncacheable write access	Privileged, except when using STRT

Chapter 9

Preload Engine

The design can include a *Preload Engine* (PLE). The PLE loads selected regions of memory into L2. This chapter describes the PLE. It contains:

- *About the Preload Engine* on page 9-2
- *PLE control register descriptions* on page 9-3
- *PLE operations* on page 9-4.

9.1 About the Preload Engine

If implemented, the PLE loads selected regions of memory into L2. Use the MCRR preload channel operation to program the PLE. Dedicated events monitor the behavior of the memory region. Additional L2C-310 events can also monitor PLE behavior.

The preload operation parameters enter the PLE FIFO which includes:

- programmed parameters:
 - base address
 - length of stride
 - number of blocks.
- a valid bit
- an NS state bit
- a *Translation Table Base* (TTB) address
- an *Address Space Identifier* (ASID) value.

Preload blocks can span multiple page entries. Programmed entries can still be valid in case of context switches.

The Preload Engine handles cache line preload requests in the same way as a standard PLD request except that it uses its own TTB and ASID parameters. If there is a translation abort, the preload request is ignored and the Preload Engine issues the next request.

Not all the MMU settings are saved. The Domain, Tex-Remap, Primary Remap, Normal Remap, and Access Permission registers are not saved. As a consequence, a write operation in any of these registers causes a flush of the entire FIFO and of the active channel.

Additionally, for *Translation Lookaside Buffer* (TLB) maintenance operations, the maintenance operation must be applied to the FIFO entries too. This is done as follows:

On Invalidate by MVA and ASID

Invalidate all entries with a matching ASID

On Invalidate by ASID

Invalidate all entries with a matching ASID

On Invalidate by MVA all ASID

Flush the entire FIFO

On Invalidate entire TLB

Flush the entire FIFO

These rules are also applicable to the PLE active channel.

The Preload Engine defines the following MCRR instruction to use with the preload blocks.

```
MCRR p15, 0, <Rt>, <Rt2> c11; Program new PLE channel
```

The number of entries in the FIFO can be set as an RTL configuration design choice. Available sizes are:

- 16 entries
- 8 entries
- 4 entries.

9.2 PLE control register descriptions

The PLE control registers are CP15 registers, accessed when CRn is c11. See *CP15 c11 register summary* on page 4-30. The following sections describe the PLE control registers:

- *PLE ID Register* on page 4-30
- *PLE Activity Status Register* on page 4-31
- *PLE FIFO Status Register* on page 4-32
- *Preload Engine User Accessibility Register* on page 4-32
- *Preload Engine Parameters Control Register* on page 4-33.

For all CP15 c11 system control registers NSAC.PLE controls Non-secure accesses. *PLE operations* on page 9-4 shows the operations to use with these control registers.

9.3 PLE operations

The following sections describe the PLE operations:

- *Preload Engine FIFO flush operation*
- *Preload Engine pause channel operation*
- *Preload Engine resume channel operation*
- *Preload Engine kill channel operation* on page 9-5
- *PLE Program New Channel operation* on page 9-5.

For all Preload Engine operations:

- NSACR.PLE controls Non-secure execution.
- PLEUAR.EN controls User execution.
- the operations are only available in configurations where the Preload Engine is present, otherwise an Undefined Instruction exception is taken.

9.3.1 Preload Engine FIFO flush operation

The PLEFF operation characteristics are:

Purpose Flushes all PLE channels programmed previously including the PLE channel currently being executed.

To perform the PLE FIFO Flush operation, use:

MCR p15, 0, <Rt>, c11, c2, 1

<Rt> is not taken into account in this operation.

9.3.2 Preload Engine pause channel operation

The PLEPC operation characteristics are:

Purpose Pauses PLE activity.

You can perform a PLEPC operation even if no PLE channel is currently active. In this case, even if a new PLE channel is programmed afterwards, its execution does not start until after a PLE Resume Channel operation.

To perform the PLE PC operation, use:

MCR p15, 0, <Rt>, c11, c3, 0

<Rt> is not taken into account in this operation.

9.3.3 Preload Engine resume channel operation

The PLERC operation characteristics are:

Purpose Causes Preload Engine activity to resume.

If you perform a PLERC operation when the PLE is not paused, the Resume Channel operation is ignored.

To perform a PLERC operation, use:

MCR p15, 0, <Rt>, c11, c3, 1

9.3.4 Preload Engine kill channel operation

The PLEKC operation characteristics are:

Purpose Kills the PLE channel currently active.

This operation does not operate on any PLE request in the PLE FIFO.

To perform a PLEKC operation, use

MCR p15, 0, <Rt>, c11, c3, 2

9.3.5 PLE Program New Channel operation

The PLE Program new channel operation characteristics are:

Purpose Programs a new memory region to preload into L2 memory. Kills the PLE channel currently active.

Figure 9-1 shows the <Rt>. and <Rt2> bit assignments for PLE program new channel operations. Rt is the register that contains the Base address. Rt2 is the register that contains the length, stride, and number of blocks.




Figure 9-1 Program new channel operation bit assignments

Table 9-1 shows the PLE program new channel operation bit assignments.

Table 9-1 PLE program new channel operation bit assignments

Bits	Name	Description
[63:34]	Base address (VA)	This is the 32-bit Base Virtual Address of the first block of memory to preload. The address is aligned on a word boundary. That is, bits [33:32] are RAZ/WI.
[33:32]	-	RAZ/WI
[31:18]	Length	Specifies the length of the block to preload. Length is encoded as word multiples. The range is from 14'b0000000000, a single word block, to 14'b11111111111111, a 16K word block.
[17:10]	Stride	Indicates the preload stride between blocks. The preload stride is the difference between the start address of two blocks. The stride is encoded as a word multiple. The range is from 8'b00000000, contiguous blocks, to 8'b11111111, prefetch blocks every 256 words.
[9:2]	Number of blocks	Specifies the number of blocks to preload. Values range from 8'b00000000, indicating a single block preload, to 8'b11111111 indicating 256 blocks.
[1:0]	-	RAZ/WI

To program a new channel operation, use the MCRR operation:

MCRR p15, 0, <Rt>, <Rt2> c11; Program new PLE channel

Note

A newly programmed PLE entry is written to the PLE FIFO if the FIFO has available entries. In cases of FIFO overflow, the instruction silently fails, and the FIFO Overflow event signal is asserted. See Preload events in Table 11-7 on page 11-7. See *PLE FIFO Status Register* on page 4-32.

Chapter 10

Debug

This chapter describes the processor debug unit. This feature assists the development of application software, operating systems, and hardware. This chapter contains the following sections:

- *About the debug interface* on page 10-2
- *About the Cortex-A9 debug interface* on page 10-4
- *Debug register descriptions* on page 10-7
- *Debug management registers* on page 10-13
- *External debug interface* on page 10-16.

10.1 About the debug interface

The Cortex-A9 processor implements the ARMv7 debug architecture as described in the *ARM Architecture Reference Manual*. It implements the set of debug events described in the *ARM Architecture Reference Manual*.

In addition, there are:

- Cortex-A9 processor specific events. These are described in *Performance monitoring events* on page 11-7.
- system coherency events.

See *Performance monitoring* on page 2-3. See also Chapter 11 *Performance Monitoring Unit*

10.1.1 Debugging modes

Authentication signals control the debugging modes. The authentication signals configure the processor so its activity can only be debugged or traced in a certain subset of processor modes and security states. See *Authentication signals* on page 10-16.

———— Note —————

The Cortex-A9 processor only supports halt mode debugging in secure user mode when invasive debugging is enabled by the SPIDEN pin. If SPIDEN is LOW only monitor mode debugging in secure user mode is available by setting the SDR.SUIDEN bit.. That is, when **SPIDEN** is LOW, the core is not allowed to enter Halting Debug Mode even if the SDR.SUIDEN bit is set to 1. You can bypass this restriction by setting the external **SPIDEN** pin HIGH.

10.1.2 Breakpoints and watchpoints

There are:

- six breakpoints, two with Context ID comparison capability, BRP4 and BRP5. See *Breakpoint Value Registers* on page 10-7 and *Breakpoint Control Registers* on page 10-8.
- four watchpoints.

A watchpoint event is always synchronous. It has the same behavior as a synchronous data abort. The method of debug entry (DBGDSCR[5:2]) never has the value b0010.

If a synchronous abort occurs on a watchpointed access, the synchronous abort takes priority over the watchpoint.

If the abort is asynchronous and cannot be associated with the access, the exception that is taken is unpredictable.

Cache maintenance operations do not generate watchpoint events

See *Watchpoint Value Registers* on page 10-10 and *Watchpoint Control Registers* on page 10-11.

10.1.3 Asynchronous aborts

The Cortex-A9 processor ensures that all possible outstanding asynchronous data aborts have been recognized prior to entry to debug state.

10.1.4 Processor interfaces

The Cortex-A9 processor has the following interfaces to the debug, performance monitor:

Debug registers

This interface is Baseline CP14, Extended CP14, and memory-mapped. See *CTI signals* on page A-22 and *APB interface signals* on page A-22

Performance monitor

This interface is CP15 based and memory-mapped. See *Performance monitoring* on page 2-3. See also Chapter 11 *Performance Monitoring Unit*.

10.1.5 Effects of resets on debug registers

nDBGRESET

nDBGRESET is the debug logic reset signals. This signal must be asserted during a power-on reset sequence.

Other reset signals, **nCPURESET** and **nNEONRESET**, if MPE is present, have no effect on the debug logic.

On a debug reset:

- The debug state is unchanged. That is, **DBGSCR.HALTED** is unchanged.
- The processor removes the pending halting debug events **DBGDRCR.HaltReq**.

10.2 About the Cortex-A9 debug interface

The debug interface consists of:

- a Baseline CP14 interface
- an Extended CP14 interface
- an external debug interface connected to the external debugger through a *Debug Access Port (DAP)*.

Figure 10-1 shows the Cortex-A9 debug registers interface.

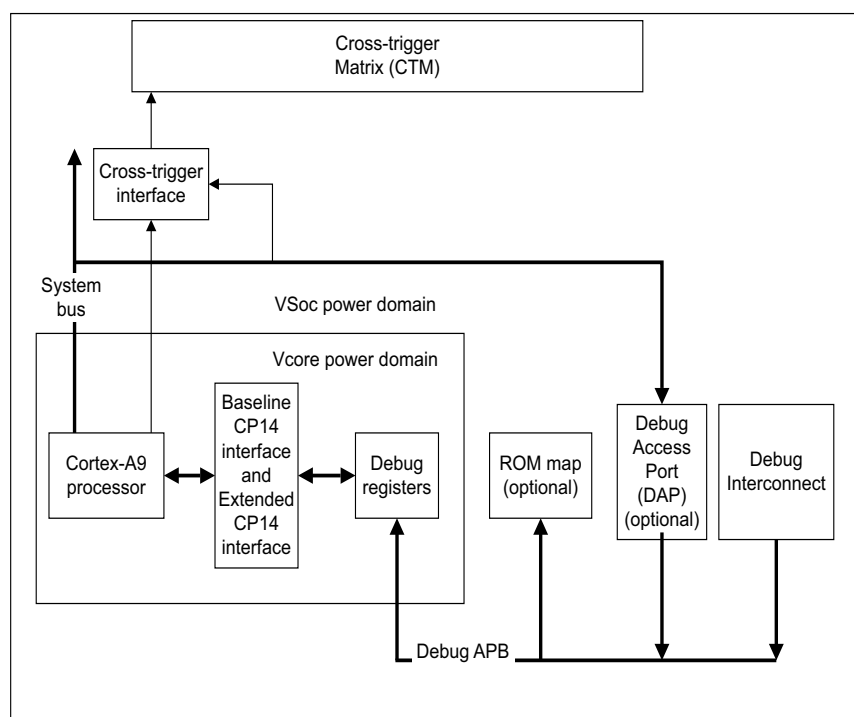


Figure 10-1 Debug registers interface and CoreSight infrastructure

10.2.1 Debug register access

You can access the debug registers:

- through the cp14 interface. The debug registers are mapped to coprocessor instructions.
- through the APB using the relevant offset, with the following exceptions:
 - DBGRAR
 - DBGSAR
 - DBGSCR-int
 - DBGTR-int.

External views of DBSCR and DBGTR are accessible through memory-mapped APB access.

Table 10-1 on page 10-5 shows the CP14 interface registers. All other registers are described in the *ARM Architecture Reference Manual*.

Table 10-1 CP14 interface registers

Register number	Offset	CP14 instruction	Access	Register name	Description
0	0x000	0, c0, c0, 0	RO	DBGDIDR ^a	- b
-	-	0, c1, c0, 0	RO	DBGDRAR ^a	-
-	-	0, c2, c0, 0	RO	DBGDSAR ^a	-
-	-	0, c0, c1, 0	RO	DBGDSCRint ^{ab}	-
5	-	0, c0, c5, 0	R	DBGDTRRXint ^a	-
	-		W	DBGDTRTXint ^a	-
6	0x018	0, c0, c6, 0	RW	DBGWFAR	Use of DBGWFAR is deprecated in the ARMv7 architecture, because watchpoints are synchronous
7	0x01C	0, c0, c7, 0	RW	DBGVCR	-
8	-	-	-	Reserved	-
9	0x024	0, c0, c9, 0	RAZ/WI	DBGECR	Not implemented
10	0x028	0, c0, c10, 0	RAZ/WI	DBGDSCCR	<i>Debug State Cache Control Register (DBGDSCCR) on page 10-7</i>
11	0x02C	0, c0, c11, 0	RAZ/WI	DBGDSMCR	Not implemented
12-31	-	-	-	Reserved	-
32	0x080	0, c0, c0, 2	RW	DBGDTRRXext	-
33	0x084	0, c0, c1, 2	WO	DBGITR	-
33	0x084	0, c0, c1, 2	RO	DBGPCSR	-
34	0x088	0, c0, c2, 2	RW	DBGDSCRext	-
35	0x08C	0, c0, c3, 2	RW	DBGDTRTXext	-
36	0x090	0, c0, c4, 2	WO	DBGDRCR	-
37-63	-	-	-	Reserved	-
64-68	0x100-0x114	0, c0, c0-c5, 4	RW	DBGBVRn	<i>Breakpoint Value Registers on page 10-7</i>
69-79	-	-	-	Reserved	-
80-85	0x140-0x154	0, c0, c0-c5, 5	RW	DBGBCRn	<i>Breakpoint Control Registers on page 10-8</i>
86-95	-	-	-	Reserved	-
96-99	0x180-0x18BC	0, c0, c0-c3, 6	RW	DBGWVRn	<i>Watchpoint Value Registers on page 10-10</i>
100-111	-	-	-	Reserved	-
112-115	0x1C0-0x1DC	0, c0, c0-c3, 7	RW	DBGWCRn	<i>Watchpoint Control Registers on page 10-11</i>
116-191	-	-	-	Reserved	-

Table 10-1 CP14 interface registers (continued)

Register number	Offset	CP14 instruction	Access	Register name	Description
192	0x300	0, c1, c0, 4	RAZ/WI	DBGOSLAR	Not implemented
193	0x304	0, c1, c1, 4	RAZ/WI	DBGOSLSR	Not implemented
194	0x308	0, c1, c2, 4	RAZ/WI	DBGOSSRR	Not implemented
195	-	-	-	Reserved	-
196	0x310	0, c1, c4, 4	RO	DBGPRCR	-
197	0x314	0, c1, c5, 4	RO	DBGPRSR	-
198-511	-	-	-	Reserved	-
512-575	0x800-0x8FC	-	-	-	PMU registers ^c
576-831	-	-	-	Reserved	-
832-895	0xD00-0xDFC	0, c6, c0, 15, 4-7	RW	Unpredictable	-
896-927	-	-	-	Reserved	-
928-959	0xE80-0xEFC0	0, c7, c0, 15, 2-3	RAZ/WI	-	-
960	0xF00	0, c7, c0, 4	RAZ/WI	DBGITCTRL	Integration Mode Control Register
961-999	0xF04-0xF9C	-	-	-	-
1000	0xFA0	0, c7, c8, 6	RW	DBGCLAIMSET	Claim Tag Set Register
1001	0xFA4	0, c7, c9, 6	RW	DBGCLAIMCLR	Claim Tag Clear Register
1002-1003	-	-	-	Reserved	-
1004	0xFB0	0, c7, c12, 6	WO	DBGLAR	Lock Access Register
1005	0xFB4	0, c7, c13, 6	RO	DBGLSR	Lock Status Register
1006	0xFB8	0, c7, c14, 6	RO	DBGAUTHSTATUS	Authentication Status Register
1007-1009	-	-	-	Reserved	-
1010	0xFC8	0, c7, c2, 7	RAZ	DBGDEVID	-
1011	0xFCC	0, c7, c3, 7	RO	DBGDEVTYPE	Device Type Register
1012-1016	0xFD0-0xFEC	0, c7, c4-8, 7	RO	PERIPHERALID	<i>CoreSight Identification Registers on page 10-14</i>
1017-1019	-	-	-	Reserved	-
1020-1023	0xFF0-0xFFC	0, c7, c12-15, 7	RO	COMPONENTID	<i>CoreSight Identification Registers on page 10-14</i>

a. Baseline CP14 interface. This register also has an external view through the memory-mapped interface and the CP14 interface.

b. Accessible in User mode if bit[12] of the DBGSCR is clear. Also accessible in privileged modes.

c. PMU registers are part of the CP15 interface. Reads from the extended CP14 interface return zero. See *CP15 c9 register summary* on page 4-28. See also Chapter 11 *Performance Monitoring Unit*.

10.3 Debug register descriptions

This section describes the debug registers.

10.3.1 Debug State Cache Control Register (DBGDSCCR)

The DSCCR controls cache behavior while the processor is in debug state. The Cortex-A9 processor does not implement any of the features of the Debug State Cache Control Register. The Debug State Cache Control Register reads as zero.

10.3.2 Breakpoint Value Registers

The *Breakpoint Value Registers* (BVRs) are registers 64-68, at offsets 0x100-0x114. Each BVR is associated with a *Breakpoint Control Register* (BCR), for example:

- BVR0 with BCR0
- BVR1 with BCR1.

This pattern continues up to BVR5 with BCR5.

A pair of breakpoint registers, BVR_n and BCR_n, is called a *Breakpoint Register Pair* (BRP_n).

Table 10-2 shows the BVRs and corresponding BCRs.

Table 10-2 BVRs and corresponding BCRs

Breakpoint Value Registers			Breakpoint Control Registers		
Register	Register number	Offset	Register	Register number	Offset
BVR0	64	0x100	BCR0	80	0x140
BVR1	65	0x104	BCR1	81	0x144
BVR2	66	0x108	BCR2	82	0x148
BVR3	66	0x10C	BCR3	83	0x14C
BVR4	67	0x110	BCR4	84	0x150
BVR5	68	0x114	BCR5	85	0x154

The breakpoint value contained in this register corresponds to either an *Instruction Virtual Address* (IVA) or a context ID. Breakpoints can be set on:

- an IVA
- a context ID value
- an IVA and context ID pair.

For an IVA and context ID pair, two BRPs must be linked. A debug event is generated when both the IVA and the context ID pair match at the same time.

Table 10-3 shows how the bit values correspond with the Breakpoint Value Registers functions.

Table 10-3 Breakpoint Value Registers bit functions

Bits	Name	Description
[31:0]	-	Breakpoint value. The reset value is 0.

Note

- Only BRP4 and BRP5 support context ID comparison.
- BVR0[1:0], BVR1[1:0], BVR2[1:0], and BVR3[1:0] are Should Be Zero or Preserved on writes and Read As Zero on reads because these registers do not support context ID comparisons.
- The context ID value for a BVR to match with is given by the contents of the CP15 Context ID Register.

10.3.3 Breakpoint Control Registers

The BCR is a read and write register that contains the necessary control bits for setting:

- breakpoints
- linked breakpoints.

Figure 10-2 shows the bit arrangement of the BCRs.




Figure 10-2 Breakpoint Control Registers bit assignments

Table 10-4 shows how the bit values correspond with the Breakpoint Control Registers functions.

Table 10-4 Breakpoint Control Registers bit assignments

Bits	Name	Description
[31:29]	-	RAZ on reads, SBZP on writes.
[28:24]	Breakpoint address mask	Breakpoint address mask. RAZ/WI b00000 = No mask
[23]	-	RAZ on reads, SBZP on writes.
[22:20]	M	Meaning of BVR: b000 = Instruction virtual address match b001 = Linked instruction virtual address match b010 = Unlinked context ID b011 = Linked context ID b100 = Instruction virtual address mismatch b101 = Linked instruction virtual address mismatch b11x = Reserved.

Note

BCR0[21], BCR1[21], BCR2[21], and BCR3[21] are RAZ on reads because these registers do not have context ID comparison capability.

Table 10-4 Breakpoint Control Registers bit assignments (continued)

Bits	Name	Description
[19:16]	Linked BRP	<p>Linked BRP number. The binary number encoded here indicates another BRP to link this one with.</p> <hr/> <p style="text-align: center;">Note</p> <ul style="list-style-type: none"> • 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. <hr/>
[15:14]	Secure state access control	<p>Secure state access control. This field enables the breakpoint to be conditional on the security state of the processor.</p> <p>b00 = Breakpoint matches in both Secure and Non-secure state b01 = Breakpoint only matches in Non-secure state b10 = Breakpoint only matches in Secure state b11 = Reserved.</p>
[13:9]	-	RAZ on reads, SBZP on writes.
[8:5]	Byte address select	<p>Byte address select. For breakpoints programmed to match an IVA, you must write a word-aligned address to the BVR. You can then use this field to program the breakpoint so it hits only if you access certain byte addresses.</p> <p>If you program the BRP for IVA match:</p> <p>b0000 = The breakpoint never hits b0011 = The breakpoint hits if any of the two bytes starting at address BVR & 0xFFFFFFFFC +0 is accessed b1100 = The breakpoint hits if any of the two bytes starting at address BVR & 0xFFFFFFFFC +2 is accessed b1111 = The breakpoint hits if any of the four bytes starting at address BVR & 0xFFFFFFFFC +0 is accessed.</p> <p>If you program the BRP for IVA mismatch, the breakpoint hits where the corresponding IVA breakpoint does not hit, that is, the range of addresses covered by an IVA mismatch breakpoint is the negative image of the corresponding IVA breakpoint.</p> <p>If you program the BRP for context ID comparison, this field must be set to b1111. Otherwise, breakpoint and watchpoint debug events might not be generated as expected.</p> <hr/> <p style="text-align: center;">Note</p> <p>Writing a value to BCR[8:5] where BCR[8] is not equal to BCR[7], or BCR[6] is not equal to BCR[5], has Unpredictable results.</p> <hr/>
[4:3]	-	RAZ on reads, SBZP on writes.
[2:1]	SP	<p>Supervisor access control. The breakpoint can be conditioned on the mode of the processor.</p> <p>b00 = User, System, or Supervisor b01 = Privileged b10 = User b11 = Any.</p>
[0]	B	<p>Breakpoint enable:</p> <p>0 = Breakpoint disabled, reset value 1 = Breakpoint enabled.</p>

Table 10-5 shows the meaning of the BVR bits.

Table 10-5 Meaning of BVR bits [22:20]

BVR[22:20]	Meaning
b000	The corresponding BVR[31:2] is compared against the IVA bus and the state of the processor against this BCR. It generates a breakpoint debug event on a joint IVA and state match.
b001	The corresponding BVR[31:2] is compared against the IVA bus and the state of the processor against this BCR. This BRP is linked with the one indicated by BCR[19:16] linked BRP field. They generate a breakpoint debug event on a joint IVA, context ID, and state match.
b010	The corresponding BVR[31:0] is compared against CP15 Context ID Register, c13 and the state of the processor against this BCR. 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, BCR[8:5] must be set to b1111. Otherwise, it is Unpredictable whether a breakpoint debug event is generated.
b011	The corresponding BVR[31:0] is compared against CP15 Context ID Register, c13. This BRP links another BRP (of the BCR[21:20]=b01 type), or WRP (with WCR[20]=b1). They generate a breakpoint or watchpoint debug event on a joint IVA or DVA and context ID match. For this BRP, BCR[8:5] must be set to b1111, BCR[15:14] must be set to b00, and BCR[2:1] must be set to b11. Otherwise, it is Unpredictable whether a breakpoint debug event is generated.
b100	The corresponding BVR[31:2] and BCR[8:5] are compared against the IVA bus and the state of the processor against this BCR. It generates a breakpoint debug event on a joint IVA mismatch and state match.
b101	The corresponding BVR[31:2] and BCR[8:5] are compared against the IVA bus and the state of the processor against this BCR. This BRP is linked with the one indicated by BCR[19:16] linked BRP field. It generates a breakpoint debug event on a joint IVA mismatch, state and context ID match.
b11x	Reserved. The behavior is Unpredictable.

10.3.4 Watchpoint Value Registers

The *Watchpoint Value Registers* (WVRs) are registers 96-99, at offsets 0x180-0x18C. Each WVR is associated with a *Watchpoint Control Register* (WCR), for example:

- WVR0 with WCR0
- WVR1 with WCR1.

This pattern continues up to WVR3 with WCR3.

Table 10-6 shows the WVRs and corresponding WCRs.

Table 10-6 WVRs and corresponding WCRs

Watchpoint Value Registers			Watchpoint Control Registers		
Register	Register number	Offset	Register	Register number	Offset
WVR0	96	0x180	WCR0	112	0x1C0
WVR1	97	0x184	WCR1	113	0x1C4
WVR2	98	0x188	WCR2	114	0x1C8
WVR3	99	0x18C	WCR3	115	0x1DC

A pair of watchpoint registers, WVR_n and WCR_n, is called a *Watchpoint Register Pair* (WRP_n).

The watchpoint value contained in the WVR always corresponds to a *Data Virtual Address* (DVA) and can be set either on:

- a DVA
- a DVA and context ID pair.

For a DVA and context ID pair, a WRP and a BRPs with context ID comparison capability must be linked. A debug event is generated when both the DVA and the context ID pair match simultaneously. Table 10-7 shows how the bit values correspond with the Watchpoint Value Registers functions.

Table 10-7 Watchpoint Value Registers bit functions

Bits	Name	Description
[31:2]	-	Watchpoint address
[1:0]	-	RAZ on reads, SBZP on writes

10.3.5 Watchpoint Control Registers

The WCRs contain the necessary control bits for setting:

- watchpoints
- linked watchpoints.

Figure 10-3 shows the bit arrangement of the Watchpoint Control Registers.




Figure 10-3 Watchpoint Control Registers bit assignments

Table 10-8 shows how the bit values correspond with the Watchpoint Control Registers functions.

Table 10-8 Watchpoint Control Registers bit assignments

Bits	Name	Description
[31:29]	-	RAZ on reads, SBZP on writes.
[28:24]	Watchpoint address mask	Watchpoint address mask.
[23:21]	-	RAZ on reads, SBZP on writes.
[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.

Table 10-8 Watchpoint Control Registers bit assignments (continued)

Bits	Name	Description
[15:14]	Secure state access control	Secure state access control. This field enables the watchpoint to be conditioned on the security state of the processor. b00 = Watchpoint matches in both Secure and Non-secure state b01 = Watchpoint only matches in Non-secure state b10 = Watchpoint only matches in Secure state b11 = Reserved.
[13]	-	RAZ on reads, SBZP on writes.
[12:9]	-	RAZ/WI
[8:5]	Byte address select	Byte address select. The WVR 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.
[4:3]	L/S	Load/store access. The watchpoint can be conditioned to the type of access being done. b00 = Reserved b01 = Load, load exclusive, or swap b10 = Store, store exclusive or swap b11 = Either. SWP and SWPB trigger a watchpoint on b01, b10, or b11. A load exclusive instruction triggers a watchpoint on b01 or b11. A store exclusive instruction triggers a watchpoint on b10 or b11 only if it passes the local monitor within the processor. ^a
[2:1]	SP	Privileged access control. The watchpoint can be conditioned to the privilege of the access being done: b00 = Reserved b01 = Privileged, match if the processor does a privileged access to memory b10 = User, match only on nonprivileged 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, reset value 1 = Watchpoint enabled.

- a. A store exclusive can generate an MMU fault or cause the processor to take a data watchpoint exception regardless of the state of the local monitor.

10.4 Debug management registers

The Management registers define the standardized set of registers that is implemented by all CoreSight components. These registers are described in this section. The cp14 interface must be used to access these registers.

Table 10-9 shows the contents of the Management registers for the Cortex-A9 debug unit.

Table 10-9 Debug management registers

Offset	Register number	Access	Mnemonic	Description
0xD00-0xDFC	832-895	RO	-	<i>Processor ID Registers.</i>
0xE00-0xEF0	854-956	-	-	RAZ.
0xF00	960	RW	ITCTRL	-
0xF04-0xF9C	961-999	RAZ	-	Reserved for Management Register expansion.
0xFA0	1000	RW	CLAIMSET	-
0xFA4	1001	RW	CLAIMCLR	-
0xFA8-0xFBC	1002-1003	-	-	RAZ.
0xFB0	1004	WO	LOCKACCESS	-
0xFB4		RO	LOCKSTATUS	-
0xFB8		RO	AUTHSTATUS	-
0xFBC-0xFC4	1007-1009	-	-	RAZ.
0xFC8	1010	RO	DEVID	Device Identifier.
0xFCC	1011	RO	DEVTYPE	-
0xFD0-0xFFC	1012-1023	R	-	<i>CoreSight Identification Registers on page 10-14.</i>

10.4.1 Processor ID Registers

The Processor ID Registers are read-only registers that return the same values as the corresponding CP15 ID Code Register and Feature ID Register.

Table 10-10 shows the offset value, register number, mnemonic, and description that are associated with each Process ID Register.

Table 10-10 Processor Identifier Registers

Offset (hex)	Register number	Mnemonic		Register value	Description
0xD00	832	CPUID	RO	0x80000000	ID Code Register ^a
0xD04	833	CTYPR	RO	0x80038003	Cache Type Register
0xD08	834	-	RAZ	-	-
0xD0C	835	TTYPR	RO	0x00000400	TLB Type Register
0xD10-0xD1C	836-839	-	-	-	Reserved

Table 10-10 Processor Identifier Registers (continued)

Offset (hex)	Register number	Mnemonic		Register value	Description
0xD20	840	ID_PFR0	RO	0x00001231	Processor Feature Register 0
0xD24	841	ID_PFR1	RO	0x00000011	Processor Feature Register 1
0xD28	842	ID_DFR0	RO	0x00010444	Debug Feature Register 0
0xD2C	843	ID_AFR0	RAZ	-	Auxiliary Feature Register 0
0xD30	844	ID_MMFR0	RO	0x00100103	Memory Model Feature Register 0
0xD34	845	ID_MMFR1	RO	0x20000000	Memory Model Feature Register 1
0xD38	846	ID_MMFR2	RO	0x01230000	Memory Model Feature Register 2
0xD3C	847	ID_MMFR3	RO	0x00002111	Memory Model Feature Register 3
0xD40	848	ID_ISAR0	RO	0x00101111	Instruction Set Attribute Register 0
0xD44	849	ID_ISAR1	RO	0x13112111	Instruction Set Attribute Register 1
0xD48	850	ID_ISAR2	RO	0x21232041	Instruction Set Attribute Register 2
0xD4C	851	ID_ISAR3	RO	0x11112131	Instruction Set Attribute Register 3
0xD50	852	ID_ISAR4	RO	0x00011142	Instruction Set Attribute Register 4
0xD54	853	ID_ISAR5	RAZ	-	Instruction Set Attribute Register 5

- a. For uniprocessor versions = 0x80000000
 For multiprocessor versions = 0xC000n0m
 n = **CLUSTERID** input
 m = CPU number (0x0 for CPU0, 0x1 for CPU1, 0x2 for CPU2, and 0x3 for CPU3).

10.4.2 CoreSight Identification Registers

The 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 required by all CoreSight components. Only bits [7:0] of each register are used.

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 10-11 shows the offset value, register number, and description that are associated with each Peripheral Identification Register.

Table 10-11 Peripheral Identification Registers

Offset (hex)	Register number	Value	Description
0xFD0	1012	0x04	Peripheral Identification Register 4
0xFD4	1013	-	Reserved
0xFD8	1014	-	Reserved
0xFDC	1015	-	Reserved

Table 10-11 Peripheral Identification Registers (continued)

Offset (hex)	Register number	Value	Description
0xFE0	1016	0x09	Peripheral Identification Register 0
0xFE4	1017	0xBC	Peripheral Identification Register 1
0xFE8	1018	0x0B	Peripheral Identification Register 2
0xFEC	1019	0x00	Peripheral Identification Register 3

Table 10-12 shows the offset value, register number, and value that are associated with each Component Identification Register.

Table 10-12 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

10.5 External debug interface

The system can access memory-mapped debug registers through the Cortex-A9 APB slave port.

This APB slave interface supports 32-bits wide data, stalls, slave-generated aborts, and eleven address bits [12:2] mapping 2x4KB of memory. Bit[12] of **PADDRDBG[12:0]** selects which of the components is accessed:

- Use **PADDRDBG[12] = 0** to access the debug area of the Cortex-A9 processor. See Table 10-1 on page 10-5 for debug resources memory mapping.
- Use **PADDRDBG[12] = 1** to access the Performance Monitoring Unit (PMU) area of the Cortex-A9 processor. See Chapter 11 *Performance Monitoring Unit* for PMU resources memory mapping.

The **PADDRDBG31** signal indicates to the processor the source of the access.

See Appendix A *Signal Descriptions* for a complete list of the external debug signals.

Figure 10-4 shows the external debug interface signals.

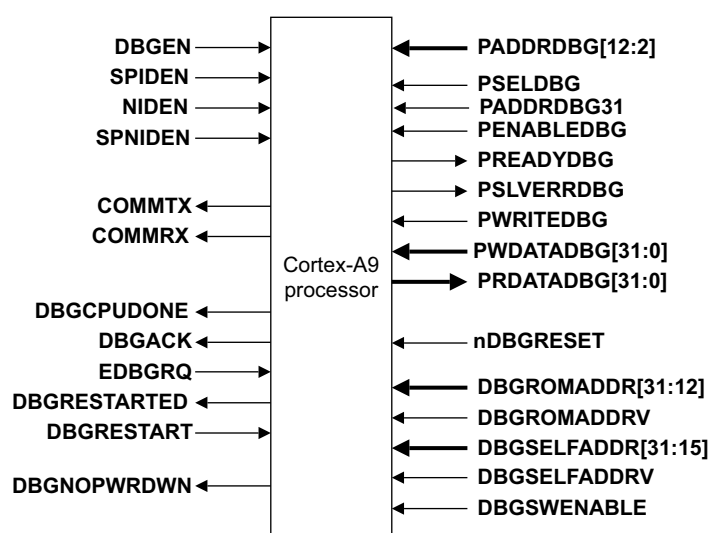


Figure 10-4 External debug interface signals

10.5.1 Authentication signals

Table 10-13 shows a list of the valid combinations of authentication signals along with their associated debug permissions.

Table 10-13 Authentication signal restrictions

SPIDEN	DBGEN ^a	SPNIDEN	NIDEN	Secure ^b invasive debug permitted	Non-secure invasive debug permitted	Secure non-invasive debug permitted	Non-secure non-invasive debug permitted
0	0	0	0	No	No	No	No
0	0	0	1	No	No	No	Yes
0	0	1	0	No	No	No	No
0	0	1	1	No	No	Yes	Yes

Table 10-13 Authentication signal restrictions (continued)

SPIDEN	DBGEN ^a	SPNIDEN	NIDEN	Secure ^b invasive debug permitted	Non-secure invasive debug permitted	Secure non-invasive debug permitted	Non-secure non-invasive debug permitted
0	1	0	0	No	Yes	No	Yes
0	1	0	1	No	Yes	No	Yes
0	1	1	0	No	Yes	Yes	Yes
0	1	1	1	No	Yes	Yes	Yes
1	0	0	0	No	No	No	No
1	0	0	1	No	No	Yes	Yes
1	0	1	0	No	No	No	No
1	0	1	1	No	No	Yes	Yes
1	1	0	0	Yes	Yes	Yes	Yes
1	1	0	1	Yes	Yes	Yes	Yes
1	1	1	0	Yes	Yes	Yes	Yes
1	1	1	1	Yes	Yes	Yes	Yes

- a. When **DBGEN** 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.
- b. Invasive debug is defined as those operations that affect the behavior of the core. For example, taking a breakpoint is defined as invasive debug but performance counters and trace are noninvasive.

10.5.2 Changing the authentication signals

The **NIDEN**, **DBGEN**, **SPIDEN**, and **SPNIDEN** input signals are either tied off to some fixed value or controlled by some external device.

If software running on the Cortex-A9 processor 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 **DSB**.
3. Poll the **DSCR** or Authentication Status Register to check whether the processor has already detected the changed value of these signals. This is required because the system might not issue the signal change to the processor until several cycles after the **DSB** completes.
4. Perform an **ISB**, an Exception entry, or Exception exit.

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 processor through the **ITR** while in debug state.

The relevant combinations of the **DBGEN**, **NIDEN**, **SPIDEN**, and **SPNIDEN** values can be determined by polling **DSCR**[17:16], **DSCR**[15:14], or the Authentication Status Register.

10.5.3 Debug APB interface

Table 10-14 shows the PMU register names and corresponding addresses on the Debug APB interface.

Table 10-14 PMU register names and Debug APB interface addresses

PMU register name	Debug APB Address
PMU event counter 0	0x000
PMU event counter 1	0x004
PMU event counter 2	0x008
PMU event counter 3	0x00C
PMU event counter 4	0x010
PMU event counter 5	0x014
pmcentr	0x07C
pmevtyper0	0x400
pmevtyper1	0x404
pmevtyper2	0x408
pmevtyper3	0x40C
pmevtyper4	0x410
pmevtyper5	0x414
pmcntenset	0xC00
pmcentenclr	0xC20
pmintenset	0xC40
pmintenclr	0xC60
pmovsr	0xC80
pmswinc	0xCA0
pmcr	0xE04
pmuserenr	0xE08

10.5.4 External debug request interface

The following sections describe the external debug request interface signals:

- *EDBGRQ* on page 10-19
- *DBGACK* on page 10-19
- *DBGCPUDONE* on page 10-19
- *COMMRX* and *COMMTX* on page 10-19
- *Memory mapped accesses, DBGROMADDR, and DBGSELFADDR* on page 10-20.

EDBGRQ

This signal generates a halting debug event, that is, it requests the processor to enter debug state. When this occurs, the DSCR[5:2] method of debug entry bits are set to b0100. When **EDBGRQ** is asserted, it must be held until **DBGACK** is asserted. Failure to do so leads to Unpredictable behavior of the processor.

DBGACK

The processor asserts **DBGACK** to indicate that the system has entered debug state. It serves as a handshake for the **EDBGRQ** signal. The **DBGACK** signal is also driven HIGH when the debugger sets the DSCR[10] DbgAck bit to 1.

DBGCPUDONE

DBGCPUDONE is asserted when the core has completed a *Data Synchronization Barrier* (DSB) as part of the entry procedure to debug state.

The processor asserts **DBGCPUDONE** only after it has completed all Non-debug state memory accesses. Therefore the system can use **DBGCPUDONE** as an indicator that all memory accesses issued by the processor result from operations performed by a debugger.

Figure 10-5 shows the Cortex-A9 connections specific to debug request and restart and the CoreSight pins.

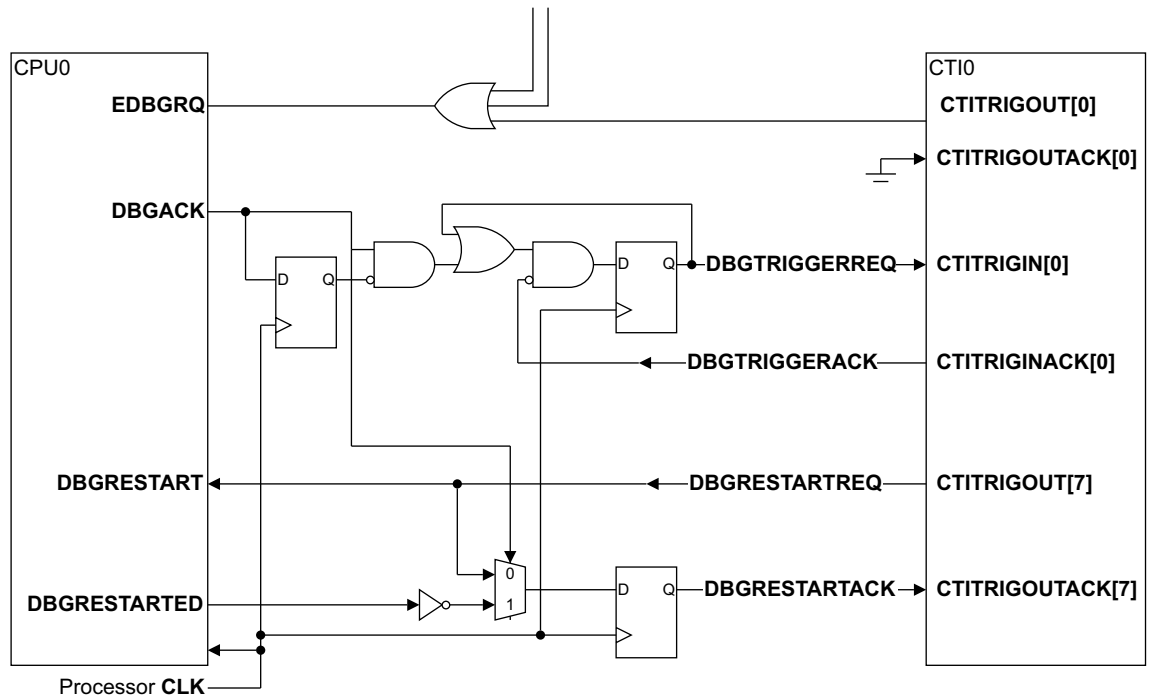


Figure 10-5 Debug request restart-specific connections

COMMRX and COMMTX

The **COMMRX** and **COMMTX** output signals enable interrupt-driven communications over the DTR. By connecting these signals to an interrupt controller, software using the debug communications channel can be interrupted whenever there is new data on the channel or when the channel is clear for transmission.

COMMRX is asserted when the CP14 DTR has data for the processor to read, and it is deasserted when the processor reads the data. Its value is equal to the DBGDSCR[30] DTRRX full flag.

COMMTX is asserted when the CP14 is ready for write data, and it is deasserted when the processor writes the data. Its value equals the inverse of the DBGDSCR[29] DTRTX full flag.

Memory mapped accesses, DBGROMADDR, and DBGSELFADDR

Cortex-A9 processors have a memory-mapped debug interface. Cortex-A9 processors can access the debug and PMU registers by executing load and store instructions going through the AXI bus.

DBGROMADDR gives the base address for the ROM table which locates the physical addresses of the debug components.

DBGSELFADDR gives the offset from the ROM table to the physical addresses of the registers owned by the processor itself.

Chapter 11

Performance Monitoring Unit

This chapter describes the *Performance Monitoring Unit* (PMU) and the registers that it can use. It contains the following sections:

- *About the Performance Monitoring Unit* on page 11-2
- *PMU management registers* on page 11-3
- *Performance monitoring events* on page 11-7.

11.1 About the Performance Monitoring Unit

The Cortex-A9 processor PMU provides six counters to gather statistics on the operation of the processor and memory system. Each counter can count any of the 58 events available in the Cortex-A9 processor.

The PMU counters, and their associated control registers, are accessible from the internal CP15 interface as well as from the Debug APB interface. Table 11-1 shows the mappings of the PMU registers.

Table 11-1 Performance monitoring instructions and Debug APB mapping

Debug APB interface mapping	CP15 instruction	Access	Reset	Name
0x000	0, c9, c13, 2	RW	-	PMXEVNTR0
0x004	0, c9, c13, 2	RW	-	PMXEVNTR1
0x008	0, c9, c13, 2	RW	-	PMXEVNTR2
0x00C	0, c9, c13, 2	RW	-	PMXEVNTR3
0x010	0, c9, c13, 2	RW	-	PMXEVNTR4
0x014	0, c9, c13, 2	RW	-	PMXEVNTR5
0x07C	0, c9, c13, 0	RW	-	PMCCNTR
0x400	0, c9, c13, 1	RW	-	PMXEVTYPER0
0x404	0, c9, c13, 1	RW	-	PMXEVTYPER1
0x408	0, c9, c13, 1	RW	-	PMXEVTYPER2
0x40C	0, c9, c13, 1	RW	-	PMXEVTYPER3
0x410	0, c9, c13, 1	RW	-	PMXEVTYPER4
0x414	0, c9, c13, 1	RW	-	PMXEVTYPER5
0xC00	0, c9, 0 c12, 1	RW	-	PMCNTENSET
0xC20	0, c9, c12, 2	RW	-	PMCNTENCLR
0xC40	0, c9, c14, 1	RW	-	PMINTENSET
0xC60	0, c9, c14, 2	RW	-	PMINTENCLR
0xC80	0, c9, c12, 3	RW	-	PMOVSr
0xCA0	0, c9, c12, 4	WO	-	PMSWINC
0xE04	0, c9, c12, 0	RW	0x41093000	PMCR
0xE08	0, c9, c14, 0	RW ^a	0x00000000	PMUSERENR
-	0, c9, c12, 5	RW	-	PMSELR

a. Read only in user mode.

11.2 PMU management registers

The PMU management registers define the standardized set of registers that is implemented by all CoreSight components. These registers are described in this section. The cp14 interface must be used to access these registers.

Table 11-2 shows the contents of the Management registers for the Cortex-A9 debug unit.

Table 11-2 PMU Management registers

Offset	Register number	Access	Mnemonic	Description
0xD00-0xDFC	832-895	RO	-	<i>Processor ID Registers.</i>
0xE00-0xEF0	854-956	-	-	RAZ.
0xF00	960	RW	ITCTRL	-
0xF04-0xF9C	961-999	RAZ	-	Reserved for Management Register expansion.
0xFA0	1000	RW	CLAIMSET	-
0xFA4	1001	RW	CLAIMCLR	-
0xFA8-0xFBC	1002-1003	-	-	RAZ.
0xFB0	1004	WO	LOCKACCESS	-
0xFB4		RO	LOCKSTATUS	-
0xFB8		RO	AUTHSTATUS	-
0xFBC-0xFC4	1007-1009	-	-	RAZ.
0xFC8	1010	RO	DEVID	Device Identifier.
0xFCC	1011	RO	DEVTYPE	-
0xFD0-0xFFC	1012-1023	R	-	<i>CoreSight Identification Registers on page 11-4.</i>

11.2.1 Processor ID Registers

The Processor ID Registers are read-only registers that return the same values as the corresponding CP15 ID Code Register and Feature ID Register.

Table 11-3 shows the offset value, register number, mnemonic, and description that are associated with each Processor ID Register.

Table 11-3 Processor Identifier Registers

Offset (hex)	Register number	Mnemonic	Access	Register value	Description
0xD00	832	CPUID	RO	0x80000000	ID Code Register ^a
0xD04	833	CTYPR	RO	0x80038003	Cache Type Register
0xD08	834	-	RAZ	-	-
0xD0C	835	TTYPR	RO	0x00000400	TLB Type Register
0xD10-0xD1C	836-839	-	-	-	Reserved

Table 11-3 Processor Identifier Registers (continued)

Offset (hex)	Register number	Mnemonic	Access	Register value	Description
0xD20	840	ID_PFR0	RO	0x00001231	Processor Feature Register 0
0xD24	841	ID_PFR1	RO	0x00000011	Processor Feature Register 1
0xD28	842	ID_DFR0	RO	0x00010444	Debug Feature Register 0
0xD2C	843	ID_AFR0	RAZ	-	Auxiliary Feature Register 0
0xD30	844	ID_MMFR0	RO	0x00100103	Memory Model Feature Register 0
0xD34	845	ID_MMFR1	RO	0x20000000	Memory Model Feature Register 1
0xD38	846	ID_MMFR2	RO	0x01230000	Memory Model Feature Register 2
0xD3C	847	ID_MMFR3	RO	0x00002111	Memory Model Feature Register 3
0xD40	848	ID_ISAR0	RO	0x00101111	Instruction Set Attribute Register 0
0xD44	849	ID_ISAR1	RO	0x13112111	Instruction Set Attribute Register 1
0xD48	850	ID_ISAR2	RO	0x21232041	Instruction Set Attribute Register 2
0xD4C	851	ID_ISAR3	RO	0x11112131	Instruction Set Attribute Register 3
0xD50	852	ID_ISAR4	RO	0x00011142	Instruction Set Attribute Register 4
0xD54	853	ID_ISAR5	RAZ	-	Instruction Set Attribute Register 5

- a. For uniprocessor versions = 0x80000000
 For multiprocessor versions = 0xC000n0m
 n = CLUSTERID input
 m = CPU number (0x0 for CPU0, 0x1 for CPU1, 0x2 for CPU2, and 0x3 for CPU3).

11.2.2 CoreSight Identification Registers

The 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 required by all CoreSight components. Only bits [7:0] of each register are used.

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 11-4 shows the offset value, register number, value, and description that are associated with each Peripheral Identification Register.

Table 11-4 Peripheral Identification Registers

Offset (hex)	Register number	Value	Description
0xFD0	1012	0x04	Peripheral Identification Register 4
0xFD4	1013	-	Reserved
0xFD8	1014	-	Reserved
0xFDC	1015	-	Reserved

Table 11-4 Peripheral Identification Registers (continued)

Offset (hex)	Register number	Value	Description
0xFE0	1016	0xA0	Peripheral Identification Register 0
0xFE4	1017	0xB9	Peripheral Identification Register 1
0xFE8	1018	0x0B	Peripheral Identification Register 2
0xFEC	1019	0x00	Peripheral Identification Register 3

Table 11-5 shows the offset value, register number, and value that are associated with each Component Identification Register.

Table 11-5 Component Identification Registers

Offset (hex)	Register number	Value	Description
0xFF0	1020	0x00	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

11.2.3 PMU APB interface

Table 11-6 shows the PMU register names and corresponding addresses on the APB interface.

Table 11-6 PMU register names and APB addresses

PMU register name	Debug APB Address
PMU event counter 0	0x000
PMU event counter 1	0x004
PMU event counter 2	0x008
PMU event counter 3	0x00C
PMU event counter 4	0x010
PMU event counter 5	0x014
pmccntr	0x07C
pmevtyper0	0x400
pmevtyper1	0x404
pmevtyper2	0x408
pmevtyper3	0x40C
pmevtyper4	0x410
pmevtyper5	0x414
pmcntenset	0xC00

Table 11-6 PMU register names and APB addresses (continued)

PMU register name	Debug APB Address
pmcntenclr	0xC20
pmintenset	0xC40
pmintenclr	0xC60
pmovsr	0xC80
pmswinc	0xCA0
pmcr	0xE04
pmuserenr	0xE08

11.3 Performance monitoring events

The Cortex-A9 processor implements the architectural events described in the *ARM Architecture Reference Manual*, with the exception of:

- 0x08** Memory-reading instruction architecturally executed
- 0x0E** Procedure return, other than exception return, architecturally executed.

For events and the corresponding **PMUEVENT** signals, see Table A-18 on page A-14.

The PMU provides an additional set of Cortex-A9 specific events.

11.3.1 Cortex-A9 specific events

Table 11-7 shows the Cortex-A9 specific events. In the value column of Table 11-7 Precise means the event is counted precisely. Events related to stalls and speculative instructions appear as Approximate entries in this column.

Table 11-7 Cortex-A9 specific events

Event	Description	Value
0x40	Java bytecode execute ^a Counts the number of Java bytecodes being decoded, including speculative ones.	Approximate
0x41	Software Java bytecode executed. ^a Counts the number of software java bytecodes being decoded, including speculative ones.	Approximate
0x42	Jazelle backward branches executed ^a . Counts the number of Jazelle taken branches being executed. This includes the branches that are flushed because of a previous load/store which aborts late.	Approximate
0x50	Coherent linefill miss ^b Counts the number of coherent linefill requests performed by the Cortex-A9 processor which also miss in all the other Cortex-A9 processors, meaning that the request is sent to the external memory.	Precise
0x51	Coherent linefill hit ^b Counts the number of coherent linefill requests performed by the Cortex-A9 processor which hit in another Cortex-A9 processor, meaning that the linefill data is fetched directly from the relevant Cortex-A9 cache.	Precise
0x60	Instruction cache dependent stall cycles Counts the number of cycles where the processor is ready to accept new instructions, but does not receive any because of the instruction side not being able to provide any and the instruction cache is currently performing at least one linefill.	Approximate
0x61	Data cache dependent stall cycles Counts the number of cycles where the core has some instructions that it cannot issue to any pipeline, and the Load Store unit has at least one pending linefill request, and no pending TLB requests.	Approximate
0x62	Main TLB miss stall cycles Counts the number of cycles where the processor is stalled waiting for the completion of translation table walks from the main TLB. The processor stalls can be because of the instruction side not being able to provide the instructions, or to the data side not being able to provide the necessary data, because of them waiting for the main TLB translation table walk to complete.	Approximate
0x63	STREX passed Counts the number of STREX instructions architecturally executed and passed.	Precise

Table 11-7 Cortex-A9 specific events (continued)

Event	Description	Value
0x64	STREX failed Counts the number of STREX instructions architecturally executed and failed.	Precise
0x65	Data eviction Counts the number of eviction requests because of a linefill in the data cache.	Precise
0x66	Issue does not dispatch any instruction Counts the number of cycles where the issue stage does not dispatch any instruction because it is empty or cannot dispatch any instructions.	Precise
0x67	Issue is empty Counts the number of cycles where the issue stage is empty.	Precise
0x68	Instructions coming out of the core renaming stage Counts the number of instructions going through the Register Renaming stage. This number is an approximate number of the total number of instructions speculatively executed, and even more approximate of the total number of instructions architecturally executed. The approximation depends mainly on the branch misprediction rate. The renaming stage can handle two instructions in the same cycle so the event is two bits long: <ul style="list-style-type: none"> • b00 no instructions coming out of the core renaming stage • b01 one instruction coming out of the core renaming stage • b10 two instructions coming out of the core renaming stage. See Table A-17 on page A-14 for a description of how these values map to the PMUEVENT bus bits.	Approximate
0x6E	Predictable function returns Counts the number of procedure returns whose condition codes do not fail, excluding all returns from exception. This count includes procedure returns which are flushed because of a previous load/store which aborts late. Only the following instructions are reported: <ul style="list-style-type: none"> • BX R14 • MOV PC, LR • POP {...,pc} • LDR pc, [sp], #offset. The following instructions are not reported: <ul style="list-style-type: none"> • LDMIA R9!, {...,PC} (ThumbEE state only) • LDR PC, [R9], #offset (ThumbEE state only) • BX R0 (Rm != R14) • MOV PC, R0 (Rm != R14) • LDM SP, {...,PC} (writeback not specified) • LDR PC, [SP, #offset] (wrong addressing mode). 	Approximate
0x70	Main execution unit instructions Counts the number of instructions being executed in the main execution pipeline of the processor, the multiply pipeline and arithmetic logic unit pipeline. The counted instructions are still speculative.	Approximate
0x71	Second execution unit instructions Counts the number of instructions being executed in the processor second execution pipeline (ALU). The counted instructions are still speculative.	Approximate
0x72	Load/Store Instructions Counts the number of instructions being executed in the Load/Store unit. The counted instructions are still speculative.	Approximate

Table 11-7 Cortex-A9 specific events (continued)

Event	Description	Value
0x73	<p>Floating-point instructions</p> <p>Counts the number of Floating-point instructions going through the Register Rename stage. Instructions are still speculative in this stage.</p> <p>Two floating-point instructions can be renamed in the same cycle so the event is two bits long:</p> <p>0b00 no floating-point instruction renamed</p> <p>0b01 one floating-point instruction renamed</p> <p>0b10 two floating-point instructions renamed.</p> <p>See Table A-17 on page A-14 for a description of how these values map to the PMUEVENT bus bits.</p>	Approximate
0x74	<p>NEON instructions</p> <p>Counts the number of NEON instructions going through the Register Rename stage. Instructions are still speculative in this stage.</p> <p>Two NEON instructions can be renamed in the same cycle so the event is two bits long:</p> <p>0b00 no NEON instruction renamed</p> <p>0b01 one NEON instruction renamed</p> <p>0b10 two NEON instructions renamed.</p> <p>See Table A-17 on page A-14 for a description of how these values map to the PMUEVENT bus bits.</p>	Approximate
0x80	<p>Processor stalls because of PLDs</p> <p>Counts the number of cycles where the processor is stalled because PLD slots are all full.</p>	Approximate
0x81	<p>Processor stalled because of a write to memory</p> <p>Counts the number of cycles when the processor is stalled and the data side is stalled too because it is full and executing writes to the external memory.</p>	Approximate
0x82	<p>Processor stalled because of instruction side main TLB miss</p> <p>Counts the number of stall cycles because of main TLB misses on requests issued by the instruction side.</p>	Approximate
0x83	<p>Processor stalled because of data side main TLB miss</p> <p>Counts the number of stall cycles because of main TLB misses on requests issued by the data side.</p>	Approximate
0x84	<p>Processor stalled because of instruction micro TLB miss</p> <p>Counts the number of stall cycles because of micro TLB misses on the instruction side. This event does not include main TLB miss stall cycles that are already counted in the corresponding main TLB event.</p>	Approximate
0x85	<p>Processor stalled because of data micro TLB miss</p> <p>Counts the number of stall cycles because of micro TLB misses on the data side. This event does not include main TLB miss stall cycles that are already counted in the corresponding main TLB event.</p>	Approximate
0x86	<p>Processor stalled because of DMB</p> <p>Counts the number of stall cycles because of the execution of a DMB memory barrier. This includes all DMB instructions being executed, even speculatively.</p>	Approximate
0x8A	<p>Integer clock enabled</p> <p>Counts the number of cycles during which the integer core clock is enabled.</p>	Approximate
0x8B	<p>Data Engine clock enabled</p> <p>Counts the number of cycles during which the Data Engine clock is enabled.</p>	Approximate
0x90	<p>ISB instructions</p> <p>Counts the number of ISB instructions architecturally executed.</p>	Precise

Table 11-7 Cortex-A9 specific events (continued)

Event	Description	Value
0x91	DSB instructions Counts the number of DSB instructions architecturally executed.	Precise
0x92	DMB instructions Counts the number of DMB instructions speculatively executed.	Approximate
0x93	External interrupts Counts the number of external interrupts executed by the processor.	Approximate
0xA0	PLE cache line request completed. ^c	Precise
0xA1	PLE cache line request skipped. ^c	Precise
0xA2	PLE FIFO flush. ^c	Precise
0xA3	PLE request completed. ^c	Precise
0xA4	PLE FIFO overflow. ^c	Precise
0xA5	PLE request programmed. ^c	Precise

- a. Only when the design implements the Jazelle Extension. Otherwise reads as 0.
- b. For use with Cortex-A9 multiprocessor variants.
- c. Active only when the PLE is present. Otherwise reads as 0.

Appendix A

Signal Descriptions

This appendix lists and describes the Cortex-A9 signals. It contains the following sections:

- *Clock and clock control signals* on page A-2
- *Resets and reset control* on page A-3
- *Interrupts* on page A-4
- *Configuration signals* on page A-5
- *Standby and Wait For Event signals* on page A-6
- *Power management signals* on page A-7
- *AXI interfaces* on page A-8
- *Performance monitoring signals* on page A-14
- *Exception flags signal* on page A-17
- *Parity signal* on page A-18.
- *MBIST interface* on page A-19
- *Scan test signal* on page A-20.
- *External Debug interface* on page A-21
- *PTM interface signals* on page A-24.

A.1 Clock and clock control signals

The Cortex-A9 processor has a single externally generated global clock. Table A-1 shows the clock and clock control signal.

Table A-1 Clock and clock control signals for Cortex-A9

Name	I/O	Source	Description
CLK	I	Clock controller	Global clock. See <i>Clocking and resets</i> on page 2-6.
MAXCLKLATENCY[2:0]	I	Implementation-specific static value	Controls dynamic clock gating delays. This pin is sampled during reset of the processor. See <i>Dynamic high level clock gating</i> on page 2-8

A.2 Resets and reset control

Table A-2 shows the reset and reset control signals.

Table A-2 Cortex-A9 processor reset signals

Name	I/O	Source	Description
nCPURESET	I	Reset controller	Cortex-A9 processor reset.
nDBGRESET	I		Cortex-A9 processor debug logic reset.
NEONCLKOFF^a	I		MPE SIMD logic clock control 0 = Do not cut MPE SIMD logic clock 1 = Cut MPE SIMD logic clock.
nNEONRESET^a	I		Cortex-A9 MPE SIMD logic reset.

a. Only if the MPE is present.

See *Reset* on page 2-6.

A.3 Interrupts

Table A-3 shows the interrupt line signals.

Table A-3 Interrupt line signals

Name	I/O	Source	Description
nFIQ	I	Interrupt sources	Cortex-A9 processor FIQ request input line. Active-LOW fast interrupt request: 0 = Activate fast interrupt 1 = Do not activate fast interrupt. The processor treats the nFIQ input as level sensitive.
nIRQ	I	Interrupt sources	Cortex-A9 processor IRQ request input line. Active-LOW interrupt request: 0 = Activate interrupt 1 = Do not activate interrupt. The processor treats the nIRQ input as level sensitive.

A.4 Configuration signals

Table A-4 shows the configuration signals only sampled during reset of the processor.

Table A-4 Configuration signals

Name	I/O	Source	Description
CFGEND	I	System configuration control	Controls the state of EE bit in the SCTLR at reset: 0 = EE bit is LOW 1 = EE bit is HIGH
CFGNMFI	I		Configures fast interrupts to be nonmaskable: 0 = Clear the NMFI bit in the CP15 c1 Control Register 1 = Set the NMFI bit in the CP15 c1 Control Register.
TEINIT	I		Default exception handling state: 0 = ARM 1 = Thumb. It sets the SCTLR.TE bit at reset.
VINITHI	I		Controls the location of the exception vectors at reset: 0 = Start exception vectors at address 0x00000000 1 = Start exception vectors at address 0xFFFF0000. It sets the SCTLR.V bit.

Table A-5 shows the CP15SDISABLE signal.

Table A-5 CP15SDISABLE signal

Name	I/O	Source	Description
CP15SDISABLE	I	Security controller	Disables write access to some system control processor registers in Secure state: 0 = Not enabled 1 = Enabled. See <i>System Control Register</i> on page 4-15.

A.5 Standby and Wait For Event signals

Table A-6 shows standby and wait for event signals.

Table A-6 Standby and wait for event signals

Name	I/O	Source or destination	Description
EVENTI	I	External coherent agent	Event input for Cortex-A9 processor wake-up from WFE state.
EVENTO	O		Event output. This signal is active HIGH for one processor clock cycle when one SEV instruction is executed.
STANDBYWFI	O	Power controller	Indicates if the processor is in WFI state: 0 = Processor not in wait for event state 1 = Processor in wait for event state.
STANDBYWFE	O		Indicates if the processor is in WFE state: 0 = Processor not in wait for event state 1 = Processor in wait for event state.

See *Standby modes* on page 2-11.

A.6 Power management signals

Table A-7 shows the power management signals.

Table A-7 Power management signals

Name	I/O	Source	Description
CPURAMCLAMP	I	Power controller	Activates the CPU RAM interface clamps: 0 = Clamps not active 1 = Clamps active.
NEONCLAMP ^a	I		Activates the Cortex-A9 MPE SIMD logic clamps: 0 = Clamps not active 1 = Clamps active.

a. Only if the MPE is present.

See *Power management* on page 2-10.

A.7 AXI interfaces

In Cortex-A9 designs there can be two AXI master ports. The following sections describe the AXI interfaces:

- *AXI Master0 signals data accesses*
- *AXI Master1 signals instruction accesses* on page A-11.

A.7.1 AXI Master0 signals data accesses

The following sections describe the AXI Master0 interface signals used for data read and write accesses:

- *Write address signals for AXI Master0*
- *Write data channel signals* on page A-9
- *Write response channel signals* on page A-10
- *Read data channel signals* on page A-10
- *Read data channel signals* on page A-11
- *AXI Master0 Clock enable signals* on page A-11.

Write address signals for AXI Master0

Table A-8 shows the AXI write address signals for AXI Master0.

Table A-8 AXI-AW signals for AXI Master0

Name	I/O	Source or destination	Description
AWADDRM0[31:0]	O	AXI system devices	Address.
AWBURSTM0[1:0]	O		Burst type = b01, INCR incrementing burst.
AWCACHEM0[3:0]	O		Cache type giving additional information about cacheable characteristics, determined by the memory type and Outer cache policy for the memory region.
AWIDM0[1:0]	O		Request ID

Table A-8 AXI-AW signals for AXI Master0 (continued)

Name	I/O	Source or destination	Description
AWLENM0[3:0]	O	AXI system devices	The number of data transfers that can occur within each burst..
AWLOCKM0[1:0]	O		Lock type.
AWPROTM0[2:0]	O		Protection Type.
AWREADYM0	I		Address ready.
AWSIZEM0[1:0]	O		Data transfer size: b000 = 8-bit transfer b001 = 16-bit transfer b010 = 32-bit transfer b011 = 64-bit transfer.
AWUSERM0[8:0]	O		[8] early BRESP . Used with L2C-310. [7] full line of write zeros. Used with the L2C-310. [6] clean eviction. [5] level 1 eviction. [4:1] memory type and Inner cache policy. b0000 = Strongly-ordered. b0001 = Device b0011 = Normal Memory Non-Cacheable. b0110 = Write-Through. b0111 = Write-Back no Write-Allocate. b1111 = Write-Back Write-Allocate. [0] shared.
AWVALIDM0	O		Address valid.

Write data channel signals

Table A-9 shows the AXI write data signals for AXI Master0.

Table A-9 AXI-W signals for AXI Master0

Name	I/O	Source or destination	Description
WDATAM0[63:0]	O	AXI system devices	Write data
WIDM0[1:0]	O		Write ID
WLASTM0	O		Write last indication
WREADYM0	I		Write ready
WSTRBM0[7:0]	O		Write byte lane strobe
WVALIDM0	O		Write valid

Write response channel signals

Table A-10 shows the AXI write response signals for AXI Master0.

Table A-10 AXI-B signals for AXI Master0

Name	I/O	Source or destination	Description
BIDM0[1:0]	I	AXI system devices	Response ID
BREADYM0	O		Response ready
BRESPM0[1:0]	I		Write response
BVALIDM0	I		Response valid

Read data channel signals

Table A-11 shows the AXI read address signals for AXI Master0.

Table A-11 AXI-AR signals for AXI Master0

Name	I/O	Source or destination	Description
ARADDRM0[31:0]	O	AXI system devices	Address.
ARBURSTM0[1:0]	O		Burst type: b01 = INCR incrementing burst b10 = WRAP Wrapping burst.
ARCAHEM0[3:0]	O		Cache type giving additional information about cacheable characteristics.
ARIDM0[1:0]	O		Request ID
ARLENM0[3:0]	O		The number of data transfers that can occur within each burst.
ARLOCKM0[1:0]	O		Lock type.
ARPROTM0[2:0]	O		Protection Type
ARREADYM0	I		Address ready.
AR SizEM0[1:0]	O	AXI system devices	Burst size: b000 = 8-bit transfer b001 = 16-bit transfer b010 = 32-bit transfer b011 = 64-bit transfer.
ARUSERM0[4:0]	O		[4:1] memory type and Inner cache policy b0000 = Strongly-ordered b0001 = Device b0011 = Normal Memory Non-Cacheable b0110 = Write-Through b0111 = Write-Back no Write-Allocate b1111 = Write-Back Write-Allocate. [0] shared.
ARVALIDM0	O		Address valid.

Read data channel signals

Table A-12 shows the AXI read data signals for AXI Master0.

Table A-12 AXI-R signals for AXI Master0

Name	I/O	Source or destination	Description
RVALIDM0	I	AXI system devices	Read valid
RDATAM0[63:0]	I		Read data
RRESPM0[1:0]	I		Read response
RLASTM0	I		Read Last indication
RIDM0[1:0]	I		Read ID
RREADYM0	O		Read ready

AXI Master0 Clock enable signals

This section describes the AXI Master0 clock enable signals. Table A-13 shows the AXI Master0 clock enable signal.

Table A-13 AXI Master0 clock enable signal

Name	I/O	Source	Description
ACLKENM0	I	Clock controller	Clock enable for the AXI bus that enables the AXI interface to operate at integer ratios of the system clock. See <i>Clocking and resets</i> on page 2-6.

A.7.2 AXI Master1 signals instruction accesses

The following sections describe the AXI Master1 interface signals, which are used for instruction accesses:

- *Read data channel signals* on page A-12
- *Read data channel signals* on page A-13
- *AXI Master1 Clock enable signals* on page A-13.

Read data channel signals

Table A-14 shows the AXI read address signals for AXI Master1.

Table A-14 AXI-AR signals for AXI Master1

Name	I/O	Destination	Description
ARADDRM1[31:0]	O	AXI system devices	Address.
ARBURSTM1[1:0]	O		Burst type: b01 = INCR incrementing burst b10 = WRAP Wrapping burst.
ARCACHEM1[3:0]	O		Cache type giving additional information about cacheable characteristics.
ARIDM1[5:0]	O		Request ID.
ARLENM1[3:0]	O		The number of data transfers that can occur within each burst.
ARLOCKM1[1:0]	O		Lock type: b00 = Normal access.
ARPROTM1[2:0]	O		Protection Type.
ARREADYM1	I		Address ready.
ARSIZEM1[1:0]	O	AXI system devices	Burst size: b000 = 8-bit transfer b001 = 16-bit transfer b010 = 32-bit transfer b011 = 64-bit transfer.
ARUSERM1[4:0]	O		[4:1] = Inner attributes b0000 = Strongly-ordered b0001 = Device b0011 = Normal Memory Non-Cacheable b0110 = Write-Through b0111 = Write-Back no Write-Allocate b1111 = Write-Back Write-Allocate. [0] = Shared.
ARVALIDM1	O		Address valid.

Read data channel signals

Table A-15 shows the AXI read data signals for AXI Master1.

Table A-15 AXI-R signals for AXI Master1

Name	I/O	Source or destination	Description
RVALIDM1	I	AXI system devices	Read valid
RDATAM1[63:0]	I		Read data
RRESPM1[1:0]	I		Read response
RLASTM1	I		Read Last indication
RIDM1[5:0]	I		Read ID
RREADYM1	O		Read ready

AXI Master1 Clock enable signals

This section describes the AXI Master1 clock enable signals. Table A-16 shows the AXI Master1 clock enable signals.

Table A-16 AXI Master1 clock enable signal

Name	I/O	Source	Description
ACLKENM1	I	Clock controller	Clock enable for the AXI bus that enables the AXI interface to operate at integer ratios of the system clock. See <i>Clocking and resets</i> on page 2-6.

See Chapter 8 *Level 2 Memory Interface*.

A.8 Performance monitoring signals

Table A-17 shows the performance monitoring signals.

Table A-17 Performance monitoring signals

Name	I/O	Destination	Description
PMUEVENT[57:0]	O	PTM or external monitoring unit	Performance Monitoring Unit event bus. See Table A-18.
PMUIRQ	O		Performance Monitoring Unit interrupt signal.
PMUSECURE	O		Gives the status of the Cortex-A9 processor 0 = In Non-secure state 1 = In Secure state. This signal does not provide input to CoreSight Trace delivery infrastructure.
PMUPRIV	O		Gives the status of the Cortex-A9 processor 0 = In user mode 1 = In privileged mode. This signal does not provide input to CoreSight Trace delivery infrastructure.

Table A-18 gives the correlation between **PMUEVENT** signals and their event numbers.

Table A-18 Event signals and event numbers

Name	Event number	Description
PMUEVENT[0]	0x00	Software increment
PMUEVENT[1]	0x01	Instruction cache miss
PMUEVENT[2]	0x02	Instruction micro TLB miss
PMUEVENT[3]	0x03	Data cache miss
PMUEVENT[4]	0x04	Data cache access
PMUEVENT[5]	0x05	Data micro TLB miss
PMUEVENT[6]	0x06	Data read
PMUEVENT[7]	0x07	Data writes
-	0x08	Unused ^a
PMUEVENT[8]	0x68	b00 = No instructions renamed
PMUEVENT[9]		b01 = Nne instruction renamed b10 = Two instructions renamed.
PMUEVENT[10]	0x09	Exception taken
PMUEVENT[11]	0x0A	Exception returns
PMUEVENT[12]	0x0B	Write context id
PMUEVENT[13]	0x0C	Software change of PC
PMUEVENT[14]	0x0D	Immediate branch
-	0x0E	Unused ^b
PMUEVENT[15]	0x6E	Predictable function return ^b

Table A-18 Event signals and event numbers (continued)

Name	Event number	Description
PMUEVENT[16]	0x0F	Unaligned
PMUEVENT[17]	0x10	Branch mispredicted or not predicted
Not exported	0x11	Cycle count
PMUEVENT[18]	0x12	Predictable branches
PMUEVENT[19]	0x40	Java bytecode
PMUEVENT[20]	0x41	Software Java bytecode
PMUEVENT[21]	0x42	Jazelle backward branch
PMUEVENT[22]	0x50	Coherent linefill miss ^c
PMUEVENT[23]	0x51	Coherent linefill hit ^c
PMUEVENT[24]	0x60	Instruction cache dependent stall
PMUEVENT[25]	0x61	Data cache dependent stall
PMUEVENT[26]	0x62	Main TLB miss stall
PMUEVENT[27]	0x63	STREX passed
PMUEVENT[28]	0x64	STREX failed
PMUEVENT[29]	0x65	Data eviction
PMUEVENT[30]	0x66	Issue does not dispatch any instruction
PMUEVENT[31]	0x67	Issue is empty
PMUEVENT[32]	0x70	Main Execution Unit pipe
PMUEVENT[33]	0x71	Second Execution Unit pipe
PMUEVENT[34]	0x72	Load/Store pipe
PMUEVENT[35]	0x73	b00 = No floating-point instruction renamed
PMUEVENT[36]		b01 = One floating-point instruction renamed
		b10 = Two floating-point instructions renamed
PMUEVENT[37]	0x74	b00 = No NEON instruction renamed
PMUEVENT[38]		b01 = One NEON instruction renamed
		b10 = Two NEON instructions renamed
PMUEVENT[39]	0x80	PLD stall
PMUEVENT[40]	0x81	Write stall
PMUEVENT[41]	0x82	Instruction main TLB miss stall
PMUEVENT[42]	0x83	Data main TLB miss stall
PMUEVENT[43]	0x84	Instruction micro TLB miss stall
PMUEVENT[44]	0x85	Data micro TLB miss stall
PMUEVENT[45]	0x86	DMB stall
PMUEVENT[46]	0x8A	Integer core clock enabled

Table A-18 Event signals and event numbers (continued)

Name	Event number	Description
PMUEVENT[47]	0x8B	Data Engine clock enabled
PMUEVENT[48]	0x90	ISB
PMUEVENT[49]	0x91	DSB
PMUEVENT[50]	0x92	DMB
PMUEVENT[51]	0x93	External interrupt
PMUEVENT[52]	0xA0	PLE cache line request completed
PMUEVENT[53]	0xA1	PLE cache line request skipped
PMUEVENT[54]	0xA2	PLE FIFO Flush
PMUEVENT[55]	0xA3	PLE request completed
PMUEVENT[56]	0xA4	PLE FIFO Overflow
PMUEVENT[57]	0xA5	PLE request programmed

- a. Not generated by Cortex-A9 processors. Replaced by the similar event 0x68.
- b. Not generated by Cortex-A9 processors. Replaced by the similar event 0x6E.
- c. Used in multiprocessor configurations

See *Cortex-A9 specific events* on page 11-7.

A.9 Exception flags signal

Table A-19 shows the **DEFLAGS** signal.

Table A-19 DEFLAGS signal

Name	I/O	Destination	Description
DEFLAGS[6:0]	O	Exception monitoring unit	<p>Data Engine output flags. Only implemented if the Cortex-A9 processor includes a Data Engine, either an MPE or FPU.</p> <p>If the DE is MPE:</p> <ul style="list-style-type: none"> • Bit[6] gives the value of FPSCR[27] • Bit[5] gives the value of FPSCR[7] • Bits[4:0] give the value of FPSCR[4:0]. <p>If the DE is FPU:</p> <ul style="list-style-type: none"> • Bit[6] is zero. • Bit[5] gives the value of FPSCR[7] • Bits[4:0] give the value of FPSCR[4:0].

For additional information on the FPSCR, see the *Cortex-A9 Floating-Point Unit (FPU) Technical Reference Manual* and the *Cortex-A9 NEON® Media Processing Engine Technical Reference Manual*.

A.10 Parity signal

Table A-20 shows the parity signal. This signal is present only if parity is defined. See *Parity error support* on page 7-11.

Table A-20 Parity signal

Name	I/O	Destination	Description
PARITYFAIL[7:0]	O	Parity monitoring device	Parity output pin from the RAM arrays: 0 no parity fail 1 parity fail Bit [7] BTAC parity error Bit [6] GHB parity error Bit [5] Instruction tag RAM parity error Bit [4] Instruction data RAM parity error Bit [3] main TLB parity error Bit [2] D outer RAM parity error Bit [1] Data tag RAM parity error Bit [0] Data data RAM parity error.

A.11 MBIST interface

Table A-21 shows the MBIST interface signals. These signals are present only when the BIST interface is present.

Table A-21 MBIST interface signals

Name	I/O	Source	Description
MBISTADDR[10:0]	I	MBIST controller	MBIST address bus.
MBISTARRAY[19:0]	I		MBIST arrays used for testing RAMs.
MBISTENABLE	I		MBIST test enable
MBISTWRITEEN	I		Global write enable.
MBISTREADEN	I		Global read enable.

The size of some MBIST signals depends on whether the implementation has parity support or not. Table A-22 shows these signals with parity support implemented.

Table A-22 MBIST signals with parity support implemented

Name	I/O	Source or destination	Description
MBISTBE[32:0]	I	MBIST controller	MBIST write enable
MBISTINDATA[71:0]	I		MBIST data in
MBISTOUTDATA[71:0]	O		MBIST data out

Table A-23 shows these signals without parity support implemented.

Table A-23 MBIST signals without parity support implemented

Name	I/O	Source/Destination	Description
MBISTBE[25:0]	I	MBIST controller	MBIST write enable
MBISTINDATA[63:0]	I		MBIST data in
MBISTOUTDATA[63:0]	O		MBIST data out

See the *Cortex-A9 r0p0 MBIST TRM* for a description of MBIST.

A.12 Scan test signal

Table A-24 lists the scan test signal.

Table A-24 Scan test signal

Name	I/O	Destination	Description
SE	I	DFT controller	Scan enable: 0 = Not enabled 1 = Enabled.

A.13 External Debug interface

The following sections describe the external debug interface signals:

- *Authentication interface*
- *APB interface signals* on page A-22
- *CTI signals* on page A-22
- *Miscellaneous debug interface signals* on page A-23.

A.13.1 Authentication interface

Table A-25 shows the authentication interface signals.

Table A-25 Authentication interface signals

Name	I/O	Source	Description
DBGEN	I	Security controller	Invasive debug enable: 0 = Not enabled 1 = Enabled.
NIDEN	I		Noninvasive debug enable: 0 = Not enabled 1 = Enabled.
SPIDEN	I		Secure privileged invasive debug enable: 0 = Not enabled 1 = Enabled.
SPNIDEN	I		Secure privileged noninvasive debug enable: 0 = Not enabled 1 = Enabled.

A.13.2 APB interface signals

Table A-26 shows the APB interface signals.

Table A-26 APB interface signals

Name	I/O	Source or destination	Description
PENABLEDBG	I	CoreSight APB devices	APB clock enable.
PRDATADBG[31:0]	O		APB read data bus.
PSELDBG	I		Debug registers select: 0 = Debug registers not selected 1 = Debug registers selected.
PSLVERRDBG	O		APB slave error signal.
PWRITEDBG	I		APB Read/Write signal.
PADDRDBG[12:2]	I		Programming address.
PADDRDBG31	I		APB address bus bit [31]: 0 = Not an external debugger access 1 = External debugger access.
PREADYDBG	O		APB slave ready. An APB slave can assert PREADY to extend a transfer.
PWDATADBG[31:0]	I		APB write data.

A.13.3 CTI signals

Table A-27 shows the CTI signals.

Table A-27 CTI signals

Name	I/O	Source or destination	Description
EDBGRQ	I	External debugger or CoreSight interconnect	External debug request: 0 = No external debug request 1 = External debug request. The processor treats the EDBGRQ input as level-sensitive. The EDBGRQ input must be asserted until the processor asserts DBGACK .
DBGACK	O		Debug acknowledge signal.
DBGCPUDONE	O		Indicates that all memory accesses issued by the Cortex-A9 processor result from operations performed by a debugger. Active HIGH.
DBGRESTART	I		Causes the core to exit from Debug state. It must be held HIGH until DBGRESTARTED is deasserted. 0 = Not enabled 1 = Enabled.
DBGRESTARTED	O		Used with DBGRESTART to move between Debug state and Normal state. 0 = Not enabled 1 = Enabled.

A.13.4 Miscellaneous debug interface signals

Table A-28 shows the miscellaneous debug interface signals.

Table A-28 Miscellaneous debug signals

Name	I/O	Source or destination	Description
COMMRX	O	Debug comms channel	Communications channel receive. Receive portion of Data Transfer Register full flag: 0 = Empty 1 = Full.
COMMTX	O	Debug comms channel	Communications channel transmit. Transmit portion of Data Transfer Register full flag: 0 = Empty 1 = Full.
DBGNOPWRDWN	O	Debugger	Debugger has requested the Cortex-A9 processor is not powered down.
DBGSWENABLE	I	External debugger	When LOW only the external debug agent can modify debug registers. 0 = Not enabled. 1 = Enabled.
DBGROMADDR[31:12]	I	System configuration	Specifies bits [31:12] of the ROM table physical address. If the address cannot be determined tie this signal off to zero.
DBGROMADDRV	I		Valid signal for DBGROMADDR . If the address cannot be determined tie this signal LOW.
DBGSELFADDR[31:15]	I		Specifies bits [31:15] of the two's complement signed offset from the ROM table physical address to the physical address where the debug registers are memory-mapped. If the offset cannot be determined tie this signal off to zero.
DBGSELFADDRV	I		Valid signal for DBGSELFADDR . If the offset cannot be determined tie this signal LOW.

See Chapter 10 *Debug*.

A.14 PTM interface signals

Table A-29 shows the PTM interface signals. These signals are present only if the PTM interface is present.

In the Input/Output column I indicates an input from the PTM interface to the Cortex-A9 processor. O indicates an output from the Cortex-A9 processor to the PTM. All these signals are in the Cortex-A9 clock domain.

Table A-29 PTM interface signals

Name	I/O	Source or destination	Description
WPTCOMMIT[1:0]	O	PTM device	Number of waypoints committed this cycle. It is valid to indicate a valid waypoint and commit it in the same cycle.
WPTCONTEXTID[31:0]	O		Context ID for the waypoint. This signal must be true regardless of the condition code of the waypoint. If the core Context ID has not been set, then WPTCONTEXTID[31:0] must report 0.
WPTENABLE	I		Enable waypoint.
WPTEXCEPTIONTYPE[3:0]	O		Exception type: b0001 = Halting debug-mode b0010 = Secure Monitor b0100 = Imprecise Data Abort b0101 = T2EE trap b1000 = Reset b1001 = UNDEF b1010 = SVC b1011 = Prefetch abort/software breakpoint b1100 = Precise data abort/software watchpoint b1110 = IRQ b1111 = FIQ.
WPTFLUSH	O		Waypoint flush signal.
WPTLINK	O		The waypoint is a branch and updates the link register. Only HIGH if WPTTYPE is a direct branch or an indirect branch.

Table A-29 PTM interface signals (continued)

Name	I/O	Source or destination	Description
WPTPC[31:0]	O	PTM device	Waypoint last executed address indicator. This is the base Link Register in the case of an exception. Equal to 0 if the waypoint is reset exception.
WPTT32LINK	O		Indicates the size of the last executed address when in Thumb state: 0 = 16-bit instruction 1 = 32-bit instruction.
WPTTAKEN	O		The waypoint passed its condition codes. The address is still used, irrespective of the value of this signal. Must be set for all waypoints except branch.
WPTTARGETJBIT	O		J bit for waypoint destination.
WPTTARGETPC[31:0]	O		Waypoint target address. Bit [1] must be zero if the T bit is zero. Bit [0] must be zero if the J bit is zero. The value is zero if WPTTYPE is either prohibited or debug.
WPTTARGETTBIT	O		T bit for waypoint destination
WPTTRACEPROHIBITED	O	PTM device	Trace is prohibited for the current waypoint target. Indicates entry to prohibited region. No more waypoints are traced until trace can resume. This signal must be permanently asserted if NIDEN and DBGEN are both LOW, after the in-flight waypoints have exited the core. Either an exception or a serial branch is required to ensure that changes to the inputs have been sampled. Only one WPTVALID cycle must be seen with WPTTRACEPROHIBITED set. Trace stops with this waypoint and the next waypoint seen is an Isync packet. See the <i>CoreSight PTM Architecture Specification</i> for a description of the packets used in trace.
WPTTYPE[2:0]	O		Waypoint Type. b000 = Direct branch b001 = Indirect branch b010 = Exception b011 = DMB/DSB/ISB b100 = Debug entry b101 = Debug exit b110 = Invalid b111 = Invalid. Debug Entry must be followed by Debug Exit. ————— Note ————— Debug exit does not reflect the execution of an instruction.

Table A-29 PTM interface signals (continued)

Name	I/O	Source or destination	Description
WPTVALID	O	PTM device	Waypoint is confirmed as valid.
WPTnSECURE	O		Instructions following the current waypoint are executed in Non-secure state. An instruction is in Non-secure state if the NS bit is set and the processor is not in secure monitor mode. See <i>About system control</i> on page 4-2 for information about security extensions.
WPTFIFOEMPTY	O		There are no speculative waypoints in the PTM interface FIFO.

See *Interfaces* on page 2-4.

Appendix B

Instruction Cycle Timings

This chapter describes the cycle timings of integer instructions on Cortex-A9 processors. It contains the following sections:

- *About instruction cycle timing* on page B-2
- *Data-processing instructions* on page B-3
- *Load and store instructions* on page B-4
- *Multiplication instructions* on page B-7
- *Branch instructions* on page B-8
- *Serializing instructions* on page B-9.

B.1 About instruction cycle timing

This chapter provides information to estimate how much execution time particular code sequences require. The complexity of the Cortex-A9 processor makes it impossible to calculate precise timing information manually. The timing of an instruction is often affected by other concurrent instructions, memory system activity, and additional events outside the instruction flow. Detailed descriptions of all possible instruction interactions and all possible events taking place in the processor is beyond the scope of this document.

B.2 Data-processing instructions

Table B-1 shows the execution unit cycle time for data-processing instructions.

Table B-1 shows the following cases:

no shift on source registers

For example, ADD r0, r1, r2

shift by immediate source register

For example, ADD r0, r1, r2 LSL #2

shift by register

For example, ADD r0, r1, r2 LSL r3.

Table B-1 Data-processing instructions cycle timings

Instruction	No shift	Shift by	
		Constant	Register
MOV	1	1	2
AND, EOR, SUB, RSB, ADD, ADC, SBC, RSC, CMN, ORR, BIC, MVN, TST, TEQ, CMP	1	2	3
QADD, QSUB, QADD8, QADD16, QSUB8, QSUB16, SHADD8, SHADD16, SHSUB8, SHSUB16, UQADD8, UQADD16, UQSUB8, UQSUB16, UHADD8, UHADD16, UHSUB8, UHSUB16, QASX, QSAX, SHASX, SHSAX, UQASX, UQSAX, UHASX, UHSAX	2	-	-
QDADD, QDSUB, SSAT, USAT	3	-	-
PKHBT, PKHTB	1	2	-
SSAT16, USAT16, SADD8, SADD16, SSUB8, SSUB16, UADD8, UADD16, USUB8, USUB16, SASX, SSAX, UASX, USAX	1	-	-
SXTAB, SXTAB16, SXTAH, UXTAB, UXTAB16, UXTAH	3	-	-
SXTB, STXB16, SXTH, UXTB, UTXB16, UXTH	2	-	-
BFC, BFI, UBFX, SBFX	2	-	-
CLZ, MOVT, MOVW, RBIT, REV, REV16, REVSH, MRS	1	-	-
MSR not modifying mode or control bits See <i>Serializing instructions</i> on page B-9.	1	-	-

B.3 Load and store instructions

Load and store instructions are classed as:

- single load and store instructions such as LDR instructions
- load and store multiple instructions such as LDM instructions.

For load multiple and store multiple instructions, the number of registers in the register list usually determines the number of cycles required to execute a load or store instruction.

The Cortex-A9 processor has special paths that immediately forward data from a load instruction to a subsequent data processing instruction in the execution units.

This path is used when the following conditions are met:

- the data-processing instruction is one of: SUB, RSB, ADD, ADC, SBC, RSC, CMN, MVN, or CMP
- the forwarded source register is not part of a shift operation.

Table B-2 shows cycle timing for single load and store operations. The result latency is the latency of the first loaded register.

Table B-2 Single load and store operation cycle timings

Instruction cycles	AGU cycles	Result latency	
		Fast forward cases	other cases
LDR ,[reg] LDR ,[reg imm] LDR ,[reg reg] LDR ,[reg reg LSL #2]	1	2	3
LDR ,[reg reg LSL reg] LDR ,[reg reg LSR reg] LDR ,[reg reg ASR reg] LDR ,[reg reg ROR reg] LDR ,[reg reg,RRX]	1	3	4
LDRB ,[reg] LDRB ,[reg imm] LDRB ,[reg reg] LDRB ,[reg reg LSL #2] LDRH ,[reg] LDRH ,[reg imm] LDRH ,[reg reg] LDRH ,[reg reg LSL #2]	2	3	4
LDRB ,[reg reg LSL reg] LDRB ,[reg reg ASR reg] LDRB ,[reg reg LSL reg] LDRB ,[reg reg ASR reg] LDRH ,[reg reg LSL reg] LDRH ,[reg reg ASR reg] LDRH ,[reg reg LSL reg] LDRH ,[reg reg ASR reg]	2	4	5

The Cortex-A9 processor can load or store two 32-bit registers in each cycle. However, to access 64 bits, the address must be 64-bit aligned.

This scheduling is done in the *Address Generation Unit* (AGU). The number of cycles required by the AGU to process the load multiple or store multiple operations depends on the length of the register list and the 64-bit alignment of the address. The resulting latency is the latency of the first loaded register. Table B-3 shows the cycle timings for load multiple operations.

Table B-3 Load multiple operations cycle timings

Instruction	AGU cycles to process the instruction		Resulting latency	
	Address aligned on a 64-bit boundary		Fast forward case	Other cases
	Yes	No		
LDM ,{1 register}	1	1	2	3
LDM ,{2 registers} LDRD RFE	1	2	2	3
LDM ,{3 registers}	2	2	2	3
LDM ,{4 registers}	2	3	2	3
LDM ,{5 registers}	3	3	2	3
LDM ,{6 registers}	3	4	2	3
LDM ,{7 registers}	4	4	2	3
LDM ,{8 registers}	4	5	2	3
LDM ,{9 registers}	5	5	2	3
LDM ,{10 registers}	5	6	2	3
LDM ,{11 registers}	6	6	2	3
LDM ,{12 registers}	6	7	2	3
LDM ,{13 registers}	7	7	2	3
LDM ,{14 registers}	7	8	2	3
LDM ,{15 registers}	8	8	2	3
LDM ,{16 registers}	8	9	2	3

Table B-4 shows the cycle timings of store multiple operations.

Table B-4 Store multiple operations cycle timings

Instruction	AGU cycles	
	Aligned on a 64-bit boundary	
	Yes	No
STM, {1 register}	1	1
STM, {2 registers} STRD SRS	1	2
STM, {3 registers}	2	2
STM, {4 registers}	2	3
STM, {5 registers}	3	3
STM, {6 registers}	3	4
STM, {7 registers}	4	4
STM, {8 registers}	4	5
STM, {9 registers}	5	5
STM, {10 registers}	5	6
STM, {11 registers}	6	6
STM, {12 registers}	6	7
STM, {13 registers}	7	7
STM, {14 registers}	7	8
STM, {15 registers}	8	8
STM, {16 registers}	8	9

B.4 Multiplication instructions

Table B-4 on page B-6 shows the cycle timings for multiplication instructions.

Table B-5 Multiplication instruction cycle timings

Instruction	Cycles	Result latency
MUL(S), MLA(S)	2	4
SMULL(S), UMULL(S), SMLAL(S), UMLAL(S)	3	4 for the first written register 5 for the second written register
SMULxy, SMLAxy, SMULWy, SMLAWy	1	3
SMLALxy	2	3 for the first written register 4 for the second written register
SMUAD, SMUADX, SMLAD, SMLADX, SMUSD, SMUSDx, SMLSD, SMLSDx	1	3
SMMUL, SMMULR, SMMLA, SMMLAR, SMMLS, SMMLSR	2	4
SMLALD, SMLALDX, SMLSLD, SMLDLDX	2	3 for the first written register 4 for the second written register
UMAAL	3	4 for the first written register 5 for the second written register

B.5 Branch instructions

Branch instructions have different timing characteristics:

- Branch instructions to immediate locations do not consume execution unit cycles.
- Data-processing instructions to the PC register are processed in the execution units as standard instructions. See *Data-processing instructions* on page B-3.
- Load instructions to the PC register are processed in the execution units as standard instructions. See *Load and store instructions* on page B-4.

Also, see *About the L1 instruction side memory system* on page 7-5 for some information on dynamic branch prediction.

B.6 Serializing instructions

Out of order execution is not always possible. Some instructions are serializing. Serializing instructions force the processor to complete all modifications to flags and general-purpose registers by previous instructions before the next instruction is executed.

This section lists timings for serializing instructions

B.6.1 Serializing instructions

The following exception entry instructions are serializing:

- SVC
- SMC
- BKPT
- instructions that take the prefetch abort handler.
- instructions that take the Undefined Instruction exception handler,

The following instructions that modify mode or program control are serializing:

- MSR CPSR when they modify control or mode bits
- Data processing to PC with the S bit set (for example, `MOVS pc, r14`)
- `LDM pc ^.`
- CPS
- SETEND
- RFE.

The following instructions are serializing:

- all MCR to cp14 or cp15 except ISB and DMB.
- MRC p14 for debug registers
- WFE, WFI, SEV
- CLREX
- DSB.

In the r1p0 implementation DMB waits for all previous LDR/STR instructions to finish, not for all instructions to finish.

The following instruction, which modifies the SPSR, is serializing:

- MSR SPSR.

Appendix C

Revisions

This appendix describes the technical changes between released issues of this book.s

Table C-1 Issue A

Change	Location
First release	-

Table C-2 Differences between issue A and issue B

Change	Location
Clarified Load/Store Unit and address generation	Figure 1-1 on page 1-2.
Changed fast loop mode to small loop mode	<ul style="list-style-type: none"> • Figure 1-1 on page 1-2 • Small loop mode on page 1-3 • <i>Instruction cache features</i> on page 7-2 • About power consumption control on page 12-6.
Changed branch prediction to dynamic branch prediction	<ul style="list-style-type: none"> • <i>Features</i> on page 1-6 • <i>About the L1 instruction side memory system</i> on page 7-5 • <i>Branch instructions</i> on page B-8.
Changed L1 cache coherency to L1 data cache coherency	<i>Cortex-A9 variants</i> on page 1-4.
Corrected Processor Feature Register 0 reset value	Table 4-29 on page 4-46.

Table C-2 Differences between issue A and issue B (continued)

Change	Location
Made PMSWINC descriptions consistent	<ul style="list-style-type: none"> • Table 4-29 on page 4-46 • <i>Software Increment Register</i> on page 4-100.
Updated MIDR bits[3:0] from 0 to 1	Table 4-1 on page 4-5.
Corrected ID_MMFR3 [23:20] bit value to 0x1	Table 4-42 on page 4-50.
Corrected AFE bit description	Table 4-51 on page 4-62.
Corrected Auxiliary Control Register bit field	<ul style="list-style-type: none"> • Table 4-52 on page 4-66 • Figure 4-36 on page 4-66.
Corrected S parameter values	Set/Way format on page 4-83.
Made descriptions of bits[11], [10], and [8] consistent with table	Figure 4-41 on page 4-87.
Corrected description of event 0x68 , architecturally removed.	Table 4-80 on page 4-123.
Corrected TLB lockdown entries number from 8 to 4	c10, TLB Lockdown Register on page 4-134.
Corrected A, I, and F bit descriptions	c12, Interrupt Status Register on page 4-147.
Changed number of micro TLB entries from 8 to 32	<i>Micro TLB</i> on page 6-4.
Removed repeated information about cache types	<i>Micro TLB</i> on page 6-4.
Amended IRGN bits description from TTBCR to TTBR0/TTRBR1	<i>Main TLB</i> on page 6-4.
Added note about invalidating the caches and BTAC before use	<i>About the L1 memory system</i> on page 7-2.
Added parity support scheme information section	<i>Parity error support</i> on page 7-11.
Listed and described L2 master interfaces, M0 and M1	<i>About the Cortex-A9 L2 interface</i> on page 8-2.
Added Cross reference to DBSCR external description . Extended Footnote to include reference to the DBSCR external view	Table 10-1 on page 10-5.
Corrected DBGDSCR description with the addition of internal and external view descriptions.	<i>CPI4 c1, Debug Status and Control Register (DBGDSCR)</i> on page 8-9.
Re-ordered and extended MOE bits descriptions	Table 8-2 on page 8-10.

Table C-2 Differences between issue A and issue B (continued)

Change	Location
Added more cross-references from Table 10-1	<ul style="list-style-type: none"> • <i>Debug State Cache Control Register (DBGDSCCR)</i> on page 8-8 • <i>CP14 c1, Debug Status and Control Register (DBGDSCR)</i> on page 8-9 • <i>Device Power-down and Reset Status Register (DBGPRSR)</i> on page 8-27 • <i>Integration Mode Control Register (DBGITCTRL)</i> on page 8-45 • <i>Claim Tag Clear Register (DBGCLAIMCLR)</i> on page 8-47 • <i>Lock Access Register (DBGLAR)</i> on page 8-48 • <i>Lock Status Register (DBGLSR)</i> on page 8-49 • <i>Authentication Status Register (DBGAUTHSTATUS)</i> on page 8-49 • <i>Device Type Register (DBGDEVTYPE)</i> on page 8-50.
Corrected Table 10-1 footnotes	Table 10-1 on page 10-5.
Corrected byte address field entries	Table 10-8 on page 10-11.
Corrected interrupt signal descriptions	Table A-3 on page A-4.
Extended AXI USER descriptions	<ul style="list-style-type: none"> • Table A-8 on page A-8 • Table A-11 on page A-10 • Table A-14 on page A-12.

Table C-3 Differences between issue B and issue C

Change	Location
Removed 2.8.1 LE and BE-8 accesses on a 64-bit wide bus .	-
Removed Chapter 4 Unaligned and Mixed-Endian Data Access Support .	-
Removed the power management signal BISTSCLAMP .	-
Added dynamic high level clock gating .	<i>Dynamic high level clock gating</i> on page 2-9
Updated TLB information .	Table 1-1 on page 1-10, Table 4-10 on page 4-15, Table 4-37 on page 4-44
Shortened ID_MMF3[15:12] description .	<i>Memory Model Features Register 3</i> on page 4-49
Updated ACTLR to include reference to PL310 optimizations.	<i>Auxiliary Control Register</i> on page 4-64
Added information about a second replacement strategy. Selection done by SCTLRR.RR bit.	<i>System Control Register</i> on page 4-15
Extended event information .	<i>Cortex-A9 specific events</i> on page 4-32

Table C-3 Differences between issue B and issue C (continued)

Change	Location
Added DEFLAGS[6:0]	<i>DEFLAGS[6:0]</i> on page 4-37, <i>Performance monitoring signals</i> on page A-14
Added Power Control Register description	<i>Power Control Register</i> on page 4-63
Added PL310 optimizations to L2 memory interface description	<i>Optimized accesses to the L2 memory interface</i> on page 8-7
Added watchpoint address masking	<i>Watchpoint Control Registers</i> on page 10-11
Added debug request restart diagram.	<i>Effects of resets on debug registers</i> on page 10-3
Added CPUCLKOFF information	Table A-4 on page A-5, <i>Unregistered signals</i> on page B-3
Added DECLKOFF information	Table A-4 on page A-5, <i>Unregistered signals</i> on page B-3
Added MAXCLKLATENCY[2:0] information	<i>Configuration signals</i> on page A-5
Extended PMUEVENT bus description	<i>Performance monitoring signals</i> on page A-14
Added PMUSECURE and PMUPRIV	<i>Performance monitoring signals</i> on page A-14
Updated description of serializing behavior of DMB	<i>Serializing instructions</i> on page B-9

Table C-4 Differences between issue C and issue D

Change	Location
Included <i>Preload Engine (PE)</i> in block diagram	Figure 1-1 on page 1-2
Amended interrupt signals	
Clarified Data Engine options	<i>Data Engine</i> on page 1-2
Clarified system design components	<i>System design components</i> on page 1-3
Clarified <i>Compliance</i>	<i>Compliance</i> on page 1-5
Added PE to features	<i>Features</i> on page 1-6
Included PE and PE FIFO size in configurable options	<i>Configurable options</i> on page 1-8
Clarified NEON SIMD and FPU options	Table 1-1 on page 1-8
Added <i>Test Features</i> section	<i>Test features</i> on page 1-9
2.1.3 PTM interface reworded	<i>Performance monitoring</i> on page 2-3
2.1.5 Virtualization of interrupts added	<i>Virtualization of interrupts</i> on page 2-3
Included NEON SIMD clock gating in power control description	<i>Power Control Register</i> on page 2-9

Table C-4 Differences between issue C and issue D (continued)

Change	Location
Replaced nDERESET with nNEONRESET	<i>Reset modes</i> on page 2-7
Added nWDRESET	
Added nPERIPHRESET	
Changed voltage domain boundaries and description	Figure 2-4 on page 2-14
2.5.4 Date Engine logic reset replaced	<i>MPE SIMD logic reset</i> on page 2-8
Cortex-A9 input signals DECLAMP removed, level shifters reference removed	<i>Communication to the power management controller</i> on page 2-13
Table 3-1 J and T bit encoding removed	-
The Jazelle extension on page 3-3 moved	<i>The Jazelle Extension</i> on page 3-7
NEON technology on page 3-4 renamed and rewritten	<i>Advanced SIMD architecture</i> on page 3-4
3.4 Processor operating states removed	-
3.5 Data types removed	-
Multiprocessing Extensions section added	<i>Multiprocessing Extensions</i> on page 3-6
3.6 Memory formats renamed and moved	<i>Memory model</i> on page 3-8
3.8 Security extensions overview renamed and moved	<i>Security Extensions architecture</i> on page 3-5
Removed content, tables and figures from 4.1 that duplicates <i>ARM Architecture Reference Manual</i> material	<i>About system control</i> on page 4-2
4.2 Duplicates of <i>ARM Architecture Reference Manual</i> material removed, section renamed	<i>Register summary</i> on page 4-3
4.3 Duplicates of <i>ARM Architecture Reference Manual</i> material removed, section renamed	<i>Register descriptions</i> on page 4-8
Footnote e removed	Table 4-8 on page 4-15
Preload Engine registers added	<i>CP15 c11 register summary</i> on page 4-30
-	<i>PLE ID Register</i> on page 4-30
-	<i>PLE Activity Status Register</i> on page 4-31
-	<i>PLE FIFO Status Register</i> on page 4-32
-	<i>Preload Engine User Accessibility Register</i> on page 4-32
-	<i>Preload Engine Parameters Control Register</i> on page 4-33
4.4 CP14 Jazelle registers and 4.5 CP14 Jazelle register descriptions in a new chapter	Chapter 5 <i>Jazelle DBX registers</i>
Chapter 5 Memory Management Unit, 5.6 MMU software-accessible registers section removed	-
Level 1 Memory System chapter, Cortex-A9 cache policies section removed	-

Table C-4 Differences between issue C and issue D (continued)

Change	Location
Prefetch hint to the L2 memory interface, description re-written and extended	<i>Prefetch hint to the L2 memory interface</i> on page 8-7
Clarifications of BRESP and cache controller behavior	<i>Early BRESP</i> on page 8-7
Write full line of zeros, signal name corrected to AWUSERM0[7]	<i>Write full line of zeros</i> on page 8-8
Speculative coherent requests section added	<i>Speculative coherent requests</i> on page 8-8
Removed sentence about tying unused bits of PARITYFAIL HIGH	<i>Parity error support</i> on page 7-11
Added PE description	Chapter 9 <i>Preload Engine</i>
Added PMU description	Chapter 11 <i>Performance Monitoring Unit</i>
<i>Debug chapter</i> ; About debug systems removed	-
<i>Debug chapter</i> ; Debugging modes removed	-
Duplicates of <i>ARM Architecture Reference Manual</i> material removed	-
External debug interface, description of PADDRDBG[12:0] added	<i>External debug interface</i> on page 10-16
Debug APB interface section added	<i>Debug APB interface</i> on page 10-18
Amended and extended signals descriptions, source destination column added	Appendix A <i>Signal Descriptions</i>
PMUEVENT[46] description corrected	Table A-17 on page A-14
PMUEVENT[47] description corrected	
Removed <i>AC Characteristics</i> Appendix	-

No differences between issue D and issue E.

Differences between issue D and issue F.

Table C-5 Differences between issue D and issue F

Change	Location
PL310 renamed L2C-310	Throughout the book
VFPv3 corrected to VFPv3 D-32	<i>Media Processing Engine</i> on page 1-2
Cortex-A9 FPU hardware description rewritten for clarity	<i>Floating-Point Unit</i> on page 1-2
SCU description extended	<i>Cortex-A9 variants</i> on page 1-4
Dynamic branch prediction description added	<i>Dynamic branch prediction</i> on page 2-3
Final paragraph removed	<i>Energy efficiency features</i> on page 2-10
WFI/WFE corrected to Standby	Table 2-2 on page 2-10
Renamed and rewritten for clarity	<i>Standby modes</i> on page 2-11
Dormant mode clamping information removed	<i>Dormant mode</i> on page 2-12
IEM support renamed and rewritten	<i>Power domains</i> on page 2-13
Repeated material removed	<i>About the programmers model</i> on page 3-2

Table C-5 Differences between issue D and issue F (continued)

Change	Location
Debug register description corrected	Table 4-1 on page 4-3
Main ID Register values for r2p1 and r2p2 added	Table 4-2 on page 4-8
Debug register name corrected	Table 4-2 on page 4-8
Descriptions clarified and footnote added.	Table 4-4 on page 4-11
Purpose description extended	<i>Cache Size Identification Register</i> on page 4-11
System Control Register value corrected. Footnotes amended.	Table 4-8 on page 4-15
Bit[17] function corrected	Table 4-9 on page 4-16
Footnote d corrected	Table 4-33 on page 4-36
Purpose description extended	<i>Power Control Register</i> on page 4-36
Configurations description corrected	<i>Configuration Base Address Register</i> on page 4-38
Chapter renamed	Chapter 5 <i>Jazelle DBX registers</i>
6.1 application specific corrected to address space specific	<i>About the MMU</i> on page 6-2
Unified Main TLB description clarified	<i>Memory Management Unit</i> on page 6-2
Duplicate information about page sizes removed	
ASID description corrected and extended. Cross-reference added.	
TLB match process duplicate information about page sizes removed	<i>TLB match process</i> on page 6-4
Synchronous and asynchronous aborts incorrect cross-reference removed	<i>Synchronous and asynchronous aborts</i> on page 6-8
Cache features cross-reference corrected	<i>Cache features</i> on page 7-2
Implementation information removed	
Return stack predictions ARM or Thumb state replaced by instruction state	<i>Return stack predictions</i> on page 7-7
DSB section added	<i>About DSB</i> on page 7-9
AXI master 0 interface attributes corrections to values	Table 8-1 on page 8-2
Debug chapter moved to before PMU chapter	
Figure redrawn	Figure 10-1 on page 10-4
Corrections to bit format	Table 10-1 on page 10-5
Footnote about CLUSTERID values added	Table 10-10 on page 10-13
Value column added	Table 10-11 on page 10-14
DBGCPUDONE description extended	<i>DBGCPUDONE</i> on page 10-19
PMU management registers section added	<i>PMU management registers</i> on page 11-3
Signal descriptions extended.	<i>Configuration signals</i> on page A-5

Table C-5 Differences between issue D and issue F (continued)

Change	Location
Signal descriptions extended, information repeated from AXI removed	Table A-8 on page A-8
AWBURSTM0[1:0]	
AWLENM0[3:0]	
AWLOCKM0[1:0]	
Signal descriptions extended, information repeated from AXI removed	Table A-11 on page A-10
ARLENM0[3:0]	
ARLOCKM0[1:0]	
Title changed	<i>AXI Master1 signals instruction accesses</i> on page A-11
Information repeated from AXI removed	Table A-14 on page A-12
ARLENM1[3:0]	
PMUEVENT[46] and PMUEVENT[47] corrected	Table A-17 on page A-14
Introduction reduced. Note about DSB behavior added.	<i>Serializing instructions</i> on page B-9

Glossary

This glossary describes some of the terms used in ARM manuals. Where terms can have several meanings, the meaning presented here is intended.

Abort A mechanism that indicates to a core that the value associated with a memory access is invalid. An abort can be caused by the external or internal memory system as a result of attempting to access invalid instruction or data memory. An abort is classified as either a Prefetch or Data Abort, and an internal or External Abort.

See also Data Abort, External Abort and Prefetch Abort.

Abort model An abort model is the defined behavior of an ARM processor in response to a Data Abort exception. Different abort models behave differently with regard to load and store instructions that specify base register write-back.

Addressing modes A mechanism, shared by many different instructions, for generating values used by the instructions. For four of the ARM addressing modes, the values generated are memory addresses (the traditional role of an addressing mode). A fifth addressing mode generates values to be used as operands by data-processing instructions.

Advanced eXtensible Interface (AXI)

A bus protocol that supports separate address/control and data phases, unaligned data transfers using byte strobes, burst-based transactions with only start address issued, separate read and write data channels to enable low-cost DMA, ability to issue multiple outstanding addresses, out-of-order transaction completion, and easy addition of register stages to provide timing closure. The AXI protocol also includes optional extensions to cover signaling for low-power operation.

AXI is targeted at high performance, high clock frequency system designs and includes a number of features that make it very suitable for high speed sub-micron interconnect.

Advanced High-performance Bus (AHB)

A bus protocol with a fixed pipeline between address/control and data phases. It only supports a subset of the functionality provided by the AMBA AXI protocol. The full AMBA AHB protocol specification includes a number of features that are not commonly required for master and slave IP developments and ARM Limited recommends only a subset of the protocol is usually used. This subset is defined as the AMBA AHB-Lite protocol.

See also Advanced Microcontroller Bus Architecture and AHB-Lite.

Advanced Microcontroller Bus Architecture (AMBA)

A family of protocol specifications that describe a strategy for the interconnect. AMBA is the ARM open standard for on-chip buses. It is an on-chip bus specification that describes a strategy for the interconnection and management of functional blocks that make up a *System-on-Chip* (SoC). It aids in the development of embedded processors with one or more CPUs or signal processors and multiple peripherals. AMBA complements a reusable design methodology by defining a common backbone for SoC modules.

Advanced Peripheral Bus (APB)

A simpler bus protocol than AXI and AHB. It is designed for use with ancillary or general-purpose peripherals such as timers, interrupt controllers, UARTs, and I/O ports. Connection to the main system bus is through a system-to-peripheral bus bridge that helps to reduce system power consumption.

AHB

See Advanced High-performance Bus.

AHB Access Port (AHB-AP)

An optional component of the DAP that provides an AHB interface to a SoC.

AHB-AP

See AHB Access Port.

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 by using an AMBA AXI protocol interface.

Aligned

A data item stored at an address that is divisible by the number of bytes that defines the data size is said to be aligned. Aligned words and halfwords have addresses that are divisible by four and two respectively. The terms word-aligned and halfword-aligned therefore stipulate addresses that are divisible by four and two respectively.

AMBA

See Advanced Microcontroller Bus Architecture.

Advanced Trace Bus (ATB)

A bus used by trace devices to share CoreSight capture resources.

APB

See Advanced Peripheral Bus.

Architecture

The organization of hardware and/or software that characterizes a processor and its attached components, and enables devices with similar characteristics to be grouped together when describing their behavior, for example, Harvard architecture, instruction set architecture, ARMv6 architecture.

ARM instruction

A word that specifies an operation for an ARM processor to perform. ARM instructions must be word-aligned.

ARM state

A processor that is executing ARM (32-bit) word-aligned instructions is operating in ARM state.

ATB

See Advanced Trace Bus.

ATB bridge

A synchronous ATB bridge provides a register slice to facilitate timing closure through the addition of a pipeline stage. It also provides a unidirectional link between two synchronous ATB domains.

An asynchronous ATB bridge provides a unidirectional link between two ATB domains with asynchronous clocks. It is intended to support connection of components with ATB ports residing in different clock domains.

ATPG

See Automatic Test Pattern Generation.

Automatic Test Pattern Generation (ATPG)

The process of automatically generating manufacturing test vectors for an ASIC design, using a specialized software tool.


AXI

See Advanced eXtensible Interface.

AXI channel order and interfaces

The block diagram shows:

- the order that AXI channel signals are described in
- the master and slave interface conventions for AXI components.

**AXI terminology**

The following AXI terms are general. They apply to both masters and slaves:

Active read transaction

A transaction where the read address has transferred, but the last read data has not yet transferred.

Active transfer

A transfer where the **xVALID**¹ handshake has asserted, but **xREADY** has not yet asserted.

Active write transaction

A transaction where the write address or leading write data has transferred, but the write response has not yet transferred.

Completed transfer

A transfer where the **xVALID/xREADY** handshake is complete.

Payload

The non-handshake signals in a transfer.

Transaction

An entire burst of transfers, comprising an address, one or more data transfers and a response transfer (writes only).

1. The letter x in the signal name denotes an AXI channel as follows:

AW	Write address channel.
W	Write data channel.
B	Write response channel.
AR	Read address channel.
R	Read data channel.

Transmit	An initiator driving the payload and asserting the relevant xVALID signal.
Transfer	A single exchange of information. That is, with one xVALID/xREADY handshake.

The following AXI terms are master interface attributes. To obtain optimum performance, they must be specified for all components with an AXI master interface:

Combined issuing capability

The maximum number of active transactions that a master interface can generate. This is specified instead of write or read issuing capability for master interfaces that use a combined storage for active write and read transactions.

Read ID capability

The maximum number of different **ARID** values that a master 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 a master interface can generate.

Write ID capability

The maximum number of different **AWID** values that a master 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 that the master interface is capable of transmitting data for. This is counted from the earliest transaction.

Write issuing capability

The maximum number of active write transactions that a master interface can generate.

The following AXI terms are slave interface attributes. To obtain optimum performance, they must be specified for all components with an AXI slave interface

Combined acceptance capability

The maximum number of active transactions that a slave interface can accept. This is specified instead of write or read acceptance capability for slave interfaces that use a combined storage for active write and read transactions.

Read acceptance capability

The maximum number of active read transactions that a slave interface can accept.

Read data reordering depth

The number of active read transactions that a slave interface can transmit data for. This is counted from the earliest transaction.

	<p>Write acceptance capability</p> <p>The maximum number of active write transactions that a slave interface can accept.</p>
	<p>Write interleave depth</p> <p>The number of active write transactions that the slave interface can receive data for. This is counted from the earliest transaction.</p>
Banked registers	Those physical registers whose use is defined by the current processor mode. The banked registers are r8 to r14.
Base register	A register specified by a load or store instruction that is used to hold the base value for the instruction's address calculation. 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	Updating the contents of the base register used in an instruction target address calculation so that the modified address is changed to the next higher or lower sequential address in memory. This means that it is not necessary to fetch the target address for successive instruction transfers and enables faster burst accesses to sequential memory.
Beat	Alternative word for an individual transfer within a burst. For example, an INCR4 burst comprises four beats. <i>See also</i> Burst.
BE-8	Big-endian view of memory in a byte-invariant system. <i>See also</i> BE-32, LE, Byte-invariant and Word-invariant.
BE-32	Big-endian view of memory in a word-invariant system. <i>See also</i> BE-8, LE, Byte-invariant and Word-invariant.
Big-endian	Byte ordering scheme where bytes of decreasing significance in a data word are stored at increasing addresses in memory. <i>See also</i> Little-endian and Endianness.
Big-endian memory	Memory where: <ul style="list-style-type: none"> • a byte or halfword at a word-aligned address is the most significant byte or halfword within the word at that address • a byte at a halfword-aligned address is the most significant byte within the halfword at that address. <i>See also</i> Little-endian memory.
Block address	An address that comprises a tag, an index, and a word field. The tag bits identify the way that contains the matching cache entry for a cache hit. The index bits identify the set being addressed. The word field contains the word address that can be used to identify specific words, halfwords, or bytes within the cache entry. <i>See also</i> Cache terminology diagram on the last page of this glossary.

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. Processors can contain several shift registers to enable you to access selected parts of the device.
Branch prediction	The process of predicting if conditional branches are to be taken or not in pipelined processors. Successfully predicting if branches are to be taken enables the processor to prefetch the instructions following a branch before the condition is fully resolved. Branch prediction can be done in software or by using custom hardware. Branch prediction techniques are categorized as static, where the prediction decision is decided before run time, and dynamic, where the prediction decision can change during program execution.
Breakpoint	A breakpoint is a mechanism provided by debuggers to identify an instruction that program execution is to be halted at. Breakpoints are inserted by the programmer to enable inspection of register contents, memory locations, variable values at fixed points in the program execution to test that the program is operating correctly. Breakpoints are removed after the program is successfully tested. <i>See also</i> Watchpoint.
Burst	A group of transfers to consecutive addresses. Because the addresses are consecutive, there is no requirement to supply an address for any of the transfers after the first one. This increases the speed that the group of transfers can occur at. Bursts over AHB buses are controlled using the HBURST signals to specify if transfers are single, four-beat, eight-beat, or 16-beat bursts, and to specify how the addresses are incremented. <i>See also</i> Beat.
Byte	An 8-bit data item.
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. Multi-word accesses are expected to be word-aligned. <i>See also</i> Word-invariant.
Byte lane strobe	An AHB signal, HBSTRB , that is used for unaligned or mixed-endian data accesses to determine the byte lanes that are active in a transfer. One bit of HBSTRB corresponds to eight bits of the data bus.
Cache	A block of on-chip or off-chip fast access memory locations, situated between the processor and main memory, used for storing and retrieving copies of often used instructions and/or data. This is done to greatly reduce the average speed of memory accesses and so to increase processor performance. <i>See also</i> Cache terminology diagram on the last page of this glossary.
Cache contention	When the number of frequently-used memory cache lines that use a particular cache set exceeds the set-associativity of the cache. In this case, main memory activity increases and performance decreases.
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	<p>The basic unit of storage in a cache. It is always a power of two words in size (usually four or eight words), and is required to be aligned to a suitable memory boundary.</p> <p><i>See also</i> Cache terminology diagram on the last page of this glossary.</p>
Cache line index	<p>The number associated with each cache line in a cache way. Within each cache way, the cache lines are numbered from 0 to (set associativity) -1.</p> <p><i>See also</i> Cache terminology diagram on the last page of this glossary.</p>
Cache lockdown	<p>To fix a line in cache memory so that it cannot be overwritten. Enables critical instructions and/or data to be loaded into the cache so that the cache lines containing them are not subsequently reallocated. This ensures that all subsequent accesses to the instructions/data concerned are cache hits, and therefore complete as quickly as possible.</p>
Cache miss	<p>A memory access that cannot be processed at high speed because the instruction/data it addresses is not in the cache and a main memory access is required.</p>
Cache set	<p>A cache set is a group of cache lines (or blocks). A set contains all the ways that can be addressed with the same index. The number of cache sets is always a power of two.</p> <p><i>See also</i> Cache terminology diagram on the last page of this glossary.</p>
Cache way	<p>A group of cache lines (or blocks). It is 2 to the power of the number of index bits in size.</p> <p><i>See also</i> Cache terminology diagram on the last page of this glossary.</p>
Clean	<p>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.</p> <p><i>See also</i> Dirty.</p>
Clock gating	<p>Gating a clock signal for a macrocell with a control signal and using the modified clock that results to control the operating state of the macrocell.</p>
Clocks Per Instruction (CPI)	<p><i>See</i> Cycles Per Instruction (CPI).</p>
Coherency	<p><i>See</i> Memory coherency.</p>
Cold reset	<p>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 enable the system to be used again. In other cases, only a warm reset is required.</p> <p><i>See also</i> Warm reset.</p>
Communications channel	<p>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 Debug Comms Channel. In an ARMv6 compliant core, the communications channel includes the Data Transfer Register, some bits of the Data Status and Control Register, and the external debug interface controller, such as the DBGJTAG controller in the case of the JTAG interface.</p>
Condition field	<p>A four-bit field in an instruction that specifies a condition under which the instruction can execute.</p>
Conditional execution	<p>If the condition code flags indicate that the corresponding condition is true when the instruction starts executing, it executes normally. Otherwise, the instruction does nothing.</p>

Context	The environment that each process operates in for a multitasking operating system. In ARM processors, this is limited to mean the Physical Address range that it can access in memory and the associated memory access permissions.
Control bits	The bottom eight bits of a Program Status Register (PSR). The control bits change when an exception arises and can be altered by software only when the processor is in a privileged mode.
Coprocessor	A processor that supplements the main processor. It carries out additional functions that the main processor cannot perform. Usually used for floating-point math calculations, signal processing, or memory management.
Core reset	<i>See</i> Warm reset.
CPI	<i>See</i> Cycles per instruction.
CPSR	<i>See</i> Current Program Status Register
Current Program Status Register (CPSR)	The register that holds the current operating processor status.
Cycles Per instruction (CPI)	Cycles per instruction (or clocks per instruction) is a measure of the number of computer instructions that can be performed in one clock cycle. This figure of merit can be used to compare the performance of different CPUs that implement the same instruction set against each other. The lower the value, the better the performance.
Data Abort	An indication from a memory system to a core that it must halt execution of an attempted illegal memory access. A Data Abort is attempting to access invalid data memory. <i>See also</i> Abort, External Abort, and Prefetch Abort.
Data cache	A block of on-chip fast access memory locations, situated between the processor and main memory, used for storing and retrieving copies of often used data. This is done to greatly reduce the average speed of memory accesses and so to increase processor performance.
DBGTAP	<i>See</i> Debug Test Access Port.
Debug Access Port (DAP)	A TAP block that acts as an AMBA (AHB or AHB-Lite) master for access to a system bus. The DAP is the term used to encompass a set of modular blocks that support system wide debug. The DAP is a modular component, intended to be extendable to support optional access to multiple systems such as memory mapped AHB and CoreSight APB through a single debug interface.
Debugger	A debugging system that includes a program, used to detect, locate, and correct software faults, together with custom hardware that supports software debugging.
Direct-mapped cache	A one-way set-associative cache. Each cache set consists of a single cache line, so cache lookup selects and checks a single cache line.
Dirty	A cache line in a write-back cache that has been modified while it is in the cache is said to be dirty. A cache line is marked as dirty by setting the dirty bit. If a cache line is dirty, it must be written to memory on a cache miss because the next level of memory contains data that has not been updated. The process of writing dirty data to main memory is called cache cleaning. <i>See also</i> Clean.
DNM	<i>See</i> Do Not Modify.

Do Not Modify (DNM)	In Do Not Modify fields, the value must not be altered by software. DNM fields read as Unpredictable values, and must only be written with the same value read from the same field on the same processor. DNM fields are sometimes followed by RAZ or RAO in parentheses to show the way the bits must read for future compatibility, but programmers must not rely on this behavior.
Doubleword	A 64-bit data item. The contents are taken as being an unsigned integer unless otherwise stated.
Doubleword-aligned	A data item having a memory address that is divisible by eight.
EmbeddedICE logic	An on-chip logic block that provides TAP-based debug support for ARM processor cores. It is accessed through the TAP controller on the ARM core using the JTAG interface.
EmbeddedICE-RT	The JTAG-based hardware provided by debuggable ARM processors to aid debugging in real-time.
Endianness	Byte ordering. The scheme that determines the order that successive bytes of a data word are stored in, in memory. An aspect of the system's memory mapping. <i>See also</i> Little-endian and Big-endian
Exception	A fault or error event that is considered serious enough to require that program execution is interrupted. Examples include attempting to perform an invalid memory access, external interrupts, and undefined instructions. When an exception occurs, normal program flow is interrupted and execution is resumed at the corresponding exception vector. This contains the first instruction of the interrupt handler to deal with the exception.
Exception service routine	<i>See</i> Interrupt handler.
Exception vector	<i>See</i> Interrupt vector.
Exponent	The component of a floating-point number that normally signifies the integer power to which two is raised in determining the value of the represented number.
External Abort	An indication from an external memory system to a core that it must halt execution of an attempted illegal memory access. An External Abort is caused by the external memory system as a result of attempting to access invalid memory. <i>See also</i> Abort, Data Abort and Prefetch Abort.
Flat address mapping	A system of organizing memory where each Physical Address contained within the memory space is the same as its corresponding Virtual Address.
Front of queue pointer	Pointer to the next entry to be written to in the write buffer.
Fully-associative cache	A cache that has only one cache set that consists of the entire cache. The number of cache entries is the same as the number of cache ways. <i>See also</i> Direct-mapped cache.
Halfword	A 16-bit data item.

Halting debug-mode	<p>One of two mutually exclusive debug modes. In Halting debug-mode a <i>debug event</i>, such as a breakpoint or watchpoint, causes the processor to enter a special Debug state. In Debug state the processor is controlled through the external debug interface. This interface also provides access to all processor state, coprocessor state, memory and input/output locations.</p> <p><i>See also</i> Monitor debug-mode.</p>
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.
Host	A computer that provides data and other services to another computer. Especially, a computer providing debugging services to a target being debugged.
IEEE 754 standard	<i>IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Std. 754-1985.</i> The standard that defines data types, correct operation, exception types and handling, and error bounds for floating-point systems. Most processors are built in compliance with the standard in either hardware or a combination of hardware and software.
IEM	<i>See</i> Intelligent Energy Manager.
IGN	<i>See</i> Ignore.
Ignore (IGN)	Must ignore memory writes.
Illegal instruction	An instruction that is architecturally Undefined.
Implementation-defined	Means that the behavior is not architecturally defined, but must be defined and documented by individual implementations.
Implementation-specific	Means that the behavior is not architecturally defined, and does not have to be documented by individual implementations. Used when there are a number of implementation options available and the option chosen does not affect software compatibility.
Imprecise tracing	<p>A filtering configuration where instruction or data tracing can start or finish earlier or later than expected. Most cases cause tracing to start or finish later than expected.</p> <p>For example, if TraceEnable is configured to use a counter so that tracing begins after the fourth write to a location in memory, the instruction that caused the fourth write is not traced, although subsequent instructions are. This is because the use of a counter in the TraceEnable configuration always results in imprecise tracing.</p>
Index	<i>See</i> Cache index.
Index register	A register specified in some load or store instructions. The value of this register is used as an offset to be added to or subtracted from the base register value to form the virtual address, which is sent to memory. Some addressing modes optionally enable the index register value to be shifted prior to the addition or subtraction.
Instruction cache	A block of on-chip fast access memory locations, situated between the processor and main memory, used for storing and retrieving copies of often used instructions. This is done to greatly reduce the average speed of memory accesses and so to increase processor performance.
Instruction cycle count	The number of cycles that an instruction occupies the Execute stage of the pipeline for.
Intelligent Energy Manager (IEM)	A technology that enables dynamic voltage scaling and clock frequency variation to be used to reduce power consumption in a device.

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	A program that control of the processor is passed to when an interrupt occurs.
Interrupt vector	One of a number of fixed addresses in low memory, or in high memory if high vectors are configured, that contains the first instruction of the corresponding interrupt handler.
Invalidate	To mark a cache line as being not valid by clearing the valid bit. This must be done whenever the line does not contain a valid cache entry. For example, after a cache flush all lines are invalid.
Joint Test Action Group (JTAG)	The name of the organization that developed standard IEEE 1149.1. This standard defines a boundary-scan architecture used for in-circuit testing of integrated circuit devices. It is commonly known by the initials JTAG.
JTAG	<i>See</i> Joint Test Action Group.
LE	Little endian view of memory in both byte-invariant and word-invariant systems. <i>See also</i> Byte-invariant, Word-invariant.
Line	<i>See</i> Cache line.
Little-endian	Byte ordering scheme where bytes of increasing significance in a data word are stored at increasing addresses in memory. <i>See also</i> Big-endian and Endianness.
Little-endian memory	Memory where: <ul style="list-style-type: none"> • a byte or halfword at a word-aligned address is the least significant byte or halfword within the word at that address • a byte at a halfword-aligned address is the least significant byte within the halfword at that address. <i>See also</i> Big-endian memory.
Load/store architecture	A processor architecture where data-processing operations only operate on register contents, not directly on memory contents.
Load Store Unit (LSU)	The part of a processor that handles load and store transfers.
LSU	<i>See</i> Load Store Unit.
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) plus application-specific logic.
Memory bank	One of two or more parallel divisions of interleaved memory, usually one word wide, that enable reads and writes of multiple words at a time, rather than single words. All memory banks are addressed simultaneously and a bank enable or chip select signal determines the bank that is accessed for each transfer. Accesses to sequential word addresses cause accesses to sequential banks. This enables the delays associated with accessing a bank to occur during the access to its adjacent bank, speeding up memory transfers.
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 there are multiple possible physical locations that are involved, such as a system that has main memory, a write buffer and a cache.

Memory Management Unit (MMU)

Hardware that controls caches and access permissions to blocks of memory, and translates virtual addresses to physical addresses.

Memory Protection Unit (MPU)

Hardware that controls access permissions to blocks of memory. Unlike an MMU, an MPU does not translate virtual addresses to physical addresses.

Microprocessor

See Processor.

Miss

See Cache miss.

MMU

See Memory Management Unit.

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, this enables vital system interrupts to continue to be serviced while normal program execution is suspended.

See also Halt mode.

MPU

See Memory Protection Unit.

VA

See Modified Virtual Address.

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.

Power-on reset

See Cold reset.

Prefetching

In pipelined processors, the process of fetching instructions from memory to fill up the pipeline before the preceding instructions have finished executing. Prefetching an instruction does not mean that the instruction must be executed.

Prefetch Abort

An indication from a memory system to a core that it must halt execution of an attempted illegal memory access. 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.

Processor

A processor is the circuitry in a computer system required to process data using the computer instructions. It is an abbreviation of microprocessor. A clock source, power supplies, and main memory are also required to create a minimum complete working computer system.

Physical Address (PA)

The MMU performs a translation on *Modified Virtual Addresses* (VA) to produce the *Physical Address* (PA) that is given to AXI to perform an external access. The PA is also stored in the data cache to avoid the necessity for address translation when data is cast out of the cache.

Read

Reads are defined as memory operations that have the semantics of a load. That is, the ARM instructions LDM, LDRD, LDC, LDR, LDRT, LDRSH, LDRH, LDRSB, LDRB, LDRBT, LDREX, RFE, STREX, SWP, and SWPB, and the Thumb instructions LDM, LDR, LDRSH, LDRH, LDRSB, LDRB, and POP. Java bytecodes that are accelerated by hardware can cause a number of reads to occur, according to the state of the Java stack and the implementation of the Java hardware acceleration.

RealView ICE

A system for debugging embedded processor cores using a JTAG interface.

Region

A partition of instruction or data memory space.

Remapping	Changing the address of physical memory or devices after the application has started executing. This is typically done to enable RAM to replace ROM when the initialization has been completed.
Reserved	A field in a control register or instruction format is reserved if the field is to be defined by the implementation, or produces Unpredictable results if the contents of the field are not zero. These fields are reserved for use in future extensions of the architecture or are implementation-specific. All reserved bits not used by the implementation must be written as 0 and read as 0.
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.
SBO	<i>See</i> Should Be One.
SBZ	<i>See</i> Should Be Zero.
SBZP	<i>See</i> Should Be Zero or Preserved.
Scan chain	A 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. Processors can contain several shift registers to enable you to access selected parts of the device.
SCREG	The currently selected scan chain number in an ARM TAP controller.
Set	<i>See</i> Cache set.
Set-associative cache	In a set-associative cache, lines can only be placed in the cache in locations that correspond to the modulo division of the memory address by the number of sets. If there are n ways in a cache, the cache is termed n -way set-associative. The set-associativity can be any number greater than or equal to 1 and is not restricted to being a power of two.
Should Be One (SBO)	Should be written as 1 (or all 1s for bit fields) by software. Writing a 0 produces Unpredictable results.
Should Be Zero (SBZ)	Should be written as 0 (or all 0s for bit fields) by software. Writing a 1 produces Unpredictable results.
Should Be Zero or Preserved (SBZP)	Should be written as 0 (or all 0s for bit fields) by software, or preserved by writing the same value back that has been previously read from the same field on the same processor.
SPSR	<i>See</i> Saved Program Status Register
Standard Delay Format (SDF)	The format of a file that contains timing information to the level of individual bits of buses and is used in SDF back-annotation. An SDF file can be generated in a number of ways, but most commonly from a delay calculator.
Synchronization primitive	The memory synchronization primitive instructions are those instructions that are used to ensure memory synchronization. That is, the LDREX, STREX, SWP, and SWPB instructions.

Tag	<p>The upper portion of a block address used to identify a cache line within a cache. The block address from the CPU is compared with each tag in a set in parallel to determine if the corresponding line is in the cache. If it is, it is said to be a cache hit and the line can be fetched from cache. If the block address does not correspond to any of the tags, it is said to be a cache miss and the line must be fetched from the next level of memory.</p> <p><i>See also</i> Cache terminology diagram on the last page of this glossary.</p>
TAP	<i>See</i> Test access port.
Test Access Port (TAP)	<p>The collection of four mandatory and one optional terminals that form the input/output and control interface to a JTAG boundary-scan architecture. The mandatory terminals are TDI, TDO, TMS, and TCK. The optional terminal is TRST. This signal is required in ARM cores because it is used to reset the debug logic.</p>
Thumb instruction	A halfword that specifies an operation for an ARM processor in Thumb state to perform. Thumb instructions must be halfword-aligned.
Thumb state	A processor that is executing Thumb (16-bit) halfword aligned instructions is operating in Thumb state.
TLB	<i>See</i> Translation Look-aside Buffer.
Translation Lookaside Buffer (TLB)	A cache of recently used page table entries that avoid the overhead of translation table walking on every memory access. Part of the Memory Management Unit.
Translation table	A table, held in memory, that contains data that defines the properties of memory areas of various fixed sizes.
Translation table walk	The process of doing a full translation table lookup. It is performed automatically by hardware.
Trap	An exceptional condition in a VFP coprocessor that has the respective exception enable bit set in the FPSCR register. The user trap handler is executed.
Undefined	Indicates an instruction that generates an Undefined instruction trap. <i>See the ARM Architecture Reference Manual</i> for more details on ARM exceptions.
UNP	<i>See</i> Unpredictable.
Unpredictable	For reads, the data returned when reading from this location is unpredictable. It can have any value. For writes, writing to this location causes unpredictable behavior, or an unpredictable change in device configuration. Unpredictable instructions must not halt or hang the processor, or any part of the system.
Unsupported values	Specific data values that are not processed by the VFP coprocessor hardware but bounced to the support code for completion. These data can include infinities, NaNs, subnormal values, and zeros. An implementation is free to select which of these values is supported in hardware fully or partially, or requires assistance from support code to complete the operation. Any exception resulting from processing unsupported data is trapped to user code if the corresponding exception enable bit for the exception is set.
VA	<i>See</i> Virtual Address.
Victim	A cache line, selected to be discarded to make room for a replacement cache line that is required as a result of a cache miss. The method used to select the victim for eviction is processor-specific. A victim is also known as a cast out.

Virtual Address (VA)	<p>The MMU uses its translation tables to translate a Virtual Address into a Physical Address. The processor executes code at the Virtual Address, possibly located elsewhere in physical memory.</p> <p><i>See also</i> Modified Virtual Address, and Physical Address.</p>
Warm reset	<p>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.</p>
Watchpoint	<p>A watchpoint is a mechanism provided by debuggers to halt program execution when the data contained by a particular memory address is changed. Watchpoints are inserted by the programmer to enable inspection of register contents, memory locations, and variable values when memory is written to test that the program is operating correctly. Watchpoints are removed after the program is successfully tested. <i>See also</i> Breakpoint.</p>
Way	<p><i>See</i> Cache way.</p>
WB	<p><i>See</i> Write-back.</p>
Word	<p>A 32-bit data item.</p>
Word-invariant	<p>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 that are given unaligned addresses is instruction-specific, and is in general not the expected behavior for an unaligned access. It is recommended 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 use only aligned word memory accesses.</p> <p><i>See also</i> Byte-invariant.</p>
Write	<p>Writes are defined as operations that have the semantics of a store. That is, the ARM instructions SRS, STM, STRD, STC, STRT, STRH, STRB, STRBT, STREX, SWP, and SWPB, and the Thumb instructions STM, STR, STRH, STRB, and PUSH. Java bytecodes that are accelerated by hardware can cause a number of writes to occur, according to the state of the Java stack and the implementation of the Java hardware acceleration.</p>
Write-back (WB)	<p>In a write-back cache, data is only written to main memory when it is forced out of the cache on line replacement following a cache miss. Otherwise, writes by the processor only update the cache. (Also known as copyback).</p>
Write buffer	<p>A block of high-speed memory, arranged as a FIFO buffer, between the data cache and main memory, whose purpose is to optimize stores to main memory.</p>
Write completion	<p>The memory system indicates to the processor that a write has been completed at a point in the transaction where the memory system is able to guarantee that the effect of the write is visible to all processors in the system. This is not the case if the write is associated with a memory synchronization primitive, or is to a Device or Strongly-ordered region. In these cases the memory system might only indicate completion of the write when the access has affected the state of the target, unless it is impossible to distinguish between having the effect of the write visible and having the state of target updated.</p>

This stricter requirement for some types of memory ensures that any side-effects of the memory access can be guaranteed by the processor to have taken place. You can use this to prevent the starting of a subsequent operation in the program order until the side-effects are visible.

Write-through (WT)

In a write-through cache, data is written to main memory at the same time as the cache is updated.

WT

See Write-through.

Cache terminology diagram

The diagram illustrates the following cache terminology:

- block address
- cache line
- cache set
- cache way
- index
- tag.

