



Arm® Neoverse™ N1

Software Optimization Guide

Non-Confidential

Copyright © 2018-2023 Arm Limited (or its affiliates).
All rights reserved.

Issue 4.0

PJDOC-466751330-9707

Arm® Neoverse™ N1

Software Optimization Guide

Copyright © 2018-2023 Arm Limited (or its affiliates). All rights reserved.

Confidential Proprietary Notice

This document is **NON-CONFIDENTIAL** and any use by you is subject to the terms of the agreement between you and Arm or the terms of the agreement between you and the party authorised by Arm to disclose this document to you.

This document is protected by copyright and other related rights and the practice or implementation of the information contained in this document may be protected by one or more patents or pending patent applications. No part of this document may be reproduced in any form by any means without the express prior written permission of Arm. **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.**

Your access to the information in this document is conditional upon your acceptance that you will not use or permit others to use the information: (i) for the purposes of determining whether implementations infringe any third party patents; (ii) for developing technology or products which avoid any of Arm's intellectual property; or (iii) as a reference for modifying existing patents or patent applications or creating any continuation, continuation in part, or extension of existing patents or patent applications; or (iv) for generating data for publication or disclosure to third parties, which compares the performance or functionality of the Arm technology described in this document with any other products created by you or a third party, without obtaining Arm's prior written consent.

THIS DOCUMENT IS PROVIDED "AS IS". ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, third party patents, copyrights, trade secrets, or other rights.

This document may include technical inaccuracies or typographical errors.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

This document consists solely of commercial items. You shall be responsible for ensuring that any use, duplication or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word "partner" in reference to Arm's customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

If any of the provisions contained in these terms conflict with any of the provisions of any click through or signed written agreement covering this document with Arm, then the click through or signed written agreement prevails over and supersedes the conflicting provisions of these terms. This

document may be translated into other languages for convenience, and you agree that if there is any conflict between the English version of this document and any translation, the terms of the English version of the Agreement shall prevail.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners. Please follow Arm's trademark usage guidelines at <http://www.arm.com/company/policies/trademarks>.

Copyright © 2018-2023 Arm Limited (or its affiliates). All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

LES-PRE-20348

Confidentiality Status

This document is Non-Confidential. This document may only be used and distributed 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

1 Introduction	6
1.1 Product revision status	6
1.2 Intended audience	6
1.3 Conventions	6
1.3.1 Glossary	6
1.3.2 Typographical conventions	7
1.4 Additional reading	8
1.5 Feedback	8
1.5.1 Feedback on this product	8
1.5.2 Feedback on content	8
2 About this document	9
2.1 Scope	9
2.2 Pipeline overview	9
3 Instruction characteristics	11
3.1 Instruction tables	11
3.2 Legend for reading the utilized pipelines	11
3.3 Branch instructions	11
3.4 Arithmetic and logical instructions	12
3.5 Move and shift instructions	15
3.6 Divide and multiply instructions	15
3.7 Saturating and parallel arithmetic instructions	18
3.8 Miscellaneous data-processing instructions	19
3.9 Load instructions	20
3.10 Store instructions	25
3.11 FP data processing instructions	28
3.12 FP miscellaneous instructions	30
3.13 FP load instructions	32
3.14 FP store instructions	34
3.15 ASIMD integer instructions	36
3.16 ASIMD floating-point instructions	42
3.17 ASIMD miscellaneous instructions	47

3.18 ASIMD load instructions	50
3.19 ASIMD store instructions	54
3.20 Cryptography extensions	57
3.21 CRC	58
4 Special considerations	59
4.1 Dispatch constraints	59
4.2 Dispatch stall	59
4.3 Optimizing general-purpose register spills and fills	59
4.4 Optimizing memory copy	59
4.5 Load/Store alignment	60
4.6 Store to Load Forwarding	61
4.7 AES encryption/decryption	61
4.8 Region based fast forwarding	61
4.9 Branch instruction alignment	63
4.10 FPCR self-synchronization	63
4.11 Special register access	63
4.12 Register forwarding hazards	65
4.13 IT blocks	66

1 Introduction

1.1 Product revision status

The *rm**pn* identifier indicates the revision status of the product described in this book, for example, r1p2, where:

- rm* Identifies the major revision of the product, for example, r1.
- pn* Identifies the minor revision or modification status of the product, for example, p2.

1.2 Intended audience

This document is for system designers, system integrators, and programmers who are designing or programming a System-on-Chip (SoC) that uses an Arm core.

1.3 Conventions

The following subsections describe conventions used in Arm documents.

1.3.1 Glossary

The Arm Glossary is a list of terms used in Arm documentation, together with definitions for those terms. The Arm Glossary does not contain terms that are industry standard unless the Arm meaning differs from the generally accepted meaning.




See the [Arm® Glossary](#) for more information.

1.3.1.1 Terms and Abbreviations

This document uses the following terms and abbreviations.

Term	Meaning
ALU	Arithmetic and Logical Unit
ASIMD	Advanced SIMD
VFP	Vector Floating Point
MOP	Macro-Operation
μOP	Micro-Operation

1.3.2 Typographical conventions

Convention	Use
<i>italic</i>	Introduces special terminology, denotes 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.
Monospace bold	Denotes language keywords when used outside example code.
<i>monospace italic</i>	Denotes arguments to monospace text where the argument is to be replaced by a specific value.
monospace <u>underline</u>	Denotes a permitted abbreviation for a command or option. You can enter the underlined text instead of the full command or option name.
<and>	Encloses replaceable terms for assembler syntax where they appear in code or code fragments. For example: <code>MRC p15, 0, <Rd>, <CRn>, <CRm>, <Opcode_2></code>
SMALL CAPITALS	Used in body text for a few terms that have specific technical meanings, that are defined in the Arm® Glossary. For example, IMPLEMENTATION DEFINED, IMPLEMENTATION SPECIFIC, UNKNOWN, and UNPREDICTABLE.
	Caution
	Warning
	Note

1.4 Additional reading

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

Table 1: Arm publications

Document name	Document ID	Licensee only Y/N
Arm® Architecture Reference Manual, Armv8, for Armv8-A architecture profile	DDI 0487	N
Arm® Neoverse™ N1 Technical Reference Manual	100616	N

1.5 Feedback

1.5.1 Feedback on this product

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

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

1.5.2 Feedback on content

If you have comments on content, send an e-mail to errata@arm.com and give:

- The title: Arm® Neoverse™ N1 Software Optimization Guide.
- The number: PJDOC-466751330-9707.
- If applicable, the page number(s) to which your comments refer.
- A concise explanation of your comments.

Arm also welcomes general suggestions for additions and improvements.

Note:



Arm tests the PDF only in Adobe Acrobat and Acrobat Reader, and cannot guarantee the quality of the represented document when used with any other PDF reader.

2 About this document

This document contains a guide to the Neoverse N1 micro-architecture with a view to aiding software optimization.

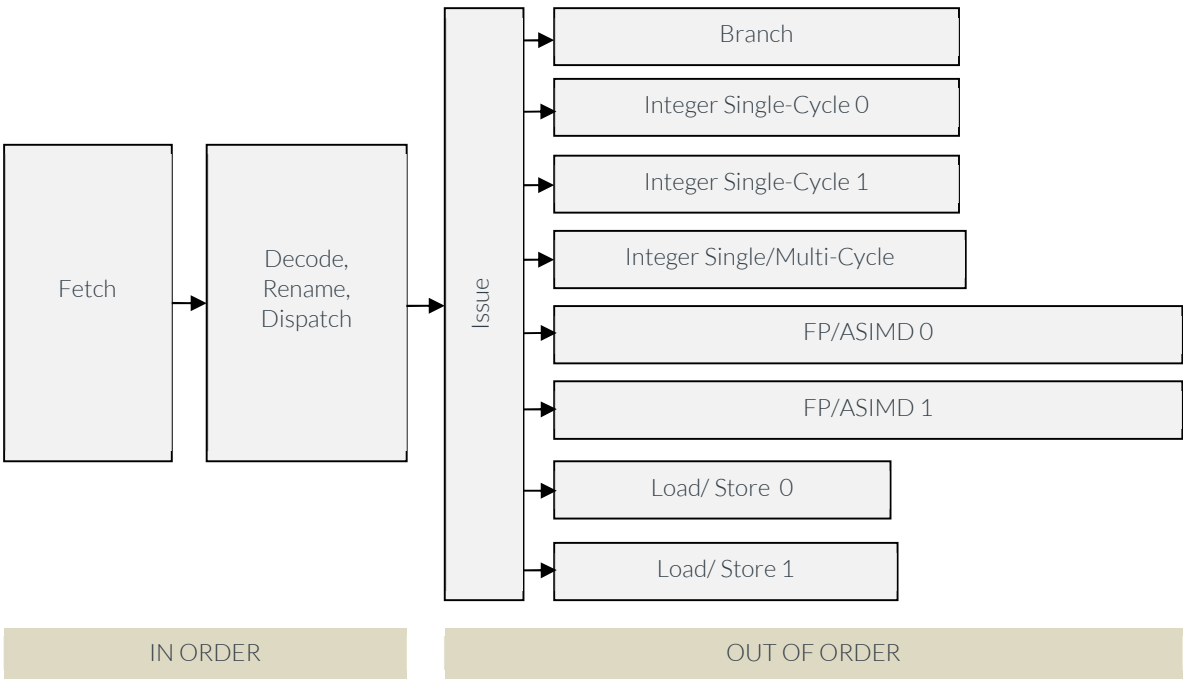
2.1 Scope

This document provides high-level information about the Neoverse N1 pipeline, instruction performance characteristics, and special performance considerations. This information is intended to aid people who are optimizing software and compilers for Neoverse N1. For a more complete description of the Neoverse N1 processor, please refer to the Arm® Neoverse™ N1 Technical Reference Manual.

2.2 Pipeline overview

The following diagram describes the high-level Neoverse N1 instruction processing pipeline. Instructions are first fetched, then decoded into internal macro-operations (Mops). From there, the Mops proceed through register renaming and dispatch stages. A Mop can be split further into two Uops at dispatch stage. Once dispatched, uops wait for their operands and issue out-of-order to one of eight execution pipelines. Each execution pipeline can accept and complete one uop per cycle.

Figure 1: Neoverse N1 pipeline



The execution pipelines support different types of operations, as shown in the following table.

Table 2: Neoverse N1 operations

Instruction groups	Instructions
Branch	Branch μ OPs
Integer Single-Cycle 0/1	Integer ALU μ OPs
Integer Single/Multi-cycle 0/1	Integer shift-ALU, multiply, divide, CRC and sum-of-absolute-differences μ OPs
Load/Store Address Generation 0/1	Load, Store address generation and special memory μ OPs
FP/ASIMD-0	ASIMD ALU, ASIMD misc, ASIMD integer multiply, FP convert, FP misc, FP add, FP multiply, FP divide, FP sqrt, crypto μ OPs, store data μ OPs
FP/ASIMD-1	ASIMD ALU, ASIMD misc, FP misc, FP add, FP multiply, ASIMD shift μ OPs, store data μ OPs, crypto μ OPs.

3 Instruction characteristics

3.1 Instruction tables

This chapter describes high-level performance characteristics for most Arm v8.2-A A32, T32 and A64 instructions. A series of tables summarize the effective execution latency and throughput (instruction bandwidth per cycle), pipelines utilized, and special behaviours associated with each group of instructions. Utilized pipelines correspond to the execution pipelines described in chapter 2.

In the tables below, Exec Latency is defined as the minimum latency seen by an operation dependent on an instruction in the described group.

In the tables below, Execution Throughput is defined as the maximum throughput (in instructions per cycle) of the specified instruction group that can be achieved in the entirety of the Neoverse N1 microarchitecture.

3.2 Legend for reading the utilized pipelines

Table 3: Neoverse N1 pipeline names and symbols

Pipeline name	Symbol used in tables
Branch	B
Integer single Cycle 0/1	S
Integer single Cycle 0/1 and single/multicycle	I
Integer single/multicycle	M
Integer single Cycle 1 and Integer multicycle	D
Load/Store 0/1	L
FP/ASIMD 0/1	V
FP/ASIMD 0	V0
FP/ASIMD 1	V1

3.3 Branch instructions

Table 4: AArch64 Branch instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Branch, immed	B	1	1	B	-
Branch, register	BR, RET	1	1	B	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Branch and link, immed	BL	1	1	I, B	-
Branch and link, register (reg != lr)	BLR	2	1	I, B	-
Branch and link, register (reg == lr)	BLR	2	1	I, B	-
Compare and branch	CBZ, CBNZ, TBZ, TBNZ	1	1	B	-

Table 5: AArch32 Branch instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Branch, immed	B	1	1	B	-
Branch, register	BX	1	1	B	-
Branch and link, immed	BL, BLX	1	1	B	-
Branch and link, register (reg != lr)	BLX	1	1	I, B	-
Branch and link, register (reg == lr)	BLX	2	1	I, B	-
Compare and branch	CBZ, CBNZ	1	1	B	-

3.4 Arithmetic and logical instructions

Table 6: AArch64 Arithmetic and logical instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Arithmetic, basic	ADD{S}, ADC{S}, SUB{S}, SBC{S}	1	3	I	
Arithmetic, extend and shift	ADD{S}, SUB{S}	2	1	M	
Arithmetic, LSL shift, shift <= 4	ADD{S}, SUB{S}	1	3	I	

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Arithmetic, LSR/ASR/ROR shift or LSL shift > 4	ADD{S}, SUB{S}	2	1	M	
Conditional compare	CCMN, CCMP	1	3	I	
Conditional select	CSEL, CSINC, CSINV, CSNEG	1	3	I	
Logical, basic	AND{S}, BIC{S}, EON, EOR, ORN, ORR	1	3	I	
Logical, shift, no flagset	AND, BIC, EON, EOR, ORN, ORR	1	3	I	
Logical, shift, flagset	ANDS, BICS	2	1	M	

Table 7: AArch32 Arithmetic and logical instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ALU, basic	ADD{S}, ADC{S}, ADR, AND{S}, BIC{S}, CMN, CMP, EOR{S}, ORN{S}, ORR{S}, RSB{S}, RSC{S}, SUB{S}, SBC{S}, TEQ, TST	1	3	I	
ALU, shift by register, unconditional	(same as ALU, basic)	2	1	M	
ALU, shift by register, conditional	(same as ALU, basic)	2	1	I, M	

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Arithmetic, LSL shift by immed, shift ≤ 4 , unconditional	ADD{S}, ADC{S}, RSB{S}, RSC{S}, SUB{S}, SBC{S}	1	3	I	
Arithmetic, LSL shift by immed, shift ≤ 4 , conditional	ADD{S}, ADC{S}, RSB{S}, RSC{S}, SUB{S}, SBC{S}	1	1	M	
Arithmetic, LSR/ASR/ROR shift by immed or LSL shift by immed > 4	ADD{S}, ADC{S}, RSB{S}, RSC{S}, SUB{S}, SBC{S}	2	1	M	
Logical, shift by immed, noflagset	AND, BIC, EOR, ORN, ORR	1	3	I	
Logical, shift by immed, flagset	AND{S}, BIC{S}, EOR{S}, ORN{S}, ORR{S}	2	1	M	
Test/Compare, shift by immed	CMN, CMP, TEQ, TST	2	1	M	
Branch forms		+1	1	+B	1

Note:



Branch forms are possible when the instruction destination register is the PC. For those cases, an additional branch μ OP is required. This adds 1 cycle to the latency.

3.5 Move and shift instructions

Table 8: AArch32 Move and shift instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
Move, basic	MOV{S}, MOVW, MVN{S}	1	3	I	
Move, shift by immed, no setflags	ASR, LSL, LSR, ROR, RRX, MVN	1	3	I	
Move, shift by immed, setflags	ASRS, LSLS, LSRS, RORS, RRXS, MVNS	2	1	M	
Move, shift by register, no setflags, unconditional	ASR, LSL, LSR, ROR, RRX, MVN	1	3	I	
Move, shift by register, no setflags, conditional	ASR, LSL, LSR, ROR, RRX, MVN	2	3/2	I	
Move, shift by register, setflags	ASRS, LSLS, LSRS, RORS, RRXS, MVNS	2	1	M	
Move, top	MOVT	1	3	I	
Move, branch forms		+1	1	+B	

Note:



Branch forms are possible when the instruction destination register is the PC. For those cases, an additional branch μ OP is required. This adds 1 cycle to the latency.

3.6 Divide and multiply instructions

Table 9: AArch64 Divide and multiply instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
Divide, W-form	SDIV, UDIV	5 to 12	1/12 to 1/5	M	1
Divide, X-form	SDIV, UDIV	5 to 20	1/20 to 1/5	M	1

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Multiply accumulate, W-form	MADD, MSUB	2(1)	1	M	2
Multiply accumulate, X-form	MADD, MSUB	4(3)	1/3	M	2, 4
Multiply accumulate long	SMADDL, SMSUBL, UMADDL, UMSUBL	2(1)	1	M	2
Multiply high	SMULH, UMULH	5(3)	1/4	M	5

Table 10: AArch32 Divide and multiply instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Divide	SDIV, UDIV	5 to 12	1/12 to 1/5	M	1
Multiply	MUL, SMULBB, SMULBT, SMULTB, SMULTT, SMULWB, SMULWT, SMMUL{R}, SMUAD{X}, SMUSD{X}	2	1	M	
Multiply accumulate, conditional	MLA, MLS, SMLABB, SMLABT, SMLATB, SMLATT, SMLAWB, SMLAWT, SMLAD{X}, SMLSD{X}, SMMLA{R}, SMMLS{R}	3	1	M, I	
Multiply accumulate, unconditional	MLA, MLS, SMLABB, SMLABT, SMLATB, SMLATT, SMLAWB, SMLAWT, SMLAD{X}, SMLSD{X}, SMMLA{R}, SMMLS{R}	2(1)	1	M	2
Multiply accumulate long, conditional	UMAAL	4	1	I, M	

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
Multiply accumulate long, unconditional	UMAAL	3	1	I, M	
Multiply accumulate long	SMLAL, SMLALBB, SMLALBT, SMLALTB, SMLALTT, SMLALD{X}, SMLSLD{X}, UMLAL	3	1	M, I	
Multiply long, all setflag, conditional and no setflag	SMULL, UMULL	3	1	M, I	
Multiply long, unconditional and no setflag	SMULL, UMULL	2	1	M	
(Multiply, setflags forms)		+1	(Same as above)	+I	3

Note:



1. Integer divides are performed using an iterative algorithm and block any subsequent divide operations until complete. Early termination is possible, depending upon the data values.
2. Multiply-accumulate pipelines support late-forwarding of accumulate operands from similar μ OPs, allowing a typical sequence of multiply-accumulate μ OPs to issue one every N cycles (accumulate latency N shown in parentheses).
3. Multiplies that set the condition flags require an additional integer μ OP.
4. X-form multiply accumulates stall the multiplier pipeline for 2 extra cycles.
5. Multiply high operations stall the multiplier pipeline for N extra cycles before any other type M μ op can be issued to that pipeline, with N shown in parentheses.

3.7 Saturating and parallel arithmetic instructions

Table 11: AArch32 Saturating and parallel arithmetic instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Parallel arith, unconditional	SADD16, SADD8, SSUB16, SSUB8, UADD16, UADD8, USUB16, USUB8	2	1	M	-
Parallel arith, conditional	SADD16, SADD8, SSUB16, SSUB8, UADD16, UADD8, USUB16, USUB8	2(4)	3/5	M, I	1
Parallel arith with exchange, unconditional	SASX, SSAX, UASX, USAX	3	1	I, M	
Parallel arith with exchange, conditional	SASX, SSAX, UASX, USAX	3(5)	3/5	I, M	1
Parallel halving arith	SHADD16, SHADD8, SHSUB16, SHSUB8, UHADD16, UHADD8, UHSUB16, UHSUB8	2	1	M	
Parallel halving arith with exchange	SHASX, SHSAX, UHASX, UHSAX	3	1	I, M	
Parallel saturating arith	QADD16, QADD8, QSUB16, QSUB8, UQADD16, UQADD8, UQSUB16, UQSUB8	2	2	M	
Parallel saturating arith with exchange	QASX, QSAX, UQASX, UQSAX	3	1	I, M	
Saturate	SSAT, SSAT16, USAT, USAT16	2	1	M	
Saturating arith	QADD, QSUB	2	1	M	

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Saturating doubling arith	QDADD, QDSUB	3	1	I, M	

Note:



1. Branch forms are possible Conditional GE-setting instructions require three extra uops and two additional cycles to conditionally update the GE field (GE latency shown in parentheses).

3.8 Miscellaneous data-processing instructions

Table 12: AArch64 Miscellaneous data-processing instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
Address generation	ADR, ADRP	1	3	I	-
Bitfield extract, one reg	EXTR	1	3	I	-
Bitfield extract, two regs	EXTR	3	1	I, M	-
Bitfield move, basic	SBFM, UBFM	1	3	I	-
Bitfield move, insert	BFM	2	1	M	-
Count leading	CLS, CLZ	1	3	I	-
Move immed	MOVN, MOVK, MOVZ	1	3	I	-
Reverse bits/bytes	RBIT, REV, REV16, REV32	1	3	I	-
Variable shift	ASRV, LSLV, LSRV, RORV	1	3	I	-

Table 13: AArch32 Miscellaneous data-processing instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Bit field extract	SBFX, UBFX	1	3	I	-
Bit field insert/clear	BFI, BFC	2	1	M	-
Count leading zeros	CLZ	1	3	I	-
Pack halfword	PKH	2	1	M	-
Reverse bits/bytes	RBIT, REV, REV16, REVSH	1	3	I	-
Select bytes, unconditional	SEL	1	3	I	-
Select bytes, conditional	SEL	2	3/2	I	-
Sign/zero extend, normal	SXTB, SXTH, UXTB, UXTH	1	3	I	-
Sign/zero extend, parallel	SXTB16, UXTB16	2	1	M	-
Sign/zero extend and add, normal	SXTAB, SXTAH, UXTAB, UXTAH	2	1	M	-
Sign/zero extend and add, parallel	SXTAB16, UXTAB16	4	1/2	M	-
Sum of absolute differences	USAD8, USADA8	2	1	M	-

3.9 Load instructions

Table 14: AArch64 Load instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Load register, literal	LDR, LDRSW, PRFM	4	2	L	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Load register, unscaled immed	LDUR, LDURB, LDURH, LDURSB, LDURSH, LDURSW, PRFUM	4	2	L	
Load register, immed post-index	LDR, LDRB, LDRH, LDRSB, LDRSH, LDRSW	4	2	L, I	
Load register, immed pre-index	LDR, LDRB, LDRH, LDRSB, LDRSH, LDRSW	4	2	L, I	
Load register, immed unprivileged	LDTR, LDTRB, LDTRH, LDTRSB, LDTRSH, LDTRSW	4	2	L	
Load register, unsigned immed	LDR, LDRB, LDRH, LDRSB, LDRSH, LDRSW, PRFM	4	2	L	
Load register, register offset, basic	LDR, LDRB, LDRH, LDRSB, LDRSH, LDRSW, PRFM	4	2	L	
Load register, register offset, scale by 4/8	LDR, LDRSW, PRFM	4	2	L	
Load register, register offset, scale by 2	LDRH, LDRSH	5	2	I, L	
Load register, register offset, extend	LDR, LDRB, LDRH, LDRSB, LDRSH, LDRSW, PRFM	4	2	L	
Load register, register offset, extend, scale by 4/8	LDR, LDRSW, PRFM	4	2	L	

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Load register, register offset, extend, scale by 2	LDRH, LDRSH	5	2	I, L	
Load pair, signed immed offset, normal, W-form	LDP, LDNP	4	2	L	
Load pair, signed immed offset, normal, X-form	LDP, LDNP	4	1	L	
Load pair, signed immed offset, signed words, base != SP	LDPSW	5	1	I, L	
Load pair, signed immed offset, signed words, base = SP	LDPSW	5	1	I, L	
Load pair, immed post-index, normal	LDP	4	1	L, I	
Load pair, immed post-index, signed words	LDPSW	5	1	I, L	
Load pair, immed pre-index, normal	LDP	4	1	L, I	
Load pair, immed pre-index, signed words	LDPSW	5	1	I, L	

Table 15: AArch32 Load instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Load, immed offset	LDR{T}, LDRB{T}, LDRD, LDRH{T}, LDRSB{T}, LDRSH{T}	4	2	L	1,2

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Load, register offset, plus	LDR, LDRB, LDRD, LDRH, LDRSB, LDRSH	4	2	L	1,2
Load, register offset, minus	LDR, LDRB, LDRD, LDRH, LDRSB, LDRSH	5	2	I, L	1,2
Load, scaled register offset, plus, LSL2	LDR, LDRB	4	2	L	1
Load, scaled register offset, other	LDR, LDRB, LDRH, LDRSB, LDRSH	5	2	I, L	1
Load, immed pre-indexed	LDR, LDRB, LDRD, LDRH, LDRSB, LDRSH	4	2	L, I	1,2
Load, register pre-indexed, shift Rm, plus and minus	LDR, LDRB, LDRH, LDRSB, LDRSH	5	2	I, L, M	3
Load, register pre-indexed	LDRD	4	2	L, I	-
Load, register pre-indexed, cond	LDRD	5	1 1/2	L, I	-
Load, scaled register pre-indexed, plus, LSL2	LDR, LDRB	4	2	L, I	1
Load, scaled register pre-indexed, unshifted	LDR, LDRB	4	2	L, I	-
Load, immed post-indexed	LDR{T}, LDRB{T}, LDRD, LDRH{T}, LDRSB{T}, LDRSH{T}	4	2	L, I	1,2

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Load, register post-indexed	LDR, LDRB, LDRH{T}, LDRSB{T}, LDRSH{T}	5	2	I, L	-
Load, register post-indexed	LDRD	4	2	L, I	-
Load, register post-indexed	LDRT, LDRBT	5	2	I, L	-
Load, scaled register post-indexed	LDR, LDRB	4	2	L, M	3
Load, scaled register post-indexed	LDRT, LDRBT	4	2	L, M	3
Preload, immed offset	PLD, PLDW	4	2	L	-
Preload, register offset, plus	PLD, PLDW	4	2	L	-
Preload, register offset, minus	PLD, PLDW	5	2	I, L	-
Preload, scaled register offset, plus LSL2	PLD, PLDW	5	2	I, L	-
Preload, scaled register offset, other	PLD, PLDW	5	2	I, L	-
Load multiple, no writeback, base reg not in list	LDMIA, LDMIB, LDMDA, LDMDB	N	2/R	L	1, 4, 5
Load multiple, no writeback, base reg in list	LDMIA, LDMIB, LDMDA, LDMDB	1+ N	2/R	I, L	1, 4, 5
Load multiple, writeback	LDMIA, LDMIB, LDMDA, LDMDB, POP	1+N	2/R	L, I	1, 4, 5
(Load, all branch forms)	-	+1	-	+ B	6

Note:

1. Condition loads have an extra μ OP which goes down pipeline I and have 1 cycle extra latency compared to their unconditional counterparts.



2. The throughput of conditional LDRD is 1 as compared to a throughput of 2 for unconditional LDRD.
3. The address update op for addressing forms which use reg scaled reg, or reg extend goes down pipeline 'I' if the shift is LSL where the shift value is less than or equal to 4.
4. N is $\text{floor}[(\text{num_reg}+3)/4]$.
5. R is $\text{floor}[(\text{num_reg}+1)/2]$.
6. Branch forms are possible when the instruction destination register is the PC. For those cases, an additional branch μOP is required. This adds 1 cycle to the latency.

3.10 Store instructions

The following tables describes performance characteristics for standard store instructions. Stores μOPs are split into address and data μOPs . Once executed, stores are buffered and committed in the background.

Table 16: AArch64 Store instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Store register, unscaled immed	STUR, STURB, STURH	1	2	L, D	-
Store register, immed post-index	STR, STRB, STRH	1	2	L, D	-
Store register, immed pre-index	STR, STRB, STRH	1	2	L, D	-
Store register, immed unprivileged	STTR, STTRB, STTRH	1	2	L, D	-
Store register, unsigned immed	STR, STRB, STRH	1	2	L, D	-
Store register, register offset, basic	STR, STRB, STRH	1	2	L, D	-
Store register, register offset, scaled by 4/8	STR	1	2	L, D	-
Store register, register offset, scaled by 2	STRH	2	3/2	I, L, D	-
Store register, register offset, extend	STR, STRB, STRH	1	2	L, D	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Store register, register offset, extend, scale by 4/8	STR	1	2	L, D	-
Store register, register offset, extend, scale by 1	STRH	2	3/2	I, L, D	-
Store pair, immed offset, W-form	STP, STNP	1	2	L, D	-
Store pair, immed offset, X-form	STP, STNP	1	1	L, D	-
Store pair, immed post-index, W-form	STP	1	1	L, D	-
Store pair, immed post-index, X-form	STP	1	1	L, D	-
Store pair, immed pre-index, W-form	STP	1	1	L, D	-
Store pair, immed pre-index, X-form	STP	1	1	L, D	-

Table 17: AArch32 Store instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Store, immed offset	STR{T}, STRB{T}, STRD, STRH{T}	1	2	L, D	-
Store, register offset, plus	STR, STRB, STRD, STRH	1	2	L, D	-
Store, register offset, minus	STR, STRB, STRD, STRH	1	2	L, D	-
Store, register offset, no shift, plus	STR, STRB	1	2	L, D	-
Store, scaled register offset, plus LSL2	STR, STRB	1	2	L, D	-
Store, scaled register offset, other	STR, STRB	2	3/2	I, L, D	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Store, scaled register offset, minus	STR, STRB	2	3/2	I, L, D	-
Store, immed pre-indexed	STR, STRB, STRD, STRH	1	3/2	I, L, D	-
Store, register pre-indexed, plus, no shift	STR, STRB, STRD, STRH	1	3/2	L, D	-
Store, register pre-indexed, minus	STR, STRB, STRD, STRH	2	1	I, L, D	-
Store, scaled register pre-indexed, plus LSL2	STR, STRB	1	3/2	L, D	-
Store, scaled register pre-indexed, other	STR, STRB	2	1	I, L, D, M	1
Store, immed post-indexed	STR{T}, STRB{T}, STRD, STRH{T}	1	3/2	L, D	-
Store, register post-indexed	STRH{T}, STRD	1	3/2	L, D	-
Store, register post-indexed	STR{T}, STRB{T}	1	3/2	L, D	-
Store, scaled register post-indexed	STR{T}, STRB{T}	1	3/2	L, D	-
Store multiple, no writeback	STMIA, STMIB, STMDA, STMDB	N	1/N	L, D	2
Store multiple, writeback	STMIA, STMIB, STMDA, STMDB, PUSH	N	1/N	L, D	2

Note:



1. The address update op for addressing forms which use reg scaled reg, or reg extend goes down pipeline 'I' if the shift is LSL where the shift value is less than or equal to 4.
2. For store multiple instructions, $N = \text{floor}((\text{num_regs} + 3) / 4)$.

3.11 FP data processing instructions

Table 18: AArch64 FP data processing instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
FP absolute value	FABS	2	2	V	-
FP arithmetic	FADD, FSUB	2	2	V	-
FP compare	FCCMP{E}, FCMP{E}	2	1	V0	-
FP divide, H-form	FDIV	7	4/7	V0	1
FP divide, S-form	FDIV	7 to 10	4/9 to 4/7	V0	1
FP divide, D-form	FDIV	7 to 15	1/7 to 2/7	V0	1
FP min/max	FMIN, FMINNM, FMAX, FMAXNM	2	2	V	-
FP multiply	FMUL, FNMUL	3	2	V	2
FP multiply accumulate	FMADD, FMSUB, FNMADD, FNMSUB	4 (2)	2	V	3
FP negate	FNEG	2	2	V	-
FP round to integral	FRINTA, FRINTI, FRINTM, FRINTN, FRINTP, FRINTX, FRINTZ	3	1	V0	-
FP select	FCSEL	2	2	V	-
FP square root, H-form	FSQRT	7	4/7	V0	1
FP square root, S-form	FSQRT	7 to 10	4/9 to 4/7	V0	1
FP square root, D-form	FSQRT	7 to 17	1/8 to 2/7	V0	1

Table 19: AArch32 FP data processing instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
VFP absolute value	VABS	2	2	V	-
VFP arith	VADD, VSUB	2	2	V	-
VFP compare, unconditional	VCMP, VCMPE	2	1	VO	-
VFP compare, conditional	VCMP, VCMPE	4	1	V, VO	-
VFP convert	VCVT{R}, VCVTB, VCVTT, VCVTA, VCVTM, VCVTN, VCVTP	3	1	VO	-
VFP divide, H-form	VDIV	7	4/7	VO	1
VFP divide, S-form	VDIV	7 to 10	4/9 to 4/7	VO	1
VFP divide, D-form	VDIV	7 to 15	1/7 to 2/7	VO	1
VFP max/min	VMAXNM, VMINNM	2	2	V	-
VFP multiply	VMUL, VMUL	3	2	V	2
VFP multiply accumulate (chained)	VMLA, VMLS, VNMLA, VNMLS	5 (2)	2	V	3
VFP multiply accumulate (fused)	VFMA, VFMS, VFNMA, VFNMS	4 (2)	2	V	3
VFP negate	VNEG	2	2	V	-
VFP round to integral	VRINTA, VRINTM, VRINTN, VRINTP, VRINTR, VRINTX, VRINTZ	3	1	VO	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
VFP select	VSELEQ, VSELGE, VSELGT, VSELVS	2	2	V	-
VFP square root, H-form	VSQRT	7	4/7	V0	1
VFP square root, S-form	VSQRT	7 to 10	4/9 to 4/7	V0	1
VFP square root, D-form	VSQRT	7 to 17	1/8 to 2/7	V0	1

Note:



1. FP divide and square root operations are performed using an iterative algorithm and block subsequent similar operations to the same pipeline until complete.
2. FP multiply-accumulate pipelines support late forwarding of the result from FP multiply μ OPs to the accumulate operands of an FP multiply-accumulate μ OP. The latter can potentially be issued 1 cycle after the FP multiply μ OP has been issued.
3. FP multiply-accumulate pipelines support late-forwarding of accumulate operands from similar μ OPs, allowing a typical sequence of multiply-accumulate μ OPs to issue one every N cycles (accumulate latency N shown in parentheses).

3.12 FP miscellaneous instructions

Table 20: AArch64 FP miscellaneous instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
FP convert, from vec to vec reg	FCVT, FCVTXN	3	1	V0	-
FP convert, from gen to vec reg	SCVTF, UCVTF	6	1	M, V0	-
FP convert, from vec to gen reg	FCVTAS, FCVTAU, FCVTMS, FCVTMU, FCVTNS, FCVTNU, FCVTPS, FCVTPU, FCVTZS, FCVTZU	4	1	V0, V1	-
FP move, immed	FMOV	2	2	V	-
FP move, register	FMOV	2	2	V	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
FP transfer, from gen to vec reg	FMOV	3	1	M	-
FP transfer, from vec to gen reg	FMOV	2	1	V1	-

Table 21: AArch32 FP miscellaneous instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
VFP move, immed	VMOV	2	2	V	-
VFP move, register	VMOV	2	2	V	-
VFP move, insert	VINS	2	2	V	-
VFP move, extraction	VMOVX	2	2	V	-
VFP transfer, core to vfp, single reg to S-reg, cond	VMOV	5	1	M, V	-
VFP transfer, core to vfp, single reg to S-reg, uncond	VMOV	3	1	M	-
VFP transfer, core to vfp, single reg to upper/lower half of D-reg	VMOV	5	1	M, V	-
VFP transfer, core to vfp, 2 regs to 2 S-reg, cond	VMOV	6	1/2	M, V	-
VFP transfer, core to vfp, 2 regs to 2 S-reg, uncond	VMOV	4	1/2	M	-
VFP transfer, core to vfp, 2 regs to D-reg, cond	VMOV	5	1	M, V	-
VFP transfer, core to vfp, 2 regs to D-reg, uncond	VMOV	3	1	M	-
VFP transfer, vfp S-reg or upper/lower half of vfp D-reg to core reg, cond	VMOV	3	1	V1, I	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
VFP transfer, vfp S-reg or upper/lower half of vfp D-reg to core reg, uncond	VMOV	2	1	V1	-
VFP transfer, vfp 2 S-regs or D-reg to 2 core regs, cond	VMOV	3	1	V1, I	
VFP transfer, vfp 2 S-regs or D-reg to 2 core regs, uncond	VMOV	2	1	V1	

3.13 FP load instructions

The latencies shown assume the memory access hits in the Level 1 Data Cache. Compared to standard loads, an extra cycle is required to forward results to FP/ASIMD pipelines.

Table 22: AArch64 FP load instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
Load vector reg, literal, S/D/Q forms	LDR	-	2	L	-
Load vector reg, unscaled immed	LDUR	5	2	L	-
Load vector reg, immed post-index	LDR	5	2	L, I	-
Load vector reg, immed pre-index	LDR	5	2	L, I	-
Load vector reg, unsigned immed	LDR	5	2	L, I	-
Load vector reg, register offset, basic	LDR	5	2	L, I	-
Load vector reg, register offset, scale, S/D-form	LDR	5	2	L, I	-
Load vector reg, register offset, scale, H/Q-form	LDR	6	2	I, L	-
Load vector reg, register offset, extend	LDR	5	2	L, I	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Load vector reg, register offset, extend, scale, S/D-form	LDR	5	2	L, I	-
Load vector reg, register offset, extend, scale, H/Q-form	LDR	6	2	I, L	-
Load vector pair, immed offset, S/D-form	LDP, LDNP	5	1	L, I	-
Load vector pair, immed offset, Q-form	LDP, LDNP	7	1	L	-
Load vector pair, immed post-index, S/D-form	LDP	5	1	I, L	-
Load vector pair, immed post-index, Q-form	LDP	7	1	L, I	-
Load vector pair, immed pre-index, S/D-form	LDP	5	1	I, L	-
Load vector pair, immed pre-index, Q-form	LDP	7	1	L, I	-

Table 23: AArch32 FP load instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
FP load, register	VLDR	4	2	L	1
FP load multiple, S form	VLDmia, VLDmDB, VPOP	N	2/R	L	1, 2, 3
FP load multiple, D form	VLDmia, VLDmDB, VPOP	N + 2	1/R	L, V	1, 2, 3
(FP load, writeback forms)	-	(1)	-	+ I	4

Note:



1. Condition loads have an extra μOP which goes down pipeline V and have 2 cycle extra latency compared to their unconditional counterparts.
2. N is $\text{floor}[(\text{num_reg}+3)/4]$.
3. R is $\text{floor}[(\text{num_reg}+1)/2]$.

4. Writeback forms of load instructions require an extra μ OP to update the base address. This update is typically performed in parallel with or prior to the load μ OP (update latency shown in parentheses).

3.14 FP store instructions

Stores MOPs are split into store address and store data μ OPs. Once executed, stores are buffered and committed in the background.

Table 24: AArch64 FP store instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Store vector reg, unscaled immed, B/H/S/D-form	STUR	2	2	L, I	-
Store vector reg, unscaled immed, Q-form	STUR	2	1	L, I	-
Store vector reg, immed post-index, B/H/S/D-form	STR	2	2	L, V	-
Store vector reg, immed post-index, Q-form	STR	2	1	L, V	-
Store vector reg, immed pre-index, B/H/S/D-form	STR	2	2	L, V	-
Store vector reg, immed pre-index, Q-form	STR	2	1	L, V	-
Store vector reg, unsigned immed, B/H/S/D-form	STR	2	2	L, V	-
Store vector reg, unsigned immed, Q-form	STR	2	1	L, V	-
Store vector reg, register offset, basic, B/H/S/D-form	STR	2	2	L, V	-
Store vector reg, register offset, basic, Q-form	STR	2	1	L, V	-
Store vector reg, register offset, scale, H-form	STR	2	2	I, L, V	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Store vector reg, register offset, scale, S/D-form	STR	2	2	L, V	-
Store vector reg, register offset, scale, Q-form	STR	2	1	I, L, V	-
Store vector reg, register offset, extend, B/H/S/D-form	STR	2	2	L, V	-
Store vector reg, register offset, extend, Q-form	STR	2	1	L, V	-
Store vector reg, register offset, extend, scale, H-form	STR	2	2	I, L, V	-
Store vector reg, register offset, extend, scale, S/D-form	STR	2	2	L, V	-
Store vector reg, register offset, extend, scale, Q-form	STR	2	1	I, L, V	-
Store vector pair, immed offset, S-form	STP, STNP	2	2	L, V	-
Store vector pair, immed offset, D-form	STP, STNP	2	1	L, V	-
Store vector pair, immed offset, Q-form	STP, STNP	3	1/2	L, V	-
Store vector pair, immed post-index, S-form	STP	2	1	L, V	-
Store vector pair, immed post-index, D-form	STP	2	1	L, V	-
Store vector pair, immed post-index, Q-form	STP	3	1	L, V	-
Store vector pair, immed pre-index, S-form	STP	2	1	L, V	-
Store vector pair, immed pre-index, D-form	STP	2	1	L, V	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Store vector pair, immed pre-index, Q-form	STP	3	1/2	L, V	-

Table 25: AArch32 FP store instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
FP store, immed offset	VSTR	2	2	L, I	-
FP store multiple, S-form	VSTMIA, VSTMDB, VPUSH	N+1	2/R	L, V	1, 3
FP store multiple, D-form	VSTMIA, VSTMDB, VPUSH	P + 1	1/R	L, V	2, 3
(FP store, writeback forms)	-	(1)	-	+ I	4

Note:



1. For store multiple instructions, $N = \text{floor}((\text{num_regs} + 3) / 4)$.
2. For store multiple instructions, $P = \text{floor}((\text{num_regs} + 1) / 2)$.
3. $R = \text{floor}[(\text{num_regs} + 1) / 2]$.
4. Writeback forms of store instructions require an extra μOP to update the base address. This update is typically performed in parallel with or prior to the store μOP (update latency shown in parentheses).

3.15 ASIMD integer instructions

Table 26: AArch64 ASIMD integer instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD absolute diff	SABD, UABD	2	2	V	-
ASIMD absolute diff accum	SABA, UABA	4(1)	1	V1	2

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD absolute diff accum long	SABAL(2), UABAL(2)	4(1)	1	V1	2
ASIMD absolute diff long	SABDL(2), UABDL(2)	2	2	V	-
ASIMD arith, basic	ABS, ADD, NEG, SADDL(2), SADDW(2), SHADD, SHSUB, SSUBL(2), SSUBW(2), SUB, UADDL(2), UADDW(2), UHADD, UHSUB, USUBL(2), USUBW(2)	2	2	V	-
ASIMD arith, complex	ADDHN(2), RADDHN(2), RSUBHN(2), SQABS, SQADD, SQNEG, SQSUB, SRHADD, SUBHN(2), SUQADD, UQADD, UQSUB, URHADD, USQADD	2	2	V	-
ASIMD arith, pair-wise	ADDP, SADDLP, UADDLP	2	2	V	-
ASIMD arith, reduce, 4H/4S	ADDV, SADDLV, UADDLV	3	1	V1	-
ASIMD arith, reduce, 8B/8H	ADDV, SADDLV, UADDLV	5	1	V1, V	-
ASIMD arith, reduce, 16B	ADDV, SADDLV, UADDLV	6	1/2	V1	-
ASIMD compare	CMEQ, CMGE, CMGT, CMHI, CMHS, CMLE, CMLT, CMTST	2	2	V	-
ASIMD dot product	SDOT, UDOT	2	2	V	-
ASIMD logical	AND, BIC, EOR, MOV, MVN, ORN, ORR, NOT	2	2	V	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD max/min, basic and pairwise	SMAX, SMAXP, SMIN, SMINP, UMAX, UMAXP, UMIN, UMINP	2	2	V	-
ASIMD max/min, reduce, 4H/4S	SMAXV, SMINV, UMAXV, UMINV	3	1	V1	-
ASIMD max/min, reduce, 8B/8H	SMAXV, SMINV, UMAXV, UMINV	5	1	V1, V	-
ASIMD max/min, reduce, 16B	SMAXV, SMINV, UMAXV, UMINV	6	1/2	V1	-
ASIMD multiply, D-form	MUL, SQDMULH, SQRDMULH	4	1	V0	-
ASIMD multiply, Q-form	MUL, SQDMULH, SQRDMULH	5	1/2	V0	-
ASIMD multiply accumulate, D-form	MLA, MLS	4(1)	1	V0	1
ASIMD multiply accumulate, Q-form	MLA, MLS	5(2)	1/2	V0	1
ASIMD multiply accumulate high, D-form	SQRDMLAH, SQRDMLSH	4	1	V0	-
ASIMD multiply accumulate high, Q-form	SQRDMLAH, SQRDMLSH	5	1/2	V0	-
ASIMD multiply accumulate long	SMLAL(2), SMLSL(2), UMLAL(2), UMLSL(2)	4(1)	1	V0	1
ASIMD multiply accumulate saturating long	SQDMLAL(2), SQDMLSL(2)	4	1	V0	-
ASIMD multiply/multiply long (8x8) polynomial, D-form	PMUL, PMULL(2)	3	1	V0	3

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD multiply/multiply long (8x8) polynomial, Q-form	PMUL, PMULL(2)	4	1/2	V0	3
ASIMD multiply long	SMULL(2), UMULL(2), SQDMULL(2)	4	1	V0	-
ASIMD pairwise add and accumulate long	SADALP, UADALP	4(1)	1	V1	2
ASIMD shift accumulate	SSRA, SRSRA, USRA, URSRA	4(1)	1	V1	2
ASIMD shift by immed, basic	SHL, SHLL(2), SHRN(2), SSHLL(2), SSHR, SXTL(2), USHLL(2), USHR, UXTL(2)	2	1	V1	-
ASIMD shift by immed and insert, basic	SLI, SRI	2	1	V1	-
ASIMD shift by immed, complex	RSHRN(2), SQRSHRN(2), SQRSHRUN(2), SQSHL{U}, SQSHRN(2), SQSHRUN(2), RSHR, UQRSHRN(2), UQSHL, UQSHRN(2), URSR	4	1	V1	-
ASIMD shift by register, basic	SSHL, USHL	2	1	V1	-
ASIMD shift by register, complex	SRSHL, SQRSHL, SQSHL, URSHL, UQRSHL, UQSHL	4	1	V1	-

Table 27: AArch32 ASIMD integer instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD absolute diff	VABD	2	2	V	-
ASIMD absolute diff accum	VABA	4(1)	1	V1	2
ASIMD absolute diff accum long	VABAL	4(1)	1	V1	2
ASIMD absolute diff long	VABDL	2	2	V	-
ASIMD arith, basic	VADD, VADDL, VADDW, VNEG, VSUB, VSUBL, VSUBW	2	2	V	-
ASIMD arith, complex	VABS, VADDHN, VHADD, VHSUB, VQABS, VQADD, VQNEG, VQSUB, VRADDHN, VRHADD, VRSUBHN, VSUBHN	2	2	V	-
ASIMD arith, pair-wise	VPADD, VPADDL	2	2	V	-
ASIMD compare	VCEQ, VCGE, VCGT, VCLE, VTST	2	1	V	-
ASIMD dot product	VSDOT, VUDOT	2	2	V	-
ASIMD logical	VAND, VBIC, VMVN, VORR, VORN, VEOR	2	1	V	-
ASIMD max/min	VMAX, VMIN, VPMAX, VPMIN	2	1	V	-
ASIMD multiply, D-form	VMUL, VQDMULH, VQRDMULH	4	1	V0	-
ASIMD multiply, Q-form	VMUL, VQDMULH, VQRDMULH	5	1/2	V0	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD multiply accumulate, D-form	VMLA, VMLS	4(1)	1	V0	1
ASIMD multiply accumulate, Q-form	VMLA, VMLS	5(2)	1/2	V0	1
ASIMD multiply accumulate long	VMLAL, VMLS�	4(1)	1	V0	1
ASIMD multiply accumulate high, D-form	VQRDMLAH, VQRDMLSH	4	1	V0	-
ASIMD multiply accumulate high, Q-form	VQRDMLAH, VQRDMLSH	5	1/2	V0	-
ASIMD multiply accumulate saturating long	VQDMLAL, VQDMLSL	4	1	V0	-
ASIMD multiply/multiply long (8x8) polynomial, D-form	VMUL (.P8), VMULL (.P8)	3	1	V0	-
ASIMD multiply (8x8) polynomial, Q-form	VMUL (.P8)	4	1/2	V0	-
ASIMD multiply long	VMULL (.S, .L), VQDMULL	4	1	V0	-
ASIMD pairwise add and accumulate	VPADAL	4(1)	1	V1	1
ASIMD shift accumulate	VSRA, VRSRA	4(1)	1	V1	1
ASIMD shift by immed, basic	VMOVL, VSHL, VSHLL, VSHR, VSHRN	2	1	V1	-
ASIMD shift by immed and insert, basic	VSLI, VSRI	2	1	V1	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD shift by immed, complex	VQRSHRN, VQRSHRUN, VQSHL{U}, VQSHRN, VQSHRUN, VRSHR, VRSHRN	4	1	V1	-
ASIMD shift by register, basic	VSHL	2	1	V1	-
ASIMD shift by register, complex	VQRSHL, VQSHL, VRSHL	4	1	V1	-

Note:



1. Multiply-accumulate pipelines support late-forwarding of accumulate operands from similar μ OPs, allowing a typical sequence of integer multiply-accumulate μ OPs to issue one every cycle or one every other cycle (accumulate latency shown in parentheses).
2. Other accumulate pipelines also support late-forwarding of accumulate operands from similar μ OPs, allowing a typical sequence of such μ OPs to issue one every cycle (accumulate latency shown in parentheses).
3. This category includes instructions of the form “PMULL Vd.8H, Vn.8B, Vm.8B” and “PMULL2 Vd.8H, Vn.16B, Vm.16B”.

3.16 ASIMD floating-point instructions

Table 28: AArch64 ASIMD floating point instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
ASIMD FP absolute value/difference	FABS, FABD	2	2	V	-
ASIMD FP arith, normal	FABD, FADD, FSUB, FADDP	2	2	V	-
ASIMD FP compare	FACGE, FACGT, FCMEQ, FCMGE, FCMGT, FCMLE, FCMLT	2	2	V	-
ASIMD FP convert, long (F16 to F32)	FCVTL(2)	4	1/2	V0	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD FP convert, long (F32 to F64)	FCVTL(2)	3	1	V0	-
ASIMD FP convert, narrow (F32 to F16)	FCVTN(2)	4	1/2	V0	-
ASIMD FP convert, narrow (F64 to F32)	FCVTN(2), FCVTXN(2)	3	1	V0	-
ASIMD FP convert, other, D-form F32 and Q-form F64	FCVTAS, FCVTAU, FCVTMS, FCVTMU, FCVTNS, FCVTNU, FCVTPS, FCVTPU, FCVTZS, FCVTZU, SCVTF, UCVTF	3	1	V0	-
ASIMD FP convert, other, D-form F16 and Q-form F32	FCVTAS, FCVTAU, FCVTMS, FCVTMU, FCVTNS, FCVTNU, FCVTPS, FCVTPU, FCVTZS, FCVTZU, SCVTF, UCVTF	4	1/2	V0	-
ASIMD FP convert, other, Q-form F16	FCVTAS, FCVTAU, FCVTMS, FCVTMU, FCVTNS, FCVTNU, FCVTPS, FCVTPU, FCVTZS, FCVTZU, SCVTF, UCVTF	6	1/4	V0	-
ASIMD FP divide, D-form, F16	FDIV	7	1/7	V0	3
ASIMD FP divide, D-form, F32	FDIV	7 to 10	2/9 to 2/7	V0	3
ASIMD FP divide, Q-form, F16	FDIV	10 to 13	1/13 to 1/10	V0	3
ASIMD FP divide, Q-form, F32	FDIV	7 to 10	1/9 to 1/7	V0	3

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD FP divide, Q-form, F64	FDIV	7 to 15	1/14 to 1/7	V0	3
ASIMD FP max/min, normal	FMAX, FMAXNM, FMIN, FMINNM	2	2	V	-
ASIMD FP max/min, pairwise	FMAXP, FMAXNMP, FMINP, FMINNMP	2	2	V	-
ASIMD FP max/min, reduce	FMAXV, FMAXNMV, FMINV, FMINNMV	5	2	V	-
ASIMD FP max/min, reduce, Q-form F16	FMAXV, FMAXNMV, FMINV, FMINNMV	8	2/3	V	-
ASIMD FP multiply	FMUL, FMULX	3	2	V	2
ASIMD FP multiply accumulate	FMLA, FMLS	4 (2)	2	V	1
ASIMD FP multiply accumulate long	FMLAL(2), FMLSL(2)	5(2)	2	V	1
ASIMD FP negate	FNEG	2	2	V	-
ASIMD FP round, D-form F32 and Q-form F64	FRINTA, FRINTI, FRINTM, FRINTN, FRINTP, FRINTX, FRINTZ	3	1	V0	-
ASIMD FP round, D-form F16 and Q-form F32	FRINTA, FRINTI, FRINTM, FRINTN, FRINTP, FRINTX, FRINTZ	4	1/2	V0	-
ASIMD FP round, Q-form F16	FRINTA, FRINTI, FRINTM, FRINTN, FRINTP, FRINTX, FRINTZ	6	1/4	V0	-
ASIMD FP square root, D-form, F16	FSQRT	7	1/7	V0	3

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD FP square root, D-form, F32	FSQRT	7 to 10	2/9 to 2/7	V0	3
ASIMD FP square root, Q-form, F16	FSQRT	11 to 13	1/13 to 1/11	V0	3
ASIMD FP square root, Q-form, F32	FSQRT	7 to 10	1/9 to 1/7	V0	3
ASIMD FP square root, Q-form, F64	FSQRT	7 to 17	1/16 to 1/7	V0	3

Table 29: AArch32 ASIMD floating point instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD FP absolute value	VABS	2	2	V	-
ASIMD FP arith	VABD, VADD, VPADD, VSUB	2	2	V	-
ASIMD FP compare	VACGE, VACGT, VACLE, VACLT, VCEQ, VCGE, VCGT, VCLE	2	2	V	-
ASIMD FP convert, integer, D-form	VCVT, VCVTA, VCVTM, VCVTN, VCVTP	3	1	V0	-
ASIMD FP convert, integer, Q-form	VCVT, VCVTA, VCVTM, VCVTN, VCVTP	4	1/2	V0	-
ASIMD FP convert, fixed, D-form	VCVT	3	1	V0	-
ASIMD FP convert, fixed, Q-form	VCVT	4	1/2	V0	-
ASIMD FP convert, half-precision	VCVT	4	1/2	V0	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD FP max/min	VMAX, VMIN, VPMAX, VPMIN, VMAXNM, VMINNM	2	2	V	-
ASIMD FP multiply	VMUL, VNMUL	3	2	V	2
ASIMD FP chained multiply accumulate	VMLA, VMLS	5(2)	2	V	1
ASIMD FP fused multiply accumulate	VFMA, VFMS	4(2)	2	V	1
ASIMD FP fused multiply accumulate long	VFMAL(2), VFMSL(2)	4(2)	2	V	1
ASIMD FP negate	VNEG	2	2	V	
ASIMD FP round to integral, D-form	VRINTA, VRINTM, VRINTN, VRINTP, VRINTX, VRINTZ	3	1	VO	-
ASIMD FP round to integral, Q-form	VRINTA, VRINTM, VRINTN, VRINTP, VRINTX, VRINTZ	4	1/2	VO	-

Note:



1. ASIMD multiply-accumulate pipelines support late-forwarding of accumulate operands from similar μ OPs, allowing a typical sequence of floating-point multiply-accumulate μ OPs to issue one every N cycles (accumulate latency N shown in parentheses).
2. ASIMD multiply-accumulate pipelines support late forwarding of the result from ASIMD FP multiply μ OPs to the accumulate operands of an ASIMD FP multiply-accumulate μ OP. The latter can potentially be issued 1 cycle after the ASIMD FP multiply μ OP has been issued.
3. ASIMD divide and square root operations are performed using an iterative algorithm and block subsequent similar operations to the same pipeline until complete.

3.17 ASIMD miscellaneous instructions

Table 30: AArch64 ASIMD miscellaneous instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD bit reverse	RBIT	2	2	V	-
ASIMD bitwise insert	BIF, BIT, BSL	2	2	V	
ASIMD count	CLS, CLZ, CNT	2	2	V	
ASIMD duplicate, gen reg	DUP	3	1	M	
ASIMD duplicate, element	DUP	2	2	V	
ASIMD extract	EXT	2	2	V	
ASIMD extract narrow	XTN	2	2	V	
ASIMD extract narrow, saturating	SQXTN(2), SQXTUN(2), UQXTN(2)	4	1	V1	
ASIMD insert, element to element	INS	2	2	V	
ASIMD move, FP immed	FMOV	2	2	V	
ASIMD move, integer immed	MOVI, MVNI	2	2	V	
ASIMD reciprocal estimate, D-form F32 and F64	FRECPE, FRECPX, FRSQRTE, URECPE, URSQRTE	3	1	V0	
ASIMD reciprocal estimate, D-form F16 and Q-form F32	FRECPE, FRECPX, FRSQRTE, URECPE, URSQRTE	4	1/2	V0	
ASIMD reciprocal estimate, Q-form F16	FRECPE, FRECPX, FRSQRTE, URECPE, URSQRTE	6	1/4	V0	

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD reciprocal step	FRECPS, FRSQRTS	4	2	V	
ASIMD reverse	REV16, REV32, REV64	2	2	V	
ASIMD table lookup, 1 or 2 table regs	TBL	2	2	V	
ASIMD table lookup, 3 table regs	TBL	4	1/2	V	
ASIMD table lookup, 4 table regs	TBL	4	2/3	V	
ASIMD table lookup extension, 1 table reg	TBX	2	2	V	
ASIMD table lookup extension, 2 table reg	TBX	4	1/2	V	
ASIMD table lookup extension, 3 table reg	TBX	6	2/3	V	
ASIMD table lookup extension, 4 table reg	TBX	6	2/5	V	
ASIMD transfer, element to gen reg	UMOV, SMOV	2	1	V1	
ASIMD transfer, gen reg to element	INS	5	1	M, V	
ASIMD transpose	TRN1, TRN2	2	2	V	
ASIMD unzip/zip	UZP1, UZP2, ZIP1, ZIP2	2	2	V	

Table 31: AArch32 ASIMD miscellaneous instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD bitwise insert	VBIF, VBIT, VBSL	2	2	V	-
ASIMD count	VCLS, VCLZ, VCNT	2	2	V	-
ASIMD duplicate, core reg	VDUP	3	1	M	-
ASIMD duplicate, scalar	VDUP	2	2	V	-
ASIMD extract	VEXT	2	2	V	-
ASIMD move, immed	VMOV	2	2	V	-
ASIMD move, register	VMOV	2	2	V	-
ASIMD move, narrowing	VMOVN	2	2	V	-
ASIMD move, saturating	VQMOVN, VQMOVUN	4	1	V1	-
ASIMD reciprocal estimate, D-form	VRECPE, VRSQRTE	3	1	V0	-
ASIMD reciprocal estimate, Q-form	VRECPE, VRSQRTE	4	1/2	V0	-
ASIMD reciprocal step	VRECPS, VRSQRTS	5	2	V	-
ASIMD reverse	VREV16, VREV32, VREV64	2	2	V	-
ASIMD swap	VSWP	4	2/3	V	-
ASIMD table lookup, 1 or 2 table regs	VTBL	2	2	V	-
ASIMD table lookup, 3 table regs	VTBL	4	1/2	V	-
ASIMD table lookup, 4 table regs	VTBL	4	2/3	V	-
ASIMD table lookup extension, 1 reg	VTBX	2	2	V	-
ASIMD table lookup extension, 2 table reg	VTBX	4	1/2	V	-
ASIMD table lookup extension, 3 table reg	VTBX	6	2/3	V	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD table lookup extension, 4 table reg	VTBX	6	2/5	V	-
ASIMD transfer, scalar to core reg, word	VMOV	2	1	V1	-
ASIMD transfer, scalar to core reg, byte/hword	VMOV	3	1	V1, I	-
ASIMD transfer, core reg to scalar	VMOV	5	1	M, V	-
ASIMD transpose	VTRN	4	2/3	V	-
ASIMD unzip/zip	VUZP, VZIP	4	2/3	V	-

3.18 ASIMD load instructions

The latencies shown assume the memory access hits in the Level 1 Data Cache. Compared to standard loads, an extra cycle is required to forward results to FP/ASIMD pipelines.

Table 32: AArch64 ASIMD load instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipeline s	Notes
ASIMD load, 1 element, multiple, 1 reg, D-form	LD1	5	2	L	-
ASIMD load, 1 element, multiple, 1 reg, Q-form	LD1	5	2	L	-
ASIMD load, 1 element, multiple, 2 reg, D-form	LD1	5	1	L	-
ASIMD load, 1 element, multiple, 2 reg, Q-form	LD1	5	1	L	-
ASIMD load, 1 element, multiple, 3 reg, D-form	LD1	6	2/3	L	-
ASIMD load, 1 element, multiple, 3 reg, Q-form	LD1	6	2/3	L	-
ASIMD load, 1 element, multiple, 4 reg, D-form	LD1	6	1/2	L	-
ASIMD load, 1 element, multiple, 4 reg, Q-form	LD1	6	1/2	L	-
ASIMD load, 1 element, one lane, B/H/S	LD1	7	2	L, V	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD load, 1 element, one lane, D	LD1	7	2	L, V	-
ASIMD load, 1 element, all lanes, D-form, B/H/S	LD1R	7	2	L, V	-
ASIMD load, 1 element, all lanes, D-form, D	LD1R	7	2	L, V	-
ASIMD load, 1 element, all lanes, Q-form	LD1R	7	2	L, V	-
ASIMD load, 2 element, multiple, D-form, B/H/S	LD2	7	1	L, V	-
ASIMD load, 2 element, multiple, Q-form, B/H/S	LD2	7	1	L, V	-
ASIMD load, 2 element, multiple, Q-form, D	LD2	7	1	L, V	-
ASIMD load, 2 element, one lane, B/H	LD2	7	1	L, V	-
ASIMD load, 2 element, one lane, S	LD2	7	1	L, V	-
ASIMD load, 2 element, one lane, D	LD2	7	1	L, V	-
ASIMD load, 2 element, all lanes, D-form, B/H/S	LD2R	7	1	L, V	-
ASIMD load, 2 element, all lanes, D-form, D	LD2R	7	1	L, V	-
ASIMD load, 2 element, all lanes, Q-form	LD2R	7	1	L, V	-
ASIMD load, 3 element, multiple, D-form, B/H/S	LD3	8	1/2	L, V	-
ASIMD load, 3 element, multiple, Q-form, B/H/S	LD3	8	1/2	L, V	-
ASIMD load, 3 element, multiple, Q-form, D	LD3	8	1/2	L, V	-
ASIMD load, 3 element, one lane, B/H	LD3	7	1/2	L, V	-
ASIMD load, 3 element, one lane, S	LD3	7	1/2	L, V	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD load, 3 element, one lane, D	LD3	7	1/2	L, V	-
ASIMD load, 3 element, all lanes, D-form, B/H/S	LD3R	7	1/2	L, V	-
ASIMD load, 3 element, all lanes, D-form, D	LD3R	7	1/2	L, V	-
ASIMD load, 3 element, all lanes, Q-form, B/H/S	LD3R	7	1/2	L, V	-
ASIMD load, 3 element, all lanes, Q-form, D	LD3R	7	1/2	L, V	-
ASIMD load, 4 element, multiple, D-form, B/H/S	LD4	8	2/7	L, V	-
ASIMD load, 4 element, multiple, Q-form, B/H/S	LD4	10	1/5	L, V	-
ASIMD load, 4 element, multiple, Q-form, D	LD4	10	1/5	L, V	-
ASIMD load, 4 element, one lane, B/H	LD4	8	1/2	L, V	-
ASIMD load, 4 element, one lane, S	LD4	8	1/2	L, V	-
ASIMD load, 4 element, one lane, D	LD4	8	1/2	L, V	-
ASIMD load, 4 element, all lanes, D-form, B/H/S	LD4R	8	1/2	L, V	-
ASIMD load, 4 element, all lanes, D-form, D	LD4R	8	1/2	L, V	-
ASIMD load, 4 element, all lanes, Q-form, B/H/S	LD4R	8	1/2	L, V	-
ASIMD load, 4 element, all lanes, Q-form, D	LD4R	8	1/2	L, V	-
(ASIMD load, writeback form)	-	(1)	-	+ I	1

Table 33: AArch32 ASIMD load instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD load, 1 element, multiple, 1 reg	VLD1	5	2	L	-
ASIMD load, 1 element, multiple, 2 reg	VLD1	5	2	L	-
ASIMD load, 1 element, multiple, 3 reg	VLD1	5	1	L	-
ASIMD load, 1 element, multiple, 4 reg	VLD1	5	1	L	-
ASIMD load, 1 element, one lane	VLD1	7	2	L, V	-
ASIMD load, 1 element, all lanes, 1 reg	VLD1	7	2	L, V	-
ASIMD load, 1 element, all lanes, 2 reg	VLD1	7	2/3	L, V	-
ASIMD load, 2 element, multiple, 2 reg	VLD2	7	2/3	L, V	-
ASIMD load, 2 element, multiple, 4 reg	VLD2	8	1/2	L, V	-
ASIMD load, 2 element, one lane, size 32	VLD2	7	1	L, V	-
ASIMD load, 2 element, one lane, size 8/16	VLD2	7	1	L, V	-
ASIMD load, 2 element, all lanes	VLD2	7	1	L, V	-
ASIMD load, 3 element, multiple, 3 reg	VLD3	8	2/3	L, V	-
ASIMD load, 3 element, one lane, size 32	VLD3	8	2/3	L, V	-
ASIMD load, 3 element, one lane, size 8/16	VLD3	8	2/3	L, V	-
ASIMD load, 3 element, all lanes	VLD3	8	2/3	L, V	-
ASIMD load, 4 element, multiple, 4 reg	VLD4	8	1/2	L, V	-
ASIMD load, 4 element, one lane, size 32	VLD4	8	1/2	L, V	-

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD load, 4 element, one lane, size 8/16	VLD4	8	1/2	L, V	-
ASIMD load, 4 element, all lanes	VLD4	8	1/2	L, V	-
(ASIMD load, writeback form)	-	(1)	-	+I	1

Note:



1. Writeback forms of load instructions require an extra μ OP to update the base address. This update is typically performed in parallel with the load μ OP (update latency shown in parentheses).

3.19 ASIMD store instructions

Stores MOPs are split into store address and store data μ OPs. Once executed, stores are buffered and committed in the background.

Table 34: AArch64 ASIMD store instructions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD store, 1 element, multiple, 1 reg, D-form	ST1	2	2	L, V	-
ASIMD store, 1 element, multiple, 1 reg, Q-form	ST1	2	1	L, V	-
ASIMD store, 1 element, multiple, 2 reg, D-form	ST1	2	1	L, V	-
ASIMD store, 1 element, multiple, 2 reg, Q-form	ST1	3	1/2	L, V	-
ASIMD store, 1 element, multiple, 3 reg, D-form	ST1	3	2/3	L, V	-
ASIMD store, 1 element, multiple, 3 reg, Q-form	ST1	4	1/3	L, V	-
ASIMD store, 1 element, multiple, 4 reg, D-form	ST1	3	1/2	L, V	-
ASIMD store, 1 element, multiple, 4 reg, Q-form	ST1	5	1/4	L, V	-
ASIMD store, 1 element, one lane, B/H/S	ST1	4	1	V, L	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD store, 1 element, one lane, D	ST1	4	1	V, L	-
ASIMD store, 2 element, multiple, D-form, B/H/S	ST2	4	1	V, L	-
ASIMD store, 2 element, multiple, Q-form, B/H/S	ST2	5	1/2	V, L	-
ASIMD store, 2 element, multiple, Q-form, D	ST2	5	1/2	V, L	-
ASIMD store, 2 element, one lane, B/H/S	ST2	4	1	V, L	-
ASIMD store, 2 element, one lane, D	ST2	4	1	V, L	-
ASIMD store, 3 element, multiple, D-form, B/H/S	ST3	5	1/2	V, L	-
ASIMD store, 3 element, multiple, Q-form, B/H/S	ST3	6	1/3	V, L	-
ASIMD store, 3 element, multiple, Q-form, D	ST3	6	1/3	V, L	-
ASIMD store, 3 element, one lane, B/H	ST3	4	1/2	V, L	-
ASIMD store, 3 element, one lane, S	ST3	4	1/2	V, L	-
ASIMD store, 3 element, one lane, D	ST3	5	1/2	V, L	-
ASIMD store, 4 element, multiple, D-form, B/H/S	ST4	7	1/3	V, L	-
ASIMD store, 4 element, multiple, Q-form, B/H/S	ST4	9	1/6	V, L	-
ASIMD store, 4 element, multiple, Q-form, D	ST4	6	1/4	V, L	-
ASIMD store, 4 element, one lane, B/H	ST4	5	-	V, L	-
ASIMD store, 4 element, one lane, S	ST4	-	2/3	V, L	-
ASIMD store, 4 element, one lane, D	ST4	-	-	V, L	-

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
(ASIMD store, writeback form)	-	(1)	-	Add I	1

Table 35: AArch32 ASIMD store instructions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
ASIMD store, 1 element, multiple, 1 reg	VST1	2	2	L, V	-
ASIMD store, 1 element, multiple, 2 reg	VST1	2	2	L, V	-
ASIMD store, 1 element, multiple, 3 reg	VST1	3	2/3	L, V	-
ASIMD store, 1 element, multiple, 4 reg	VST1	3	1/2	L, V	-
ASIMD store, 1 element, one lane	VST1	4	2	V, L	-
ASIMD store, 2 element, multiple, 2 reg	VST2	4	1	V, L	-
ASIMD store, 2 element, multiple, 4 reg	VST2	5	1/2	V, L	-
ASIMD store, 2 element, one lane	VST2	4	2	V, L	-
ASIMD store, 3 element, multiple, 3 reg	VST3	5	2/3	V, L	-
ASIMD store, 3 element, one lane, size 32	VST3	4	1	V, L	-
ASIMD store, 3 element, one lane, size 8/16	VST3	4	1	V, L	-
ASIMD store, 4 element, multiple, 4 reg	VST4	8	1/2	V, L	-
ASIMD store, 4 element, one lane, size 32	VST4	7	2	V, L	-
ASIMD store, 4 element, one lane, size 8/16	VST4	7	2	V, L	-
(ASIMD store, writeback form)	-	(1)	-	+I	1



Note:

1. Writeback forms of store instructions require an extra μ OP to update the base address. This update is typically performed in parallel with the store μ OP (update latency shown in parentheses).

3.20 Cryptography extensions

Table 36: AArch64 Cryptography extensions

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Crypto AES ops	AESD, AESE, AESIMC, AESMC	2	1	V0	-
Crypto polynomial (64x64) multiply long	PMULL (2)	2	1	V0	-
Crypto SHA1 hash acceleration op	SHA1H	2	1	V0	-
Crypto SHA1 hash acceleration ops	SHA1C, SHA1M, SHA1P	4	1	V0	-
Crypto SHA1 schedule acceleration ops	SHA1SU0, SHA1SU1	2	1	V0	-
Crypto SHA256 hash acceleration ops	SHA256H, SHA256H2	4	1	V0	-
Crypto SHA256 schedule acceleration ops	SHA256SU0, SHA256SU1	2	1	V0	-

Table 37: AArch32 Cryptography extensions

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Crypto AES ops	AESD, AESE, AESIMC, AESMC	2	1	V0	1

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
Crypto polynomial (64x64) multiply long	VMULL.P64	2	1	VO	-
Crypto SHA1 hash acceleration op	SHA1H	2	1	VO	-
Crypto SHA1 hash acceleration ops	SHA1C, SHA1M, SHA1P	4	1	VO	-
Crypto SHA1 schedule acceleration ops	SHA1SU0, SHA1SU1	2	1	VO	-
Crypto SHA256 hash acceleration ops	SHA256H, SHA256H2	4	1	VO	-
Crypto SHA256 schedule acceleration ops	SHA256SU0, SHA256SU1	2	1	VO	-

Note:

1. Adjacent AESE/AESMC instruction pairs and adjacent AESD/AESIMC instruction pairs will exhibit the performance characteristics described in Section 4.6.



3.21 CRC

Table 38: AArch64 CRC

Instruction Group	AArch64 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
CRC checksum ops	CRC32, CRC32C	2	1	M	1

Table 39: AArch32 CRC

Instruction Group	AArch32 Instructions	Exec Latency	Execution Throughput	Utilized Pipelines	Notes
CRC checksum ops	CRC32, CRC32C	2	1	M	1

Note:

1. CRC execution supports late forwarding of the result from a producer μ OP to a consumer μ OP. This results in a 1 cycle reduction in latency as seen by the consumer.



4 Special considerations

4.1 Dispatch constraints

Dispatch of uops from the in-order portion to the out-of-order portion of the microarchitecture includes a number of constraints. It is important to consider these constraints during code generation in order to maximize the effective dispatch bandwidth and subsequent execution bandwidth of Neoverse N1.

The dispatch stage can process up to 4 Mops per cycle and dispatch up to 8 uops per cycle, with the following limitations on the number of uops of each type that may be simultaneously dispatched.

- Up to 2 uops utilizing B pipeline
- Up to 4 uops utilizing S pipelines
- Up to 2 uops utilizing M pipeline
- Up to 2 uops utilizing each of the V pipelines.
- Up to 2 uops utilizing each of the L pipelines

In the event there are more uops available to be dispatched in a given cycle than can be supported by the constraints above, uops will be dispatched in oldest to youngest age-order to the extent allowed by the above.

4.2 Dispatch stall

In the event of a V-pipeline μ OP containing more than 1 quad-word register source, a portion or all of which was previously written as one or multiple single words, that μ OP will stall in dispatch for three cycles. This stall occurs only on the first such instance, and subsequent consumers of the same register will not experience this stall.

4.3 Optimizing general-purpose register spills and fills

Register transfers between general-purpose registers (GPR) and ASIMD registers (VPR) are lower latency than reads and writes to the cache hierarchy, thus it is recommended that GPR registers be filled/spilled to the VPR rather to memory, when possible.

4.4 Optimizing memory copy

To achieve maximum throughput for memory copy (or similar loops), one should do the following.

- Unroll the loop to include multiple load and store operations per iteration, minimizing the overheads of looping.
- Use discrete, non-writeback forms of load and store instructions while interleaving them.
- Align stores on 16B boundary wherever possible.

- The following examples show the recommended instruction sequence for a long memory copy in AArch64 state

For forward copies:

```

Loop_start:
    SUBS    X2,X2,#96
    LDP     Q3,Q4,[x1,#0]
    STP     Q3,Q4,[x0,#0]
    LDP     Q3,Q4,[x1,#32]
    STP     Q3,Q4,[x0,#32]
    LDP     Q3,Q4,[x1,#64]
    STP     Q3,Q4,[x0,#64]
    ADD     X1,X1,#96
    ADD     X0,X0,#96
    BGT     Loop_start
  
```

For backward copies:

```

Loop_start:
    SUBS    X2,X2,#96
    LDP     Q4,Q3,[x1,#-32]
    STR     Q3,.[x0,#-16]
    STR     Q4,.[x0,#-32]
    LDP     Q4,Q3,[x1,#-64]
    STR     Q3,[x0,#-48]
    STR     Q4,[x0,#-64]
    LDP     Q4,Q3,[x1,#-96]
    STR     Q3,[x0,#-80]
    STR     Q4,[x0,#-96]
    SUB     X1,X1,#96
    SUB     X0,X0,#96
    BGT     Loop_start
  
```

A recommended copy routine for AArch32 would look like the sequence above but would use LDRD/STRD instructions. Avoid load-/store-multiple instruction encodings (such as LDM and STM).

4.5 Load/Store alignment

The Armv8.2-A architecture allows many types of load and store accesses to be arbitrarily aligned. The Neoverse N1 handles most unaligned accesses without performance penalties. However, there are cases which reduce bandwidth or incur additional latency, as described below.

- Load operations that cross a cache-line (64-byte) boundary.
- Quad-word load operations that are not 4B aligned.
- Store operations that cross a 16B boundary.

Copyright © 2018-2023 Arm Limited (or its affiliates). All rights reserved.

Non-Confidential

4.6 Store to Load Forwarding

Under most circumstances, the Neoverse N1 processor implements an optimization known as store to load forwarding (STLF). This optimization allows load operations to a memory address which has recently been written to forward directly from an internal processor data structure, reducing load latency times.

On Neoverse N1, this optimization is effective where data of size Z bytes is stored at an address X , and read back at address $X + Y$, where Y is 0 or a multiple of $Z/2$.

For example, if a four byte store is written at address $0x100$ followed by a two byte load from address $0x102$, STLF will be effective. However, if the same store is followed by a two byte load starting at $0x101$, STLF will not trigger.

Consequently, stores and loads that are aligned to addresses that match the data size they write or read will more frequently satisfy the STLF conditions.

To enable the store to load forwarding optimization, it is recommended that data expected to be read back with close temporal locality is written and read at element aligned addresses.

4.7 AES encryption/decryption

Neoverse N1 can issue one AESE/AESMC/AESD/AESIMC instruction every cycle (fully pipelined) with an execution latency of two cycles. This means encryption or decryption for at least two data chunks should be interleaved for maximum performance:

```
AESE  data0, key0
AESMC data0, data0
AESE  data1, key0
AESMC data1, data1
AESE  data0, key1
AESMC data0, data0
AESE  data1, key1
AESMC data1, data1
...
```

Pairs of dependent AESE/AESMC and AESD/AESIMC instructions are higher performance when they are adjacent in the program code and both instructions use the same destination register.

4.8 Region based fast forwarding

The forwarding logic in the V pipelines is optimized to provide optimal latency for instructions which are expected to commonly forward to one another. The effective latency of FP and ASIMD instructions as described in section 3 is increased by one cycle if the producer and consumer instructions are not part of the same forwarding region. These optimized forwarding regions are defined in the following table.

Table 40: Optimized forwarding regions

Region	Instruction Types	Notes
1	ASIMD ALU, ASIMD shift, ASIMD/ scalar insert and move, ASIMD abs/cmp/max/min and the ASIMD miscellaneous instructions in table 30 and table 31.	1
2	FP multiply, FP multiply-accumulate, FP compare, FP add/sub and the ASIMD miscellaneous instructions in table 30 and table 31.	1,2,3
3	Crypto SHA1/SHA256.	

Notes:

1. Reciprocal step and estimate instructions are excluded from this region.
2. ASIMD extract narrow, saturating instructions are excluded from this region.
3. ASIMD miscellaneous instructions can only be consumers of this region.

The following instructions are not a part of any region:

- FP div/sqrt
- FP convert and rounding
- ASIMD integer mul/mac
- ASIMD reduction.

In addition to the regions mentioned in the table above, all floating point and ASIMD instructions can fast forward to FP and ASIMD stores.

More special notes about the forwarding region in table 40:

- Fast forwarding will not occur in AArch32 mode if the consuming register's width is greater than that of the producer.
- Element sources used by FP multiply and multiply-accumulate operations cannot be consumers.
- Complex ASIMD shift by immediate/register and shift accumulate instructions cannot be producers (see section 3.14) in region 1.
- ASIMD extract narrow, saturating instructions cannot be producers (see section 3.16) in region 1.
- ASIMD absolute difference accumulate and pairwise add and accumulate instructions cannot be producers (see section 3.14) in region 1.
- For FP producer-consumer pairs, the precision of the instructions should match (single, double or half) in region 2.
- Pair-wise FP instructions cannot be producers or consumers in region 2.

It is not advisable to interleave instructions belonging to different regions. Also, certain instructions can only be producers or consumers in a particular region but not both (see notes for table 40). For example, the code below interleaves producers and consumers from regions 1 and 2. This will result in an additional latency of 1 cycle as seen by FMUL.

FSUB v27.2s, v28.2s, v20.2s – Region 2

FADD v20.2s, v28.2s, v20.2s – Region 2

MOV v27.s[1], v20.s[1] - Region 2 producer but not a region 2 consumer

FMUL v26.2s, v27.2s, v6.2s – Region 2

4.9 Branch instruction alignment

Branch instruction and branch target instruction alignment and density can affect performance. For best-case performance, consider the following guidelines.

- Avoid placing more than 4 branch instructions within an aligned 32-byte instruction memory region.
- When possible, a branch and its target should be located within the same 2M aligned memory region.

Consider aligning subroutine entry points and branch targets to 32B boundaries, within the bounds of the code-density requirements of the program. This will ensure that the subsequent fetch can maximize bandwidth following the taken branch by bringing in all useful instructions

For loops which comprise 32 or fewer instruction bytes, it is preferred that the loop be located entirely within a single aligned 32-byte instruction memory region.

4.10 FPCR self-synchronization

Programmers and compiler writers should note that writes to the FPCR register are self-synchronizing, i.e. its effect on subsequent instructions can be relied upon without an intervening context synchronizing operation.

4.11 Special register access

The Neoverse N1 performs register renaming for general purpose registers to enable speculative and out-of-order instruction execution. But most special-purpose registers are not renamed. Instructions that read or write non-renamed registers are subjected to one or more of the following additional execution constraints.

- Non-Speculative Execution – Instructions may only execute non-speculatively.
- In-Order Execution – Instructions must execute in-order with respect to other similar instructions or in some cases all instructions.
- Flush Side-Effects – Instructions trigger a flush side-effect after executing for synchronization.

The table below summarizes various special-purpose register read accesses and the associated execution constraints or side-effects.

Table 41: Special-purpose register read accesses

Register Read	Non-Speculative	In-Order	Flush Side-Effect	Notes
APSR	Yes	Yes	No	3

Copyright © 2018-2023 Arm Limited (or its affiliates). All rights reserved.

Non-Confidential

Register Read	Non-Speculative	In-Order	Flush Side-Effect	Notes
CurrentEL	No	Yes	No	-
DAIF	No	Yes	No	-
DLR_ELO	No	Yes	No	-
DSPSR_ELO	No	Yes	No	-
ELR_*	No	Yes	No	-
FPCR	No	Yes	No	-
FPSCR	Yes	Yes	No	2
FPSR	Yes	Yes	No	2
NZCV	No	No	No	1
SP_*	No	No	No	1
SPSel	No	Yes	No	-
SPSR_*	No	Yes	No	-

Note:

1. The NZCV and SP registers are fully renamed.
2. FPSR/FPSCR reads must wait for all prior instructions that may update the status flags to execute and retire.
3. APSR reads must wait for all prior instructions that may set the Q bit to execute and retire.

The table below summarizes various special-purpose register write accesses and the associated execution constraints or side-effects.

Table 42: Special-purpose register write accesses

Register Write	Non-Speculative	In-Order	Flush Side-Effect	Notes
APSR	Yes	Yes	No	4
DAIF	Yes	Yes	No	-
DLR_ELO	Yes	Yes	No	-
DSPSR_ELO	Yes	Yes	No	-
ELR_*	Yes	Yes	No	-
FPCR	Yes	Yes	Maybe	2
FPSCR	Yes	Yes	Maybe	2, 3
FPSR	Yes	Yes	No	3
NZCV	No	No	No	1

Register Write	Non-Speculative	In-Order	Flush Side-Effect	Notes
SP_*	No	No	No	1
SPSel	Yes	Yes	Yes	-
SPSR_*	Yes	Yes	No	-

Note:



1. The NZCV and SP registers are fully renamed.
2. If the FPCR/FPSCR write is predicted to change the control field values, it will introduce a barrier which prevents subsequent instructions from executing. If the FPCR/FPSCR write is predicted to not change the control field values, it will execute without a barrier but trigger a flush if the values change.
3. FPSR/FPSCR writes must stall at dispatch if another FPSR/FPSCR write is still pending.
4. APSR writes that set the Q bit will introduce a barrier which prevents subsequent instructions from executing until the write completes.

4.12 Register forwarding hazards

The Armv8-A architecture allows FP/ASIMD instructions to read and write 32-bit S-registers. In AArch32, Each S-register corresponds to one half (upper or lower) of an overlaid 64-bit D-register. A Q register in turn consists of two overlaid D register. Register forwarding hazards may occur when one μ OP reads a Q-register operand that has recently been written with one or more S-register result. Consider the following scenario.

```
VADD    S0, S1, S2
VADD    Q6, Q5, Q0
```

The first instruction writes S0, which correspond to the lowest part of Q0. The second instruction then requires Q0 as an input operand. In this scenario, there is a RAW dependency between the first and the second instructions. In most cases, Neoverse N1 performs slightly worse in such situations.

Neoverse N1 is able to avoid this register-hazard condition for certain cases. The following rules describe the conditions under which a register-hazard can occur.

- The producer writes an S-register (not a D[x] scalar)
- The consumer reads an overlapping Q-register (not as a D[x] scalar)
- The consumer is a FP/ASIMD μ OP (not a store or MOV μ OP)

To avoid unnecessary hazards, it is recommended that the programmer use D[x] scalar writes when populating registers prior to ASIMD operations. For example, either of the following instruction forms would safely prevent a subsequent hazard.

```
VLD1.32 D0[x], [address]
VADD    Q1, Q0, Q2F
```

4.13 IT blocks

The Armv8-A architecture performance deprecates some uses of the IT instruction in such a way that software may be written using multiple naïve single instruction IT blocks. It is preferred that software instead generate multi instruction IT blocks rather than single instruction blocks.