

IA-64 Application Instruction Set Architecture Guide

Revision 1.0



HEWLETT®
PACKARD

THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Information in this document is provided in connection with Intel/Hewlett-Packard products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel/Hewlett-Packard's Terms and Conditions of Sale for such products, Intel/Hewlett-Packard assumes no liability whatsoever, and Intel/Hewlett-Packard disclaims any express or implied warranty, relating to sale and/or use of Intel/Hewlett-Packard products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel/Hewlett-Packard products are not intended for use in medical, life saving, or life sustaining applications.

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

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

IA-64 processors may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copyright © Intel Corporation / Hewlett-Packard Company, 1999

*Third-party brands and names are the property of their respective owners.

Table of Contents

1	About the IA-64 Application ISA Guide	1-1
1.1	Overview of IA-64 Application Instruction Set Architecture (ISA) Guide	1-1
1.2	Terminology	1-1
2	Introduction to the IA-64 Processor Architecture	2-1
2.1	IA-64 Operating Environments	2-1
2.2	Instruction Set Transition Model Overview	2-2
2.3	PA-RISC Compatibility	2-2
2.4	IA-64 Instruction Set Features	2-2
2.5	Instruction Level Parallelism	2-3
2.6	Compiler to Processor Communication	2-3
2.7	Speculation	2-3
2.7.1	Control Speculation	2-3
2.7.2	Data Speculation	2-4
2.8	Predication	2-4
2.9	Register Stack	2-5
2.10	Branching	2-5
2.11	Register Rotation	2-5
2.12	Floating-point Architecture	2-5
2.13	Multimedia Support	2-6
3	IA-64 Execution Environment	3-1
3.1	Application Register State	3-1
3.1.1	Reserved and Ignored Registers	3-1
3.1.2	General Registers	3-3
3.1.3	Floating-Point Registers	3-3
3.1.4	Predicate Registers	3-3
3.1.5	Branch Registers	3-3
3.1.6	Instruction Pointer	3-3
3.1.7	Current Frame Marker	3-4
3.1.8	Application Registers	3-4
3.1.9	Performance Monitor Data Registers (PMD)	3-8
3.1.10	User Mask (UM)	3-8
3.1.11	Processor Identification Registers	3-8
3.2	Memory	3-10
3.2.1	Application Memory Addressing Model	3-10
3.2.2	Addressable Units and Alignment	3-10
3.2.3	Byte Ordering	3-10
3.3	Instruction Encoding Overview	3-11
3.4	Instruction Sequencing	3-12
4	IA-64 Application Programming Model	4-1
4.1	Register Stack	4-1
4.1.1	Register Stack Operation	4-1
4.1.2	Register Stack Instructions	4-3
4.2	Integer Computation Instructions	4-3
4.2.1	Arithmetic Instructions	4-3
4.2.2	Logical Instructions	4-4
4.2.3	32-Bit Addresses and Integers	4-4
4.2.4	Bit Field and Shift Instructions	4-4
4.2.5	Large Constants	4-5
4.3	Compare Instructions and Predication	4-5
4.3.1	Predication	4-6
4.3.2	Compare Instructions	4-6
4.3.3	Compare Types	4-6
4.3.4	Predicate Register Transfers	4-8
4.4	Memory Access Instructions	4-8

4.4.1	Load Instructions	4-9
4.4.2	Store Instructions	4-9
4.4.3	Semaphore Instructions	4-10
4.4.4	Control Speculation	4-10
4.4.5	Data Speculation	4-12
4.4.6	Memory Hierarchy Control and Consistency	4-16
4.4.7	Memory Access Ordering	4-18
4.5	Branch Instructions.	4-19
4.5.1	Modulo-Scheduled Loop Support	4-20
4.5.2	Branch Prediction Hints	4-22
4.6	Multimedia Instructions	4-23
4.6.1	Parallel Arithmetic	4-23
4.6.2	Parallel Shifts.	4-24
4.6.3	Data Arrangement	4-24
4.7	Register File Transfers	4-24
4.8	Character Strings and Population Count	4-25
4.8.1	Character Strings	4-25
4.8.2	Population Count	4-26
5	IA-64 Floating-point Programming Model	5-1
5.1	Data Types and Formats	5-1
5.1.1	Real Types	5-1
5.1.2	Floating-point Register Format	5-1
5.1.3	Representation of Values in Floating-point Registers	5-2
5.2	Floating-point Status Register	5-4
5.3	Floating-point Instructions	5-6
5.3.1	Memory Access Instructions	5-6
5.3.2	Floating-Point Register to/from General Register Transfer Instructions	5-11
5.3.3	Arithmetic Instructions	5-12
5.3.4	Non-Arithmetic Instructions	5-13
5.3.5	Floating-point Status Register (FPSR) Status Field Instructions	5-14
5.3.6	Integer Multiply and Add Instructions	5-14
5.4	Additional IEEE Considerations	5-15
5.4.1	Definition of SNaNs, QNaNs, and Propagation of NaNs	5-15
5.4.2	IEEE Standard Mandated Operations Deferred to Software	5-15
5.4.3	Additions beyond the IEEE Standard	5-15
6	IA-64 Instruction Reference	6-1
6.1	Instruction Page Conventions	6-1
6.2	Instruction Descriptions	6-2
A	Instruction Sequencing Considerations	A-1
A.1	RAW Ordering Exceptions	A-2
A.2	WAW Ordering Exceptions.	A-3
A.3	WAR Ordering Exceptions	A-3
B	IA-64 Pseudo-Code Functions	B-1
C	IA-64 Instruction Formats	C-1
C.1	Format Summary	C-2
C.2	A-Unit Instruction Encodings.	C-8
C.2.1	Integer ALU	C-8
C.2.2	Integer Compare	C-10
C.2.3	Multimedia	C-13
C.3	I-Unit Instruction Encodings	C-16
C.3.1	Multimedia and Variable Shifts	C-16
C.3.2	Integer Shifts	C-20
C.3.3	Test Bit	C-22
C.3.4	Miscellaneous I-Unit Instructions	C-23
C.3.5	GR/BR Moves	C-24
C.3.6	GR/Predicate/IP Moves	C-24
C.3.7	GR/AR Moves (I-Unit).	C-25

C.3.8	Sign/Zero Extend/Compute Zero Index	C-25
C.4	M-Unit Instruction Encodings	C-26
C.4.1	Loads and Stores.	C-26
C.4.2	Line Prefetch	C-38
C.4.3	Semaphores	C-39
C.4.4	Set/Get FR.	C-39
C.4.5	Speculation and Advanced Load Checks	C-40
C.4.6	Cache/Synchronization/RSE/ALAT.	C-41
C.4.7	GR/AR Moves (M-Unit)	C-42
C.4.8	Miscellaneous M-Unit Instructions	C-42
C.4.9	Memory Management	C-43
C.5	B-Unit Instruction Encodings	C-45
C.5.1	Branches	C-45
C.5.2	Nop	C-48
C.5.3	Miscellaneous B-Unit Instructions	C-49
C.6	F-Unit Instruction Encodings	C-49
C.6.1	Arithmetic	C-51
C.6.2	Parallel Floating-point Select	C-52
C.6.3	Compare and Classify	C-52
C.6.4	Approximation.	C-53
C.6.5	Minimum/Maximum and Parallel Compare.	C-54
C.6.6	Merge and Logical.	C-54
C.6.7	Conversion.	C-55
C.6.8	Status Field Manipulation	C-55
C.6.9	Miscellaneous F-Unit Instructions	C-56
C.7	X-Unit Instruction Encodings	C-56
C.7.1	Miscellaneous X-Unit Instructions	C-56
C.7.2	Move Long Immediate ₆₄	C-57
C.8	Immediate Formation.	C-57

List of Figures

Figure 3-1.	Application Register Model	3-2
Figure 3-2.	Frame Marker Format	3-4
Figure 3-3.	RSC Format	3-6
Figure 3-4.	BSP Register Format.	3-6
Figure 3-5.	BSPSTORE Register Format	3-6
Figure 3-6.	RNAT Register Format	3-6
Figure 3-7.	PFS Format	3-7
Figure 3-8.	Epilog Count Register Format	3-8
Figure 3-9.	User Mask Format	3-8
Figure 3-10.	CPUID Registers 0 and 1 – Vendor Information	3-9
Figure 3-11.	CPUID Register 2 – Processor Serial Number	3-9
Figure 3-12.	CPUID Register 3 – Version Information	3-9
Figure 3-13.	CPUID Register 4 – General Features/Capability Bits	3-9
Figure 3-14.	Little-endian Loads	3-10
Figure 3-15.	Big-endian Loads	3-11
Figure 3-16.	Bundle Format	3-11
Figure 4-1.	Register Stack Behavior on Procedure Call and Return.	4-2
Figure 4-2.	Data Speculation Recovery Using ld.c	4-13
Figure 4-3.	Data Speculation Recovery Using chk.a.	4-14
Figure 4-4.	Memory Hierarchy.	4-16
Figure 4-5.	Allocation Paths Supported in the Memory Hierarchy	4-17
Figure 5-1.	Floating-point Register Format	5-2
Figure 5-2.	Floating-point Status Register Format.	5-4
Figure 5-3.	Floating-point Status Field Format	5-4
Figure 5-4.	Memory to Floating-point Register Data Translation – Single Precision	5-7
Figure 5-5.	Memory to Floating-point Register Data Translation – Double Precision	5-8
Figure 5-6.	Memory to Floating-point Register Data Translation – Double Extended, Integer and Fill	5-9
Figure 5-7.	Floating-point Register to Memory Data Translation	5-10
Figure 5-8.	Spill/Fill and Double-Extended (80-bit) Floating-point Memory Formats	5-11
Figure 6-1.	Add Pointer	6-4
Figure 6-2.	Stack Frame	6-5
Figure 6-3.	Operation of br.ctop and br.cexit	6-10
Figure 6-4.	Operation of br.wtop and br.wexit.	6-11
Figure 6-5.	Deposit Example.	6-27
Figure 6-6.	Extract Example	6-28
Figure 6-7.	Floating-point Merge Negative Sign Operation	6-49
Figure 6-8.	Floating-point Merge Sign Operation	6-49
Figure 6-9.	Floating-point Merge Sign and Exponent Operation	6-49
Figure 6-10.	Floating-point Mix Left	6-52
Figure 6-11.	Floating-point Mix Right	6-52
Figure 6-12.	Floating-point Mix Left-Right.	6-52
Figure 6-13.	Floating-point Pack	6-63
Figure 6-14.	Floating-point Merge Negative Sign Operation	6-73
Figure 6-15.	Floating-point Merge Sign Operation	6-73
Figure 6-16.	Floating-point Merge Sign and Exponent Operation	6-73
Figure 6-17.	Floating-point Swap	6-95
Figure 6-18.	Floating-point Swap Negate Left or Right.	6-95
Figure 6-19.	Floating-point Sign Extend Left	6-97
Figure 6-20.	Floating-point Sign Extend Right	6-97
Figure 6-21.	Function of getf.exp	6-99
Figure 6-22.	Function of getf.sig	6-99
Figure 6-23.	Mix Example.	6-114
Figure 6-24.	Mux1 Operation (8-bit elements)	6-127
Figure 6-25.	Mux2 Examples (16-bit elements).	6-128

Figure 6-26.	Pack Operation	6-132
Figure 6-27.	Parallel Add Examples	6-134
Figure 6-28.	Parallel Average Example	6-137
Figure 6-29.	Parallel Average with Round Away from Zero Example.	6-138
Figure 6-30.	Parallel Average Subtract Example.	6-140
Figure 6-31.	Parallel Compare Example	6-142
Figure 6-32.	Parallel Maximum Example	6-144
Figure 6-33.	Parallel Minimum Example.	6-145
Figure 6-34.	Parallel Multiply Operation	6-146
Figure 6-35.	Parallel Multiply and Shift Right Operation	6-147
Figure 6-36.	Parallel Sum of Absolute Difference Example	6-149
Figure 6-37.	Parallel Shift Left Example	6-150
Figure 6-38.	Parallel Subtract Example.	6-154
Figure 6-39.	Function of setf.exp	6-158
Figure 6-40.	Function of setf.sig	6-158
Figure 6-41.	Shift Left and Add Pointer	6-162
Figure 6-42.	Shift Right Pair	6-164
Figure 6-43.	Unpack Operation.	6-178
Figure C-1.	Bundle FormatC-1

List of Tables

Table 2-1.	IA-64 Processor Operating Environments	2-1
Table 3-1.	Reserved and Ignored Registers and Fields	3-2
Table 3-2.	Frame Marker Field Description.	3-4
Table 3-3.	Application Registers	3-5
Table 3-4.	RSC Field Description	3-6
Table 3-5.	PFS Field Description	3-7
Table 3-6.	User Mask Field Descriptions	3-8
Table 3-7.	CPUID Register 3 Fields.	3-9
Table 3-8.	Relationship Between Instruction Type and Execution Unit Type	3-11
Table 3-9.	Template Field Encoding and Instruction Slot Mapping	3-12
Table 4-1.	Architectural Visible State Related to the Register Stack	4-3
Table 4-2.	Register Stack Management Instructions	4-3
Table 4-3.	Integer Arithmetic Instructions	4-4
Table 4-4.	Integer Logical Instructions	4-4
Table 4-5.	32-bit Pointer and 32-bit Integer Instructions	4-4
Table 4-6.	Bit Field and Shift Instructions	4-5
Table 4-7.	Instructions to Generate Large Constants	4-5
Table 4-8.	Compare Instructions	4-6
Table 4-9.	Compare Type Function	4-7
Table 4-10.	Compare Outcome with NaT Source Input	4-7
Table 4-11.	Instructions and Compare Types Provided	4-7
Table 4-12.	Memory Access Instructions.	4-9
Table 4-13.	State Relating to Memory Access	4-10
Table 4-14.	State Related to Control Speculation	4-12
Table 4-15.	Instructions Related to Control Speculation	4-12
Table 4-16.	State Relating to Data Speculation.	4-16
Table 4-17.	Instructions Relating to Data Speculation	4-16
Table 4-18.	Locality Hints Specified by Each Instruction Class	4-17
Table 4-19.	Memory Hierarchy Control Instructions and Hint Mechanisms.	4-18
Table 4-20.	Memory Ordering Rules.	4-19
Table 4-21.	Memory Ordering Instructions.	4-19
Table 4-22.	Branch Types	4-20
Table 4-23.	State Relating to Branching	4-20
Table 4-24.	Instructions Relating to Branching	4-20
Table 4-25.	Instructions that Modify RRBs	4-21
Table 4-26.	Whether Prediction Hint on Branches	4-22
Table 4-27.	Sequential Prefetch Hint on Branches	4-22
Table 4-28.	Predictor Deallocation Hint	4-22
Table 4-29.	Parallel Arithmetic Instructions	4-23
Table 4-30.	Parallel Shift Instructions	4-24
Table 4-31.	Parallel Data Arrangement Instructions	4-24
Table 4-32.	Register File Transfer Instructions.	4-25
Table 4-33.	String Support Instructions	4-26
Table 5-1.	IEEE Real-Type Properties	5-1
Table 5-2.	Floating-point Register Encodings	5-2
Table 5-3.	Floating-point Status Register Field Description	5-4
Table 5-4.	Floating-point Status Register's Status Field Description	5-4
Table 5-5.	Floating-point Rounding Control Definitions	5-5
Table 5-6.	Floating-point Computation Model Control Definitions	5-5
Table 5-7.	Floating-point Memory Access Instructions.	5-6
Table 5-8.	Floating-point Register Transfer Instructions	5-11
Table 5-9.	General Register (Integer) to Floating-point Register Data Translation.	5-12
Table 5-10.	Floating-point Register to General Register (Integer) Data Translation.	5-12
Table 5-11.	Floating-point Instruction Status Field Specifier Definition.	5-12

Table 5-12.	Floating-point Arithmetic Instructions	5-12
Table 5-13.	Floating-point Pseudo-Operations	5-13
Table 5-14.	Non-Arithmetic Floating-point Instructions	5-13
Table 5-15.	FPSR Status Field Instructions	5-14
Table 5-16.	Integer Multiply and Add Instructions	5-14
Table 6-1.	Instruction Page Description	6-1
Table 6-2.	Instruction Page Font Conventions	6-1
Table 6-3.	Register File Notation	6-1
Table 6-4.	C Syntax Differences	6-2
Table 6-5.	Branch Types	6-8
Table 6-6.	Branch Whether Hint	6-11
Table 6-7.	Sequential Prefetch Hint	6-11
Table 6-8.	Branch Cache Deallocation Hint	6-12
Table 6-9.	ALAT Clear Completer	6-16
Table 6-10.	Comparison Types	6-19
Table 6-11.	64-bit Comparison Relations for Normal and unc Compares.	6-20
Table 6-12.	64-bit Comparison Relations for Parallel Compares	6-20
Table 6-13.	Immediate Range for 32-bit Compares	6-22
Table 6-14.	Memory Compare and Exchange Size	6-24
Table 6-15.	Compare and Exchange Semaphore Types	6-24
Table 6-16.	Result Ranges for <i>czx</i>	6-26
Table 6-17.	Specified <i>pc</i> Mnemonic Values	6-30
Table 6-18.	<i>sf</i> Mnemonic Values	6-30
Table 6-19.	Floating-point Class Relations	6-37
Table 6-20.	Floating-point Classes.	6-37
Table 6-21.	Floating-point Comparison Types	6-40
Table 6-22.	Floating-point Comparison Relations.	6-40
Table 6-23.	Fetch and Add Semaphore Types	6-45
Table 6-24.	Floating-point Parallel Comparison Results	6-66
Table 6-25.	Floating-point Parallel Comparison Relations	6-66
Table 6-26.	<i>sz</i> Completers	6-101
Table 6-27.	Load Types	6-101
Table 6-28.	Load Hints.	6-102
Table 6-29.	<i>fsz</i> Completers.	6-105
Table 6-30.	FP Load Types	6-105
Table 6-31.	<i>lftype</i> Mnemonic Values	6-110
Table 6-32.	<i>lfhint</i> Mnemonic Values.	6-110
Table 6-33.	Indirect Register File Mnemonics.	6-122
Table 6-34.	Mux Permutations for 8-bit Elements.	6-127
Table 6-35.	Pack Saturation Limits	6-132
Table 6-36.	Parallel Add Saturation Completers	6-134
Table 6-37.	Parallel Add Saturation Limits	6-134
Table 6-38.	Pcmp Relations	6-142
Table 6-39.	PMPYSHR Shift Options	6-147
Table 6-40.	Parallel Subtract Saturation Completers	6-154
Table 6-41.	Parallel Subtract Saturation Limits	6-154
Table 6-42.	Store Types	6-166
Table 6-43.	Store Hints	6-166
Table 6-44.	<i>xsz</i> Mnemonic Values	6-171
Table 6-45.	Test Bit Relations for Normal and unc tbits	6-173
Table 6-46.	Test Bit Relations for Parallel tbits	6-173
Table 6-47.	Test NaT Relations for Normal and unc tnats	6-175
Table 6-48.	Test NaT Relations for Parallel tnats	6-175
Table 6-49.	Memory Exchange Size.	6-180
Table B-1.	Pseudo-Code Functions	B-1
Table C-1.	Relationship Between Instruction Type and Execution Unit Type	C-1
Table C-2.	Template Field Encoding and Instruction Slot Mapping	C-2

Table C-3.	Major Opcode Assignments	C-3
Table C-4.	Instruction Format Summary	C-4
Table C-5.	Instruction Field Color Key	C-6
Table C-6.	Instruction Field Names	C-7
Table C-7.	Special Instruction Notations	C-7
Table C-8.	Integer ALU 2-bit+1-bit Opcode Extensions	C-8
Table C-9.	Integer ALU 4-bit+2-bit Opcode Extensions	C-8
Table C-10.	Integer Compare Opcode Extensions	C-10
Table C-11.	Integer Compare Immediate Opcode Extensions	C-10
Table C-12.	Multimedia ALU 2-bit+1-bit Opcode Extensions	C-13
Table C-13.	Multimedia ALU Size 1 4-bit+2-bit Opcode Extensions	C-14
Table C-14.	Multimedia ALU Size 2 4-bit+2-bit Opcode Extensions	C-14
Table C-15.	Multimedia ALU Size 4 4-bit+2-bit Opcode Extensions	C-15
Table C-16.	Multimedia and Variable Shift 1-bit Opcode Extensions	C-16
Table C-17.	Multimedia Max/Min/Mix/Pack/Unpack Size 1 2-bit Opcode Extensions	C-16
Table C-18.	Multimedia Multiply/Shift/Max/Min/Mix/Pack/Unpack Size 2 2-bit Opcode Extensions	C-17
Table C-19.	Multimedia Shift/Mix/Pack/Unpack Size 4 2-bit Opcode Extensions	C-17
Table C-20.	Variable Shift 2-bit Opcode Extensions	C-18
Table C-21.	Integer Shift/Test Bit/Test NaT 2-bit Opcode Extensions	C-20
Table C-22.	Deposit Opcode Extensions	C-20
Table C-23.	Test Bit Opcode Extensions	C-22
Table C-24.	Misc I-Unit 3-bit Opcode Extensions	C-23
Table C-25.	Misc I-Unit 6-bit Opcode Extensions	C-23
Table C-26.	Integer Load/Store/Semaphore/Get FR 1-bit Opcode Extensions	C-26
Table C-27.	Floating-point Load/Store/Load Pair/Set FR 1-bit Opcode Extensions	C-26
Table C-28.	Integer Load/Store Opcode Extensions	C-26
Table C-29.	Integer Load +Reg Opcode Extensions	C-27
Table C-30.	Integer Load/Store +Imm Opcode Extensions	C-27
Table C-31.	Semaphore/Get FR Opcode Extensions	C-28
Table C-32.	Floating-point Load/Store/Lfetch Opcode Extensions	C-28
Table C-33.	Floating-point Load/Lfetch +Reg Opcode Extensions	C-29
Table C-34.	Floating-point Load/Store/Lfetch +Imm Opcode Extensions	C-29
Table C-35.	Floating-point Load Pair/Set FR Opcode Extensions	C-30
Table C-36.	Floating-point Load Pair +Imm Opcode Extensions	C-30
Table C-37.	Load Hint Completer	C-30
Table C-38.	Store Hint Completer	C-30
Table C-39.	Line Prefetch Hint Completer	C-38
Table C-40.	Opcode 0 Memory Management 3-bit Opcode Extensions	C-43
Table C-41.	Opcode 0 Memory Management 4-bit+2-bit Opcode Extensions	C-44
Table C-42.	Opcode 1 Memory Management 3-bit Opcode Extensions	C-44
Table C-43.	Opcode 1 Memory Management 6-bit Opcode Extensions	C-44
Table C-44.	IP-Relative Branch Types	C-45
Table C-45.	Indirect/Miscellaneous Branch Opcode Extensions	C-46
Table C-46.	Indirect Branch Types	C-46
Table C-47.	Indirect Return Branch Types	C-46
Table C-48.	Sequential Prefetch Hint Completer	C-46
Table C-49.	Branch Whether Hint Completer	C-47
Table C-50.	Indirect Call Whether Hint Completer	C-47
Table C-51.	Branch Cache Deallocation Hint Completer	C-47
Table C-52.	Indirect Predict/Nop Opcode Extensions	C-48
Table C-53.	Miscellaneous Floating-point 1-bit Opcode Extensions	C-49
Table C-54.	Opcode 0 Miscellaneous Floating-point 6-bit Opcode Extensions	C-50
Table C-55.	Opcode 1 Miscellaneous Floating-point 6-bit Opcode Extensions	C-50
Table C-56.	Reciprocal Approximation 1-bit Opcode Extensions	C-50
Table C-57.	Floating-point Status Field Completer	C-51
Table C-58.	Floating-point Arithmetic 1-bit Opcode Extensions	C-51
Table C-59.	Fixed-point Multiply Add and Select Opcode Extensions	C-51

Table C-60.	Floating-point Compare Opcode Extensions	C-52
Table C-61.	Floating-point Class 1-bit Opcode Extensions	C-52
Table C-62.	Misc X-Unit 3-bit Opcode Extensions	C-56
Table C-63.	Misc X-Unit 6-bit Opcode Extensions	C-56
Table C-64.	Move Long 1-bit Opcode Extensions	C-57
Table C-65.	Immediate Formation	C-57

The Intel Architecture – 64-bit (IA-64) is a unique combination of innovative features, such as explicit parallelism, predication, speculation and much more. The architecture is designed to be highly scalable to fill the ever increasing performance requirements of various server and workstation market segments. The IA-64 architecture features a revolutionary 64-bit instruction set architecture (ISA) which applies a new processor architecture technology called EPIC, or Explicitly Parallel Instruction Computing. A key feature of the IA-64 architecture is IA-32 instruction set compatibility.

This document provides a comprehensive description of IA-64 architecture exposed to application software. This includes information on application level resources (registers, etc), application environment, detailed application (non-privileged) instruction description, format and encoding.

1.1 Overview of IA-64 Application Instruction Set Architecture (ISA) Guide

Chapter 1, “About the IA-64 Application ISA Guide”. Gives an overview of this guide.

Chapter 2, “Introduction to the IA-64 Processor Architecture”. Provides an overview of key features of IA-64 architecture.

Chapter 3, “IA-64 Execution Environment”. Describes the IA-64 application architectural state (registers, memory, etc).

Chapter 4, “IA-64 Application Programming Model”. Describes the IA-64 architecture from the perspective of the application programmer. IA-64 instructions are grouped into related functions and an overview of their behavior is given.

Chapter 5, “IA-64 Floating-point Programming Model”. This chapter provides a description of IA-64 floating-point registers, data types and formats and floating-point instructions.

Chapter 6, “IA-64 Instruction Reference”. Provides detailed description of IA-64 application instructions, operation, and instruction format.

Appendix A, "Instruction Sequencing Considerations". Describes the details of instruction sequencing in IA-64 architecture.

Appendix B, "IA-64 Pseudo-Code Functions". Describes pseudo-code functions used in Chapter 6, “IA-64 Instruction Reference”.

Appendix C, "IA-64 Instruction Formats". Describes the encoding and instruction format of instructions covered in Chapter 6, “IA-64 Instruction Reference”.

1.2 Terminology

The following definitions are for terms related to the IA-64 architecture and will be used in the rest of this document:

- **Instruction Set Architecture (ISA)** – Defines application and system level resources. These resources include instructions and registers.
- **IA-64 Architecture** – IA-64 defines the architectural extensions defined in the new ISA including 64-bit instruction capabilities, new performance-enhancing features and support for IA-32 instruction set.
- **IA-32 Architecture** – The 32-bit and 16-bit Intel Architecture as described in the *Intel Architecture Software Developer's Manual*.
- **IA-64 Processor** – An Intel 64-bit processor that implements both the IA-64 and the IA-32 instruction sets.

- **IA-64 System Environment** – IA-64 operating system privileged environment that supports the execution of both IA-64 and IA-32 code.
- **IA-32 System Environment** – Operating system privileged environment and resources as defined by the *Intel Architecture Software Developer's Manual*. Resources include virtual paging, control registers, debugging, performance monitoring, machine checks, and the set of privileged instructions.

2 Introduction to the IA-64 Processor Architecture

The IA-64 architecture was designed to overcome the performance limitations of traditional architectures and provide maximum headroom for the future. To achieve this, IA-64 was designed with an array of innovative features to extract greater instruction level parallelism including: speculation, predication, large register files, a register stack, advanced branch architecture, and many others. 64-bit memory addressability was added to meet the increasing large memory footprint requirements of data warehousing, E-business, and other high performance server and workstation applications.

The IA-64 architecture also provides binary compatibility with the IA-32 instruction set. IA-64 processors can run IA-32 applications on an IA-64 operating system that supports execution of IA-32 applications. IA-64 processors can run IA-32 application binaries on IA-32 legacy operating systems assuming the platform and firmware support exists in the system. The IA-64 architecture also provides the capability to support mixed IA-32 and IA-64 code execution.

The IA-64 architecture was designed with the understanding that compatibility with IA-32 and PA-RISC is a key requirement. Significant effort has been applied in the architectural definition to maximize IA-64 scalability, performance and architectural longevity. As a result, IA-64 has been designed with an array of performance optimization techniques that extend the architecture to 64-bits and enable higher performance. These features include speculation, predication, large register files, a register stack, and an advanced branch architecture and are designed to extract greater instruction level parallelism.

2.1 IA-64 Operating Environments

The IA-64 architecture supports two operating system environments:

- IA-32 System Environment: supports IA-32 32-bit operating systems, and
- IA-64 System Environment: supports IA-64 operating systems.

The architectural model also supports a mixture of IA-32 and IA-64 applications within a single IA-64 operating system. Table 2-1 defines the major operating environments supported on IA-64 processors.

Table 2-1. IA-64 Processor Operating Environments

System Environment	Application Environment	Usage
IA-32	IA-32 Instruction Set	IA-32 Protected Mode, Real Mode and Virtual 8086 Mode application and operating system environment. Compatible with IA-32 Pentium®, Pentium Pro and Pentium II processors.
IA-64	IA-32 Protected Mode	IA-32 Protected Mode applications in the IA-64 system environment, if supported by OS.
	IA-32 Real Mode	IA-32 Real Mode applications in the IA-64 system environment, if supported by OS.
	IA-32 Virtual Mode	IA-32 Virtual 86 Mode applications in the IA-64 system environment, if supported by OS.
	IA-64 Instruction Set	IA-64 Applications on IA-64 operating systems.

2.2 Instruction Set Transition Model Overview

Within the IA-64 System Environment, the processor can execute either IA-32 or IA-64 instructions at any time. Three special instructions and interruptions are defined to transition the processor between the IA-32 and the IA-64 instruction set.

- `jmpe` (IA-32 instruction) Jump to an IA-64 target instruction, and change the instruction set to IA-64.
- `br.ia` (IA-64 instruction) IA-64 branch to an IA-32 target instruction, and change the instruction set to IA-32.
- Interruptions transition the processor to the IA-64 instruction set for handling all interruption conditions.
- `rfi` (IA-64 instruction) “return from interruption” is defined to return to an IA-32 or IA-64 instruction.

The `jmpe` and `br.ia` instructions provide a low overhead mechanism to transfer control between the instruction sets. These instructions are typically incorporated into “thunks” or “stubs” that implement the required call linkage and calling conventions to call dynamic or statically linked libraries.

2.3 PA-RISC Compatibility

Binary compatibility between PA-RISC and IA-64 is handled through dynamic object code translation. This process is very efficient because there is such a high degree of correspondence between PA-RISC and IA-64 instructions. HP’s performance studies show that on average the dynamic translator only spends 1-2% of its time in translation with 98-99% of the time spent executing native code. The dynamic translator actually performs optimizations on the translated code to take advantage of IA-64’s wider instructions, and performance features such as predication, speculation and large register sets. In addition, if an application has been aggressively optimized for PA-RISC, some of the benefit of the optimizations will carry over to IA-64. In fact, an aggressively optimized PA-RISC application may actually perform faster on IA-64 using the dynamic translator than the same application recompiled at a low level of optimization on an IA-64 compiler. Of course, the best performance will result from a high level of optimization using a good native compiler.

The dynamic translator is designed to run all non-kernel intrusive code, handling both 64-bit and 32-bit instructions. This means operating systems and device drivers typically would not be supported, but all other applications will run. HP’s dynamic translator will be bundled with all versions of HP-UX sold for IA-64 systems. When HP-UX encounters code compiled for PA-RISC, it will automatically and transparently invoke the dynamic translator which will allow the code to run on IA-64 without any intervention. Correctness of the dynamic translator has been verified with the same testing regimen used to validate PA-RISC processors.

2.4 IA-64 Instruction Set Features

IA-64 incorporates architectural features which enable high sustained performance and remove barriers to further performance increases. The IA-64 architecture is based on the following principles:

- Explicit parallelism
 - Mechanisms for synergy between the compiler and the processor
 - Massive resources to take advantage of instruction level parallelism
 - 128 Integer and Floating point registers, 64 1-bit predicate registers, 8 branch registers
 - Support for many execution units and memory ports
- Features that enhance instruction level parallelism
 - Speculation (which minimizes memory latency impact).
 - Predication (which removes branches).
 - Software pipelining of loops with low overhead
 - Branch prediction to minimize the cost of branches

- Focussed enhancements for improved software performance
 - Special support for software modularity
 - High performance floating-point architecture
 - Specific multimedia instructions

The following sections highlight these important features of IA-64.

2.5 Instruction Level Parallelism

Instruction Level Parallelism (ILP) is the ability to execute multiple instructions at the same time. The IA-64 architecture allows issuing of independent instructions in bundles (three instructions per bundle) for parallel execution and can issue multiple bundles per clock. Supported by a large number of parallel resources such as large register files and multiple execution units, the IA-64 architecture enables the compiler to manage work in progress and schedule simultaneous threads of computation.

The IA-64 architecture incorporates mechanisms to take advantage of ILP. Compilers for traditional architectures are often limited in their ability to utilize speculative information because it cannot always be guaranteed to be correct. The IA-64 architecture enables the compiler to exploit speculative information without sacrificing the correct execution of an application (see “Speculation” on page 2-3). In traditional architectures, procedure calls limit performance since registers need be spilled and filled. IA-64 enables procedures to communicate register usage to the processor. This allows the processor to schedule procedure register operations even when there is a low degree of ILP. See “Register Stack” on page 2-5.

2.6 Compiler to Processor Communication

The IA-64 architecture provides mechanisms, such as instruction templates, branch hints, and cache hints to enable the compiler to communicate compile-time information to the processor. In addition, IA-64 allows compiled code to manage the processor hardware using run-time information. These communication mechanisms are vital in minimizing the performance penalties associated with branches and cache misses.

Every memory load and store in IA-64 has a 2-bit cache hint field in which the compiler encodes its prediction of the spatial and/or temporal locality of the memory area being accessed. An IA-64 processor can use this information to determine the placement of cache lines in the cache hierarchy. This leads to better utilization of the hierarchy since the relative cost of cache misses continues to grow.

2.7 Speculation

There are two types of speculation: control and data. In both control and data speculation, the compiler exposes ILP by issuing an operation early and removing the latency of this operation from the critical path. The compiler will issue an operation speculatively if it is reasonably sure that the speculation will be beneficial. To be beneficial two conditions should hold: it must be statistically frequent enough that the probability it will require recovery is small, and issuing the operation early should expose further ILP-enhancing optimization. Speculation is one of the primary mechanisms for the compiler to exploit statistical ILP by overlapping, and therefore tolerating, the latencies of operations.

2.7.1 Control Speculation

Control speculation is the execution of an operation before the branch which guards it. Consider the code sequence below:

```
if (a>b) load(ld_addr1,target1)
else load(ld_addr2, target2)
```

If the operation `load(ld_addr1,target1)` were to be performed prior to the determination of `(a>b)`, then the operation would be control speculative with respect to the controlling condition `(a>b)`. Under normal execution, the operation `load(ld_addr1,target1)` may or may not execute. If the new control speculative load causes an exception then the exception should only be serviced if `(a>b)` is true. When the compiler uses control speculation it leaves a check operation at the original location. The check verifies whether an exception has occurred and if so it branches to recovery code. The code sequence above now translates into:

```

/* off critical path */
sload(ld_addr1,target1)
sload(ld_addr2,target2)

/* other operations including uses of target1/target2 */
if (a>b) scheck(target1,recovery_addr1)
else scheck(target2, recovery_addr2)

```

2.7.2 Data Speculation

Data speculation is the execution of a memory load prior to a store that preceded it and that may potentially alias with it. Data speculative loads are also referred to as “advanced loads”. Consider the code sequence below:

```

store(st_addr,data)
load(ld_addr,target)
use(target)

```

The process of determining at compile time the relationship between memory addresses is called disambiguation. In the example above, if `ld_addr` and `st_addr` cannot be disambiguated, and if the load were to be performed prior to the store, then the load would be data speculative with respect to the store. If memory addresses overlap during execution, a data-speculative load issued before the store might return a different value than a regular load issued after the store. Therefore analogous to control speculation, when the compiler data speculates a load, it leaves a check instruction at the original location of the load. The check verifies whether an overlap has occurred and if so it branches to recovery code. The code sequence above now translates into:

```

/* off critical path */
aload(ld_addr,target)

/* other operations including uses of target */
store(st_addr,data)
acheck(target,recovery_addr)
use(target)

```

2.8 Predication

Predication is the conditional execution of instructions. Conditional execution is implemented through branches in traditional architectures. IA-64 implements this function through the use of predicated instructions. Predication removes branches used for conditional execution resulting in larger basic blocks and the elimination of associated mispredict penalties.

To illustrate, an unpredicated instruction

```
r1 = r2 + r3
```

when predicated, would be of the form

```
if (p5) r1 = r2 + r3
```

In this example `p5` is the controlling predicate that decides whether or not the instruction executes and updates state. If the predicate value is true, then the instruction updates state. Otherwise it generally behaves like a `nop`. Predicates are assigned values by compare instructions.

Predicated execution avoids branches, and simplifies compiler optimizations by converting a control dependence to a data dependence. Consider the original code:

```

if (a>b) c = c + 1
else d = d * e + f

```

The branch at `(a>b)` can be avoided by converting the code above to the predicated code:

```

pT, pF = compare(a>b)
if (pT) c = c + 1
if (pF) d = d * e + f

```

The predicate p_T is set to 1 if the condition evaluates to true, and to 0 if the condition evaluates to false. The predicate p_F is the complement of p_T . The control dependence of the instructions $c = c + 1$ and $d = d * e + f$ on the branch with the condition $(a > b)$ is now converted into a data dependence on `compare(a > b)` through predicates p_T and p_F (the branch is eliminated). An added benefit is that the compiler can schedule the instructions under p_T and p_F to execute in parallel. It is also worth noting that there are several different types of compare instructions that write predicates in different manners including unconditional compares and parallel compares.

2.9 Register Stack

IA-64 avoids the unnecessary spilling and filling of registers at procedure call and return interfaces through compiler-controlled renaming. At a call site, a new frame of registers is available to the called procedure without the need for register spill and fill (either by the caller or by the callee). Register access occurs by renaming the virtual register identifiers in the instructions through a base register into the physical registers. The callee can freely use available registers without having to spill and eventually restore the caller's registers. The callee executes an `alloc` instruction specifying the number of registers it expects to use in order to ensure that enough registers are available. If sufficient registers are not available (stack overflow), the `alloc` stalls the processor and spills the caller's registers until the requested number of registers are available.

At the return site, the base register is restored to the value that the caller was using to access registers prior to the call. Some of the caller's registers may have been spilled by the hardware and not yet restored. In this case (stack underflow), the return stalls the processor until the processor has restored an appropriate number of the caller's registers. The hardware can exploit the explicit register stack frame information to spill and fill registers from the register stack to memory at the best opportunity (independent of the calling and called procedures).

2.10 Branching

In addition to removing branches through the use of predication, several mechanisms are provided to decrease the branch misprediction rate and the cost of the remaining mispredicted branches. These mechanisms provide ways for the compiler to communicate information about branch conditions to the processor.

For indirect branches, a branch register is used to hold the target address.

Special loop-closing branches are provided to accelerate counted loops and modulo-scheduled loops. These branches provide information that allows for perfect prediction of loop termination, thereby eliminating costly mispredict penalties and a reduction of the loop overhead.

2.11 Register Rotation

Modulo scheduling of a loop is analogous to hardware pipelining of a functional unit since the next iteration of the loop starts before the previous iteration has finished. The iteration is split into stages similar to the stages of an execution pipeline. Modulo scheduling allows the compiler to execute loop iterations in parallel rather than sequentially. The concurrent execution of multiple iterations traditionally requires unrolling of the loop and software renaming of registers. IA-64 allows the renaming of registers which provide every iteration with its own set of registers, avoiding the need for unrolling. This kind of register renaming is called register rotation. The result is that software pipelining can be applied to a much wider variety of loops - both small as well as large with significantly reduced overhead.

2.12 Floating-point Architecture

IA-64 defines a floating-point architecture with full IEEE support for the single, double, and double-extended (80-bit) data types. Some extensions, such as a fused multiply and add operation, minimum and maximum functions, and a register file format with a larger range than the double-extended memory format, are also included. 128 floating-point registers are defined. Of these, 96 registers are rotating (not stacked) and can be used to modulo schedule loops compactly. Multiple floating-point status registers are provided for speculation.

IA-64 has parallel FP instructions which operate on two 32-bit single precision numbers, resident in a single floating-point register, in parallel and independently. These instructions significantly increase the single precision floating-point computation throughput and enhance the performance of 3D intensive applications and games.

2.13 Multimedia Support

IA-64 has multimedia instructions which treat the general registers as concatenations of eight 8-bit, four 16-bit, or two 32-bit elements. These instructions operate on each element in parallel, independent of the others. IA-64 multimedia instructions are semantically compatible with HP's MAX-2 multimedia technology and Intel's MMX technology instructions and Streaming SIMD Extensions instruction technology.

The architectural state consists of registers and memory. The results of instruction execution become architecturally visible according to a set of execution sequencing rules. This chapter describes the IA-64 application architectural state and the rules for execution sequencing.

3.1 Application Register State

The following is a list of the registers available to application programs (see Figure 3-1):

- **General Registers (GRs)** – General purpose 64-bit register file, GR0 – GR127. IA-32 integer and segment registers are contained in GR8 - GR31 when executing IA-32 instructions.
- **Floating-Point Registers (FRs)** – Floating-point register file, FR0 – FR127. IA-32 floating-point and multi-media registers are contained in FR8 - FR31 when executing IA-32 instructions.
- **Predicate Registers (PRs)** – Single-bit registers, used in IA-64 predication and branching, PR0 – PR63.
- **Branch Registers (BRs)** – Registers used in IA-64 branching, BR0 – BR7.
- **Instruction Pointer (IP)** – Register which holds the bundle address of the currently executing IA-64 instruction, or byte address of the currently executing IA-32 instruction.
- **Current Frame Marker (CFM)** – State that describes the current general register stack frame, and FR/PR rotation.
- **Application Registers (ARs)** – A collection of special-purpose IA-64 and IA-32 application registers.
- **Performance Monitor Data Registers (PMD)** – Data registers for performance monitor hardware.
- **User Mask (UM)** – A set of single-bit values used for alignment traps, performance monitors, and to monitor floating-point register usage.
- **Processor Identifiers (CPUID)** – Registers that describe processor implementation-dependent IA-64 features.

IA-32 application register state is entirely contained within the larger IA-64 application register set and is accessible by IA-64 instructions. IA-32 instructions cannot access the IA-64 specific register set.

3.1.1 Reserved and Ignored Registers

Registers which are not defined are either reserved or ignored. An access to a **reserved register** raises an Illegal Operation fault. A read of an **ignored register** returns zero. Software may write any value to an ignored register and the hardware will ignore the value written. In variable-sized register sets, registers which are unimplemented in a particular processor are also reserved registers. An access to one of these unimplemented registers causes a Reserved Register/Field fault.

Within defined registers, fields which are not defined are either reserved or ignored. For **reserved fields**, hardware will always return a zero on a read. Software must always write zeros to these fields. Any attempt to write a non-zero value into a reserved field will raise a Reserved register/field fault. Reserved fields may have a possible future use.

For **ignored fields**, hardware will return a 0 on a read, unless noted otherwise. Software may write any value to these fields since the hardware will ignore any value written. Except where noted otherwise some IA-32 ignored fields may have a possible future use.

Table 3-1 summarizes how the processor treats reserved and ignored registers and fields.

Table 3-1. Reserved and Ignored Registers and Fields

Type	Read	Write
Reserved register	Illegal operation fault	Illegal operation fault
Ignored register	0	value written is discarded
Reserved field	0	write of non-zero causes Reserved Reg/Field fault
Ignored field	0 (unless noted otherwise)	value written is discarded

For defined fields in registers, values which are not defined are reserved. Software must always write defined values to these fields. Any attempt to write a **reserved value** will raise a Reserved Register/Field fault. Certain registers are **read-only registers**. A write to a read-only register raises an Illegal Operation fault.

When fields are marked as **reserved**, it is essential for compatibility with future processors that software treat these fields as having a future, though unknown effect. Software should follow these guidelines when dealing with reserved fields:

- Do not depend on the state of any reserved fields. Mask all reserved fields before testing.
- Do not depend on the states of any reserved fields when storing to memory or a register.
- Do not depend on the ability to retain information written into reserved or ignored fields.
- Where possible reload reserved or ignored fields with values previously returned from the same register, otherwise load zeros.

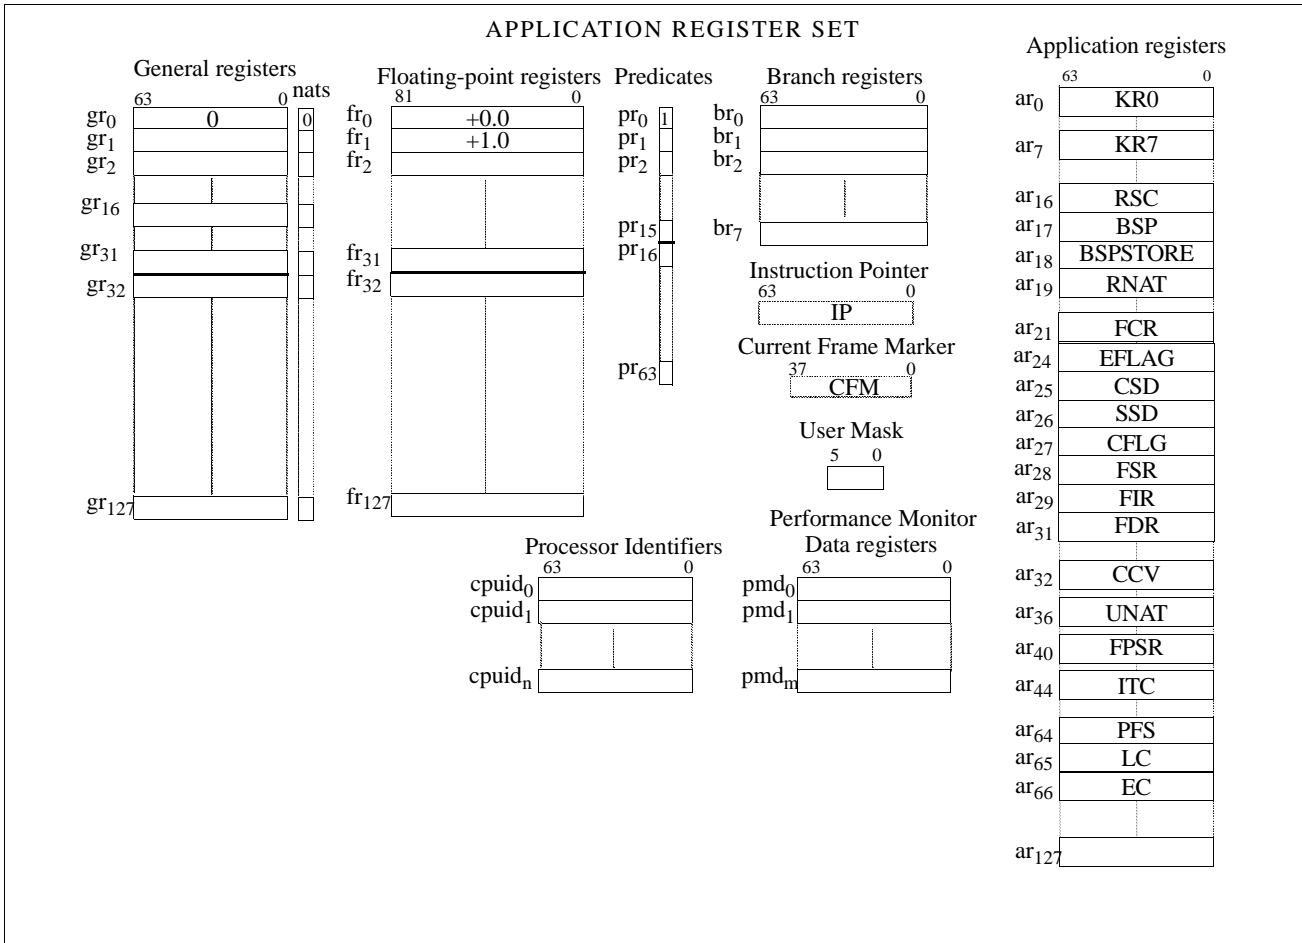


Figure 3-1. Application Register Model

3.1.2 General Registers

A set of 128 (64-bit) **general registers** provide the central resource for all integer and integer multimedia computation. They are numbered GR0 through GR127, and are available to all programs at all privilege levels. Each general register has 64 bits of normal data storage plus an additional bit, the **NaT** bit (Not a Thing), which is used to track deferred speculative exceptions.

The general registers are partitioned into two subsets. General registers 0 through 31 are termed the **static general registers**. Of these, GR0 is special in that it always reads as zero when sourced as an operand and attempting to write to GR0 causes an Illegal Operation fault. General registers 32 through 127 are termed the **stacked general registers**. The stacked registers are made available to a program by allocating a register stack frame consisting of a programmable number of local and output registers. See Chapter 4.1, “Register Stack” for a description. A portion of the stacked registers can be programmatically renamed to accelerate loops. See “Modulo-Scheduled Loop Support” on page 4-20.

General registers 8 through 31 contain the IA-32 integer, segment selector and segment descriptor registers when executing IA-32 instructions.

3.1.3 Floating-Point Registers

A set of 128 (82-bit) **floating-point registers** are used for all floating-point computation. They are numbered FR0 through FR127, and are available to all programs at all privilege levels. The floating-point registers are partitioned into two subsets. Floating-point registers 0 through 31 are termed the **static floating-point registers**. Of these, FR0 and FR1 are special. FR0 always reads as +0.0 when sourced as an operand, and FR 1 always reads as +1.0. When either of these is used as a destination, a fault is raised. Deferred speculative exceptions are recorded with a special register value called **NaTVal** (**Not a Thing Value**).

Floating-point registers 32 through 127 are termed the **rotating floating-point registers**. These registers can be programmatically renamed to accelerate loops. See “Modulo-Scheduled Loop Support” on page 4-20.

Floating-point registers 8 through 31 contain the IA-32 floating point and multi-media registers when executing IA-32 instructions.

3.1.4 Predicate Registers

A set of 64 (1-bit) **predicate registers** are used to hold the results of IA-64 compare instructions. These registers are numbered PR0 through PR63, and are available to all programs at all privilege levels. These registers are used for conditional execution of instructions.

The predicate registers are partitioned into two subsets. Predicate registers 0 through 15 are termed the **static predicate registers**. Of these, PR0 always reads as ‘1’ when sourced as an operand, and when used as a destination, the result is discarded. The static predicate registers are also used in conditional branching. See “Predication” on page 4-6.

Predicate registers 16 through 63 are termed the **rotating predicate registers**. These registers can be programmatically renamed to accelerate loops. See “Modulo-Scheduled Loop Support” on page 4-20.

3.1.5 Branch Registers

A set of 8 (64-bit) **branch registers** are used to hold IA-64 branching information. They are numbered BR 0 through BR 7, and are available to all programs at all privilege levels. The branch registers are used to specify the branch target addresses for indirect branches. For more information see “Branch Instructions” on page 4-19.

3.1.6 Instruction Pointer

The Instruction Pointer (IP) holds the address of the bundle which contains the current executing IA-64 instruction. The IP can be read directly with a `mov ip` instruction. The IP cannot be directly written, but is incremented as instructions are executed, and can be set to a new value with a `branch`. Because IA-64 instruction bundles are 16 bytes, and are 16-byte aligned, the least significant 4 bits of IP are always zero. See “Instruction Encoding Overview” on page 3-11. For IA-32 instruction set execution, IP holds the zero extended 32-bit virtual linear address of the currently executing IA-32 instruction. IA-32 instructions are byte-aligned, therefore the least significant 4 bits of IP are preserved for IA-32 instruction set execution.

3.1.7 Current Frame Marker

Each general register stack frame is associated with a frame marker. The frame marker describes the state of the IA-64 general register stack. The Current Frame Marker (CFM) holds the state of the current stack frame. The CFM cannot be directly read or written (see “Register Stack” on page 4-1).

The frame markers contain the sizes of the various portions of the stack frame, plus three Register Rename Base values (used in register rotation). The layout of the frame markers is shown in Figure 3-2 and the fields are described in Table 3-2.

On a call, the CFM is copied to the Previous Frame Marker field in the Previous Function State register (see Section 3.1.8.10). A new value is written to the CFM, creating a new stack frame with no locals or rotating registers, but with a set of output registers which are the caller’s output registers. Additionally, all Register Rename Base registers (RRBs) are set to 0. See “Modulo-Scheduled Loop Support” on page 4-20.




Figure 3-2. Frame Marker Format

Table 3-2. Frame Marker Field Description

Field	Bit Range	Description
sof	6:0	Size of stack frame
sol	13:7	Size of locals portion of stack frame
sor	17:14	Size of rotating portion of stack frame (the number of rotating registers is 8 * sor)
rrb.gr	24:18	Register Rename Base for general registers
rrb.fr	31:25	Register Rename Base for floating-point registers
rrb.pr	37:32	Register Rename Base for predicate registers

3.1.8 Application Registers

The application register file includes special-purpose data registers and control registers for application-visible processor functions for both the IA-32 and IA-64 instruction sets. These registers can be accessed by IA-64 application software (except where noted). Table 3-3 contains a list of the application registers.

Table 3-3. Application Registers

Register	Name	Description	Execution Unit Type	
AR 0-7	KR 0-7 ^a	Kernel Registers 0-7	M	
AR 8-15		Reserved		
AR 16	RSC	Register Stack Configuration Register		
AR 17	BSP	Backing Store Pointer (read-only)		
AR 18	BSPSTORE	Backing Store Pointer for Memory Stores		
AR 19	RNAT	RSE NAT Collection Register		
AR 20		Reserved		
AR 21	FCR	IA-32 Floating-point Control Register		
AR 22 – AR 23		Reserved		
AR 24	EFLAG ^b	IA-32 EFLAG register		
AR 25	CSD	IA-32 Code Segment Descriptor		
AR 26	SSD	IA-32 Stack Segment Descriptor		
AR 27	CFLG ^a	IA-32 Combined CR0 and CR4 register		
AR 28	FSR	IA-32 Floating-point Status Register		
AR 29	FIR	IA-32 Floating-point Instruction Register		
AR 30	FDR	IA-32 Floating-point Data Register		
AR 31		Reserved		
AR 32	CCV	Compare and Exchange Compare Value Register		
AR 33 – AR 35		Reserved		
AR 36	UNAT	User NAT Collection Register		
AR 37 – AR 39		Reserved		
AR 40	FPSR	Floating-Point Status Register		
AR 41 – AR 43		Reserved		
AR 44	ITC	Interval Time Counter		
AR 45 – AR 47		Reserved		
AR 48 – AR 63		Ignored		M or I
AR 64	PFS	Previous Function State		I
AR 65	LC	Loop Count Register		
AR 66	EC	Epilog Count Register		
AR 67 – AR 111		Reserved		
AR 112 – AR 127		Ignored		

a. Writes to these registers when the privilege level is not zero result in a Privileged Register fault. Reads are always allowed.

b. Some IA-32 EFLAG field writes are silently ignored if the privilege level is not zero.

Application registers can only be accessed by either a M or I execution unit. This is specified in the last column of the table. The ignored registers are for future backward-compatible extensions.

3.1.8.1 Kernel Registers (KR 0-7 – AR 0-7)

Eight user-visible IA-64 64-bit data kernel registers are provided to convey information from the operating system to the application. These registers can be read at any privilege level but are writable only at the most privileged level. KR0 - KR2 are also used to hold additional IA-32 register state when the IA-32 instruction set is executing.

3.1.8.2 Register Stack Configuration Register (RSC – AR 16)

The Register Stack Configuration (RSC) Register is a 64-bit register used to control the operation of the IA-64 Register Stack Engine (RSE). The RSC format is shown in Figure 3-3 and the field description is contained in Table 3-4. Instruc-

tions that modify the RSC can never set the privilege level field to a more privileged level than the currently executing process.



Figure 3-3. RSC Format

Table 3-4. RSC Field Description

Field	Bit Range	Description			
mode	1:0	RSE mode – controls how aggressively the RSE saves and restores register frames. Eager and intensive settings are hints and can be implemented as lazy.			
		Bit Pattern	RSE Mode	Bit 1: eager loads	Bit 0: eager stores
		00	enforced lazy	disabled	disabled
		10	load intensive	enabled	disabled
		01	store intensive	disabled	enabled
11	eager	enabled	enabled		
pl	3:2	RSE privilege level – loads and stores issued by the RSE are at this privilege level			
be	4	RSE endian mode – loads and stores issued by the RSE use this byte ordering (0: little endian; 1: big endian)			
loadrs	29:16	RSE load distance to tear point – value used in the <code>loadrs</code> instruction for synchronizing the RSE to a tear point			
rv	15:5, 63:30	Reserved			

3.1.8.3 RSE Backing Store Pointer (BSP – AR 17)

The RSE Backing Store Pointer is a 64-bit read-only register (Figure 3-4). It holds the address of the location in memory which is the save location for GR 32 in the current stack frame.




Figure 3-4. BSP Register Format

3.1.8.4 RSE Backing Store Pointer for Memory Stores (BSPSTORE – AR 18)

The RSE Backing Store Pointer for memory stores is a 64-bit register (Figure 3-5). It holds the address of the location in memory to which the RSE will spill the next value.




Figure 3-5. BSPSTORE Register Format

3.1.8.5 RSE NAT Collection Register (RNAT – AR 19)

The RSE NaT Collection Register is a 64-bit register (Figure 3-6) used by the RSE to temporarily hold NaT bits when it is spilling general registers. Bit 63 always reads as zero and ignores all writes.




Figure 3-6. RNAT Register Format

3.1.8.6 Compare and Exchange Value Register (CCV – AR 32)

The Compare and Exchange Value Register is a 64-bit register that contains the compare value used as the third source operand in the IA-64 `cmpxchg` instruction.

3.1.8.7 User NAT Collection Register (UNAT – AR 36)

The User NaT Collection Register is a 64-bit register used to temporarily hold NaT bits when saving and restoring general registers with the IA-64 `ld8.fill` and `st8.spill` instructions.

3.1.8.8 Floating-Point Status Register (FPSR – AR 40)

The floating-point status register (FPSR) controls traps, rounding mode, precision control, flags, and other control bits for IA-64 floating point instructions. FPSR does not control or reflect the status of IA-32 floating point instructions. For more details on the FPSR, see Section 5.2.

3.1.8.9 Interval Time Counter (ITC – AR 44)

The Interval Time Counter (ITC) is a 64-bit register which counts up at a fixed relationship to the processor clock frequency. Applications can directly sample the ITC for time-based calculations and performance measurements. System software can secure the interval time counter from non-privileged IA-64 access. When secured, a read of the ITC at any privilege level other than the most privileged causes a Privileged Register fault. The ITC can be written only at the most privileged level. The IA-32 Time Stamp Counter (TSC) is equivalent to ITC. ITC can directly be read by the IA-32 `rdtsc` (read time stamp counter) instruction. System software can secure the ITC from non-privileged IA-32 access. When secured, an IA-32 read of the ITC at any privilege level other than the most privileged raises an IA-32_Exception(GPfault).

3.1.8.10 Previous Function State (PFS – AR 64)

The IA-64 Previous Function State register (PFS) contains multiple fields: Previous Frame Marker (pfm), Previous Epilog Count (pec), and Previous Privilege Level (ppl). Figure 3-7 diagrams the PFS format and Table 3-5 describes the PFS fields. These values are copied automatically on a call from the CFM register, Epilog Count Register (EC) and PSR.cpl (Current Privilege Level in the Processor Status Register) to accelerate procedure calling.

When an IA-64 `br.call` is executed, the CFM, EC, and PSR.cpl are copied to the PFS and the old contents of the PFS are discarded. When an IA-64 `br.ret` is executed, the PFS is copied to the CFM and EC. PFS.ppl is copied to PSR.cpl, unless this action would increase the privilege level.

The PFS.pfm has the same layout as the CFM (see Section 3.1.7), and the PFS.pec has the same layout as the EC (see Section 3.1.8.12).




Figure 3-7. PFS Format

Table 3-5. PFS Field Description

Field	Bit Range	Description
pfm	37:0	Previous Frame Marker
pec	57:52	Previous Epilog Count
ppl	63:62	Previous Privilege Level
rv	51:38, 61:58	Reserved

3.1.8.11 Loop Count Register (LC – AR 65)

The Loop Count register (LC) is a 64-bit register used in IA-64 counted loops. LC is decremented by counted-loop-type branches.

3.1.8.12 Epilog Count Register (EC – AR 66)

The Epilog Count register (EC) is a 6-bit register used for counting the final (epilog) stages in IA-64 modulo-scheduled loops. See “Modulo-Scheduled Loop Support” on page 4-20. A diagram of the EC register is shown in Figure 3-8.




Figure 3-8. Epilog Count Register Format

3.1.9 Performance Monitor Data Registers (PMD)

A set of performance monitoring registers can be configured by privileged software to be accessible at all privilege levels. Performance monitor data can be directly sampled from within the application. The operating system is allowed to secure user-configured performance monitors. Secured performance counters return zeros when read, regardless of the current privilege level. The performance monitors can only be written at the most privileged level. Performance monitors can be used to gather performance information for both IA-32 and IA-64 instruction set execution.

3.1.10 User Mask (UM)

The user mask is a subset of the Processor Status Register and is accessible to IA-64 application programs. The user mask controls memory access alignment, byte-ordering and user-configured performance monitors. It also records the modification state of IA-64 floating-point registers. Figure 3-9 show the user mask format and Table 3-6 describes the user mask fields.




Figure 3-9. User Mask Format

Table 3-6. User Mask Field Descriptions

Field	Bit Range	Description
rv	0	Reserved
be	1	IA-64 Big-endian memory access enable (controls loads and stores but not RSE memory accesses) 0: accesses are done little-endian 1: accesses are done big-endian This bit is ignored for IA-32 data memory accesses. IA-32 data references are always performed little-endian.
up	2	User performance monitor enable for IA-32 and IA-64 instruction set execution 0: user performance monitors are disabled 1: user performance monitors are enabled
ac	3	Alignment check for IA-32 and IA-64 data memory references 0: unaligned data memory references may cause an Unaligned Data Reference fault. 1: all unaligned data memory references cause an Unaligned Data Reference fault.
mfl	4	Lower (f2 .. f31) floating-point registers written – This bit is set to one when an IA-64 instruction that uses register f2..f31 as a target register, completes. This bit is sticky and is only cleared by an explicit write of the user mask.
mfh	5	Upper (f32 .. f127) floating-point registers written – This bit is set to one when an IA-64 instruction that uses register f32..f127 as a target register, completes. This bit is sticky and only cleared by an explicit write of the user mask.

3.1.11 Processor Identification Registers

Application level processor identification information is available in an IA-64 register file termed: CPUID. This register file is divided into a fixed region, registers 0 to 4, and a variable region, register 5 and above. The CPUID[3].number field indicates the maximum number of 8-byte registers containing processor specific information.

The CPUID registers are unprivileged and accessed using the indirect `mov (from)` instruction. All registers beyond register CPUID[3].number are reserved and raise a Reserved Register/Field fault if they are accessed. Writes are not permitted and no instruction exists for such an operation.

Vendor information is located in CPUID registers 0 and 1 and specify a vendor name, in ASCII, for the processor implementation (Figure 3-10). All bytes after the end of the string up to the 16th byte are zero. Earlier ASCII characters are placed in lower number register and lower numbered byte positions.




Figure 3-10. CPUID Registers 0 and 1 – Vendor Information

A Processor Serial Number is located in CPUID register 2. If Processor Serial Numbers are supported by the processor model and are not disabled, this register returns a 64-bit number Processor Serial Number (Figure 3-11), otherwise zero is returned. The Processor Serial Number (64-bits) must be combined with the 32-bit version information (CPUID register 3; processor archrev, family, model, and revision numbers) to form a 96-bit Processor Identifier.

The 96-bit Processor Identifier is designed to be unique.




Figure 3-11. CPUID Register 2 – Processor Serial Number

CPUID register 3 contains several fields indicating version information related to the processor implementation. Figure 3-12 and Table 3-7 specify the definitions of each field.




Figure 3-12. CPUID Register 3 – Version Information

Table 3-7. CPUID Register 3 Fields

Field	Bits	Description
number	7:0	The index of the largest implemented CPUID register (one less than the number of implemented CPUID registers). This value will be at least 4.
revision	15:8	Processor revision number. An 8-bit value that represents the revision or stepping of this processor implementation within the processor model.
model	23:16	Processor model number. A unique 8-bit value representing the processor model within the processor family.
family	31:24	Processor family number. A unique 8-bit value representing the processor family.
archrev	39:32	Architecture revision. An 8-bit value that represents the architecture revision number that the processor implements.
rv	63:40	Reserved.

CPUID register 4 provides general application level information about IA-64 features. As shown in Figure 3-13, it is a set of flag bits used to indicate if a given IA-64 feature is supported in the processor model. When a bit is one the feature is supported; when 0 the feature is not supported. This register does not contain IA-32 instruction set features. IA-32 instruction set features can be acquired by the IA-32 `cpuid` instruction. There are no defined feature bits in the current architecture. As new features are added (or removed) from future processor models the presence (or removal) of new features will be indicated by new feature bits. A value of zero in this register indicates all features defined in the first IA-64 architectural revision are implemented.




Figure 3-13. CPUID Register 4 – General Features/Capability Bits

3.2 Memory

This section describes an IA-64 application program's view of memory. This includes a description of how memory is accessed, for both 32-bit and 64-bit applications. The size and alignment of addressable units in memory is also given, along with a description of how byte ordering is handled.

3.2.1 Application Memory Addressing Model

Memory is byte addressable and is accessed with 64-bit pointers. A 32-bit pointer model without a hardware mode is supported architecturally. Pointers which are 32 bits in memory are loaded and manipulated in 64-bit registers. Software must explicitly convert 32-bit pointers into 64-bit pointers before use.

3.2.2 Addressable Units and Alignment

Memory can be addressed in units of 1, 2, 4, 8, 10 and 16 bytes.

It is recommended that all addressable units be stored on their naturally aligned boundaries. Hardware and/or operating system software may have support for unaligned accesses, possibly with some performance cost. 10-byte floating-point values should be stored on 16-byte aligned boundaries.

Bits within larger units are always numbered from 0 starting with the least-significant bit. Quantities loaded from memory to general registers are always placed in the least-significant portion of the register (loaded values are placed right justified in the target general register).

Instruction bundles (3 IA-64 instructions per bundle) are 16-byte units that are always aligned on 16-byte boundaries.

3.2.3 Byte Ordering

The UM.be bit in the User Mask controls whether loads and stores use little-endian or big-endian byte ordering for IA-64 references. When the UM.be bit is 0, larger-than-byte loads and stores are little endian (lower-addressed bytes in memory correspond to the lower-order bytes in the register). When the UM.be bit is 1, larger-than-byte loads and stores are big endian (lower-addressed bytes in memory correspond to the higher-order bytes in the register). Load byte and store byte are not affected by the UM.be bit. The UM.be bit does not affect instruction fetch, IA-32 references, or the RSE. IA-64 instructions are always accessed by the processor as little-endian units. When instructions are referenced as big-endian data, the instruction will appear reversed in a register.

Figure 3-14 shows various loads in little-endian format. Figure 3-15 shows various loads in big endian format. Stores are not shown but behave similarly.

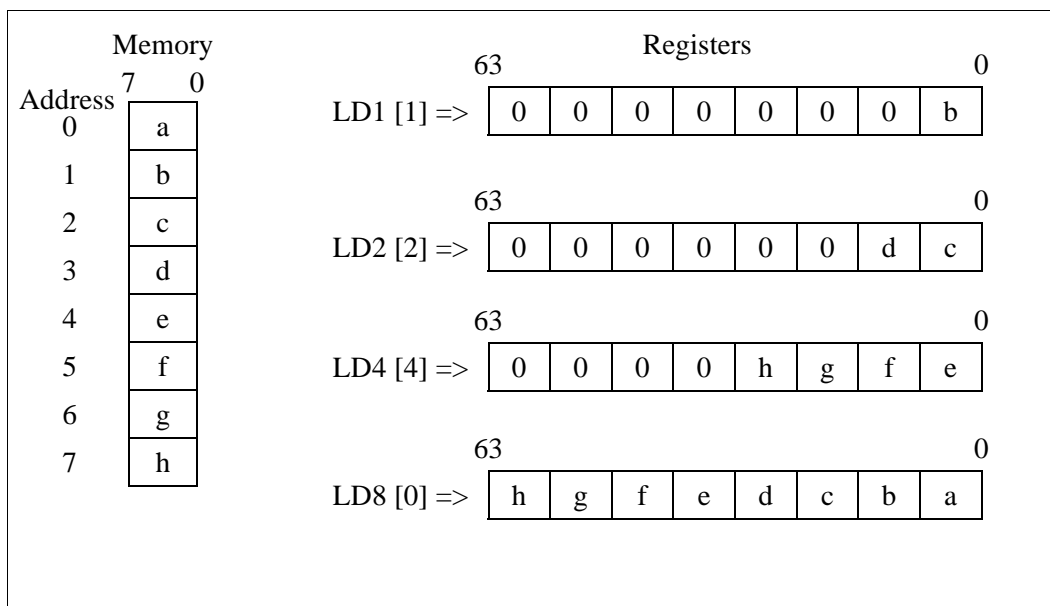


Figure 3-14. Little-endian Loads

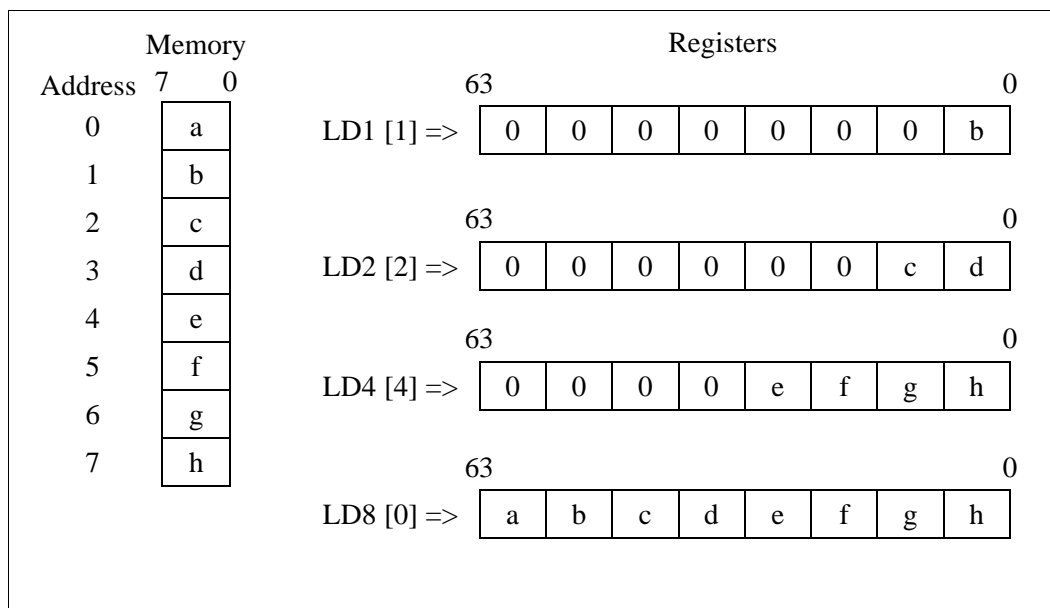


Figure 3-15. Big-endian Loads

3.3 Instruction Encoding Overview

Each IA-64 instruction is categorized into one of six types; each instruction type may be executed on one or more execution unit types. Table 3-8 lists the instruction types and the execution unit type on which they are executed:

Table 3-8. Relationship Between Instruction Type and Execution Unit Type

Instruction Type	Description	Execution Unit Type
A	Integer ALU	I-unit or M-unit
I	Non-ALU integer	I-unit
M	Memory	M-unit
F	Floating-point	F-unit
B	Branch	B-unit
L+X	Extended	I-unit

Three instructions are grouped together into 128-bit sized and aligned containers called **bundles**. Each bundle contains three 41-bit **instruction slots** and a 5-bit template field. The format of a bundle is depicted in Figure 3-16.




Figure 3-16. Bundle Format

During execution, architectural **stops** in the program indicate to the hardware that one or more instructions before the stop may have certain kinds of resource dependencies with one or more instructions after the stop. A stop is present after each slot having a double line to the right of it in Table 3-9. For example, template 00 has no stops, while template 03 has a stop after slot 1 and another after slot 2.

In addition to the location of stops, the template field specifies the mapping of instruction slots to execution unit types. Not all possible mappings of instructions to units are available. Table 3-9 indicates the defined combinations. The three rightmost columns correspond to the three instruction slots in a bundle. Listed within each column is the execution unit type controlled by that instruction slot.

Table 3-9. Template Field Encoding and Instruction Slot Mapping

Template	Slot 0	Slot 1	Slot 2
00	M-unit	I-unit	I-unit
01	M-unit	I-unit	I-unit
02	M-unit	I-unit	I-unit
03	M-unit	I-unit	I-unit
04	M-unit	L-unit	X-unit
05	M-unit	L-unit	X-unit
06			
07			
08	M-unit	M-unit	I-unit
09	M-unit	M-unit	I-unit
0A	M-unit	M-unit	I-unit
0B	M-unit	M-unit	I-unit
0C	M-unit	F-unit	I-unit
0D	M-unit	F-unit	I-unit
0E	M-unit	M-unit	F-unit
0F	M-unit	M-unit	F-unit
10	M-unit	I-unit	B-unit
11	M-unit	I-unit	B-unit
12	M-unit	B-unit	B-unit
13	M-unit	B-unit	B-unit
14			
15			
16	B-unit	B-unit	B-unit
17	B-unit	B-unit	B-unit
18	M-unit	M-unit	B-unit
19	M-unit	M-unit	B-unit
1A			
1B			
1C	M-unit	F-unit	B-unit
1D	M-unit	F-unit	B-unit
1E			
1F			

Extended instructions, used for long immediate integer, occupy two instruction slots.

3.4 Instruction Sequencing

An IA-64 program consists of a sequence of instructions and stops packed in bundles. Instruction execution is ordered as follows:

- Bundles are ordered from lowest to highest memory address. Instructions in bundles with lower memory addresses are considered to precede instructions in bundles with higher memory addresses. The byte order of each bundle in memory is little-endian (the template field is contained in byte 0 of a bundle).
- Within a bundle, instructions are ordered from instruction slot 0 to instruction slot 2 as specified in Figure 3-16 on page 3-11.

For additional details on Instruction sequencing, refer to Appendix A, “Instruction Sequencing Considerations”.

This section describes the IA-64 architectural functionality from the perspective of the application programmer. IA-64 instructions are grouped into related functions and an overview of their behavior is given. Unless otherwise noted, all immediates are sign extended to 64 bits before use. The floating-point programming model is described separately in Chapter 5, “IA-64 Floating-point Programming Model”.

The main features of the IA-64 programming model covered here are:

- General Register Stack
- Integer Computation Instructions
- Compare Instructions and Predication
- Memory Access Instructions and Speculation
- Branch Instructions and Branch Prediction
- Multimedia Instructions
- Register File Transfer Instructions
- Character Strings and Population Count

4.1 Register Stack

As described in “General Registers” on page 3-3, the general register file is divided into static and stacked subsets. The static subset is visible to all procedures and consists of the 32 registers from GR 0 through GR 31. The stacked subset is local to each procedure and may vary in size from zero to 96 registers beginning at GR 32. The register stack mechanism is implemented by renaming register addresses as a side-effect of procedure calls and returns. The implementation of this rename mechanism is not otherwise visible to application programs. The register stack is disabled during IA-32 instruction set execution.

The static subset must be saved and restored at procedure boundaries according to software convention. The stacked subset is automatically saved and restored by the Register Stack Engine (RSE) without explicit software intervention. All other register files are visible to all procedures and must be saved/restored by software according to software convention.

4.1.1 Register Stack Operation

The registers in the stacked subset visible to a given procedure are called a register stack frame. The frame is further partitioned into two variable-size areas: the local area and the output area. Immediately after a call, the size of the local area of the newly activated frame is zero and the size of the output area is equal to the size of the caller’s output area and overlays the caller’s output area.

The local and output areas of a frame can be re-sized using the `alloc` instruction which specifies immediates that determine the size of frame (`sof`) and size of locals (`sol`). (Note that in the assembly language, `alloc` specifies three operands: the size of inputs immediate; the size of locals immediate; and the size of outputs immediate. The value of `sol` is determined by adding the size of inputs immediate and the size of locals immediate; the value of `sof` is determined by adding all three immediates.) The value of `sof` specifies the size of the entire stacked subset visible to the current procedure; the value of `sol` specifies the size of the local area. The size of the output area is determined by the difference between `sof` and `sol`. The values of these parameters for the currently active procedure are maintained in the Current Frame Marker (CFM).

Reading a stacked register outside the current frame will return an undefined result. Writing a stacked register outside the current frame will cause an Illegal Operation fault.

When a call-type branch is executed, the CFM is copied to the Previous Frame Marker (PFM) field in the Previous Function State application register (PFS), and the callee's frame is created as follows:

- The stacked registers are renamed such that the first register in the caller's output area becomes GR 32 for the callee
- The size of the local area is set to zero
- The size of the callee's frame (sof_{b1}) is set to the size of the caller's output area ($sof_a - sol_a$)

Values in the output area of the caller's register stack frame are visible to the callee. This overlap permits parameter and return value passing between procedures to take place entirely in registers.

Procedure frames may be dynamically re-sized by issuing an `alloc` instruction. An `alloc` instruction causes no renaming, but only changes the size of the register stack frame and the partitioning between local and output areas. Typically, when a procedure is called, it will allocate some number of local registers for its use (which will include the parameters passed to it in the caller's output registers), plus an output area (for passing parameters to procedures it will call). Newly allocated registers (including their NaT bits) have undefined values.

When a return-type branch is executed, CFM is restored from PFM and the register renaming is restored to the caller's configuration. The PFM is procedure local state and must be saved and restored by non-leaf procedures. The CFM is not directly accessible in application programs and is updated only through the execution of calls, returns, `alloc`, and `clr-rrb`.

Figure 4-1 depicts the behavior of the register stack on a procedure call from `procA` (caller) to `procB` (callee). The state of the register stack is shown at four points: prior to the call, immediately following the call, after `procB` has executed an `alloc`, and after `procB` returns to `procA`.

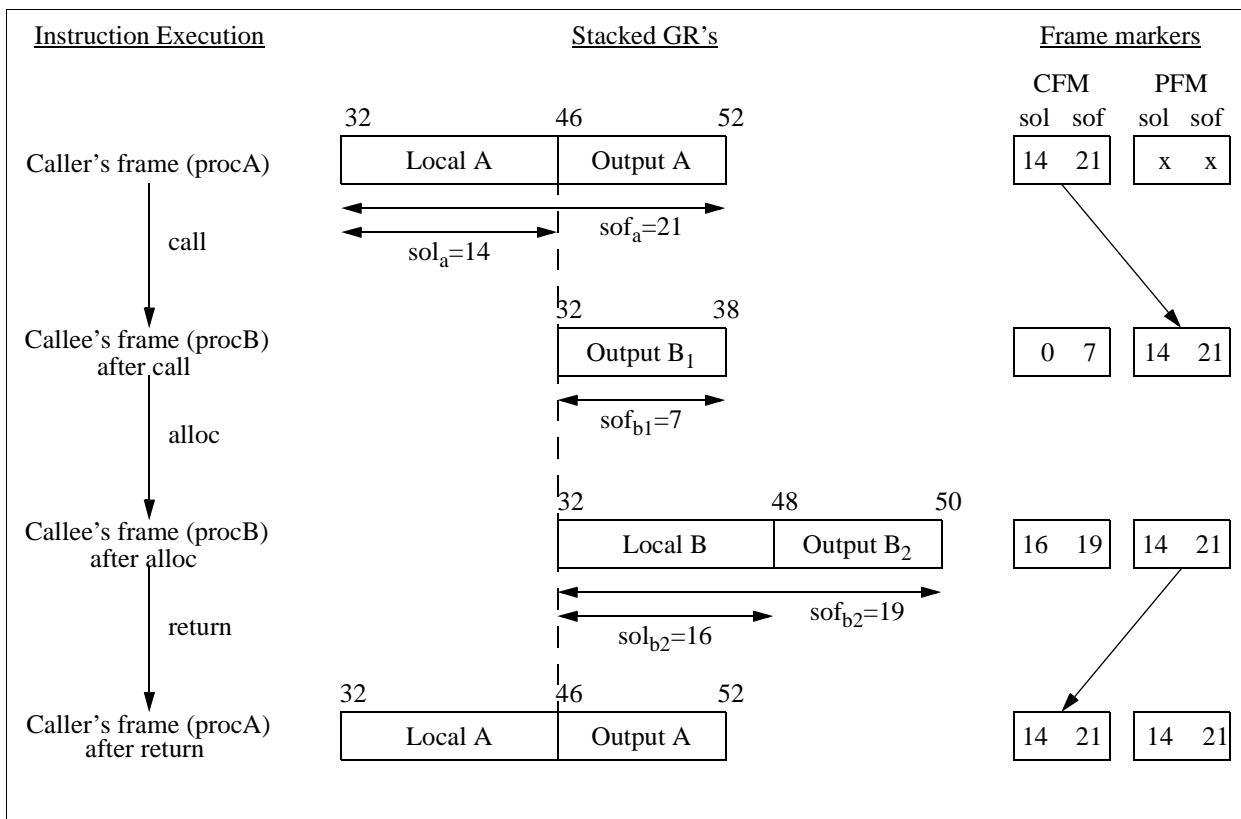


Figure 4-1. Register Stack Behavior on Procedure Call and Return

The majority of application programs need only issue `alloc` instructions and save/restore PFM in order to effectively utilize the register stack. A detailed knowledge of the RSE (Register Stack Engine) is required only by certain specialized application software such as user-level thread packages, debuggers, etc.

4.1.2 Register Stack Instructions

The `alloc` instruction is used to change the size of the current register stack frame. An `alloc` instruction must be the first instruction in an instruction group otherwise the results are undefined. An `alloc` instruction affects the register stack frame seen by all instructions in an instruction group, including the `alloc` itself. An `alloc` cannot be predicated. An `alloc` does not affect the values or NaT bits of the allocated registers. When a register stack frame is expanded, newly allocated registers may have their NaT bit set.

In addition, there are three instructions which provide explicit control over the state of the register stack. These instructions are used in thread and context switching which necessitate a corresponding switch of the backing store for the register stack.

The `flushrs` instruction is used to force all previous stack frames out to backing store memory. It stalls instruction execution until all active frames in the physical register stack up to, but not including the current frame are spilled to the backing store by the RSE. A `flushrs` instruction must be the first instruction in an instruction group; otherwise, the results are undefined. A `flushrs` cannot be predicated.

Table 4-1 lists the architectural visible state relating to the register stack. Table 4-2 summarizes the register stack management instructions. Call- and return-type branches, which affect the stack, are described in “Branch Instructions” on page 4-19.

Table 4-1. Architectural Visible State Related to the Register Stack

Register	Description
AR[PFS].pfm	Previous Frame Marker field
AR[RSC]	Register Stack Configuration application register
AR[BSP]	Backing store pointer application register
AR[BSPSTORE]	Backing store pointer application register for memory stores
AR[RNAT]	RSE NaT collection application register

Table 4-2. Register Stack Management Instructions

Mnemonic	Operation
<code>alloc</code>	Allocate register stack frame
<code>flushrs</code>	Flush register stack to backing store

4.2 Integer Computation Instructions

The integer execution units provide a set of arithmetic, logical, shift and bit-field-manipulation instructions. Additionally, they provide a set of instructions to accelerate operations on 32-bit data and pointers.

Arithmetic, logical and 32-bit acceleration instructions can be executed on both I- and M-units

4.2.1 Arithmetic Instructions

Addition and subtraction (`add`, `sub`) are supported with regular two input forms and special three input forms. The three input addition form adds one to the sum of two input registers. The three input subtraction form subtracts one from the difference of two input registers. The three input forms share the same mnemonics as the two input forms and are specified by appending a “1” as a third source operand.

Immediate forms of addition and subtraction use a register and a 15-bit immediate. The immediate form is obtained simply by specifying an immediate rather than a register as the first operand. Also, addition can be performed between a register and a 22-bit immediate; however, the source register must be GR 0, 1, 2 or 3.

A shift left and add instruction (`shladd`) shifts one register operand to the left by 1 to 4 bits and adds the result to a second register operand. Table 4-3 summarizes the integer arithmetic instructions.

Table 4-3. Integer Arithmetic Instructions

Mnemonic	Operation
add	Addition
add ...,1	Three input addition
sub	Subtraction
sub ...,1	Three input subtraction
shladd	Shift left and add

Note that an integer multiply instruction is defined which uses the floating-point registers. See “Integer Multiply and Add Instructions” on page 5-14 for details. Integer divide is performed in software similarly to floating-point divide.

4.2.2 Logical Instructions

Instructions to perform logical AND (`and`), OR (`or`), and exclusive OR (`xor`) between two registers or between a register and an immediate are defined. The `andcm` instruction performs a logical AND of a register or an immediate with the complement of another register. Table 4-4 summarizes the integer logical instructions.

Table 4-4. Integer Logical Instructions

Mnemonic	Operation
and	Logical and
or	Logical or
andcm	Logical and complement
xor	Logical exclusive or

4.2.3 32-Bit Addresses and Integers

Support for IA-64 32-bit addresses is provided in the form of add instructions that perform region bit copying. This supports the virtual address translation model. The add 32-bit pointer instruction (`addp`) adds two registers or a register and an immediate, zeroes the most significant 32-bits of the result, and copies bits 31:30 of the second source to bits 62:61 of the result. The `shladdp` instruction operates similarly but shifts the first source to the left by 1 to 4 bits before performing the add, and is provided only in the two-register form.

In addition, support for 32-bit integers is provided through 32-bit compare instructions and instructions to perform sign and zero extension. Compare instructions are described in “Compare Instructions and Predication” on page 4-5. The sign and zero extend (`sxt`, `zxt`) instructions take an 8-bit, 16-bit, or 32-bit value in a register, and produce a properly extended 64-bit result.

Table 4-5 summarizes 32-bit pointer and 32-bit integer instructions.

Table 4-5. 32-bit Pointer and 32-bit Integer Instructions

Mnemonic	Operation
addp	32-bit pointer addition
shladdp	Shift left and add 32-bit pointer
sxt	Sign extend
zxt	Zero extend

4.2.4 Bit Field and Shift Instructions

Four classes of instructions are defined for shifting and operating on bit fields within a general register: variable shifts, fixed shift-and-mask instructions, a 128-bit-input funnel shift, and special compare operations to test an individual bit within a general register. The compare instructions for testing a single bit (`tbit`), or for testing the NaT bit (`tnat`) are described in “Compare Instructions and Predication” on page 4-5.

The variable shift instructions shift the contents of a general register by an amount specified by another general register. The shift right signed (`shr`) and shift right unsigned (`shr.u`) instructions shift the contents of a register to the right with

the vacated bit positions filled with the sign bit or zeroes respectively. The shift left (`shl`) instruction shifts the contents of a register to the left.

The fixed shift-and-mask instructions (`extr`, `dep`) are generalized forms of fixed shifts. The extract instruction (`extr`) copies an arbitrary bit field from a general register to the least-significant bits of the target register. The remaining bits of the target are written with either the sign of the bit field (`extr`) or with zero (`extr.u`). The length and starting position of the field are specified by two immediates. This is essentially a shift-right-and-mask operation. A simple right shift by a fixed amount can be specified by using `shr` with an immediate value for the shift amount. This is just an assembly pseudo-op for an extract instruction where the field to be extracted extends all the way to the left-most register bit.

The deposit instruction (`dep`) takes a field from either the least-significant bits of a general register, or from an immediate value of all zeroes or all ones, places it at an arbitrary position, and fills the result to the left and right of the field with either bits from a second general register (`dep`) or with zeroes (`dep.z`). The length and starting position of the field are specified by two immediates. This is essentially a shift-left-mask-merge operation. A simple left shift by a fixed amount can be specified by using `shl` with an immediate value for the shift amount. This is just an assembly pseudo-op for `dep.z` where the deposited field extends all the way to the left-most register bit.

The shift right pair (`shrp`) instruction performs a 128-bit-input funnel shift. It extracts an arbitrary 64-bit field from a 128-bit field formed by concatenating two source general registers. The starting position is specified by an immediate. This can be used to accelerate the adjustment of unaligned data. A bit rotate operation can be performed by using `shrp` and specifying the same register for both operands.

Table 4-6 summarizes the bit field and shift instructions.

Table 4-6. Bit Field and Shift Instructions

Mnemonic	Operation
<code>shr</code>	Shift right signed
<code>shr.u</code>	Shift right unsigned
<code>shl</code>	Shift left
<code>extr</code>	Extract signed (shift right and mask)
<code>extr.u</code>	Extract unsigned (shift right and mask)
<code>dep</code>	Deposit (shift left, mask and merge)
<code>dep.z</code>	Deposit in zeroes (shift left and mask)
<code>shrp</code>	Shift right pair

4.2.5 Large Constants

A special instruction is defined for generating large constants (see Table 4-7). For constants up to 22 bits in size, the `add` instruction can be used, or the `mov` pseudo-op (pseudo-op of `add` with `GR0`, which always reads 0). For larger constants, the move long immediate instruction (`movl`) is defined to write a 64-bit immediate into a general register. This instruction occupies two instruction slots within the same bundle, and is the only such instruction.

Table 4-7. Instructions to Generate Large Constants

Mnemonic	Operation
<code>mov</code>	Move 22-bit immediate
<code>movl</code>	Move 64-bit immediate

4.3 Compare Instructions and Predication

A set of compare instructions provides the ability to test for various conditions and affect the dynamic execution of instructions. A compare instruction tests for a single specified condition and generates a boolean result. These results are written to predicate registers. The predicate registers can then be used to affect dynamic execution in two ways: as conditions for conditional branches, or as qualifying predicates for predication.

4.3.1 Predication

Predication is the conditional execution of instructions. The execution of most IA-64 instructions is gated by a qualifying predicate. If the predicate is true, the instruction executes normally; if the predicate is false, the instruction does not modify architectural state (except for the unconditional type of compare instructions, floating-point approximation instructions and while-loop branches). Predicates are one-bit values and are stored in the predicate register file. A zero predicate is interpreted as false and a one predicate is interpreted as true (predicate register PR0 is hardwired to one).

A few IA-64 instructions cannot be predicated. These instructions are: allocate stack frame (`alloc`), clear rrb (`clrrrb`), flush register stack (`flushrs`), and counted branches (`cloop`, `ctop`, `cexit`).

4.3.2 Compare Instructions

Predicate registers are written by the following instructions: general register compare (`cmp`, `cmp4`), floating-point register compare (`fcmp`), test bit and test NaT (`tbit`, `tnat`), floating-point class (`fclass`), and floating-point reciprocal approximation and reciprocal square root approximation (`frcpa`, `frsqrrta`). Most of these compare instructions (all but `frcpa` and `frsqrrta`) set two predicate registers based on the outcome of the comparison. The setting of the two target registers is described below in “Compare Types” on page 4-6. Compare instructions are summarized in Table 4-8.

Table 4-8. Compare Instructions

Mnemonic	Operation
<code>cmp</code> , <code>cmp4</code>	GR compare
<code>tbit</code>	Test bit in a GR
<code>tnat</code>	Test GR NaT bit
<code>fcmp</code>	FR compare
<code>fclass</code>	FR class
<code>frcpa</code> , <code>fprcpa</code>	Floating-point reciprocal approximation
<code>frsqrrta</code> , <code>fprsqrrta</code>	Floating-point reciprocal square root approximation

The 64-bit (`cmp`) and 32-bit (`cmp4`) compare instructions compare two registers, or a register and an immediate, for one of ten relations (e.g., `>`, `<=`). The compare instructions set two predicate targets according to the result. The `cmp4` instruction compares the least-significant 32-bits of both sources (the most significant 32-bits are ignored).

The test bit (`tbit`) instruction sets two predicate registers according to the state of a single bit in a general register (the position of the bit is specified by an immediate). The test NaT (`tnat`) instruction sets two predicate registers according to the state of the NaT bit corresponding to a general register.

The `fcmp` instruction compares two floating-point registers and sets two predicate targets according to one of eight relations. The `fclass` instruction sets two predicate targets according to the classification of the number contained in the floating-point register source.

The `frcpa` and `frsqrrta` instructions set a single predicate target if their floating-point register sources are such that a valid approximation can be produced, otherwise the predicate target is cleared.

4.3.3 Compare Types

Compare instructions can have as many as five compare types: Normal, Unconditional, AND, OR, or DeMorgan. The type defines how the instruction writes its target predicate registers based on the outcome of the comparison and on the qualifying predicate. The description of these types is contained in Table 4-9. In the table, “qp” refers to the value of the qualifying predicate of the compare and “result” refers to the outcome of the compare relation (one if the compare relation is true and zero if the compare relation is false).

Table 4-9. Compare Type Function

Compare Type	Completer	Operation	
		First Predicate Target	Second Predicate Target
Normal	<i>none</i>	if (qp) {target = result}	if (qp) {target = !result}
Unconditional	<i>unc</i>	if (qp) {target = result} else {target = 0}	if (qp) {target = !result} else {target = 0}
AND	<i>and</i> <i>andcm</i>	if (qp && !result) {target = 0} if (qp && result) {target = 0}	if (qp && !result) {target = 0} if (qp && result) {target = 0}
OR	<i>or</i> <i>orcm</i>	if (qp && result) {target = 1} if (qp && !result) {target = 1}	if (qp && result) {target = 1} if (qp && !result) {target = 1}
DeMorgan	<i>or.andcm</i> <i>and.orcm</i>	if (qp && result) {target = 1} if (qp && !result) {target = 0}	if (qp && result) {target = 0} if (qp && !result) {target = 1}

The Normal compare type simply writes the compare result to the first predicate target and the complement of the result to the second predicate target.

The Unconditional compare type behaves the same as the Normal type, except that if the qualifying predicate is 0, both predicate targets are written with 0. This can be thought of as an initialization of the predicate targets, combined with a Normal compare. Note that compare instructions with the Unconditional type modify architectural state when their qualifying predicate is false.

The AND, OR and DeMorgan types are termed “parallel” compare types because they allow multiple simultaneous compares (of the same type) to target a single predicate register. This provides the ability to compute a logical equation such as $p5 = (r4 == 0) \ || \ (r5 == r6)$ in a single cycle (assuming p5 was initialized to 0 in an earlier cycle). The DeMorgan compare type is just a combination of an OR type to one predicate target and an AND type to the other predicate target. Multiple OR-type compares (including the OR part of the DeMorgan type) may specify the same predicate target in the same instruction group. Multiple AND-type compares (including the AND part of the DeMorgan type) may also specify the same predicate target in the same instruction group.

For all compare instructions (except for *tnat* and *fclass*), if one or both of the source registers contains a deferred exception token (NaT or NaTVal – see “Control Speculation” on page 4-10), the result of the compare is different. Both predicate targets are treated the same, and are either written to 0 or left unchanged. In combination with speculation, this allows predicated code to be turned off in the presence of a deferred exception. (*fclass* behaves this way as well if NaTVal is not one of the classes being tested for.) Table 4-10 describes the behavior.

Table 4-10. Compare Outcome with NaT Source Input

Compare Type	Operation
Normal	if (qp) {target = 0}
Unconditional	target = 0
AND	if (qp) {target = 0}
OR	(not written)
DeMorgan	(not written)

Only a subset of the compare types are provided for some of the compare instructions. Table 4-11 lists the compare types which are available for each of the instructions.

Table 4-11. Instructions and Compare Types Provided

Instruction	Relation	Types Provided
<i>cmp</i> , <i>cmp4</i>	$a == b$, $a != b$, $a > 0$, $a >= 0$, $a < 0$, $a <= 0$, $0 > a$, $0 >= a$, $0 < a$, $0 <= a$ All other relations	Normal, Unconditional, AND, OR, DeMorgan
<i>tbit</i> , <i>tnat</i>	All	Normal, Unconditional Normal, Unconditional, AND, OR, DeMorgan

Table 4-11. Instructions and Compare Types Provided (Continued)

Instruction	Relation	Types Provided
<code>fcmp, fclass</code>	All	Normal, Unconditional
<code>frcpa, frsqrta,</code> <code>fprcpa, fprsqrta</code>	Not Applicable	Unconditional

4.3.4 Predicate Register Transfers

Instructions are provided to transfer between the predicate register file and a general register. These instructions operate in a “broadside” manner whereby multiple predicate registers are transferred in parallel, such that predicate register N is transferred to/from bit N of a general register.

The move to predicates instruction (`mov pr=`) loads multiple predicate registers from a general register according to a mask specified by an immediate. The mask contains one bit for each of PR 1 through PR 15 (PR 0 is hardwired to 1) and one bit for all of PR 16 through PR63 (the rotating predicates). A predicate register is written from the corresponding bit in a general register if the corresponding mask bit is 1; if the mask bit is 0 the predicate register is not modified.

The move to rotating predicates instruction (`mov pr.rot=`) copies 48 bits from an immediate value into the 48 rotating predicates (PR 16 through PR 63). The immediate value includes 28 bits, and is sign-extended. Thus PR 16 through PR 42 can be independently set to new values, and PR 43 through PR 63 are all set to either 0 or 1.

The move from predicates instruction (`mov =pr`) transfers the entire predicate register file into a general register target.

For all of these predicate register transfers, the predicate registers are accessed as though the register rename base (CFM.rrb.pr) were 0. Typically, therefore, software should clear CFM.rrb.pr before initializing rotating predicates.

4.4 Memory Access Instructions

Memory is accessed by simple load, store and semaphore instructions, which transfer data to and from general registers or floating-point registers. The memory address is specified by the contents of a general register.

Most load and store instructions can also specify base-address-register update. Base update adds either an immediate value or the contents of a general register to the address register, and places the result back in the address register. The update is done after the load or store operation, i.e., it is performed as an address post-increment.

For highest performance, data should be aligned on natural boundaries. Within a 4K-byte boundary, accesses misaligned with respect to their natural boundaries will always fault if UM.ac (alignment check bit in the User Mask register) is 1. If UM.ac is 0, then an unaligned access will succeed if it is supported by the implementation; otherwise it will cause an Unaligned Data Reference fault. All memory accesses that cross a 4K-byte boundary will cause an Unaligned Data Reference fault independent of UM.ac. Additionally, all semaphore instructions will cause an Unaligned Data Reference fault if the access is not aligned to its natural boundary, independent of UM.ac.

Accesses to memory quantities larger than a byte may be done in a big-endian or little-endian fashion. The byte ordering for all memory access instructions is determined by UM.be in the User Mask register for IA-64 memory references. All IA-32 memory references are performed little-endian.

Load, store and semaphore instructions are summarized in Table 4-12.

Table 4-12. Memory Access Instructions

Mnemonic			Operation
General	Floating-point		
	Normal	Load Pair	
ld	ldf	ldfp	Load
ld.s	ldf.s	ldfp.s	Speculative load
ld.a	ldf.a	ldfp.a	Advanced load
ld.sa	ldf.sa	ldfp.sa	Speculative advanced load
ld.c.nc, ld.c.clr	ldf.c.nc, ldf.c.clr	ldfp.c.nc, ldfp.c.clr	Check load
ld.c.clr.acq			Ordered check load
ld.acq			Ordered load
ld.bias			Biased load
ld8.fill	ldf.fill		Fill
st	stf		Store
st.rel			Ordered store
st.spill	stf.spill		Spill
cpxchg			Compare and exchange
xchg			Exchange memory and GR
fetchadd			Fetch and add

4.4.1 Load Instructions

Load instructions transfer data from memory to a general register, a floating-point register or a pair of floating-point registers.

For general register loads, access sizes of 1, 2, 4, and 8 bytes are defined. For sizes less than eight bytes, the loaded value is zero extended to 64-bits.

For floating-point loads, five access sizes are defined: single precision (4 bytes), double precision (8 bytes), double-extended precision (10 bytes), single precision pair (8 bytes), and double precision pair (16 bytes). The value(s) loaded from memory are converted into floating-point register format (see “Memory Access Instructions” on page 5-6 for details). The floating-point load pair instructions load two adjacent single or double precision numbers into two independent floating-point registers (see the `ldfp[s/d]` instruction description for restrictions on target register specifiers). The floating-point load pair instructions cannot specify base register update.

Variants of both general and floating-point register loads are defined for supporting compiler-directed control and data speculation. These use the general register NaT bits and the ALAT. See “Control Speculation” on page 4-10 and “Data Speculation” on page 4-12.

Variants are also provided for controlling the memory/cache subsystem. An ordered load can be used to force ordering in memory accesses. See “Memory Access Ordering” on page 4-18. A biased load provides a hint to acquire exclusive ownership of the accessed line. See “Memory Hierarchy Control and Consistency” on page 4-16.

Special-purpose loads are defined for restoring register values that were spilled to memory. The `ld8.fill` instruction loads a general register and the corresponding NaT bit (defined for an 8-byte access only). The `ldf.fill` instruction loads a value in floating-point register format from memory without conversion (defined for 16-byte access only). See “Register Spill and Fill” on page 4-12.

4.4.2 Store Instructions

Store instructions transfer data from a general or floating-point register to memory. Store instructions are always non-speculative. Store instructions can specify base-address-register update, but only by an immediate value. A variant is also provided for controlling the memory/cache subsystem. An ordered store can be used to force ordering in memory accesses.

Both general and floating-point register stores are defined with the same access sizes as their load counterparts. The only exception is that there are no floating-point store pair instructions.

Special purpose stores are defined for spilling register values to memory. The `st8.spill` instruction stores a general register and the corresponding NaT bit (defined for 8-byte access only). This allows the result of a speculative calculation to be spilled to memory and restored. The `stf.spill` instruction stores a floating-point register in memory in the floating-point register format without conversion. This allows register spill and restore code to be written to be compatible with possible future extensions to the floating-point register format. The `stf.spill` instruction also does not fault if the register contains a NaTVal, and is defined for 16-byte access only. See “Register Spill and Fill” on page 4-12.

4.4.3 Semaphore Instructions

Semaphore instructions atomically load a general register from memory, perform an operation and then store a result to the same memory location. Semaphore instructions are always non-speculative. No base register update is provided.

Three types of atomic semaphore operations are defined: exchange (`xchg`); compare and exchange (`cmpxchg`); and fetch and add (`fetchadd`).

The `xchg` target is loaded with the zero-extended contents of the memory location addressed by the first source and then the second source is stored into the same memory location.

The `cmpxchg` target is loaded with the zero-extended contents of the memory location addressed by the first source; if the zero-extended value is equal to the contents of the Compare and Exchange Compare Value application register (CCV), then the second source is stored into the same memory location.

The `fetchadd` instruction specifies one general register source, one general register target, and an immediate. The `fetchadd` target is loaded with the zero-extended contents of the memory location addressed by the source and then the immediate is added to the loaded value and the result is stored into the same memory location.

Table 4-13. State Relating to Memory Access

Register	Function
UM.be	User mask byte ordering
UM.ac	User mask Unaligned Data Reference fault enable
UNAT	GR NaT collection
CCV	Compare and Exchange Compare Value application register

4.4.4 Control Speculation

Special mechanisms are provided to allow for compiler-directed speculation. This speculation takes two forms, control speculation and data speculation, with a separate mechanism to support each. See also “Data Speculation” on page 4-12.

4.4.4.1 Control Speculation Concepts

Control speculation describes the compiler optimization where an instruction or a sequence of instructions is executed before it is known that the dynamic control flow of the program will actually reach the point in the program where the sequence of instructions is needed. This is done with instruction sequences that have long execution latencies. Starting the execution early allows the compiler to overlap the execution with other work, increasing the parallelism and decreasing overall execution time. The compiler performs this optimization when it determines that it is very likely that the dynamic control flow of the program will eventually require this calculation. In cases where the control flow is such that the calculation turns out not to be needed, its results are simply discarded (the results in processor registers are simply not used).

Since the speculative instruction sequence may not be required by the program, no exceptions encountered that would be visible to the program can be signalled until it is determined that the program’s control flow does require the execution of this instruction sequence. For this reason, a mechanism is provided for recording the occurrence of an exception so that it can be signalled later if and when it is necessary. In such a situation, the exception is said to be deferred. When an exception is deferred by an instruction, a special token is written into the target register to indicate the existence of a deferred exception in the program.

Deferred exception tokens are represented differently in the general and floating-point register files. In general registers, an additional bit is defined for each register called the NaT bit (Not a Thing). Thus general registers are 65 bits wide. A NaT bit equal to 1 indicates that the register contains a deferred exception token, and that its 64-bit data portion contains an implementation specific value that software cannot rely upon. In floating-point registers, a deferred exception is indi-

cated by a specific pseudo-zero encoding called the NaTVal (see “Representation of Values in Floating-point Registers” on page 5-2 for details).

4.4.4.2 Control Speculation and Instructions

Instructions are divided into two categories: speculative (instructions which can be used speculatively) and non-speculative (instructions which cannot). Non-speculative instructions will raise exceptions if they occur and are therefore unsafe to schedule before they are known to be executed. Speculative instructions defer exceptions (they do not raise them) and are therefore safe to schedule before they are known to be executed.

Loads to general and floating-point registers have both non-speculative (`ld`, `ldf`, `ldfp`) and speculative (`ld.s`, `ldf.s`, `ldfp.s`) variants. Generally, all computation instructions which write their results to general or floating-point registers are speculative. Any instruction that modifies state other than a general or floating-point register is non-speculative, since there would be no way to represent the deferred exception (there are a few exceptions).

Deferred exception tokens propagate through the program in a dataflow manner. A speculative instruction that reads a register containing a deferred exception token will propagate a deferred exception token into its target. Thus a chain of instructions can be executed speculatively, and only the result register need be checked for a deferred exception token to determine whether any exceptions occurred.

At the point in the program when it is known that the result of a speculative calculation is needed, a speculation check (`chk.s`) instruction is used. This instruction tests for a deferred exception token. If none is found, then the speculative calculation was successful, and execution continues normally. If a deferred exception token is found, then the speculative calculation was unsuccessful and must be re-done. In this case, the `chk.s` instruction branches to a new address (specified by an immediate offset in the `chk.s` instruction). Software can use this mechanism to invoke code that contains a copy of the speculative calculation (but with non-speculative loads). Since it is now known that the calculation is required, any exceptions which now occur can be signalled and handled normally.

Since computational instructions do not generally cause exceptions, the only instructions which generate deferred exception tokens are speculative loads. (IEEE floating-point exceptions are handled specially through a set of alternate status fields. See “Floating-point Status Register” on page 5-4.) Other speculative instructions propagate deferred exception tokens, but do not generate them.

4.4.4.3 Control Speculation and Compares

As stated earlier, most instructions that write a register file other than the general registers or the floating-point registers are non-speculative. The compare (`cmp`, `cmp4`, `fcmp`), test bit (`tbit`), floating-point class (`fclass`), and floating-point approximation (`frcpa`, `frsqrrta`) instructions are special cases. These instructions read general or floating-point registers and write one or two predicate registers.

For these instructions, if any source contains a deferred exception token, all predicate targets are either cleared or left unchanged, depending on the compare type (see Table 4-10 on page 4-7). Software can use this behavior to ensure that any dependent conditional branches are not taken and any dependent predicated instructions are nullified. See “Predication” on page 4-6.

Deferred exception tokens can also be tested for with certain compare instructions. The test NaT (`tnat`) instruction tests the NaT bit corresponding to the specified general register and writes two predicate results. The floating-point class (`fclass`) instruction can be used to test for a NaTVal in a floating-point register and write the result to two predicate registers. (`fclass` does not clear both predicate targets in the presence of a NaTVal input if NaTVal is one of the classes being tested for.)

4.4.4.4 Control Speculation without Recovery

A non-speculative instruction that reads a register containing a deferred exception token will raise a Register NaT Consumption fault. Such instructions can be thought of as performing a non-recoverable speculation check operation. In some compilation environments, it may be true that the only exceptions that are deferred are fatal errors. In such a program, if the result of a speculative calculation is checked and a deferred exception token is found, execution of the program is terminated. For such a program, the results of speculative calculations can be checked simply by using non-speculative instructions.

4.4.4.5 Register Spill and Fill

Special store and load instructions are provided for spilling a register to memory and preserving any deferred exception token, and for restoring a spilled register.

The spill and fill general register instructions (`st8.spill`, `ld8.fill`) are defined to save/restore a general register along with the corresponding NaT bit.

The `st8.spill` instruction writes a general register's NaT bit into the User NaT Collection application register (UNAT), and, if the NaT bit was 0, writes the register's 64-bit data portion to memory. If the register's NaT bit was 1, the UNAT is updated, but the memory update is implementation specific, and must consistently follow one of three spill behaviors:

- The `st8.spill` may not update memory with the register's 64-bit data portion, or
- The `st8.spill` may write a zero to the specified memory location, or
- The `st8.spill` may write the register's 64-bit data portion to memory, only if that implementation returns a zero into the target register of all NaT-ed speculative loads, and that implementation also guarantees that all NaT propagating instructions perform all computations as specified by the instruction pages.

Bits 8:3 of the memory address determine which bit in the UNAT register is written.

The `ld8.fill` instruction loads a general register from memory taking the corresponding NaT bit from the bit in the UNAT register addressed by bits 8:3 of the memory address. The UNAT register must be saved and restored by software. It is the responsibility of software to ensure that the contents of the UNAT register are correct while executing `st8.spill` and `ld8.fill` instructions.

The floating-point spill and fill instructions (`stf.spill`, `ldf.fill`) are defined to save/restore a floating-point register (saved as 16 bytes) without surfacing an exception if the FR contains a NaTVal (these instructions do not affect the UNAT register).

The general and floating-point spill/fill instructions allow spilling/filling of registers that are targets of a speculative instruction and may therefore contain a deferred exception token. Note also that transfers between the general and floating-point register files cause a conversion between the two deferred exception token formats.

Table 4-14 lists the state relating to control speculation. Table 4-15 summarizes the instructions related to control speculation.

Table 4-14. State Related to Control Speculation

Register	Description
NaT bits	65th bit associated with each GR indicating a deferred exception
NaTVal	Pseudo-Zero encoding for FR indicating a deferred exception
UNAT	User NaT collection application register

Table 4-15. Instructions Related to Control Speculation

Mnemonic	Operation
<code>ld.s</code> , <code>ldf.s</code> , <code>ldfp.s</code>	GR and FR speculative loads
<code>ld8.fill</code> , <code>ldf.fill</code>	Fill GR with NaT collection, fill FR
<code>st8.spill</code> , <code>stf.spill</code>	Spill GR with NaT collection, spill FR
<code>chk.s</code>	Test GR or FR for deferred exception token
<code>tnat</code>	Test GR NaT bit and set predicate

4.4.5 Data Speculation

Just as control speculative loads and checks allow the compiler to schedule instructions across control dependences, data speculative loads and checks allow the compiler to schedule instructions across some types of ambiguous data dependences. This section details the usage model and semantics of data speculation and related instructions.

4.4.5.1 Data Speculation Concepts

An ambiguous memory dependence is said to exist between a store (or any operation that may update memory state) and a load when it cannot be statically determined whether the load and store might access overlapping regions of memory. For convenience, a store that cannot be statically disambiguated relative to a particular load is said to be ambiguous relative to that load. In such cases, the compiler cannot change the order in which the load and store instructions were originally specified in the program. To overcome this scheduling limitation, a special kind of load instruction called an advanced load can be scheduled to execute earlier than one or more stores that are ambiguous relative to that load.

As with control speculation, the compiler can also speculate operations that are dependent upon the advanced load and later insert a check instruction that will determine whether the speculation was successful or not. For data speculation, the check can be placed anywhere the original non-data speculative load could have been scheduled.

Thus, a data-speculative sequence of instructions consists of an advanced load, zero or more instructions dependent on the value of that load, and a check instruction. This means that any sequence of stores followed by a load can be transformed into an advanced load followed by a sequence of stores followed by a check. The decision to perform such a transformation is highly dependent upon the likelihood and cost of recovering from an unsuccessful data speculation.

4.4.5.2 Data Speculation and Instructions

Advanced loads are available in integer (`ld.a`), floating-point (`ldf.a`), and floating-point pair (`ldfp.a`) forms. When an advanced load is executed, it allocates an entry in a structure called the Advanced Load Address Table (ALAT). Later, when a corresponding check instruction is executed, the presence of an entry indicates that the data speculation succeeded; otherwise, the speculation failed and one of two kinds of compiler-generated recovery is performed:

1. The check load instruction (`ld.c`, `ldf.c`, or `ldfp.c`) is used for recovery when the only instruction scheduled before a store that is ambiguous relative to the advanced load is the advanced load itself. The check load searches the ALAT for a matching entry. If found, the speculation was successful; if a matching entry was not found, the speculation was unsuccessful and the check load reloads the correct value from memory. Figure 4-2 shows this transformation.

Before Data Speculation	After Data Speculation
<pre>// other instructions st8 [r4] = r12 ld8 r6 = [r8];; add r5 = r6, r7;; st8 [r18] = r5</pre>	<pre>ld8.a r6 = [r8];; // advanced load // other instructions st8 [r4] = r12 ld8.c.clr r6 = [r8] // check load add r5 = r6, r7;; st8 [r18] = r5</pre>

Figure 4-2. Data Speculation Recovery Using `ld.c`

2. The advanced load check (`chk.a`) is used when an advanced load and several instructions that depend on the loaded value are scheduled before a store that is ambiguous relative to the advanced load. The advanced load check works like the speculation check (`chk.s`) in that, if the speculation was successful, execution continues inline and no recovery is necessary; if speculation was unsuccessful, the `chk.a` branches to compiler-generated recovery code. The recovery code contains instructions that will re-execute all the work that was dependent on the failed data speculative load up to the point of the check instruction. As with the check load, the success of a data speculation using an advanced load check is determined by searching the ALAT for a matching entry. This transformation is shown in Figure 4-3.

Before Data Speculation	After Data Speculation
<pre>// other instructions st8 [r4] = r12 ld8 r6 = [r8];; add r5 = r6, r7;; st8 [r18] = r5</pre>	<pre>ld8.a r6 = [r8];; // other instructions add r5 = r6, r7;; // other instructions st8 [r4] = r12 chk.a.clr r6, recover back: st8 [r18] = r5 // somewhere else in program recover: ld8 r6 = [r8];; add r5 = r6, r7 br back</pre>

Figure 4-3. Data Speculation Recovery Using `chk.a`

Recovery code may use either a normal or advanced load to obtain the correct value for the failed advanced load. An advanced load is used only when it is advantageous to have an ALAT entry reallocated after a failed speculation. The last instruction in the recovery code should branch to the instruction following the `chk.a`.

4.4.5.3 Detailed Functionality of the ALAT and Related Instructions

The ALAT is the structure that holds the state necessary for advanced loads and checks to operate correctly. The ALAT is searched in two different ways: by physical addresses and by ALAT register tags. An ALAT register tag is a unique number derived from the physical target register number and type in conjunction with other implementation-specific state. Implementation-specific state might include register stack wrap-around information to distinguish one instance of a physical register that may have been spilled by the RSE from the current instance of that register, thus avoiding the need to purge the ALAT on all register stack wrap-arounds.

IA-32 instruction set execution leaves the contents of the ALAT undefined. Software can not rely on ALAT values being preserved across an instruction set transition. On entry to IA-32 instruction set, existing entries in the ALAT are ignored.

Allocating and Checking ALAT Entries

Advanced loads perform the following actions:

1. The ALAT register tag for the advanced load is computed. (For `ldf.p.a`, a tag is computed only for the first target register.)
2. If an entry with a matching ALAT register tag exists, it is removed.
3. A new entry is allocated in the ALAT which contains the new ALAT register tag, the load access size, and a tag derived from the physical memory address.
4. The value at the address specified in the advanced load is loaded into the target register and, if specified, the base register is updated and an implicit prefetch is performed.

Since the success of a check is determined by finding a matching register tag in the ALAT, both the `chk.a` and the target register of a `ld.c` must specify the same register as their corresponding advanced load. Additionally, the check load must use the same address and operand size as the corresponding advanced load; otherwise, the value written into the target register by the check load is undefined.

An advanced load check performs the following actions:

1. It looks for a matching ALAT entry and if found, falls through to the next instruction.
2. If no matching entry is found, the `chk.a` branches to the specified address.

An implementation may choose to implement a failing advanced load check directly as a branch or as a fault where the fault-handler emulates the branch. Although the expected mode of operation is for an implementation to detect matching entries in the ALAT during checks, an implementation may fail a check instruction even when an entry with a matching ALAT register tag exists. This will be a rare occurrence but software must not assume that the ALAT does not contain the entry.

A check load checks for a matching entry in the ALAT. If no matching entry is found, it reloads the value from memory and any faults that occur during the memory reference are raised. When a matching entry is found, the target register is left unchanged.

If the check load was an ordered check load (`ld.c.clr.acq`), then it is performed with the semantics of an ordered load (`ld.acq`). ALAT register tag lookups by advanced load checks and check loads are subject to memory ordering constraints as outlined in “Memory Access Ordering” on page 4-18.

In addition to the flexibility described above, the size, organization, matching algorithm, and replacement algorithm of the ALAT are implementation dependent. Thus, the success or failure of specific advanced loads and checks in a program may change: when the program is run on different processor implementations, within the execution of a single program on the same implementation, or between different runs on the same implementation.

Invalidating ALAT Entries

In addition to entries removed by advanced loads, ALAT entry invalidations can occur implicitly by events that alter memory state or explicitly by any of the following instructions: `ld.c.clr`, `ld.c.clr.acq`, `chk.a.clr`, `invala`, `invala.e`. Events that may implicitly invalidate ALAT entries include those that change memory state or memory translation state such as:

1. The execution of stores or semaphores on other processors in the coherence domain.
2. The execution of store or semaphore instructions issued on the local processor.

When one of these events occurs, hardware checks each memory region represented by an entry in the ALAT to see if it overlaps with the locations affected by the invalidation event. ALAT entries whose memory regions overlap with the invalidation event locations are removed.

4.4.5.4 Combining Control and Data Speculation

Control speculation and data speculation are not mutually exclusive; a given load may be both control and data speculative. Both control speculative (`ld.sa`, `ldf.sa`, `ldfp.sa`) and non-control speculative (`ld.a`, `ldf.a`, `ldfp.a`) variants of advanced loads are defined for general and floating-point registers. If a speculative advanced load generates a deferred exception token then:

1. Any existing ALAT entry with the same ALAT register tag is invalidated.
2. No new ALAT entry is allocated.
3. If the target of the load was a general-purpose register, its NaT bit is set.
4. If the target of the load was a floating-point register, then NaTVal is written to the target register.

If a speculative advanced load does not generate a deferred exception, then its behavior is the same as the corresponding non-control speculative advanced load.

Since there can be no matching entry in the ALAT after a deferred fault, a single advanced load check or check load is sufficient to check both for data speculation failures and to detect deferred exceptions.

4.4.5.5 Instruction Completers for ALAT Management

To help the compiler manage the allocation and deallocation of ALAT entries, two variants of advanced load checks and check loads are provided: variants with clear (`chk.a.clr`, `ld.c.clr`, `ld.c.clr.acq`, `ldf.c.clr`, `ldfp.c.clr`) and variants with no clear (`chk.a.nc`, `ld.c.nc`, `ldf.c.nc`, `ldfp.c.nc`).

The clear variants are used when the compiler knows that the ALAT entry will not be used again and wants the entry explicitly removed. This allows software to indicate when entries are unneeded, making it less likely that a useful entry will be unnecessarily forced out because all entries are currently allocated.

For the clear variants of check load, any ALAT entry with the same ALAT register tag is invalidated independently of whether the address or size fields of the check load and the corresponding advanced load match. For `chk.a.clr`, the entry is guaranteed to be invalidated only when the instruction falls through (the recovery code is not executed). Thus, a failing `chk.a.clr` may or may not clear any matching ALAT entries. In such cases, the recovery code must explicitly invalidate the entry in question if program correctness depends on the entry being absent after a failed `chk.a.clr`.

Non-clear variants of both kinds of data speculation checks act as a hint to the processor that an existing entry should be maintained in the ALAT or that a new entry should be allocated when a matching ALAT entry doesn't exist. Such variants can be used within loops to check advanced loads which were presumed loop-invariant and moved out of the loop by the compiler. This behavior ensures that if the check load fails on one iteration, then the check load will not necessarily fail on all subsequent iterations. Whenever a new entry is inserted into the ALAT or when the contents of an entry are updated, the information written into the ALAT only uses information from the check load and does not use any residual information from a prior entry. The non-clear variant of `chk.a`, `chk.a.nc`, does not allocate entries and the 'nc' completer acts as a hint to the processor that the entry should not be cleared.

Table 4-16 and Table 4-17 summarize state and instructions relating to data speculation.

Structure	Function
ALAT	Advanced load address table

Table 4-16. State Relating to Data Speculation

Mnemonic	Operation
<code>ld.a</code> , <code>ldf.a</code> , <code>ldfp.a</code>	GR and FR advanced load
<code>st</code> , <code>st.rel</code> , <code>st8.spill</code> , <code>stf</code> , <code>stf.spill</code>	GR and FR store
<code>cmpxchg</code> , <code>fetchadd</code> , <code>xchg</code>	GR semaphore
<code>ld.c.clr</code> , <code>ld.c.clr.acq</code> , <code>ldf.c.clr</code> , <code>ldfp.c.clr</code>	GR and FR check load, clear on ALAT hit
<code>ld.c.nc</code> , <code>ldf.c.nc</code> , <code>ldfp.c.nc</code>	GR and FR check load, re-allocate on ALAT miss
<code>ld.sa</code> , <code>ldf.sa</code> , <code>ldfp.sa</code>	GR and FR speculative advanced load
<code>chk.a.clr</code> , <code>chk.a.nc</code>	GR and FR advanced load check
<code>invala</code>	Invalidate all ALAT entries
<code>invala.e</code>	Invalidate individual ALAT entry for GR or FR

Table 4-17. Instructions Relating to Data Speculation

4.4.6 Memory Hierarchy Control and Consistency

4.4.6.1 Hierarchy Control and Hints

IA-64 memory access instructions are defined to specify whether the data being accessed possesses temporal locality. In addition, memory access instructions can specify which levels of the memory hierarchy are affected by the access. This leads to an architectural view of the memory hierarchy depicted in Figure 4-4 composed of zero or more levels of cache between the register files and memory where each level may consist of two parallel structures: a temporal structure and a non-temporal structure. Note that this view applies to data accesses and not instruction accesses.

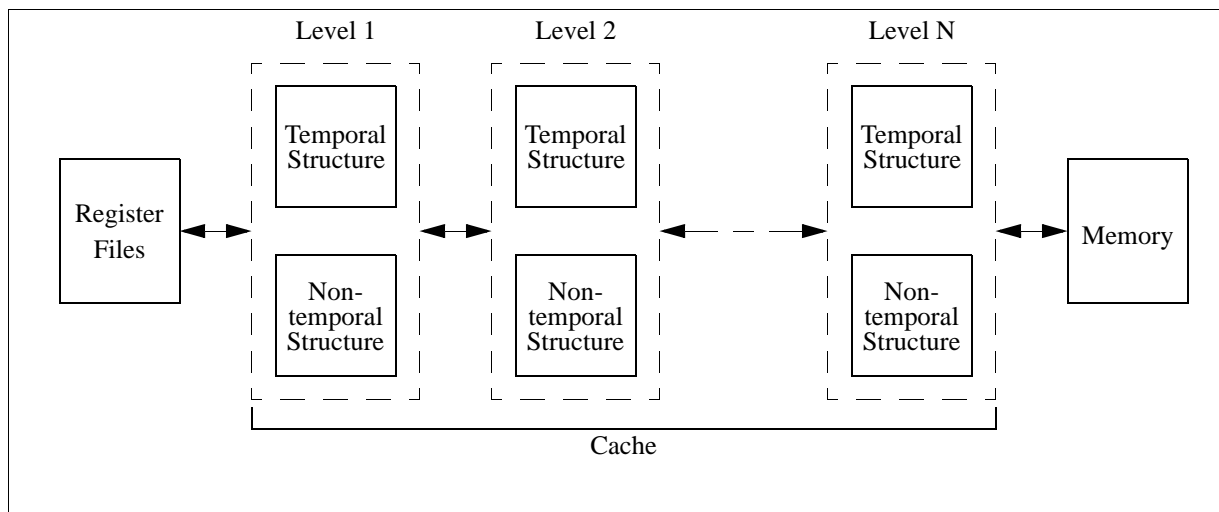


Figure 4-4. Memory Hierarchy

The temporal structures cache memory accessed with temporal locality; the non-temporal structures cache memory accessed without temporal locality. Both structures assume that memory accesses possess spatial locality. The existence of separate temporal and non-temporal structures, as well as the number of levels of cache, is implementation dependent.

Three mechanisms are defined for allocation control: locality hints; explicit prefetch; and implicit prefetch. Locality hints are specified by load, store, and explicit prefetch (`lfetch`) instructions. A locality hint specifies a hierarchy level (e.g., 1, 2, all). An access that is temporal with respect to a given hierarchy level is treated as temporal with respect to all lower (higher numbered) levels. An access that is non-temporal with respect to a given hierarchy level is treated as temporal with respect to all lower levels. Finding a cache line closer in the hierarchy than specified in the hint does not demote the line. This enables the precise management of lines using `lfetch` and then subsequent uses by `.nta` loads and stores to retain that level in the hierarchy. For example, specifying the `.nt2` hint by a prefetch indicates that the data should be cached at level 3. Subsequent loads and stores can specify `.nta` and have the data remain at level 3.

Locality hints do not affect the functional behavior of the program and may be ignored by the implementation. The locality hints available to loads, stores, and explicit prefetch instructions are given in Table 4-18. Instruction accesses are considered to possess both temporal and spatial locality with respect to level 1.

Table 4-18. Locality Hints Specified by Each Instruction Class

Mnemonic	Locality Hint	Instruction Type		
		Load	Store	lfetch, lfetch.fault
<code>none</code>	Temporal, level 1	x	x	x
<code>nt1</code>	Non-temporal, level 1	x		x
<code>nt2</code>	Non-temporal, level 2			x
<code>nta</code>	Non-temporal, all levels	x	x	x

Each locality hint implies a particular allocation path in the memory hierarchy. The allocation paths corresponding to the locality hints are depicted in Figure 4-5. The allocation path specifies the structures in which the line containing the data being referenced would best be allocated. If the line is already at the same or higher level in the hierarchy no movement occurs. Hinting that a datum should be cached in a temporal structure indicates that it is likely to be read in the near future.

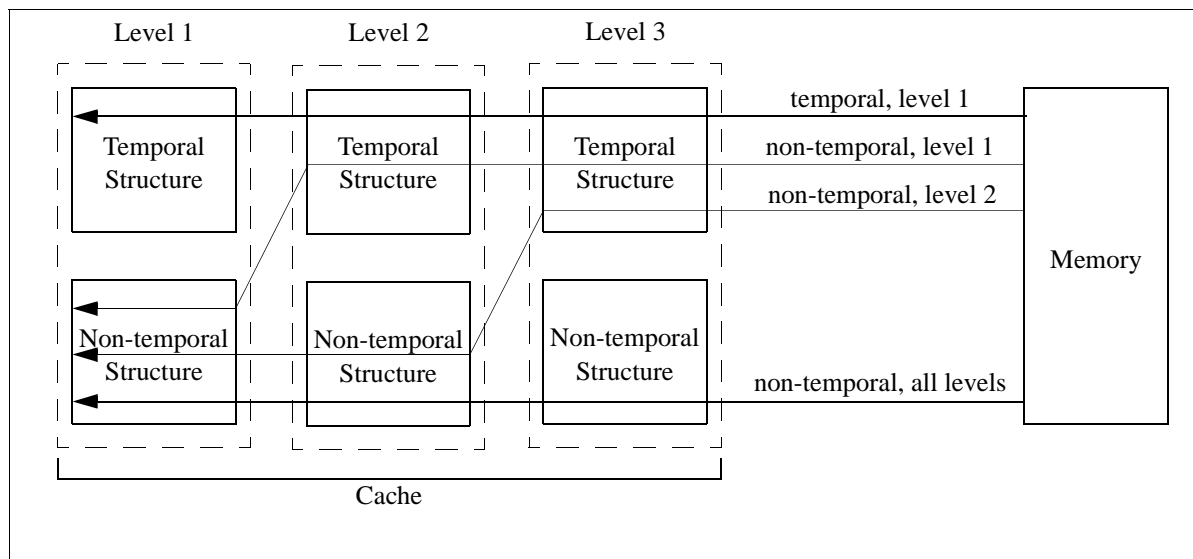


Figure 4-5. Allocation Paths Supported in the Memory Hierarchy

Explicit prefetch is defined in the form of the line prefetch instruction (`lfetch`, `lfetch.fault`). The `lfetch` instructions moves the line containing the addressed byte to a location in the memory hierarchy specified by the locality hint. If the line is already at the same or higher level in the hierarchy, no movement occurs. Both immediate and register post-increment are defined for `lfetch` and `lfetch.fault`. The `lfetch` instruction does not cause any exceptions, does not affect program behavior, and may be ignored by the implementation. The `lfetch.fault` instruction affects the memory hierarchy in exactly the same way as `lfetch` but takes exceptions as if it were a 1-byte load instruction.

Implicit prefetch is based on the address post-increment of loads, stores, `lfetch` and `lfetch.fault`. The line containing the post-incremented address is moved in the memory hierarchy based on the locality hint of the originating load, store, `lfetch` or `lfetch.fault`. If the line is already at the same or higher level in the hierarchy then no movement occurs. Implicit prefetch does not cause any exceptions, does not affect program behavior, and may be ignored by the implementation.

Another form of hint that can be provided on loads is the `ld.bias` load type. This is a hint to the implementation to acquire exclusive ownership of the line containing the addressed data. The bias hint does not affect program functionality and may be ignored by the implementation.

The `fc` instruction invalidates the cache line in all levels of the memory hierarchy above memory. If the cache line is not consistent with memory, then it is copied into memory before invalidation.

Table 4-19 summarizes the memory hierarchy control instructions and hint mechanisms.

Table 4-19. Memory Hierarchy Control Instructions and Hint Mechanisms

Mnemonic	Operation
<code>.ntl</code> and <code>.nta</code> completer on loads	Load usage hints
<code>.nta</code> completer on stores	Store usage hints
prefetch line at post-increment address on loads and stores	Prefetch hint
<code>lfetch</code> , <code>lfetch.fault</code> with <code>.ntl</code> , <code>.nt2</code> , and <code>.nta</code> hints	Prefetch line
<code>fc</code>	Flush cache

4.4.6.2 Memory Consistency

IA-64 instruction accesses made by a processor are not coherent with respect to instruction and/or data accesses made by any other processor, nor are instruction accesses made by a processor coherent with respect to data accesses made by that same processor. Therefore, hardware is not required to keep a processor's instruction caches consistent with respect to any processor's data caches, including that processor's own data caches; nor is hardware required to keep a processor's instruction caches consistent with respect to any other processor's instruction caches. Data accesses from different processors in the same coherence domain are coherent with respect to each other; this consistency is provided by the hardware. Data accesses from the same processor are subject to data dependency rules; see Section 4.4.7, "Memory Access Ordering" below.

The mechanism(s) by which coherence is maintained is implementation dependent. Separate or unified structures for caching data and instructions are not architecturally visible. Within this context there are two categories of data memory hierarchy control: allocation and flush. Allocation refers to movement towards the processor in the hierarchy (lower numbered levels) and flush refers to movement away from the processor in the hierarchy (higher numbered levels). Allocation and flush occur in line-sized units; the minimum architecturally visible line size is 32-bytes (aligned on a 32-byte boundary). The line size in an implementation may be smaller in which case the implementation will need to move multiple lines for each allocation and flush event. An implementation may allocate and flush in units larger than 32-bytes.

In order to guarantee that a write from a given processor becomes visible to the instruction stream of that same, and other, processors, the affected line(s) must be flushed to memory. Software may use the `fc` instruction for this purpose. Memory updates by DMA devices are coherent with respect to instruction and data accesses of processors. The consistency between instruction and data caches of processors with respect to memory updates by DMA devices is provided by the hardware. In case a program modifies its own instructions, the `sync.i` and `sr1z.i` instructions are used to ensure that prior coherency actions are observed by a given point in the program. Refer to the description `sync.i` on page 6-172 for an example of self-modifying code.

4.4.7 Memory Access Ordering

Memory data access ordering must satisfy read-after-write (RAW), write-after-write (WAW), and write-after-read (WAR) data dependencies to the same memory location. In addition, memory writes and flushes must observe control dependencies. Except for these restrictions, reads, writes, and flushes may occur in an order different from the specified program order. Note that no ordering exists between instruction accesses and data accesses or between any two instruction accesses. The mechanisms described below are defined to enforce a particular memory access order. In the following discussion, the terms "previous" and "subsequent" are used to refer to the program specified order. The term "visible" is used

to refer to all architecturally visible effects of performing a memory access (at a minimum this involves reading or writing memory).

Memory accesses follow one of four memory ordering semantics: unordered, release, acquire or fence. Unordered data accesses may become visible in any order. Release data accesses guarantee that all previous data accesses are made visible prior to being made visible themselves. Acquire data accesses guarantee that they are made visible prior to all subsequent data accesses. Fence operations combine the release and acquire semantics into a bi-directional fence, i.e., they guarantee that all previous data accesses are made visible prior to any subsequent data accesses being made visible.

Explicit memory ordering takes the form of a set of instructions: ordered load and ordered check load (`ld.acq`, `ld.c.clr.acq`), ordered store (`st.rel`), semaphores (`cmpxchg`, `xchg`, `fetchadd`), and memory fence (`mf`). The `ld.acq` and `ld.c.clr.acq` instructions follow acquire semantics. The `st.rel` follows release semantics. The `mf` instruction is a fence operation. The `xchg`, `fetchadd.acq`, and `cmpxchg.acq` instructions have acquire semantics. The `cmpxchg.rel`, and `fetchadd.rel` instructions have release semantics. The semaphore instructions also have implicit ordering. If there is a write, it will always follow the read. In addition, the read and write will be performed atomically with no intervening accesses to the same memory region.

Table 4-20 illustrates the ordering interactions between memory accesses with different ordering semantics. “O” indicates that the first and second reference are performed in order with respect to each other. A “-” indicates that no ordering is implied other than data dependencies (and control dependencies for writes and flushes).

Table 4-20. Memory Ordering Rules

First Reference	Second Reference			
	Fence	Acquire	Release	Unordered
fence	O	O	O	O
acquire	O	O	O	O
release	O	-	O	-
unordered	O	-	O	-

Table 4-21 summarizes memory ordering instructions related to cacheable memory.

Table 4-21. Memory Ordering Instructions

Mnemonic	Operation
<code>ld.acq</code> , <code>ld.c.clr.acq</code>	Ordered load and ordered check load
<code>st.rel</code>	Ordered store
<code>xchg</code>	Exchange memory and general register
<code>cmpxchg.acq</code> , <code>cmpxchg.rel</code>	Conditional exchange of memory and general register
<code>fetchadd.acq</code> , <code>fetchadd.rel</code>	Add immediate to memory
<code>mf</code>	Memory ordering fence

4.5 Branch Instructions

Branch instructions effect a transfer of control flow to a new address. Branch targets are bundle-aligned, which means control is always passed to the first instruction slot of the target bundle (slot 0). Branch instructions are not required to be the last instruction in an instruction group. In fact, an instruction group can contain arbitrarily many branches (provided that the normal RAW and WAW dependency requirements are met). If a branch is taken, only instructions up to the taken branch will be executed. After a taken branch, the next instruction executed will be at the target of the branch.

There are two categories of branches: IP-relative branches, and indirect branches. IP-relative branches specify their target with a signed 21-bit displacement, which is added to the IP of the bundle containing the branch to give the address of the target bundle. The displacement allows a branch reach of ± 16 MBytes and is bundle-aligned. Indirect branches use the branch registers to specify the target address.

There are several branch types, as shown in Table 4-22. The conditional branch `br` is a branch which is taken if the specified predicate is 1, and not-taken otherwise. The conditional call branch `br.call` does the same thing, and in addition, writes a link address to a specified branch register and adjusts the general register stack (see “Register Stack” on

page 4-1). The conditional return `br.ret` does the same thing as an indirect conditional branch, plus it adjusts the general register stack. Unconditional branches, calls and returns are executed by specifying PR 0 (which is always 1) as the predicate for the branch instruction.

Table 4-22. Branch Types

Mnemonic	Function	Branch Condition	Target Address
<code>br.cond</code> or <code>br</code>	Conditional branch	Qualifying predicate	IP-rel or Indirect
<code>br.call</code>	Conditional procedure call	Qualifying predicate	IP-rel or Indirect
<code>br.ret</code>	Conditional procedure return	Qualifying predicate	Indirect
<code>br.ia</code>	Invoke the IA-32 instruction set	Unconditional	Indirect
<code>br.cloop</code>	Counted loop branch	Loop count	IP-rel
<code>br.ctop</code> , <code>br.cexit</code>	Modulo-scheduled counted loop	Loop count and Epilog count	IP-rel
<code>br.wtop</code> , <code>br.wexit</code>	Modulo-scheduled while loop	Qualifying predicate and Epilog count	IP-rel

The counted loop type (CLOOP) uses the Loop Count (LC) application register. If LC is non-zero then it is decremented and the branch is taken. If LC is zero, the branch falls through. The modulo-scheduled loop type branches (CTOP, CEXIT, WTOP, WEXIT) are described in “Modulo-Scheduled Loop Support” on page 4-20. The loop type branches (CLOOP, CTOP, CEXIT, WTOP, WEXIT) are allowed only in slot 2 of a bundle. A loop type branch executed in slot 0 or 1 will cause an Illegal Operation fault.

Instructions are provided to move data between branch registers and general registers (`mov =br`, `mov br=`). Table 4-23 summarizes state and instructions relating to branching.

Table 4-23. State Relating to Branching

Register	Function
BRs	Branch registers
PRs	Predicate registers
CFM	Current Frame Marker
PFS	Previous Function State application register
LC	Loop Count application register
EC	Epilog Count application register

Table 4-24. Instructions Relating to Branching

Mnemonic	Operation
<code>br</code>	Branch
<code>mov =br</code>	Move from BR to GR
<code>mov br=</code>	Move from GR to BR

4.5.1 Modulo-Scheduled Loop Support

Support for software-pipelined loops is provided through rotating registers and loop branch types. Software pipelining of a loop is analogous to hardware pipelining of a functional unit. The loop body is partitioned into multiple “stages” with zero or more instructions in each stage. Modulo-scheduled loops have 3 phases: prolog, kernel, and epilog. During the prolog phase, new loop iterations are started each time around (filling the software pipeline). During the kernel phase, the pipeline is full. A new loop iteration is started, and another is finished each time around. During the epilog phase, no new iterations are started, but previous iterations are completed (draining the software pipeline).

A predicate is assigned to each stage to control the activation of the instructions in that stage (this predicate is called the “stage predicate”). To support the pipelining effect of stage predicates and registers in a software-pipelined loop, a fixed sized area of the predicate and floating-point register files (PR16-PR63 and FR32-FR127), and a programmable sized area of the general register file, are defined to “rotate.” The size of the rotating area in the general register file is determined by an immediate in the `alloc` instruction. This immediate must be either zero or a multiple of 8. The general register rotating area is defined to start at GR32 and overlay the local and output areas, depending on their relative sizes. The stage predi-

icates are allocated in the rotating area of the predicate register file. For counted loops, PR16 is architecturally defined to be the first stage predicate with subsequent stage predicates extending to higher predicate register numbers. For while loops, the first stage predicate may be any rotating predicate with subsequent predicates extending to higher predicate register numbers. Software is required to initialize the stage (rotating) predicates prior to entering the loop. An alloc instruction may not change the size of the rotating portion of the register stack frame unless all rotating register bases (rrb's) in the CFM are zero. All rrb's can be set to zero with the `clrrrb` instruction. The `clrrrb.pr` form can be used to clear just the rrb for the predicate registers. The `clrrrb` instruction must be the last instruction in an instruction group.

Rotation by one register position occurs when a software-pipelined loop type branch is executed. Registers are rotated towards larger register numbers in a wrap-around fashion. For example, the value in register X will be located in register X+1 after one rotation. If X is the highest addressed rotating register its value will wrap to the lowest addressed rotating register. Rotation is implemented by renaming register numbers based upon the value of a rotating register base (rrb) contained in CFM. A rrb is defined for each of the three rotating register files: CFM.rrb.gr for the general registers; CFM.rrb.fr for the floating-point registers; CFM.rrb.pr for the predicate registers. General registers only rotate when the size of the rotating region is not equal to zero. Floating-point and predicate registers always rotate. When rotation occurs, two or all three rrb's are decremented in unison. Each rrb is decremented modulo the size of their respective rotating regions (e.g., 96 for rrb.fr). The operation of the rotating register rename mechanism is not otherwise visible to software. The instructions that modify the rrb's are listed in Table 4-25.

Table 4-25. Instructions that Modify RRBs

Mnemonic	Operation
<code>clrrrb</code>	Clears all rrb's
<code>clrrrb.pr</code>	Clears rrb.pr
<code>br.call</code>	Clears all rrb's
<code>br.ret</code>	Restores CFM.rrb's from PFM.rrb's
<code>br.ctop</code> , <code>br.cexit</code> , <code>br.wtop</code> , and <code>br.wexit</code>	Decrements all rrb's

There are two categories of software-pipelined loop branch types: counted and while. Both categories have two forms: top and exit. The “top” variant is used when the loop decision is located at the bottom of the loop body. A taken branch will continue the loop while a not-taken branch will exit the loop. The “exit” variant is used when the loop decision is located somewhere other than the bottom of the loop. A not-taken branch will continue the loop and a taken branch will exit the loop. The “exit” variant is also used at intermediate points in an unrolled pipelined loop.

The branch condition of a counted loop branch is determined by the specific counted loop type (`ctop` or `cexit`), the value of the loop count application register (LC), and the value of the epilog count application register (EC). Note that the counted loop branches do not use a qualifying predicate. LC is initialized to one less than the number of iterations for the counted loop and EC is initialized to the number of stages into which the loop body has been partitioned. While LC is greater than zero, the branch direction will continue the loop, LC will be decremented, registers will be rotated (rrb's are decremented), and PR 16 will be set to 1 after rotation. (For each of the loop-type branches, PR 63 is written by the branch, and after rotation this value will be in PR 16.)

Execution of a counted loop branch with LC equal to zero signals the start of the epilog. While in the epilog and while EC is greater than one, the branch direction will continue the loop, EC will be decremented, registers will be rotated, and PR 16 will be set to 0 after rotation. Execution of a counted loop branch with LC equal to zero and EC equal to one signals the end of the loop; the branch direction will exit the loop, EC will be decremented, registers will be rotated, and PR 16 will be set to 0 after rotation. A counted loop type branch executed with both LC and EC equal to zero will have a branch direction to exit the loop. LC, EC, and the rrb's will not be modified (no rotation) and PR 63 will be set to 0. LC and EC equal to zero can occur in some types of optimized, unrolled software-pipelined loops if the target of a `cexit` branch is set to the next sequential bundle and the loop trip count is not evenly divisible by the unroll amount.

The direction of a while loop branch is determined by the specific while loop type (`wtop` or `wexit`), the value of the qualifying predicate, and the value of EC. The while loop branches do not use LC. While the qualifying predicate is one, the branch direction will continue the loop, registers will be rotated, and PR 16 will be set to 0 after rotation. While the qualifying predicate is zero and EC is greater than one, the branch direction will continue the loop, EC will be decremented, registers will be rotated, and PR 16 will be set to 0 after rotation. The qualifying predicate is one during the kernel and zero during the epilog. During the prolog, the qualifying predicate may be zero or one depending upon the scheme used to program the pipelined while loop. Execution of a while loop branch with qualifying predicate equal to zero and EC equal to one signals the end of the loop; the branch direction will exit the loop, EC will be decremented, registers will be rotated,

and PR 16 will be set to 0 after rotation. A while loop branch executed with a zero qualifying predicate and with EC equal to zero has a branch direction to exit the loop. EC and the rrb's will not be modified (no rotation) and PR 63 will be set to 0.

For while loops, the initialization of EC depends upon the scheme used to program the pipelined while loop. Often, the first valid condition for the while loop branch is not computed until several stages into the prolog. Therefore, software pipelines for while loops often have several speculative prolog stages. During these stages, the qualifying predicate can be set to zero or one depending upon the scheme used to program the loop. If the qualifying predicate is one throughout the prolog, EC will be decremented only during the epilog phase and is initialized to one more than the number of epilog stages. If the qualifying predicate is zero during the speculative stages of the prolog, EC will be decremented during this part of the prolog, and the initialization value for EC is increased accordingly.

4.5.2 Branch Prediction Hints

Information about branch behavior can be provided to the processor to improve branch prediction. This information can be encoded with branch hints as part of a branch instruction (referred to as hints). Hints do not affect the functional behavior of the program and may be ignored by the processor.

Branch instructions can provide three types of hints:

- **Whether prediction strategy:** This describes (for COND, CALL and RET type branches) how the processor should predict the branch condition. (For the loop type branches, prediction is based on LC and EC.) The suggested strategies that can be hinted are shown in Table 4-26.

Table 4-26. Whether Prediction Hint on Branches

Completer	Strategy	Operation
spnt	Static Not-Taken	Ignore this branch, do not allocate prediction resources for this branch.
sptk	Static Taken	Always predict taken, do not allocate prediction resources for this branch.
dpnt	Dynamic Not-Taken	Use dynamic prediction hardware. If no dynamic history information exists for this branch, predict not-taken.
dptk	Dynamic Taken	Use dynamic prediction hardware. If no dynamic history information exists for this branch, predict taken.

- **Sequential prefetch:** This indicates how much code the processor should prefetch at the branch target (shown in Table 4-27).

Table 4-27. Sequential Prefetch Hint on Branches

Completer	Sequential Prefetch Hint	Operation
few	Prefetch few lines	When prefetching code at the branch target, stop prefetching after a few (implementation-dependent number of) lines.
many	Prefetch many lines	When prefetching code at the branch target, prefetch more lines (also an implementation-dependent number).

- **Predictor deallocation:** This provides re-use information to allow the hardware to better manage branch prediction resources. Normally, prediction resources keep track of the most-recently executed branches. However, sometimes the most-recently executed branch is not useful to remember, either because it will not be re-visited any time soon or because a hint instruction will re-supply the information prior to re-visiting the branch. In such cases, this hint can be used to free up the prediction resources.

Table 4-28. Predictor Deallocation Hint

Completer	Operation
none	Don't deallocate
clr	Deallocate branch information

4.6 Multimedia Instructions

Multimedia instructions (see Table 4-29) treat the general registers as concatenations of eight 8-bit, four 16-bit, or two 32-bit elements. They operate on each element independently and in parallel. The elements are always aligned on their natural boundaries within a general register. Most multimedia instructions are defined to operate on multiple element sizes. Three classes of multimedia instructions are defined: arithmetic, shift and data arrangement.

4.6.1 Parallel Arithmetic

There are three forms of parallel addition and subtraction: modulo (`padd`, `psub`), signed saturation (`padd.sss`, `psub.sss`), and unsigned saturation (`padd.uuu`, `padd.uus`, `psub.uuu`, `psub.uus`). The modulo forms have the result wrap around the largest or smallest representable value in the range of the result element. In the saturating forms, results larger than the largest representable value of the range of the result element, or smaller than the smallest representable value of the range, are clamped to the largest or smallest value in the range of the result element respectively. The signed saturation form treats both sources as signed and clamps the result to the limits of a signed range. The unsigned saturation form treats one source as unsigned and clamps the result to the limits of an unsigned range. Two variants are defined that treat the second source as either signed (`.uus`) or unsigned (`.uuu`).

The parallel average instruction (`pavg`, `pavg.raz`) adds corresponding elements from each source and right shifts each result by one bit. In the simple form of the instruction, the carry out of the most-significant bit of each sum is written into the most significant bit of the result element. In the round-away-from-zero form, a 1 is added to each sum before shifting. The parallel average subtract instruction (`pavgsub`) performs a similar operation on the difference of the sources.

The parallel shift left and add instruction (`pshladd`) performs a left shift on the elements of the first source and then adds them to the corresponding elements from the second source. Signed saturation is performed on both the shift and the add operations. The parallel shift right and add instruction (`pshradd`) is similar to `pshladd`. Both of these instructions are defined for 2-byte elements only.

The parallel compare instruction (`pcomp`) compares the corresponding elements of both sources and writes all ones (if true) or all zeroes (if false) into the corresponding elements of the target according to one of two relations (`==` or `>`).

The parallel multiply right instruction (`pmpy.r`) multiplies the corresponding two even-numbered signed 2-byte elements of both sources and writes the results into two 4-byte elements in the target. The `pmpy.l` instruction performs a similar operation on odd-numbered 2-byte elements. The parallel multiply and shift right instruction (`pmpyshr`, `pmpyshr.u`) multiplies the corresponding 2-byte elements of both sources producing four 4-byte results. The 4-byte results are shifted right by 0, 7, 15, or 16 bits as specified by the instruction. The least-significant 2 bytes of the 4-byte shifted results are then stored in the target register.

The parallel sum of absolute difference instruction (`psad`) accumulates the absolute difference of corresponding 1-byte elements and writes the result in the target.

The parallel minimum (`pmin.u`, `pmin`) and the parallel maximum (`pmax.u`, `pmax`) instructions deliver the minimum or maximum, respectively, of the corresponding 1-byte or 2-byte elements in the target. The 1-byte elements are treated as unsigned values and the 2-byte elements are treated as signed values.

Table 4-29. Parallel Arithmetic Instructions

Mnemonic	Operation	1-byte	2-byte	4-byte
<code>padd</code>	Parallel modulo addition	X	X	X
<code>padd.sss</code>	Parallel addition with signed saturation	X	X	
<code>padd.uuu</code> , <code>padd.uus</code>	Parallel addition with unsigned saturation	X	X	
<code>psub</code>	Parallel modulo subtraction	X	X	X
<code>psub.sss</code>	Parallel subtraction with signed saturation	X	X	
<code>psub.uuu</code> , <code>psub.uus</code>	Parallel subtraction with unsigned saturation	X	X	
<code>pavg</code>	Parallel arithmetic average	X	X	
<code>pavg.raz</code>	Parallel arithmetic average with round away from zero	X	X	
<code>pavgsub</code>	Parallel average of a difference	X	X	
<code>pshladd</code>	Parallel shift left and add with signed saturation		X	
<code>pshradd</code>	Parallel shift right and add with signed saturation		X	

Table 4-29. Parallel Arithmetic Instructions (Continued)

Mnemonic	Operation	1-byte	2-byte	4-byte
<code>pcomp</code>	Parallel compare	x	x	x
<code>pmpy.l</code>	Parallel signed multiply of odd elements			x
<code>pmpy.r</code>	Parallel signed multiply of even elements			x
<code>pmpyshr</code>	Parallel signed multiply and shift right		x	
<code>pmpyshr.u</code>	Parallel unsigned multiply and shift right		x	
<code>psad</code>	Parallel sum of absolute difference	x		
<code>pmin</code>	Parallel minimum	x	x	
<code>pmax</code>	Parallel maximum	x	x	

4.6.2 Parallel Shifts

The parallel shift left instruction (`pshl`) individually shifts each element of the first source by a count contained in either a general register or an immediate. The parallel shift right instruction (`pshr`) performs an individual arithmetic right shift of each element of one source by a count contained in either a general register or an immediate. The `pshr.u` instruction performs an unsigned right shift. Table 4-30 summarizes the types of parallel shift instructions.

Table 4-30. Parallel Shift Instructions

Mnemonic	Operation	1-byte	2-byte	4-byte
<code>pshl</code>	Parallel shift left		x	x
<code>pshr</code>	Parallel signed shift right		x	x
<code>pshr.u</code>	Parallel unsigned shift right		x	x

4.6.3 Data Arrangement

The mix right instruction (`mix.r`) interleaves the even-numbered elements from both sources into the target. The mix left instruction (`mix.l`) interleaves the odd-numbered elements. The unpack low instruction (`unpack.l`) interleaves the elements in the least-significant 4 bytes of each source into the target register. The unpack high instruction (`unpack.h`) interleaves elements from the most significant 4 bytes. The pack instructions (`pack.sss`, `pack.uss`) convert from 32-bit or 16-bit elements to 16-bit or 8-bit elements respectively. The least-significant half of larger elements in both sources are extracted and written into smaller elements in the target register. The `pack.sss` instruction treats the extracted elements as signed values and performs signed saturation on them. The `pack.uss` instruction performs unsigned saturation. The mux instruction (`mux`) copies individual 2-byte or 1-byte elements in the source to arbitrary positions in the target according to a specified function. For 2-byte elements, an 8-bit immediate allows all possible permutations to be specified. For 1-byte elements the copy function is selected from one of five possibilities (reverse, mix, shuffle, alternate, broadcast). Table 4-31 describes the various types of parallel data arrangement instructions.

Table 4-31. Parallel Data Arrangement Instructions

Mnemonic	Operation	1-byte	2-byte	4-byte
<code>mix.l</code>	Interleave odd elements from both sources	x	x	x
<code>mix.r</code>	Interleave even elements from both sources	x	x	x
<code>mux</code>	Arbitrary copy of individual source elements	x	x	
<code>pack.sss</code>	Convert from larger to smaller elements with signed saturation		x	x
<code>pack.uss</code>	Convert from larger to smaller elements with unsigned saturation		x	
<code>unpack.l</code>	Interleave least-significant elements from both sources	x	x	x
<code>unpack.h</code>	Interleave most significant elements from both sources	x	x	x

4.7 Register File Transfers

Table 4-32 shows the instructions defined to move values between the general register file and the floating-point, branch, predicate, performance monitor, processor identification, and application register files. Several of the transfer instructions share the same mnemonic (`mov`). The value of the operand identifies which register file is accessed.

Table 4-32. Register File Transfer Instructions

Mnemonic	Operation
<code>getf.exp, getf.sig</code>	Move FR exponent or significand to GR
<code>getf.s, getf.d</code>	Move single/double precision memory format from FR to GR
<code>setf.s, setf.d</code>	Move single/double precision memory format from GR to FR
<code>setf.exp, setf.sig</code>	Move from GR to FR exponent or significand
<code>mov =br</code>	Move from BR to GR
<code>mov br=</code>	Move from GR to BR
<code>mov =pr</code>	Move from predicates to GR
<code>mov pr=, mov pr.rot=</code>	Move from GR to predicates
<code>mov ar=</code>	Move from GR to AR
<code>mov =ar</code>	Move from AR to GR
<code>sum, rum</code>	Set and reset user mask
<code>mov =pmd[...]</code>	Move from performance monitor data register to GR
<code>mov =cpuid[...]</code>	Move from processor identification register to GR
<code>mov =ip</code>	Move from Instruction Pointer

Memory access instructions only target or source the general and floating-point register files. It is necessary to use the general register file as an intermediary for transfers between memory and all other register files except the floating-point register file.

Two classes of move are defined between the general registers and the floating-point registers. The first type moves the significand or the sign/exponent (`getf.sig`, `setf.sig`, `getf.exp`, `setf.exp`). The second type moves entire single or double precision numbers (`getf.s`, `setf.s`, `getf.d`, `setf.d`). These instructions also perform a conversion between the deferred exception token formats.

Instructions are provided to transfer between the branch registers and the general registers.

Instructions are defined to transfer between the predicate register file and a general register. These instructions operate in a “broadside” manner whereby multiple predicate registers are transferred in parallel (predicate register N is transferred to and from bit N of a general register). The move to predicate instruction (`mov pr=`) transfers a general register to multiple predicate registers according to a mask specified by an immediate. The mask contains one bit for each of the static predicate registers (PR 1 through PR 15 – PR 0 is hardwired to 1) and one bit for all of the rotating predicates (PR 16 through PR63). A predicate register is written from the corresponding bit in a general register if the corresponding mask bit is set. If the mask bit is clear then the predicate register is not modified. The rotating predicates are transferred as if `CFM.rrb.pr` were zero. The actual value in `CFM.rrb.pr` is ignored and remains unchanged. The move from predicate instruction (`mov =pr`) transfers the entire predicate register file into a general register target.

The `mov =pmd[]` instruction is defined to move from a performance monitor data (PMD) register to a general register. If the operating system has not enabled reading of performance monitor data registers in user level then all zeroes are returned. The `mov =cpuid[]` instruction is defined to move from a processor identification register to a general register.

The `mov =ip` instruction is provided for copying the current value of the instruction pointer (IP) into a general register.

4.8 Character Strings and Population Count

A small set of special instructions accelerate operations on character and bit-field data.

4.8.1 Character Strings

The compute zero index instructions (`czx.l`, `czx.r`) treat the general register source as either eight 1-byte or four 2-byte elements and write the general register target with the index of the first zero element found. If there are no zero elements in the source, the target is written with a constant one higher than the largest possible index (8 for the 1-byte form, 4 for the 2-byte form). The `czx.l` instruction scans the source from left to right with the left-most element having an index of zero. The `czx.r` instruction scans from right to left with the right-most element having an index of zero. Table 4-33 summarizes the compute zero index instructions.

Table 4-33. String Support Instructions

Mnemonic	Operation	1-byte	2-byte
<code>czx.l</code>	Locate first zero element, left to right	x	x
<code>czx.r</code>	Locate first zero element, right to left	x	x

4.8.2 Population Count

The population count instruction (`popcnt`) writes the number of bits which have a value of 1 in the source register into the target register.

5 IA-64 Floating-point Programming Model

The IA-64 floating-point architecture is fully compliant with the ANSI/IEEE Standard for Binary Floating-Point Arithmetic (Std. 754-1985). There is full IEEE support for single, double, and double-extended real formats. The two IEEE methods for controlling rounding precision are supported. The first method converts results to the double-extended exponent range. The second method converts results to the destination precision. Some IEEE extensions such as fused multiply and add, minimum and maximum operations, and a register file format with a larger range than the minimum double-extended format are also included.

5.1 Data Types and Formats

Six data types are supported directly: single, double, double-extended real (IEEE real types); 64-bit signed integer, 64-bit unsigned integer, and the 82-bit floating-point register format. A “Parallel FP” format where a pair of IEEE single precision values occupy a floating-point register’s significand is also supported. A seventh data type, IEEE-style quad-precision, is supported by software routines. A future architecture extension may include additional support for the quad-precision real type.

5.1.1 Real Types

The parameters for the supported IEEE real types are summarized in Table 5-1.

Table 5-1. IEEE Real-Type Properties

	Single	Double	Double-Extended	Quad-Precision
IEEE Real-Type Parameters				
Sign	+ or –	+ or –	+ or –	+ or –
E_{\max}	+127	+1023	+16383	+16383
E_{\min}	–126	–1022	–16382	–16382
Exponent bias	+127	+1023	+16383	+16383
Precision (bits)	24	53	64	113
IEEE Memory Formats				
Total memory format width (bits)	32	64	80	128
Sign field width (bits)	1	1	1	1
Exponent field width (bits)	8	11	15	15
Significand field width (bits)	23	52	64	112

5.1.2 Floating-point Register Format

Data contained in the floating-point registers can be either integer or real type. The format of data in the floating-point registers is designed to accommodate both of these types with no loss of information.

Real numbers reside in 82-bit floating-point registers in a three-field binary format (see Figure 5-1). The three fields are:

- The 64-bit **significand** field, $b_{63} \cdot b_{62}b_{61} \dots b_1b_0$, contains the number's significant digits. This field is composed of an explicit integer bit (significand{63}), and 63 bits of fraction (significand{62:0}). For Parallel FP data, the significand field holds a pair of 32-bit IEEE single real numbers.
- The 17-bit **exponent** field locates the binary point within or beyond the significant digits (i.e., it determines the number's magnitude). The exponent field is biased by 65535 (0xFFFF). An exponent field of all ones is used to encode the

special values for IEEE signed infinity and NaNs. An exponent field of all zeros and a significand field of all zeros is used to encode the special values for IEEE signed zeros. An exponent field of all zeros and a non-zero significand field encodes the double-extended real denormals and double-extended real pseudo-denormals.

- The 1-bit **sign** field indicates whether the number is positive (sign=0) or negative (sign=1). For Parallel FP data, this bit is always 0.




Figure 5-1. Floating-point Register Format

The value of a finite floating-point number, encoded with non-zero exponent field, can be calculated using the expression:

$$(-1)^{\text{sign}} * 2^{(\text{exponent} - 65535)} * (\text{significand}\{63\}.\text{significand}\{62:0\}_2)$$

The value of a finite floating-point number, encoded with zero exponent field, can be calculated using the expression:

$$(-1)^{\text{sign}} * 2^{(-16382)} * (\text{significand}\{63\}.\text{significand}\{62:0\}_2)$$

Integers (64-bit signed/unsigned) and Parallel FP numbers reside in the 64-bit significand field. In their canonical form, the exponent field is set to 0x1003E (biased 63) and the sign field is set to 0.

5.1.3 Representation of Values in Floating-point Registers

The floating-point register encodings are grouped into classes and subclasses and listed below in Table 5-2 (shaded encodings are unsupported). The last two table entries contain the values of the constant floating-point registers, FR 0 and FR 1. The constant value in FR 1 does not change for the parallel single precision instructions or for the integer multiply accumulate instruction and would not generally be useful.

Table 5-2. Floating-point Register Encodings

Class or Subclass	Sign (1 bit)	Biased Exponent (17-bits)	Significand i.bb...bb (explicit integer bit is shown) (64-bits)
NaNs	0/1	0x1FFFF	1.000...01 through 1.111...11
Quiet NaNs	0/1	0x1FFFF	1.100...00 through 1.111...11
Quiet NaN Indefinite ^a	1	0x1FFFF	1.100...00
Signaling NaNs	0/1	0x1FFFF	1.000...01 through 1.011...11
Infinity	0/1	0x1FFFF	1.000...00
Pseudo-NaN	0/1	0x1FFFF	0.000...01 through 0.111...11
Pseudo-Infinity	0/1	0x1FFFF	0.000...00
Normalized Numbers (Floating-point Register Format Normals)	0/1	0x00001 through 0x1FFFE	1.000...00 through 1.111...11
Integers or Parallel FP (large unsigned or negative signed integers)	0	0x1003E	1.000...00 through 1.111...11
Integer Indefinite ^b	0	0x1003E	1.000...00
IEEE Single Real Normals	0/1	0x0FF81 through 0x1007E	1.000...00...(40)0s through 1.111...11...(40)0s
IEEE Double Real Normals	0/1	0x0FC01 through 0x103FE	1.000...00...(11)0s through 1.111...11...(11)0s
IEEE Double-Extended Real Normals	0/1	0x0C001 through 0x13FFE	1.000...00 through 1.111...11

Table 5-2. Floating-point Register Encodings (Continued)

Normal numbers with the same value as Double-Extended Real Pseudo-Denormals	0/1	0x0C001	1.000...00 through 1.111...11
IA-32 Stack Single Real Normals (produced when the computation model is IA-32 Stack Single)	0/1	0x0C001 through 0x13FFE	1.000...00...(40)0s through 1.111...11...(40)0s
IA-32 Stack Double Real Normals (produced when the computation model is IA-32 Stack Double)	0/1	0x0C001 through 0x13FFE	1.000...00...(11)0s through 1.111...11...(11)0s
Unnormalized Numbers (Floating-point Register Format unnormalized numbers)	0/1	0x00000	0.000...01 through 1.111...11
		0x00001 through 0x1FFFE	0.000...01 through 0.111...11
		0x00001 through 0x1FFFD	0.000...00
	1	0x1FFFE	0.000...00
Integers or Parallel FP (positive signed/unsigned integers)	0	0x1003E	0.000...00 through 0.111...11
Single Real Denormals	0/1	0x0FF81	0.000...01...(40)0s through 0.111...11...(40)0s
Double Real Denormals	0/1	0x0FC01	0.000...01...(11)0s through 0.111...11...(11)0s
Register Format Denormals	0/1	0x00001	0.000...01 through 0.111...11
Double-Extended Real Denormals	0/1	0x00000	0.000...01 through 0.111...11
Unnormal numbers with the same value as Double-Extended Real Denormals	0/1	0x0C001	0.000...01 through 0.111...11
Double-Extended Real Pseudo-Denormals (IA-32 stack and memory format)	0/1	0x00000	1.000...00 through 1.111...11
IA-32 Stack Single Real Denormals (produced when computation model is IA-32 Stack Single)	0/1	0x00000	0.000...01...(40)0s through 0.111...11...(40)0s
IA-32 Stack Double Real Denormals (produced when computation model is IA-32 Stack Double)	0/1	0x00000	0.000...01...(11)0s through 0.111...11...(11)0s
Pseudo-Zeros	0/1	0x00001 through 0x1FFFD	0.000...00
	1	0x1FFFE	0.000...00
NaTVal ^c	0	0x1FFFE	0.000...00
Zero	0/1	0x00000	0.000...00
FR 0 (positive zero)	0	0x00000	0.000...00
FR 1 (positive one)	0	0x0FFFF	1.000...00

- a. Default response on a masked real invalid operation.
- b. Default response on a masked integer invalid operation.
- c. Created by unsuccessful speculative memory operation.

All register file encodings are allowed as inputs to arithmetic operations. The result of an arithmetic operation is always the most normalized register file representation of the computed value, with the exponent range limited from Emin to Emax of the destination type, and the significand precision limited to the number of precision bits of the destination type.

Computed values, such as zeros, infinities, and NaNs that are outside these bounds are represented by the corresponding unique register file encoding. Double-extended real denormal results are mapped to the register file exponent of 0x00000 (instead of 0x0C001). Unsupported encodings (Pseudo-NaN and Pseudo-Infinities), Pseudo-zeros and Double-extended Real Pseudo-denormals are never produced as a result of an arithmetic operation.

Arithmetic on pseudo-zeros operates exactly as an equivalently signed zero, with one exception. Pseudo-zero multiplied by infinity returns the correctly signed infinity instead of an Invalid Operation Floating-Point Exception fault (and QNaN). Also, pseudo-zeros are classified as unnormalized numbers, not zeros.

5.2 Floating-point Status Register

The Floating-Point Status Register (FPSR) contains the dynamic control and status information for floating-point operations. There is one main set of control and status information (FPSR.sf0), and three alternate sets (FPSR.sf1, FPSR.sf2, FPSR.sf3). The FPSR layout is shown in Figure 5-2 and its fields are defined in Table 5-3. Table 5-4 gives the FPSR’s status field description and Figure 5-3 shows their layout.




Figure 5-2. Floating-point Status Register Format

Table 5-3. Floating-point Status Register Field Description

Field	Bits	Description
traps.vd	0	Invalid Operation Floating-Point Exception fault (IEEE Trap) disabled when this bit is set
traps.dd	1	Denormal/Unnormal Operand Floating-Point Exception fault disabled when this bit is set
traps.zd	2	Zero Divide Floating-Point Exception fault (IEEE Trap) disabled when this bit is set
traps.od	3	Overflow Floating-Point Exception trap (IEEE Trap) disabled when this bit is set
traps.ud	4	Underflow Floating-Point Exception trap (IEEE Trap) disabled when this bit is set
traps.id	5	Inexact Floating-Point Exception trap (IEEE Trap) disabled when this bit is set
sf0	18:6	Main status field
sf1	31:19	Alternate status field 1
sf2	44:32	Alternate status field 2
sf3	57:45	Alternate status field 3
rv	63:58	Reserved




Figure 5-3. Floating-point Status Field Format

Table 5-4. Floating-point Status Register’s Status Field Description

Field	Bits	Description
ftz	0	Flush-to-Zero mode
wre	1	Widest range exponent (see Table 5-6)
pc	3:2	Precision control (see Table 5-6)
rc	5:4	Rounding control (see Table 5-5)
td	6	Traps disabled ^a
v	7	Invalid Operation (IEEE Flag)

Table 5-4. Floating-point Status Register’s Status Field Description (Continued)

Field	Bits	Description
d	8	Denormal/Unnormal Operand
z	9	Zero Divide (IEEE Flag)
o	10	Overflow (IEEE Flag)
u	11	Underflow (IEEE Flag)
i	12	Inexact (IEEE Flag)

a. td is a reserved bit in the main status field, FPSR.sf0.

The Denormal/Unnormal Operand status flag is an IEEE-style sticky flag that is set if the value is used in an arithmetic instruction and in an arithmetic calculation; e.g. `unorm*NaN` doesn’t set the d flag. Canonical single/double/double-extended denormal/double-extended pseudo-denormal/register format denormal encodings are a subset of the floating-point register format unnormalized numbers.

Note that the Floating-Point Exception fault/trap occurs only if an enabled floating-point exception occurs during the processing of the instruction. Hence, setting a flag bit of a status field to 1 in software will not cause an interruption. The status fields flags are merely indications of the occurrence of floating-point exceptions.

Flush-to-Zero (FTZ) mode causes results which encounter “tininess” to be truncated to the correctly signed zero. Flush-to-Zero mode can be enabled only if Underflow is disabled. This can be accomplished by disabling all traps (FPSR.sfx.td being set to 1), or by disabling it individually (FPSR.traps.ud set to 1). If Underflow is enabled then it takes priority and Flush-to-Zero mode is ignored. Note that the software exception handler could examine the Flush-to Zero mode bit and choose to emulate the Flush-to-Zero operation when an enabled Underflow exception arises.

The FPSR.sfx.u and FPSR.sfx.i bits will be set to 1 when a result is flushed to the correctly signed zero because of Flush-to-Zero mode. If enabled, an inexact result exception is signaled.

A floating-point result is rounded based on the instruction’s `.pc` completer and the status field’s `wre`, `pc`, and `rc` control fields. The result’s significand precision and exponent range are determined as described in Table 5-6 "Floating-point Computation Model Control Definitions". If the result isn’t exact, FPSR.sfx.rc specifies the rounding direction (see Table 5-5).

Table 5-5. Floating-point Rounding Control Definitions

	Nearest (or even)	– Infinity (down)	+ Infinity (up)	Zero (truncate/chop)
FPSR.sfx.rc	00	01	10	11

Table 5-6. Floating-point Computation Model Control Definitions

Computation Model Control Fields			Computation Model Selected		
Instruction’s <code>.pc</code> Completer	FPSR.sfx’s Dynamic <code>pc</code> Field	FPSR.sfx’s Dynamic <code>wre</code> Field	Significand Precision	Exponent Range	Computational Style
<code>.s</code>	ignored	0	24 bits	8 bits	IEEE real single
<code>.d</code>	ignored	0	53 bits	11 bits	IEEE real double
<code>.s</code>	ignored	1	24 bits	17 bits	Register file range, single precision
<code>.d</code>	ignored	1	53 bits	17 bits	Register file range, double precision
<code>none</code>	00	0	24 bits	15 bits	IA-32 stack single
<code>none</code>	01	0	N.A.	N.A.	Reserved
<code>none</code>	10	0	53 bits	15 bits	IA-32 stack double
<code>none</code>	11	0	64 bits	15 bits	IA-32 double-extended
<code>none</code>	00	1	24 bits	17 bits	Register file range, single precision
<code>none</code>	01	1	N.A.	N.A.	Reserved
<code>none</code>	10	1	53 bits	17 bits	Register file range, double precision

Table 5-6. Floating-point Computation Model Control Definitions (Continued)

Computation Model Control Fields			Computation Model Selected		
Instruction's .pc Completer	FPSR.sfx's Dynamic pc Field	FPSR.sfx's Dynamic wre Field	Significant Precision	Exponent Range	Computational Style
<i>none</i>	11	1	64 bits	17 bits	Register file range, double-extended precision
not applicable ^a	ignored	ignored	24 bits	8 bits	A pair of IEEE real singles
not applicable ^b	ignored	ignored	64 bits	17 bits	Register file range, double-extended precision

a. For parallel FP instructions which have no .pc completer (e.g., fpma).

b. For non-parallel FP instructions which have no .pc completer (e.g., fmerge).

The trap disable (sfx.td) control bit allows one to easily set up a local IEEE exception trap default environment. If FPSR.sfx.td is clear (enabled), the FPSR.traps bits are used. If FPSR.sfx.td is set, the FPSR.traps bits are treated as if they are all set (disabled). Note that FPSR.sf0.td is a reserved field which returns 0 when read.

5.3 Floating-point Instructions

This section describes the IA-64 floating-point instructions.

5.3.1 Memory Access Instructions

There are floating-point load and store instructions for the single, double, double-extended floating-point real data types, and the Parallel FP or signed/unsigned integer data type. The addressing modes for floating-point load and store instructions are the same as for integer load and store instructions, except for floating-point load pair instructions which can have an implicit base-register post increment. The memory hint options for floating-point load and store instructions are the same as those for integer load and store instructions. (See “Memory Hierarchy Control and Consistency” on page 4-16.) Table 5-7 lists the types of floating-point load and store instructions. The floating-point load pair instructions require the two target registers to be odd/even or even/odd. The floating-point store instructions (*stfs*, *stfd*, *stfe*) require the value in the floating-point register to have the same type as the store for the format conversion to be correct.

Table 5-7. Floating-point Memory Access Instructions

Operations	Load to FR	Load Pair to FR	Store from FR
Single	<i>ldfs</i>	<i>ldfps</i>	<i>stfs</i>
Integer/Parallel FP	<i>ldf8</i>	<i>ldfp8</i>	<i>stf8</i>
Double	<i>ldfd</i>	<i>ldfpd</i>	<i>stfd</i>
Double-extended	<i>ldfe</i>		<i>stfe</i>
Spill/fill	<i>ldf.fill</i>		<i>stf.spill</i>

Unsuccessful speculative loads write a NaTVal into the destination register or registers (see Section 4.4.4). Storing a NaTVal to memory will cause a Register NaT Consumption fault, except for the spill instruction (*stf.spill*).

Saving and restoring floating-point registers is accomplished by the spill and fill instructions (*stf.spill*, *ldf.fill*) using a 16-byte memory container. These are the only instructions that can be used for saving and restoring the actual register contents since they do not fault on NaTVal. They save and restore all types (single, double, double-extended, register format and integer or Parallel FP) and will ensure compatibility with possible future architecture extensions.

Figure 5-4, Figure 5-5, Figure 5-6 and Figure 5-7 describe how single precision, double precision, double-extended precision, and spill/fill data is translated during transfers between floating-point registers and memory.

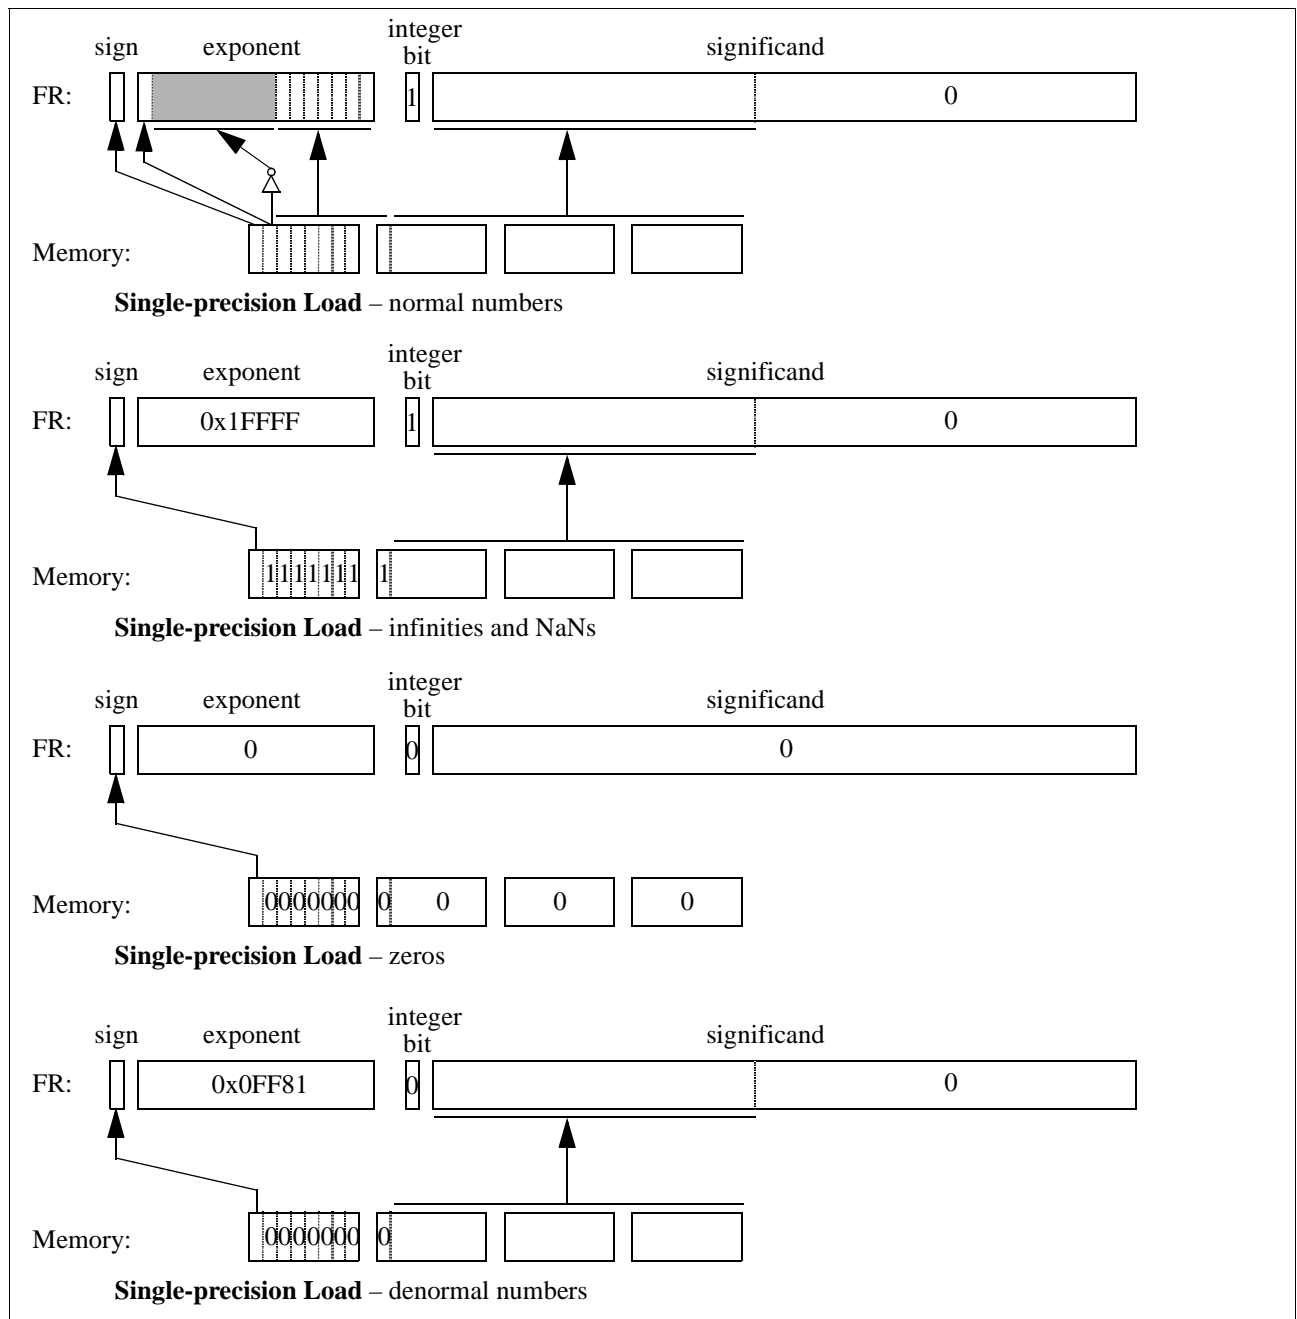


Figure 5-4. Memory to Floating-point Register Data Translation – Single Precision

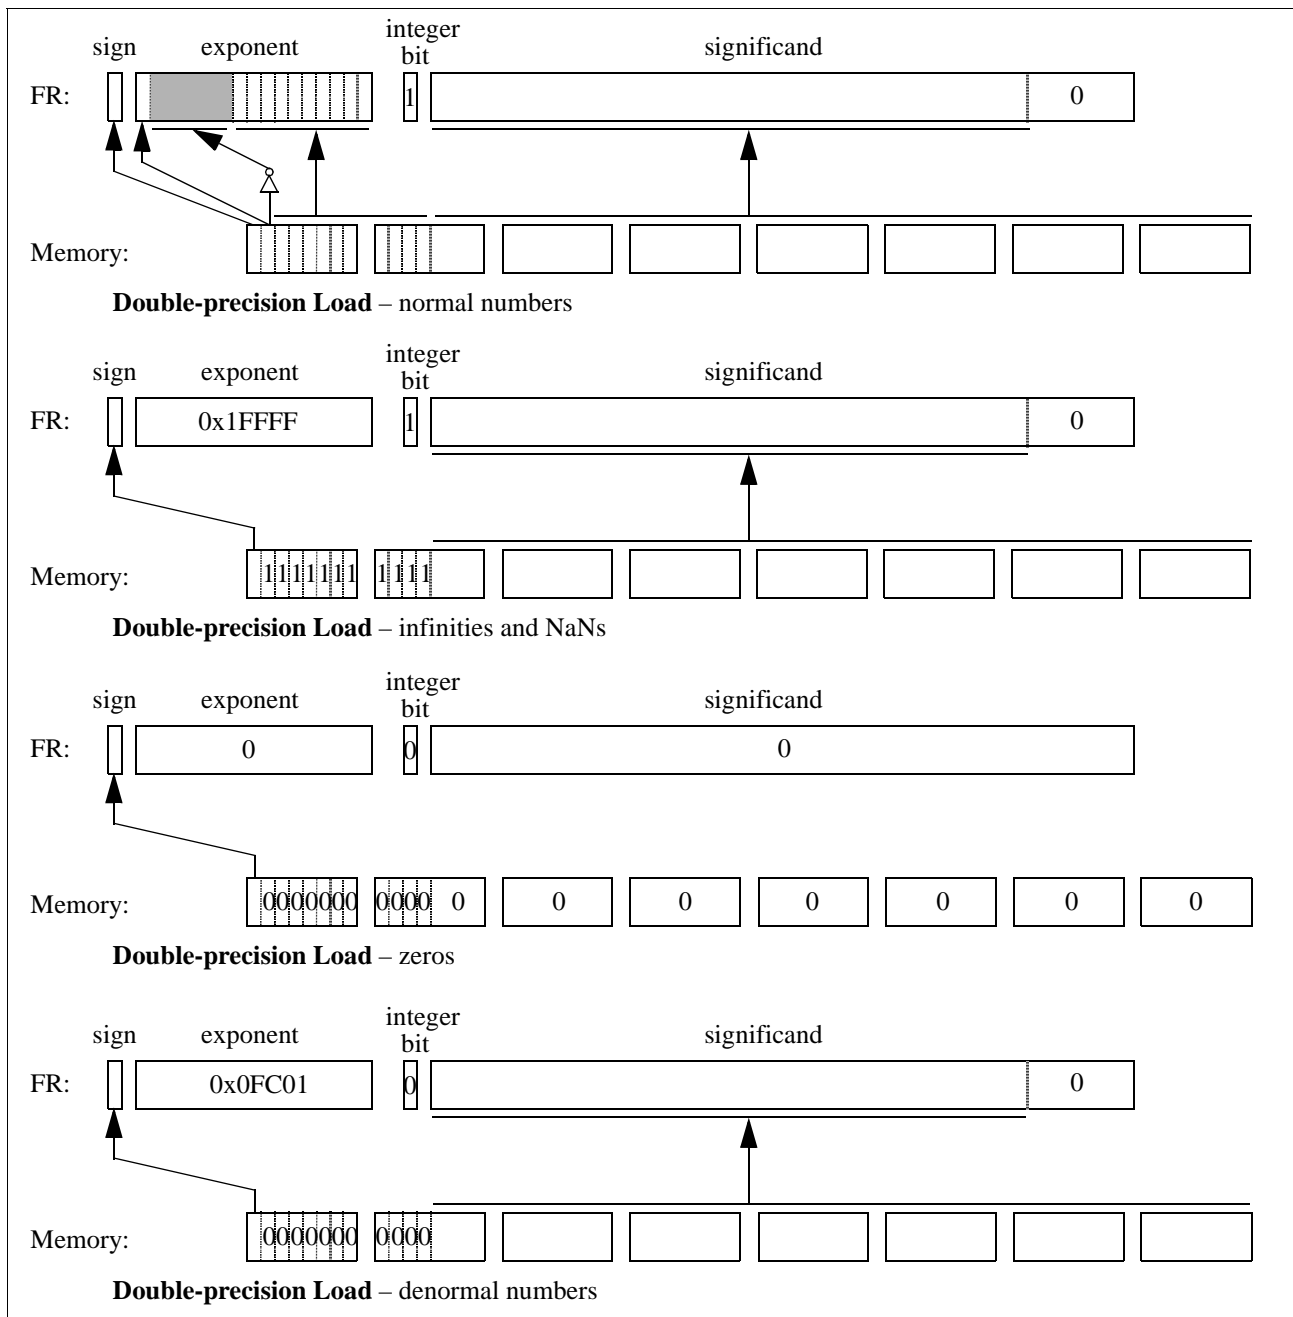


Figure 5-5. Memory to Floating-point Register Data Translation – Double Precision

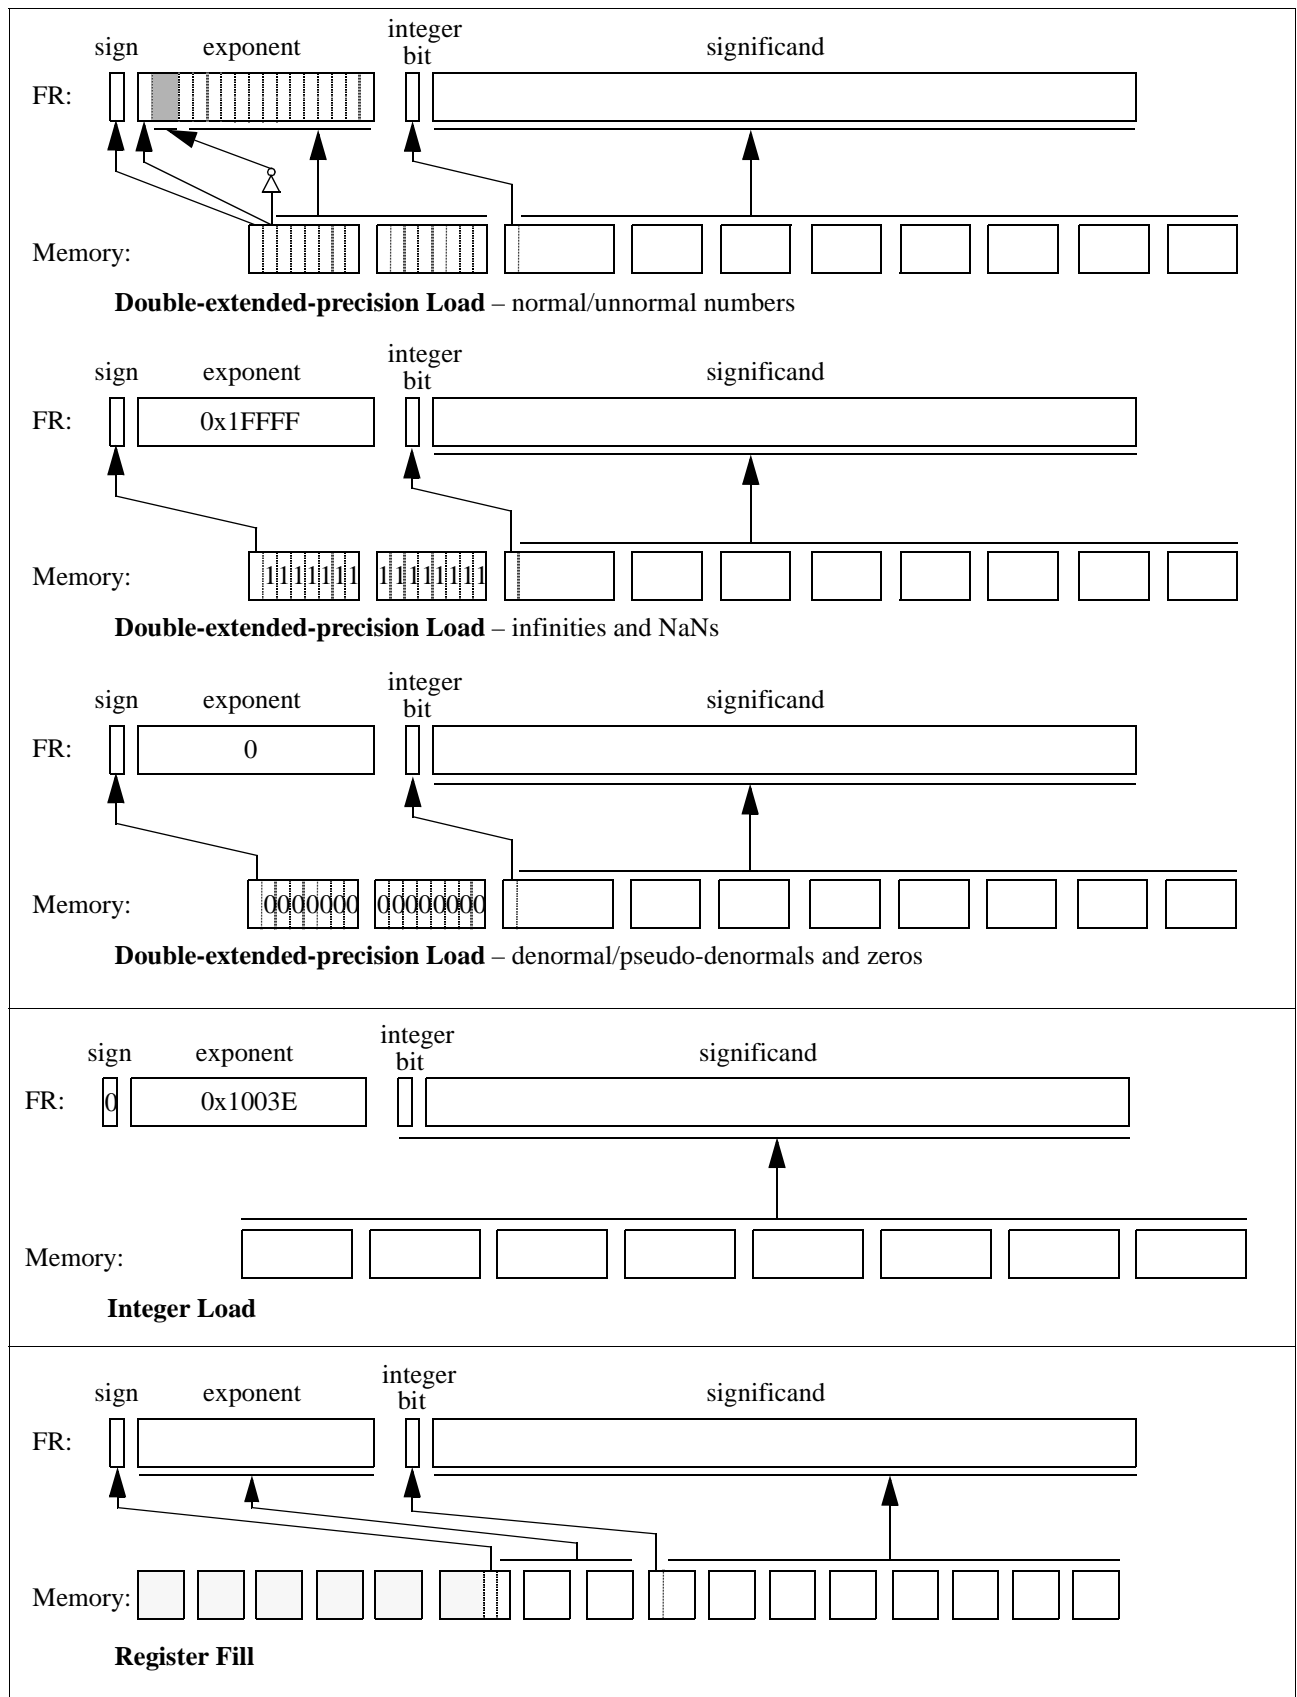


Figure 5-6. Memory to Floating-point Register Data Translation – Double Extended, Integer and Fill

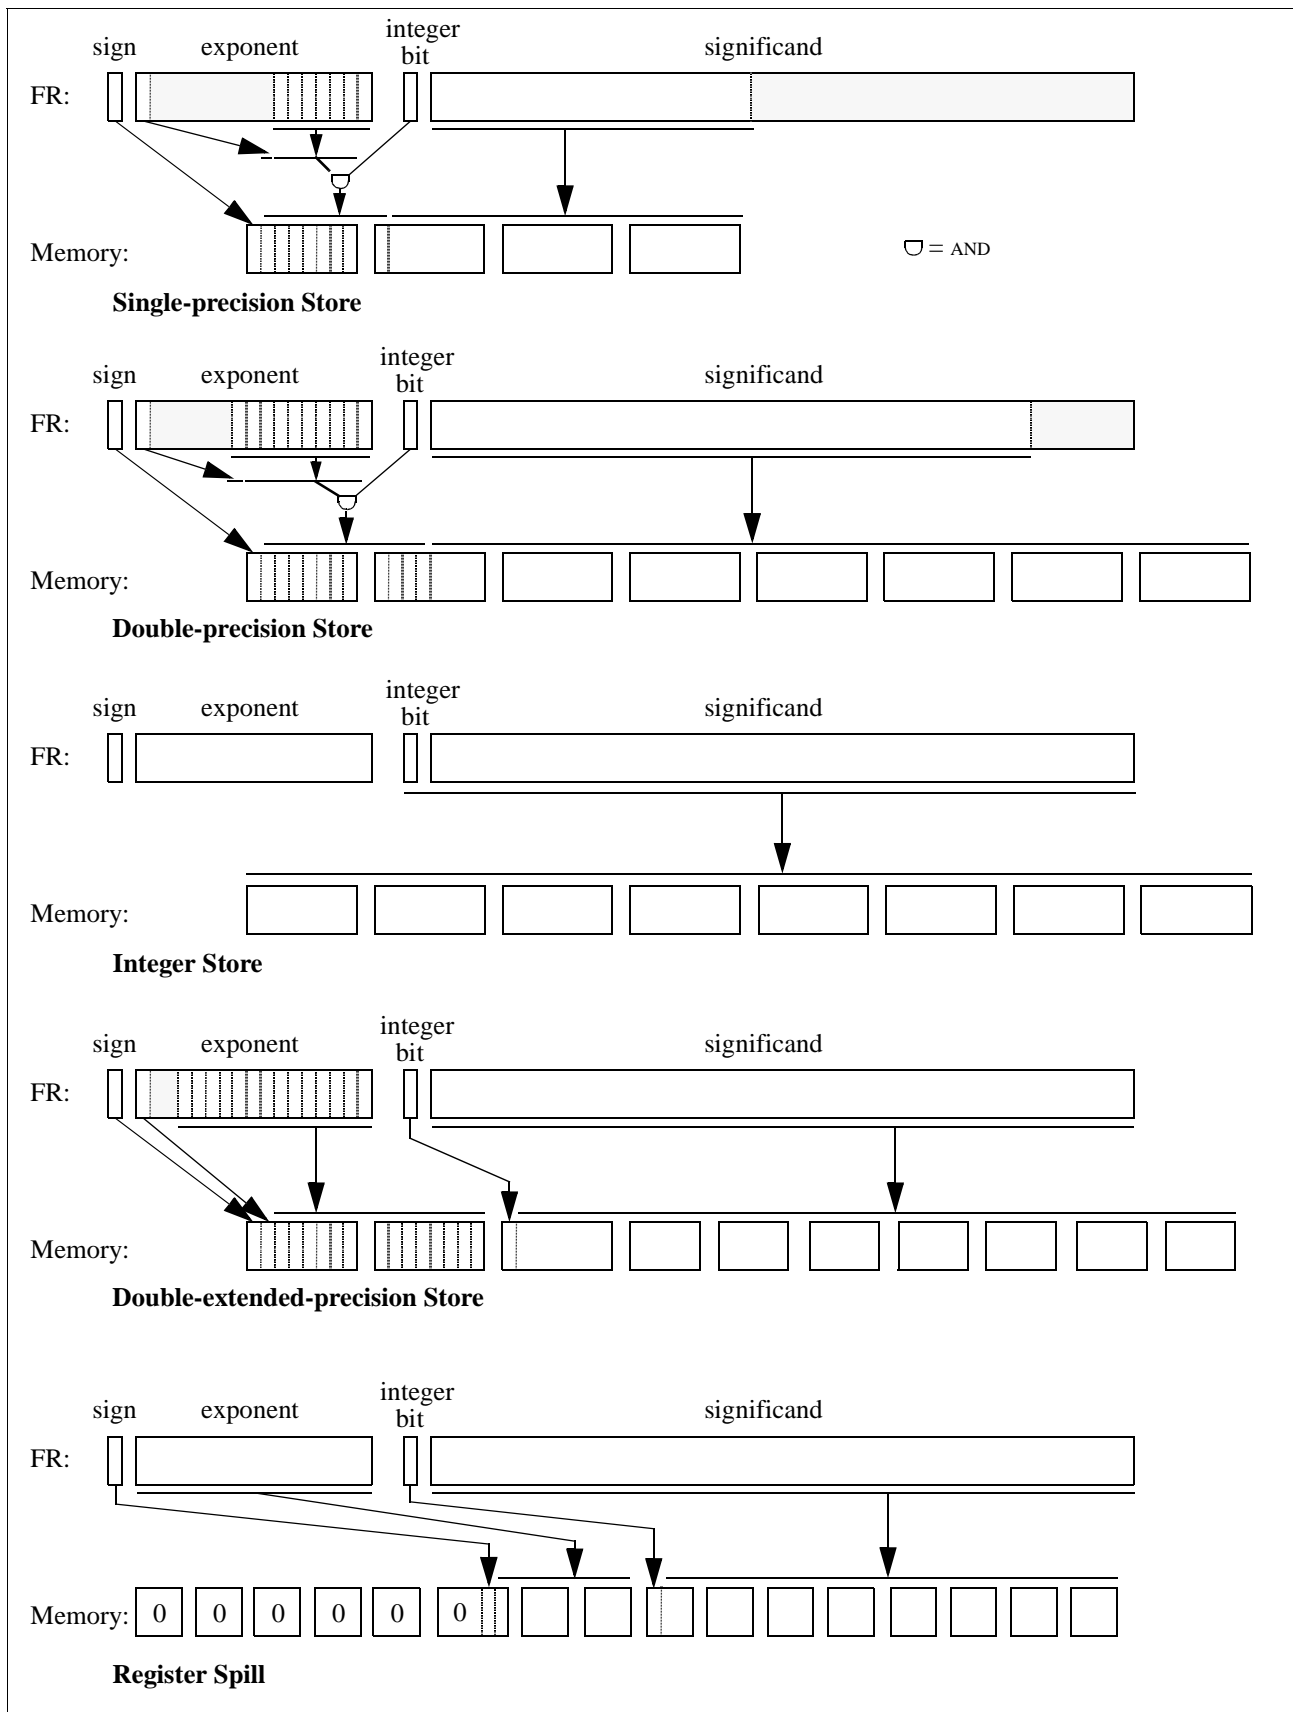


Figure 5-7. Floating-point Register to Memory Data Translation

Both little-endian and big-endian byte ordering is supported on floating-point loads and stores. For both single and double memory formats, the byte ordering is identical to the 32-bit and 64-bit integer data types (see Section 3.2.3). The byte-ordering for the spill/fill memory and double-extended formats is shown in Figure 5-8.

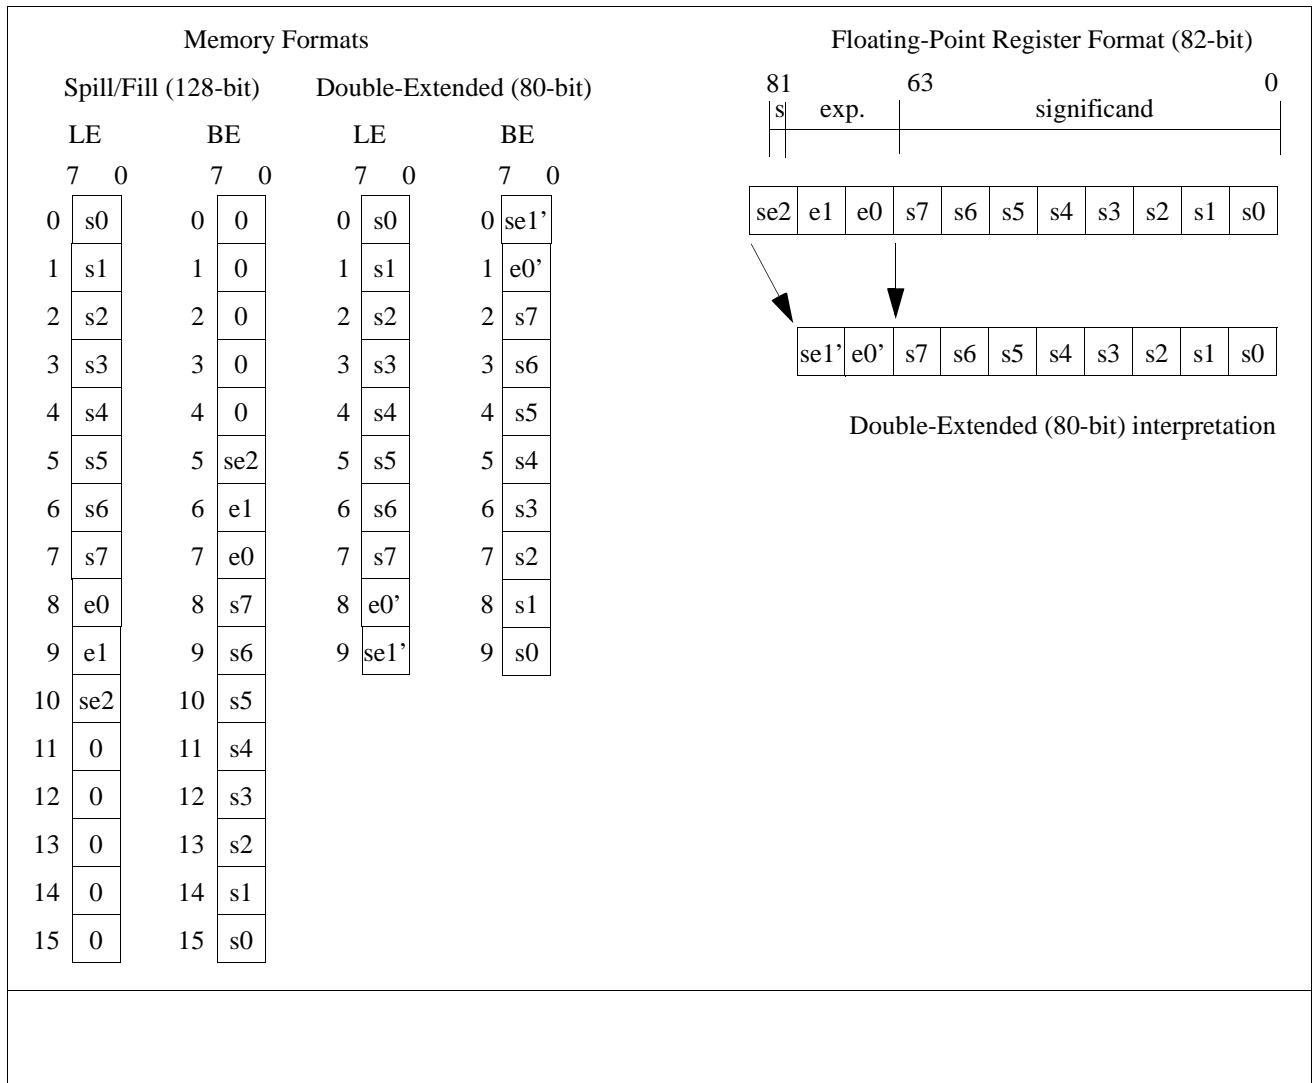


Figure 5-8. Spill/Fill and Double-Extended (80-bit) Floating-point Memory Formats

5.3.2 Floating-Point Register to/from General Register Transfer Instructions

The `setf` and `getf` instructions (see Table 5-8) transfer data between floating-point registers (FR) and general registers (GR). These instructions will translate a general register NaT to/from a floating-point register NaTVal. For all other operands, the `.s` and `.d` variants of the `setf` and `getf` instructions translate to/from FR as per Figure 5-4, Figure 5-5 and Figure 5-7. The memory representation is read from or written to the GR. The `.exp` and `.sig` variants of the `setf` and `getf` instructions operate on the sign/exponent and significand portions of a floating-point register, respectively, and their translation formats are described in Table 5-9 and Table 5-10.

Table 5-8. Floating-point Register Transfer Instructions

Operations	GR to FR	FR to GR
Single	<code>setf.s</code>	<code>getf.s</code>
Double	<code>setf.d</code>	<code>getf.d</code>
Sign and Exponent	<code>setf.exp</code>	<code>getf.exp</code>
Significand/Integer	<code>setf.sig</code>	<code>getf.sig</code>

Table 5-9. General Register (Integer) to Floating-point Register Data Translation

Class	General Register		Floating-Point Register (.sig)			Floating-Point Register (.exp)		
	NaT	Integer	Sign	Exponent	Significand	Sign	Exponent	Significand
NaT	1	ignore	NaTVal			NaTVal		
integers	0	000...00 through 111...11	0	0x1003E	integer	integer{17}	integer{16:0}	0x8000000000000000

Table 5-10. Floating-point Register to General Register (Integer) Data Translation

Class	Floating-Point Register			General Register (.sig)		General Register (.exp)	
	Sign	Exponent	Significand	NaT	Integer	NaT	Integer
NaTVal	0	0x1FFFE	0.000...00	1	0x0000000000000000	1	0x1FFFE
integers or paral- lel FP	0	0x1003E	0.000...00 through 1.111...11	0	significand	0	0x1003E
other	ignore	ignore	ignore	0	significand	0	((sign<<17) exponent)

5.3.3 Arithmetic Instructions

All of the arithmetic floating-point instructions (except `fcvt.xf` which is always exact) have a `.sf` specifier. This indicates which of the four FPSR's status fields will both control and record the status of the execution of the instruction (see Table 5-11). The status field specifies: enabled exceptions, rounding mode, exponent width, precision control, and which status field's flags to update. See "Floating-point Status Register" on page 5-4.

Table 5-11. Floating-point Instruction Status Field Specifier Definition

<code>.sf</code> Specifier	<code>.s0</code>	<code>.s1</code>	<code>.s2</code>	<code>.s3</code>
Status field	FPSR.sf0	FPSR.sf1	FPSR.sf2	FPSR.sf3

Most arithmetic floating-point instructions can specify the precision of the result statically by using a `.pc` completer, or dynamically using the `.pc` field of the FPSR status field. (see Table 5-6). Arithmetic instructions that do not have a `.pc` completer use the floating-point register file range and precision.

Table 5-12 lists the floating-point arithmetic instructions and Table 5-13 lists the pseudo-operation definitions.

Table 5-12. Floating-point Arithmetic Instructions

Operation	Normal FP Mnemonic(s)	Parallel FP Mnemonic(s)
Floating-point multiply and add	<code>fma.pc.sf</code>	<code>fpma.sf</code>
Floating-point multiply and subtract	<code>fms.pc.sf</code>	<code>fpms.sf</code>
Floating-point negate multiply and add	<code>fnma.pc.sf</code>	<code>fpnma.sf</code>
Floating-point reciprocal approximation	<code>frcpa.sf</code>	<code>fprcpa.sf</code>
Floating-point reciprocal square root approximation	<code>frsqta.sf</code>	<code>fprsqta.sf</code>
Floating-point compare	<code>fcmp.frel.fctype.sf</code>	<code>fpcmp.frel.sf</code>
Floating-point minimum	<code>fmin.sf</code>	<code>fpmin.sf</code>
Floating-point maximum	<code>fmax.sf</code>	<code>fpmax.sf</code>
Floating-point absolute minimum	<code>famin.sf</code>	<code>fpamin.sf</code>
Floating-point absolute maximum	<code>famax.sf</code>	<code>fpamax.sf</code>
Convert floating-point to signed integer	<code>fcvt.fx.sf</code> <code>fcvt.fx.trunc.sf</code>	<code>fpcvtf.fx.sf</code> <code>fpcvtf.fx.trunc.sf</code>
Convert floating-point to unsigned integer	<code>fcvt.fxu.sf</code> <code>fcvt.fxu.trunc.sf</code>	<code>fpcvtf.fxu.sf</code> <code>fpcvtf.fxu.trunc.sf</code>
Convert signed integer to floating-point	<code>fcvt.xf</code>	N.A.

Table 5-13. Floating-point Pseudo-Operations

Operation	Mnemonic	Operation Used
Floating-point multiplication (IEEE) Parallel FP multiplication	<i>fmpy.pc.sf</i> <i>fpmpy.sf</i>	<i>fma</i> , using FR 0 for addend <i>fpma</i> , using FR 0 for addend
Floating-point negate multiplication (IEEE) Parallel FP negate multiplication	<i>fnmpy.pc.sf</i> <i>fpnmpy.sf</i>	<i>fnma</i> , using FR 0 for addend <i>fpnma</i> , using FR 0 for addend
Floating-point addition (IEEE)	<i>fadd.pc.sf</i>	<i>fma</i> , using FR 1 for multiplicand
Floating-point subtraction (IEEE)	<i>fsub.pc.sf</i>	<i>fms</i> , using FR 1 for multiplicand
Floating-point negation (IEEE)	<i>fnma.pc.sf</i>	<i>fnma</i> , using FR 1 for multiplicand and FR 0 for addend
Floating-point absolute value Parallel FP absolute value	<i>fabs</i> <i>fpabs</i>	<i>fmerge.s</i> , with sign from FR 0 <i>fpmerge.s</i> , with sign from FR 0
Floating-point negate Parallel FP negate	<i>fneg</i> <i>fpneg</i>	<i>fmerge.ns</i> <i>fpmerge.ns</i>
Floating-point negate absolute value Parallel FP negate absolute value	<i>fnegabs</i> <i>fpnegabs</i>	<i>fmerge.ns</i> , with sign from FR 0 <i>fpmerge.ns</i> , with sign from FR 0
Floating-point normalization	<i>fnorm.pc.sf</i>	<i>fma</i> , using FR 1 for multiplicand and FR 0 for addend
Convert unsigned integer to floating-point	<i>fcvt.xuf.pc.sf</i>	<i>fma</i> , using FR 1 for multiplicand and FR 0 for addend

There are no pseudo-operations for Parallel FP addition, subtraction, negation or normalization since FR 1 does not contain a packed pair of single precision 1.0 values. A parallel FP addition can be performed by first forming a pair of 1.0 values in a register (using the *fpack* instruction) and then using the *fpma* instruction. Similarly, an integer add operation can be generated by first forming an integer 1 in a floating-point register and then using the *xma* instruction.

5.3.4 Non-Arithmetic Instructions

Table 5-14 lists the non-arithmetic floating-point instructions. The *fclass* instruction is used to classify the contents of a floating-point register. The *fmerge* instruction is used to merge data from two floating-point registers into one floating-point register. The *fmix*, *fsxt*, *fpack*, and *fswap* instructions are used to manipulate the Parallel FP data in the floating-point significand. The *fand*, *fandcm*, *for*, and *fxor* instructions are used to perform logical operations on the floating-point significand. The *fselect* instruction is used for conditional selects.

The non-arithmetic floating-point instructions always use the floating-point register (82-bit) precision since they do not have a *.pc* completer nor a *.sf* specifier.

Table 5-14. Non-Arithmetic Floating-point Instructions

Operation	Mnemonic(s)
Floating-point classify	<i>fclass.fcrel.fctype</i>
Floating-point merge sign Parallel FP merge sign	<i>fmerge.s</i> <i>fpmerge.s</i>
Floating-point merge negative sign Parallel FP merge negative sign	<i>fmerge.ns</i> <i>fpmerge.ns</i>
Floating-point merge sign and exponent Parallel FP merge sign and exponent	<i>fmerge.se</i> <i>fpmerge.se</i>
Floating-point mix left	<i>fmix.l</i>
Floating-point mix right	<i>fmix.r</i>
Floating-point mix left-right	<i>fmix.lr</i>
Floating-point sign-extend left	<i>fsxt.l</i>
Floating-point sign-extend right	<i>fsxt.r</i>
Floating-point pack	<i>fpack</i>
Floating-point swap	<i>fswap</i>

Table 5-14. Non-Arithmetic Floating-point Instructions (Continued)

Operation	Mnemonic(s)
Floating-point swap and negate left	<code>fswap.nl</code>
Floating-point swap and negate right	<code>fswap.nr</code>
Floating-point And	<code>fand</code>
Floating-point And Complement	<code>fandcm</code>
Floating-point Or	<code>for</code>
Floating-point Xor	<code>fxor</code>
Floating-point Select	<code>fselect</code>

5.3.5 Floating-point Status Register (FPSR) Status Field Instructions

Speculation of floating-point operations requires that the status flags be stored temporarily in one of the alternate status fields (not `FPSR.sf0`). After a speculative execution chain has been committed, a `fchkf` instruction can be used to update the normal flags (`FPSR.sf0.flags`). This operation will preserve the correctness of the IEEE flags. The `fchkf` instruction does this by comparing the flags of the status field with the `FPSR.sf0.flags` and `FPSR.traps`. If the flags of the alternate status field indicate the occurrence of an event that corresponds to an enabled floating-point exception in `FPSR.traps`, or an event that is not already registered in the `FPSR.sf0.flags` (i.e., the flag for that event in `FPSR.sf0.flags` is clear), then the `fchkf` instruction causes a Speculative Operation fault. If neither of these cases arise then the `fchkf` instruction does nothing.

The `fsetc` instruction allows bit-wise modification of a status field's control bits. The `FPSR.sf0.controls` are ANDed with a 7-bit immediate and-mask and ORed with a 7-bit immediate or-mask to produce the control bits for the status field. The `fclrf` instruction clears all of the status field's flags to zero.

Table 5-15. FPSR Status Field Instructions

Operation	Mnemonic(s)
Floating-point check flags	<code>fchkf.sf</code>
Floating-point clear flags	<code>fclrf.sf</code>
Floating-point set controls	<code>fsetc.sf</code>

5.3.6 Integer Multiply and Add Instructions

Integer (fixed-point) multiply is executed in the floating-point unit using the three-operand `xma` instructions. The operands and result of these instructions are floating-point registers. The `xma` instructions ignore the sign and exponent fields of the floating-point register, except for a NaNVal check. The product of two 64-bit source significands is added to the third 64-bit significand (zero extended) to produce a 128-bit result. The low and high versions of the instruction select the appropriate low/high 64-bits of the 128-bit result, respectively, and write it into the destination register as a canonical integer. The signed and unsigned versions of the instructions treat the input registers as signed and unsigned 64-bit integers respectively.

Table 5-16. Integer Multiply and Add Instructions

Integer Multiply and Add	Low	High
Signed	<code>xma.l</code>	<code>xma.h</code>
Unsigned	<code>xma.lu</code> (pseudo-op)	<code>xma.hu</code>

5.4 Additional IEEE Considerations

5.4.1 Definition of SNaNs, QNaNs, and Propagation of NaNs

Signaling NaNs have a zero in the most significant fractional bit of the significand. Quiet NaNs have a one in the most significant fractional bit of the significand. This definition of signaling and quiet NaNs easily preserves “NaNness” when converting between different precisions. When propagating NaNs in operations that have more than one NaN operand, the result NaN is chosen from one of the operand NaNs in the following priority based on register encoding fields: first $f4$, then $f2$, and lastly $f3$.

5.4.2 IEEE Standard Mandated Operations Deferred to Software

The following IEEE mandated operations will be implemented in software:

- String to floating-point conversion.
- Floating-point to string conversion.
- Divide (with help from `frcpa` or `fprcpa` instruction).
- Square root (with help from `frsqrrta` or `fprsqrrta` instruction).
- Remainder (with help from `frcpa` or `fprcpa` instruction).
- Floating-point to integer valued floating-point conversion.
- Correctly wrapping the exponent for single, double, and double-extended overflow and underflow values, as recommended by the IEEE standard.

5.4.3 Additions beyond the IEEE Standard

- The fused multiply and add (`fma`, `fms`, `fnma`, `fpma`, `fpms`, `fpnma`) operations enable efficient software divide, square root, and remainder algorithms.
- The extended range of the 17-bit exponent in the register file format allows simplified implementation of many basic numeric algorithms by the careful numeric programmer.
- The NaTVal is a natural extension of the IEEE concept of NaNs. It is used to support speculative execution.
- Flush-to-Zero mode is an industry standard addition.
- The minimum and maximum instructions allow the efficient execution of the common Fortran Intrinsic Functions: `MIN()`, `MAX()`, `AMIN()`, `AMAX()`; and C language idioms such as `a<b?a:b`.
- All mixed precision operations are allowed. The IEEE standard suggests that implementations allow lower precision operands to produce higher precision results; this is supported. The IEEE standard also suggests that implementations not allow higher precision operands to produce lower precision results; this suggestion is not followed.
- An IEEE style quad-precision real type that is supported in software.

This chapter describes the function of IA-64 instruction. The pages of this chapter are sorted alphabetically by assembly language mnemonic.

6.1 Instruction Page Conventions

The instruction pages are divided into multiple sections as listed in Table 6-1. The first four sections are present on all instruction pages. The last three sections are present only when necessary. Table 6-2 lists the font conventions which are used by the instruction pages.

Table 6-1. Instruction Page Description

Section Name	Contents
Format	Assembly language syntax, instruction type and encoding format
Description	Instruction function in English
Operation	Instruction function in C code
FP Exceptions	IEEE floating-point traps

Table 6-2. Instruction Page Font Conventions

Font	Interpretation
regular	(Format section) Required characters in an assembly language mnemonic
<i>italic</i>	(Format section) Assembly language field name that must be filled with one of a range of legal values listed in the Description section
code	(Operation section) C code specifying instruction behavior
<i>code_italic</i>	(Operation section) Assembly language field name corresponding to a <i>italic</i> field listed in the Format section

In the Format section, register addresses are specified using the assembly mnemonic field names given in the third column of Table 6-3. For instructions that are predicated, the Description section assumes that the qualifying predicate is true (except for instructions that modify architectural state when their qualifying predicate is false). The test of the qualifying predicate is included in the Operation section (when applicable).

In the Operation section, registers are addressed using the notation `reg[addr].field`. The register file being accessed is specified by `reg`, and has a value chosen from the second column of Table 6-3. The `addr` field specifies a register address as an assembly language field name or a register mnemonic. For the general, floating-point, and predicate register files which undergo register renaming, `addr` is the register address prior to renaming and the renaming is not shown. The `field` option specifies a named bit field within the register. If `field` is absent, then all fields of the register are accessed. The only exception is when referencing the data field of the general registers (64-bits not including the NaT bit) where the notation `GR[addr]` is used. The syntactical differences between the code found in the Operation section and standard C is listed in Table 6-4.

Table 6-3. Register File Notation

Register File	C Notation	Assembly Mnemonic	Indirect Access
Application registers	AR	ar	
Branch registers	BR	b	
CPU identification registers	CPUID	cpuid	Y

Table 6-3. Register File Notation (Continued)

Register File	C Notation	Assembly Mnemonic	Indirect Access
Floating-point registers	FR	f	
General registers	GR	r	
Performance monitor data registers	PMD	pmd	Y
Predicate registers	PR	p	

Table 6-4. C Syntax Differences

Syntax	Function
{msb:lsb}, {bit}	Bit field specifier. When appended to a variable, denotes a bit field extending from the most significant bit specified by “msb” to the least significant bit specified by “lsb” including bits “msb” and “lsb”. If “msb” and “lsb” are equal then a single bit is accessed. The second form denotes a single bit.
u>, u>=, u<, u<=	Unsigned inequality relations. Variables on either side of the operator are treated as unsigned.
u>>, u>>=	Unsigned right shift. Zeroes are shifted into the most significant bit position.
u+	Unsigned addition. Operands are treated as unsigned, and zero-extended.
u*	Unsigned multiplication. Operands are treated as unsigned.

6.2 Instruction Descriptions

The remainder of this chapter provides a description of IA-64 instruction.

Add

Format:	<i>(qp)</i> add $r_1 = r_2, r_3$	register_form	A1
	<i>(qp)</i> add $r_1 = r_2, r_3, 1$	plus1_form, register_form	A1
	<i>(qp)</i> add $r_1 = imm, r_3$	pseudo-op	
	<i>(qp)</i> adds $r_1 = imm_{14}, r_3$	imm14_form	A4
	<i>(qp)</i> addl $r_1 = imm_{22}, r_3$	imm22_form	A5

Description: The two source operands (and an optional constant 1) are added and the result placed in GR r_1 . In the register form the first operand is GR r_2 ; in the imm_14 form the first operand is taken from the sign extended imm_{14} encoding field; in the imm22_form the first operand is taken from the sign extended imm_{22} encoding field. In the imm22_form, GR r_3 can specify only GRs 0, 1, 2 and 3.

The plus1_form is available only in the register_form (although the equivalent effect in the immediate forms can be achieved by adjusting the immediate).

The immediate-form pseudo-op chooses the imm14_form or imm22_form based upon the size of the immediate operand and the value in GR r_3 .

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    if (register_form)                // register form
        tmp_src = GR[r2];
    else if (imm14_form)              // 14-bit immediate form
        tmp_src = sign_ext(imm14, 14);
    else                               // 22-bit immediate form
        tmp_src = sign_ext(imm22, 22);

    tmp_nat = (register_form ? GR[r2].nat : 0);

    if (plus1_form)
        GR[r1] = tmp_src + GR[r3] + 1;
    else
        GR[r1] = tmp_src + GR[r3];

    GR[r1].nat = tmp_nat || GR[r3].nat;
}

```

Add Pointer

Format: (qp) addp4 $r_1 = r_2, r_3$ register_form A1
 (qp) addp4 $r_1 = imm_{14}, r_3$ imm14_form A4

Description: The two source operands are added. The upper 32 bits of the result are forced to zero, and then bits {31:30} of GR r_3 are copied to bits {62:61} of the result. This result is placed in GR r_1 . In the register_form the first operand is GR r_2 ; in the imm14_form the first operand is taken from the sign extended imm_{14} encoding field.




Figure 6-1. Add Pointer

```

Operation:  if (PR[qp]) {
                check_target_register(r1);

                tmp_src = (register_form ? GR[r2] : sign_ext(imm14, 14));
                tmp_nat = (register_form ? GR[r2].nat : 0);

                tmp_res = tmp_src + GR[r3];
                tmp_res = zero_ext(tmp_res{31:0}, 32);
                tmp_res{62:61} = GR[r3]{31:30};
                GR[r1] = tmp_res;
                GR[r1].nat = tmp_nat || GR[r3].nat;
            }
  
```

Allocate Stack Frame

Format: `alloc r_1 = ar.pfs, i, l, o, r`

M34

Description: A new stack frame is allocated on the general register stack, and the Previous Function State register (PFS) is copied to GR r_1 . The change of frame size is immediate. The write of GR r_1 and subsequent instructions in the same instruction group use the new frame. This instruction cannot be predicated.

The four parameters, i (size of inputs), l (size of locals), o (size of outputs), and r (size of rotating) specify the sizes of the regions of the stack frame.




Figure 6-2. Stack Frame

The size of the frame (sof) is determined by $i + l + o$. Note that this instruction may grow or shrink the size of the current register stack frame. The size of the local region (sol) is given by $i + l$. There is no real distinction between inputs and locals. They are given as separate operands in the instruction only as a hint to the assembler about how the local registers are to be used.

The rotating registers must fit within the stack frame and be a multiple of 8 in number. If this instruction attempts to change the size of CFM.sor, and the register rename base registers (CFM.rrb.gr, CFM.rrb.fr, CFM.rrb.pr) are not all zero, then the instruction will cause a Reserved Register/Field fault.

Although the assembler does not allow illegal combinations of operands for alloc, illegal combinations can be encoded in the instruction. Attempting to allocate a stack frame larger than 96 registers, or with the rotating region larger than the stack frame, or with the size of locals larger than the stack frame, will cause an Illegal Operation fault. An alloc instruction must be the first instruction in an instruction group. Otherwise, the results are undefined.

If insufficient registers are available to allocate the desired frame alloc will stall the processor until enough dirty registers are written to the backing store. Such mandatory RSE stores may cause the data related faults listed below.

Operation:

```

tmp_sof = i + l + o;
tmp_sol = i + l;
tmp_sor = r u>> 3;
check_target_register_sof( $r_1$ , tmp_sof);
if (tmp_sof u> 96 || r u> tmp_sof || tmp_sol u> tmp_sof)
    illegal_operation_fault();
if (tmp_sor != CFM.sor &&
    (CFM.rrb.gr != 0 || CFM.rrb.fr != 0 || CFM.rrb.pr != 0))
    reserved_register_field_fault();

alat_frame_update(0, tmp_sof - CFM.sof);
rse_new_frame(CFM.sof, tmp_sof); // Make room for new registers; Mandatory RSE
                                // stores can raise faults listed below.

CFM.sof = tmp_sof;
CFM.sol = tmp_sol;
CFM.sor = tmp_sor;

GR[ $r_1$ ] = AR[PFS];
GR[ $r_1$ ].nat = 0;

```

Logical And

Format: (qp) and $r_1 = r_2, r_3$ register_form A1
 (qp) and $r_1 = imm_8, r_3$ imm8_form A3

Description: The two source operands are logically ANDed and the result placed in GR r_1 . In the register_form the first operand is GR r_2 ; in the imm8_form the first operand is taken from the imm_8 encoding field.

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    tmp_src = (register_form ? GR[r2] : sign_ext(imm8, 8));
    tmp_nat = (register_form ? GR[r2].nat : 0);

    GR[r1] = tmp_src & GR[r3];
    GR[r1].nat = tmp_nat || GR[r3].nat;
}

```


And Complement

Format:

(qp) andcm $r_1 = r_2, r_3$	register_form	A1
(qp) andcm $r_1 = imm_8, r_3$	imm8_form	A3

Description: The first source operand is logically ANDed with the 1's complement of the second source operand and the result placed in GR r_1 . In the register_form the first operand is GR r_2 ; in the imm8_form the first operand is taken from the imm_8 encoding field.

Operation:

```

if (PR[qp]) {
    check_target_register( $r_1$ );

    tmp_src = (register_form ? GR[ $r_2$ ] : sign_ext( $imm_8$ , 8));
    tmp_nat = (register_form ? GR[ $r_2$ ].nat : 0);

    GR[ $r_1$ ] = tmp_src & ~GR[ $r_3$ ];
    GR[ $r_1$ ].nat = tmp_nat || GR[ $r_3$ ].nat;
}

```

Branch

Format:	(<i>qp</i>) br. <i>btype</i> . <i>bwh</i> . <i>ph</i> . <i>dh</i> <i>target</i> ₂₅	ip_relative_form	B1
	(<i>qp</i>) br. <i>btype</i> . <i>bwh</i> . <i>ph</i> . <i>dh</i> <i>b</i> ₁ = <i>target</i> ₂₅	call_form, ip_relative_form	B3
	br. <i>btype</i> . <i>bwh</i> . <i>ph</i> . <i>dh</i> <i>target</i> ₂₅	counted_form, ip_relative_form	B2
	br. <i>ph</i> . <i>dh</i> <i>target</i> ₂₅	pseudo-op	
	(<i>qp</i>) br. <i>btype</i> . <i>bwh</i> . <i>ph</i> . <i>dh</i> <i>b</i> ₂	indirect_form	B4
	(<i>qp</i>) br. <i>btype</i> . <i>bwh</i> . <i>ph</i> . <i>dh</i> <i>b</i> ₁ = <i>b</i> ₂	call_form, indirect_form	B5
	br. <i>ph</i> . <i>dh</i> <i>b</i> ₂	pseudo-op	

Description: A branch calculation is evaluated, and either a branch is taken, or execution continues with the next sequential instruction. The execution of a branch logically follows the execution of all previous non-branch instructions in the same instruction group. On a taken branch, execution begins at slot 0.

Branches can be either IP-relative, or indirect. For IP-relative branches, the *target*₂₅ operand, in assembly, specifies a label to branch to. This is encoded in the branch instruction as a signed immediate displacement (*imm*₂₁) between the target bundle and the bundle containing this instruction (*imm*₂₁ = *target*₂₅ – IP >> 4). For indirect branches, the target address is taken from BR *b*₂.

Table 6-5. Branch Types

<i>btype</i>	Function	Branch Condition	Target Address
cond or <i>none</i>	Conditional branch	Qualifying predicate	IP-rel or Indirect
call	Conditional procedure call	Qualifying predicate	IP-rel or Indirect
ret	Conditional procedure return	Qualifying predicate	Indirect
ia	Invoke IA-32 instruction set	Unconditional	Indirect
cloop	Counted loop branch	Loop count	IP-rel
ctop, cexit	Mod-scheduled counted loop	Loop count and epilog count	IP-rel
wtop, wexit	Mod-scheduled while loop	Qualifying predicate and epilog count	IP-rel

There are two pseudo-ops for unconditional branches. These are encoded like a conditional branch (*btype* = cond), with the *qp* field specifying PR 0, and with the *bwh* hint of sptk.

The branch type determines how the branch condition is calculated and whether the branch has other effects (such as writing a link register). For the basic branch types, the branch condition is simply the value of the specified predicate register. These basic branch types are:

- **cond:** If the qualifying predicate is 1, the branch is taken. Otherwise it is not taken.
- **call:** If the qualifying predicate is 1, the branch is taken and several other actions occur:
 - The current values of the Current Frame Marker (CFM), the EC application register and the current privilege level are saved in the Previous Function State application register.
 - The caller's stack frame is effectively saved and the callee is provided with a frame containing only the caller's output region.
 - The rotation rename base registers in the CFM are reset to 0.
 - A return link value is placed in BR *b*₁.
- **return:** If the qualifying predicate is 1, the branch is taken and the following occurs:
 - CFM, EC, and the current privilege level are restored from PFS. (The privilege level is restored only if this does not increase privilege.)
 - The caller's stack frame is restored.
 - If the return lowers the privilege, and PSR.lp is 1, then a Lower-privilege Transfer trap is taken.
- **ia:** The branch is taken unconditionally, if it is not intercepted by the OS. The effect of the branch is to invoke the IA-32 instruction set (by setting PSR.is to 1) and begin processing IA-32 instructions at the virtual linear target address contained in BR *b*₂{31:0}. If the qualifying predicate is not PR 0, an Illegal Operation fault is raised.

The IA-32 target effective address is calculated relative to the current code segment, i.e. $EIP\{31:0\} = BR\ b_2\{31:0\} - CSD.base$. The IA-32 instruction set can be entered at any privilege level, provided instruction set transitions are not disabled.

Software must ensure the code segment descriptor (CSD) and selector (CS) are loaded before issuing the branch. If the target EIP value exceeds the code segment limit or has a code segment privilege violation, an IA-32_Exception(GPFault) is raised on the target IA-32 instruction. For entry into 16-bit IA-32 code, if $BR\ b_2$ is not within 64K-bytes of CSD.base a GPFault is raised on the target instruction. EFLAG.rf is unmodified until the successful completion of the first IA-32 instruction. EFLAG.rf is not cleared until the target IA-32 instruction successfully completes.

Software must issue a mf instruction before the branch if memory ordering is required between IA-32 processor consistent and IA-64 unordered memory references. The processor does not ensure IA-64-instruction-set-generated writes into the instruction stream are seen by subsequent IA-32 instruction fetches. br . ia does not perform an instruction serialization operation. The processor does ensure that prior writes (even in the same instruction group) to GRs and FRs are observed by the first IA-32 instruction. Writes to ARs within the same instruction group as br . ia are not allowed, since br . ia may implicitly reads all ARs. If an illegal RAW dependency is present between an AR write and br . ia, the first IA-32 instruction fetch and execution may or may not see the updated AR value.

IA-32 instruction set execution leaves the contents of the ALAT undefined. Software can not rely on ALAT values being preserved across an instruction set transition. On entry to IA-32 code, existing entries in the ALAT are ignored. If the register stack contains any dirty registers, an Illegal Operation fault is raised on the br . ia instruction. All registers left in the current register stack frame are left undefined during IA-32 instruction set execution. The current register stack frame is forced to zero. To flush the register file of dirty registers, the flushrs instruction must be issued in an instruction group proceeding the br . ia instruction. To enhance the performance of the instruction set transition, software can start the IA-64 register stack flush in parallel with starting the IA-32 instruction set by 1) ensuring flushrs is exactly one instruction group before the br . ia, and 2) br . ia is in the first B-slot. br . ia should always be executed in the first B-slot with a hint of “static-taken” (default), otherwise processor performance will be degraded.

Another branch type is provided for simple counted loops. This branch type uses the Loop Count application register (LC) to determine the branch condition, and does not use a qualifying predicate:

- **cloop:** If the LC register is not equal to zero, it is decremented and the branch is taken.

In addition to these simple branch types, there are four types which are used for accelerating modulo-scheduled loops. Two of these are for counted loops (which use the LC register), and two for while loops (which use the qualifying predicate). These loop types use register rotation to provide register renaming, and they use predication to turn off instructions that correspond to empty pipeline stages.

The Epilog Count application register (EC) is used to count epilog stages and, for some while loops, a portion of the prolog stages. In the epilog phase, EC is decremented each time around and, for most loops, when EC is one, the pipeline has been drained, and the loop is exited. For certain types of optimized, unrolled software-pipelined loops, the target of a br . cexit or br . wexit is set to the next sequential bundle. In this case, the pipeline may not be fully drained when EC is one, and continues to drain while EC is zero.

For these modulo-scheduled loop types, the calculation of whether the branch is taken or not depends on the kernel branch condition (LC for counted types, and the qualifying predicate for while types) and on the epilog condition (whether EC is greater than one or not).

These branch types are of two categories: top and exit. The top types (ctop and wtop) are used when the loop decision is located at the bottom of the loop body and therefore a taken branch will continue the loop while a fall through branch will exit the loop. The exit types (cexit and wexit) are used when the loop decision is located somewhere other than the bottom of the loop and therefore a fall through branch will continue the loop and a taken branch will exit the loop. The exit types are also used at intermediate points in an unrolled pipelined loop.

The modulo-scheduled loop types are:

- **ctop** and **cexit**: These branch types behave identically, except in the determination of whether to branch or not. For `br.ctop`, the branch is taken if either LC is non-zero or EC is greater than one. For `br.cexit`, the opposite is true. It is not taken if either LC is non-zero or EC is greater than one and is taken otherwise.

These branch types also use LC and EC to control register rotation and predicate initialization. During the prolog and kernel phase, when LC is non-zero, LC counts down. When `br.ctop` or `br.cexit` is executed with LC equal to zero, the epilog phase is entered, and EC counts down. When `br.ctop` or `br.cexit` is executed with LC equal to zero and EC equal to one, a final decrement of EC and a final register rotation are done. If LC and EC are equal to zero, register rotation stops. These other effects are the same for the two branch types, and are described in Figure 6-3.

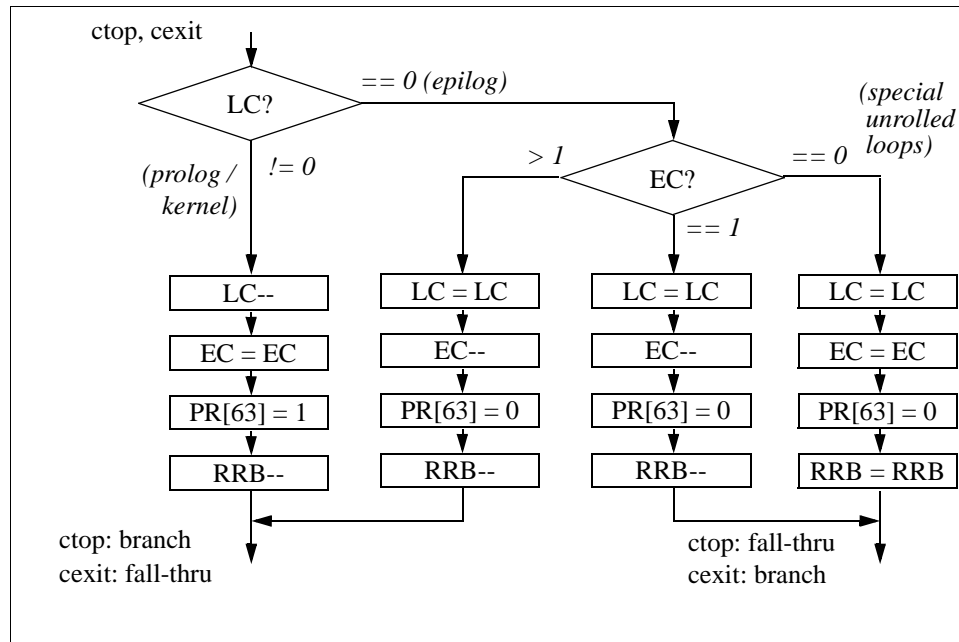


Figure 6-3. Operation of `br.ctop` and `br.cexit`

wtop and **wexit**: These branch types behave identically, except in the determination of whether to branch or not. For `br.wtop`, the branch is taken if either the qualifying predicate is one or EC is greater than one. For `br.wexit`, the opposite is true. It is not taken if either the qualifying predicate is one or EC is greater than one, and is taken otherwise.

These branch types also use the qualifying predicate and EC to control register rotation and predicate initialization. During the prolog phase, the qualifying predicate is either zero or one, depending upon the scheme used to program the loop. During the kernel phase, the qualifying predicate is one. During the epilog phase, the qualifying predicate is zero, and EC counts down. When `br.wtop` or `br.wexit` is executed with the qualifying predicate equal to zero and EC equal to one, a final decrement of EC and a final register rotation are done. If the qualifying predicate and EC are zero, register rotation stops. These other effects are the same for the two branch types, and are described in Figure 6-4.

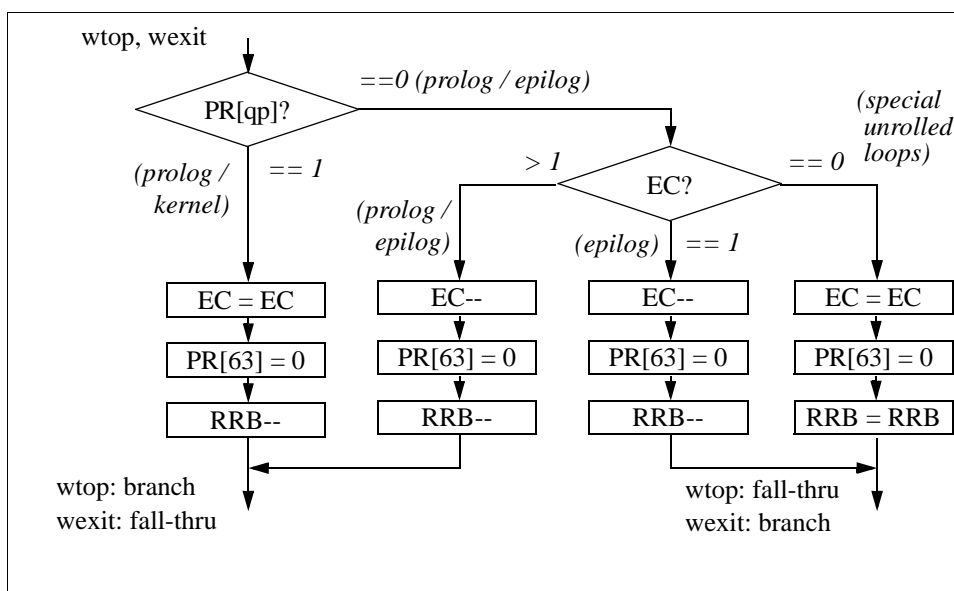


Figure 6-4. Operation of br.wtop and br.wexit

The loop-type branches (*br.cloop*, *br.ctop*, *br.cexit*, *br.wtop*, and *br.wexit*) are only allowed in instruction slot 2 within a bundle. Executing such an instruction in either slot 0 or 1 will cause an Illegal Operation fault, whether the branch would have been taken or not.

Read after Write (RAW) and Write after Read (WAR) dependency requirements are slightly different for branch instructions. Changes to BRs, PRs, and PFS by non-branch instructions are visible to a subsequent branch instruction in the same instruction group (i.e., a limited RAW is allowed for these resources). This allows for a low-latency compare-branch sequence, for example. The normal RAW requirements apply to the LC and EC application registers, and the RRBs.

Within an instruction group, a WAR dependency on PR 63 is not allowed if both the reading and writing instructions are branches. For example, a *br.wtop* or *br.wexit* may not use PR[63] as its qualifying predicate and PR[63] cannot be the qualifying predicate for any branch preceding a *br.wtop* or *br.wexit* in the same instruction group.

For dependency purposes, the loop-type branches effectively always write their associated resources, whether they are taken or not. The *cloop* type effectively always writes LC. When LC is 0, a *cloop* branch leaves it unchanged, but hardware may implement this as a re-write of LC with the same value. Similarly, *br.ctop* and *br.cexit* effectively always write LC, EC, the RRBs, and PR[63]. *br.wtop* and *br.wexit* effectively always write EC, the RRBs, and PR[63].

Values for various branch hint completers are shown in the following tables. Whether Prediction Strategy hints are shown in Table 6-6. Sequential Prefetch hints are shown in Table 6-7. Branch Cache Deallocation hints are shown in Table 6-8.

Table 6-6. Branch Whether Hint

<i>bwh</i> Completer	Branch Whether Hint
spnt	Static Not-Taken
sptk	Static Taken
dpnt	Dynamic Not-Taken
dptk	Dynamic Taken

Table 6-7. Sequential Prefetch Hint

<i>ph</i> Completer	Sequential Prefetch Hint
few or none	Few lines
many	Many lines

Table 6-8. Branch Cache Deallocation Hint

<i>dh</i> Completer	Branch Cache Deallocation Hint
<i>none</i>	Don't deallocate
<i>clr</i>	Deallocate branch information

```

Operation:  if (ip_relative_form)                               // determine branch target
            tmp_IP = IP + sign_ext((imm21 << 4), 25);
            else // indirect_form
            tmp_IP = BR[b2];

            if (btype != 'ia')                               // for IA-64 branches,
            tmp_IP = tmp_IP & ~0xf;                         // ignore bottom 4 bits of target

            lower_priv_transition = 0;

            switch (btype) {
            case 'cond':                                     // simple conditional branch
            tmp_taken = PR[qp];
            break;

            case 'call':                                     // call saves a return link
            tmp_taken = PR[qp];
            if (tmp_taken) {
                BR[b1] = IP + 16;

                AR[PFS].pfm = CFM;                          // ... and saves the stack frame
                AR[PFS].pec = AR[EC];
                AR[PFS].ppl = PSR.cpl;

                alat_frame_update(CFM.sol, 0);
                rse_preserve_frame(CFM.sol);
                CFM.sof -= CFM.sol;                          // new frame size is size of outs
                CFM.sol = 0;
                CFM.sor = 0;
                CFM.rrb.gr = 0;
                CFM.rrb.fr = 0;
                CFM.rrb.pr = 0;
            }
            break;

            case 'ret':                                     // return restores stack frame
            tmp_taken = PR[qp];
            if (tmp_taken) {
                // tmp_growth indicates the amount to move logical TOP *up*:
                // tmp_growth = sizeof(previous out) - sizeof(current frame)
                // a negative amount indicates a shrinking stack
                tmp_growth = (AR[PFS].pfm.sof - AR[PFS].pfm.sol) - CFM.sof;
                alat_frame_update(-AR[PFS].pfm.sol, 0);
                rse_fatal = rse_restore_frame(AR[PFS].pfm.sol, tmp_growth, CFM.sof);
                if (rse_fatal) {
                    CFM.sof = 0;
                    CFM.sol = 0;
                    CFM.sor = 0;
                    CFM.rrb.gr = 0;
                    CFM.rrb.fr = 0;
                    CFM.rrb.pr = 0;
                } else // normal branch return
                CFM = AR[PFS].pfm;

                rse_enable_current_frame_load();
                AR[EC] = AR[PFS].pec;
                if (PSR.cpl < AR[PFS].ppl) {                 // ... and restores privilege
                    PSR.cpl = AR[PFS].ppl;
                    lower_priv_transition = 1;
                }
            }
        }

```

```

        break;

    case 'ia':
        // switch to IA mode
        tmp_taken = 1;
        if (qp != 0)
            illegal_operation_fault();
        if (AR[BSPSTORE] != AR[BSP])
            illegal_operation_fault();
        if (PSR.di)
            disabled_instruction_set_transition_fault();
        PSR.is = 1;
        CFM.sof = 0;
        CFM.sol = 0;
        CFM.sor = 0;
        CFM.rrb.gr = 0;
        CFM.rrb.fr = 0;
        CFM.rrb.pr = 0;
        rse_invalidate_non_current_regs();

        // Note the register stack is disabled during IA-32 instruction set execution
        break;

    case 'cloop':
        // simple counted loop
        if (slot != 2)
            illegal_operation_fault();
        tmp_taken = (AR[LC] != 0);
        if (AR[LC] != 0)
            AR[LC]--;
        break;

    case 'ctop':
    case 'cexit':
        // SW pipelined counted loop
        if (slot != 2)
            illegal_operation_fault();
        if (btype == 'ctop') tmp_taken = ((AR[LC] != 0) || (AR[EC] u> 1));
        if (btype == 'cexit') tmp_taken = !((AR[LC] != 0) || (AR[EC] u> 1));
        if (AR[LC] != 0) {
            AR[LC]--;
            AR[EC] = AR[EC];
            PR[63] = 1;
            rotate_regs();
        } else if (AR[EC] != 0) {
            AR[LC] = AR[LC];
            AR[EC]--;
            PR[63] = 0;
            rotate_regs();
        } else {
            AR[LC] = AR[LC];
            AR[EC] = AR[EC];
            PR[63] = 0;
            CFM.rrb.gr = CFM.rrb.gr;
            CFM.rrb.fr = CFM.rrb.fr;
            CFM.rrb.pr = CFM.rrb.pr;
        }
        break;

    case 'wtop':
    case 'wexit':
        // SW pipelined while loop
        if (slot != 2)
            illegal_operation_fault();
        if (btype == 'wtop') tmp_taken = (PR[qp] || (AR[EC] u> 1));
        if (btype == 'wexit') tmp_taken = !(PR[qp] || (AR[EC] u> 1));
        if (PR[qp]) {
            AR[EC] = AR[EC];
            PR[63] = 0;
            rotate_regs();
        } else if (AR[EC] != 0) {
            AR[EC]--;

```

```
        PR[63] = 0;
        rotate_regs();
    } else {
        AR[EC] = AR[EC];
        PR[63] = 0;
        CFM.rrb.gr = CFM.rrb.gr;
        CFM.rrb.fr = CFM.rrb.fr;
        CFM.rrb.pr = CFM.rrb.pr;
    }
    break;
}
if (tmp_taken) {
    taken_branch = 1;
    IP = tmp_IP; // set the new value for IP
    if ((PSR.it && unimplemented_virtual_address(tmp_IP))
        || (!PSR.it && unimplemented_physical_address(tmp_IP)))
        unimplemented_instruction_address_trap(lower_priv_transition, tmp_IP);
    if (lower_priv_transition && PSR.lp)
        lower_privilege_transfer_trap();
    if (PSR.tb)
        taken_branch_trap();
}
```


Break

Format:	<i>(qp)</i> break <i>imm</i> ₂₁	pseudo-op	
	<i>(qp)</i> break.i <i>imm</i> ₂₁	i_unit_form	I19
	<i>(qp)</i> break.b <i>imm</i> ₂₁	b_unit_form	B9
	<i>(qp)</i> break.m <i>imm</i> ₂₁	m_unit_form	M37
	<i>(qp)</i> break.f <i>imm</i> ₂₁	f_unit_form	F15
	<i>(qp)</i> break.x <i>imm</i> ₆₂	x_unit_form	X1

Description: A Break Instruction fault is taken. For the i_unit_form, f_unit_form and m_unit_form, the value specified by *imm*₂₁ is zero-extended and placed in the Interruption Immediate control register (IIM).

For the b_unit_form, *imm*₂₁ is ignored and the value zero is placed in the Interruption Immediate control register (IIM).

For the x_unit_form, the lower 21 bits of the value specified by *imm*₆₂ is zero-extended and placed in the Interruption Immediate control register (IIM). The L slot of the bundle contains the upper 41 bits of *imm*₆₂.

This instruction has five forms, each of which can be executed only on a particular execution unit type. The pseudo-op can be used if the unit type to execute on is unimportant.

Operation:

```

if (PR[qp]) {
    if (b_unit_form)
        immediate = 0;
    else if (x_unit_form)
        immediate = zero_ext(imm62, 21);
    else // i_unit_form || m_unit_form || f_unit_form
        immediate = zero_ext(imm21, 21);

    break_instruction_fault(immediate);
}

```

Speculation Check

Format:	(<i>qp</i>) <i>chk.s</i> <i>r</i> ₂ , <i>target</i> ₂₅	pseudo-op	
	(<i>qp</i>) <i>chk.s.i</i> <i>r</i> ₂ , <i>target</i> ₂₅	control_form, i_unit_form, gr_form	I20
	(<i>qp</i>) <i>chk.s.m</i> <i>r</i> ₂ , <i>target</i> ₂₅	control_form, m_unit_form, gr_form	M20
	(<i>qp</i>) <i>chk.s</i> <i>f</i> ₂ , <i>target</i> ₂₅	control_form, fr_form	M21
	(<i>qp</i>) <i>chk.a.aclr</i> <i>r</i> ₁ , <i>target</i> ₂₅	data_form, gr_form	M22
	(<i>qp</i>) <i>chk.a.aclr</i> <i>f</i> ₁ , <i>target</i> ₂₅	data_form, fr_form	M23

Description: The result of a control- or data-speculative calculation is checked for success or failure. If the check fails, a branch to *target*₂₅ is taken.

In the control_form, success is determined by a NaT indication for the source register. If the NaT bit corresponding to GR *r*₂ is 1 (in the gr_form), or FR *f*₂ contains a NaTVal (in the fr_form), the check fails.

In the data_form, success is determined by the ALAT. The ALAT is queried using the general register specifier *r*₁ (in the gr_form), or the floating-point register specifier *f*₁ (in the fr_form). If no ALAT entry matches, the check fails. An implementation may optionally cause the check to fail independent of whether an ALAT entry matches.

The *target*₂₅ operand, in assembly, specifies a label to branch to. This is encoded in the instruction as a signed immediate displacement (*imm*₂₁) between the target bundle and the bundle containing this instruction (*imm*₂₁ = *target*₂₅ – IP >> 4).

The control_form of this instruction for checking general registers can be encoded on either an I-unit or an M-unit. The pseudo-op can be used if the unit type to execute on is unimportant.

For the data_form, if an ALAT entry matches, the matching ALAT entry can be optionally invalidated, based on the value of the *aclr* completer (See Table 6-9).

Table 6-9. ALAT Clear Completer

<i>aclr</i> Completer	Effect on ALAT
clr	Invalidate matching ALAT entry
nc	Don't invalidate

Note that if the *clr* value of the *aclr* completer is used and the check succeeds, the matching ALAT entry is invalidated. However, if the check fails (which may happen even if there is a matching ALAT entry), any matching ALAT entry may optionally be invalidated, but this is not required. Recovery code for data speculation, therefore, cannot rely on the absence of a matching ALAT entry.

```

Operation:  if (PR[qp]) {
            if (control_form) {
                if (fr_form && (tmp_isrkode = fp_reg_disabled(f2, 0, 0, 0)))
                    disabled_fp_register_fault(tmp_isrkode, 0);
                check_type = gr_form ? CHKS_GENERAL : CHKS_FLOAT;
                fail = (gr_form && GR[r2].nat) || (fr_form && FR[f2] == NATVAL);
            } else {
                reg_type = gr_form ? GENERAL : FLOAT;
                alat_index = gr_form ? r1 : (data_form ? f1 : f2);
                // data_form

                check_type = gr_form ? CHKA_GENERAL : CHKA_FLOAT;
                fail = !alat_cmp(reg_type, alat_index);
            }
            if (fail) {
                taken_branch = 1;
                IP = IP + sign_ext((imm21 << 4), 25);
                if ((PSR.it && unimplemented_virtual_address(IP))
                    || (!PSR.it && unimplemented_physical_address(IP)))
                    unimplemented_instruction_address_trap(0, IP);
                if (PSR.tb)
                    taken_branch_trap();
            }
            if (!fail && data_form && (aclr == 'clr'))
                alat_inval_single_entry(reg_type, alat_index);
        }

```

Clear RRB

Format: clrrrb all_form B8
clrrb.pr pred_form B8

Description: In the all_form, the register rename base registers (CFM.rrb.gr, CFM.rrb.fr, and CFM.rrb.pr) are cleared. In the pred_form, the single register rename base register for the predicates (CFM.rrb.pr) is cleared.

This instruction must be the last instruction in an instruction group, or an Illegal Operation fault is taken.

This instruction cannot be predicated.

Operation:

```
if (!followed_by_stop())
    illegal_operation_fault();

if (all_form) {
    CFM.rrb.gr = 0;
    CFM.rrb.fr = 0;
    CFM.rrb.pr = 0;
} else { // pred_form
    CFM.rrb.pr = 0;
}
```

Compare

Format:	(qp) cmp.crel.ctype $p_1, p_2 = r_2, r_3$	register_form	A6
	(qp) cmp.crel.ctype $p_1, p_2 = imm_8, r_3$	imm8_form	A8
	(qp) cmp.crel.ctype $p_1, p_2 = r0, r_3$	parallel_inequality_form	A7
	(qp) cmp.crel.ctype $p_1, p_2 = r_3, r0$	pseudo-op	

Description: The two source operands are compared for one of ten relations specified by *crel*. This produces a boolean result which is 1 if the comparison condition is true, and 0 otherwise. This result is written to the two predicate register destinations, p_1 and p_2 . The way the result is written to the destinations is determined by the compare type specified by *ctype*.

The compare types describe how the predicate targets are updated based on the result of the comparison. The normal type simply writes the compare result to one target, and the complement to the other. The parallel types update the targets only for a particular comparison result. This allows multiple simultaneous OR-type or multiple simultaneous AND-type compares to target the same predicate register.

The *unc* type is special in that it first initializes both predicate targets to 0, *independent of the qualifying predicate*. It then operates the same as the normal type. The behavior of the compare types is described in Table 6-10. A blank entry indicates the predicate target is left unchanged.

Table 6-10. Comparison Types

<i>ctype</i>	Pseudo-op of	PR[qp]==0		PR[qp]==1					
				result==0, No Source NaTs		result==1, No Source NaTs		One or More Source NaTs	
		PR[p ₁]	PR[p ₂]	PR[p ₁]	PR[p ₂]	PR[p ₁]	PR[p ₂]	PR[p ₁]	PR[p ₂]
<i>none</i>				0	1	1	0	0	0
<i>unc</i>		0	0	0	1	1	0	0	0
<i>or</i>				0	0	1	1	0	0
<i>and</i>				1	1	0	0	0	0
<i>or.andcm</i>	<i>or</i>			1	1	0	0	0	0
<i>orcm</i>	<i>and</i>			0	0	1	1	0	0
<i>andcm</i>	<i>or.andcm</i>			0	1	0	1	0	0
<i>and.orcm</i>				0	1	0	1	0	0

In the *register_form* the first operand is GR r_2 ; in the *imm8_form* the first operand is taken from the sign extended imm_8 encoding field; and in the *parallel_inequality_form* the first operand must be GR 0. The *parallel_inequality_form* is only used when the compare type is one of the parallel types, and the relation is an inequality ($>$, $>=$, $<$, $<=$). See below.

If the two predicate register destinations are the same (p_1 and p_2 specify the same predicate register), the instruction will take an Illegal Operation fault, if the qualifying predicate is set, or if the compare type is *unc*.

Of the ten relations, not all are directly implemented in hardware. Some are actually pseudo-ops. For these, the assembler simply switches the source operand specifiers and/or switches the predicate target specifiers and uses an implemented relation. For some of the pseudo-op compares in the *imm8_form*, the assembler subtracts 1 from the immediate value, making the allowed immediate range slightly different. Of the six parallel compare types, three of the types are actually pseudo-ops. The assembler simply uses the negative relation with an implemented type. The implemented relations and how the pseudo-ops map onto them are shown in Table 6-11 (for normal and *unc* type compares), and Table 6-12 (for parallel type compares).

Table 6-11. 64-bit Comparison Relations for Normal and unc Compares

<i>crel</i>	Compare Relation (<i>a rel b</i>)	Register Form is a Pseudo-op of	Immediate Form is a Pseudo-op of	Immediate Range
eq	$a == b$			-128 .. 127
ne	$a != b$	eq $p_1 \leftrightarrow p_2$	eq $p_1 \leftrightarrow p_2$	-128 .. 127
lt	$a < b$			-128 .. 127
le	$a \leq b$ signed	lt $a \leftrightarrow b$ $p_1 \leftrightarrow p_2$	lt a-1	-127 .. 128
gt	$a > b$	lt $a \leftrightarrow b$	lt a-1 $p_1 \leftrightarrow p_2$	-127 .. 128
ge	$a \geq b$	lt $p_1 \leftrightarrow p_2$	lt $p_1 \leftrightarrow p_2$	-128 .. 127
ltu	$a < b$			0 .. 127, $2^{64}-128 .. 2^{64}-1$
leu	$a \leq b$ unsigned	ltu $a \leftrightarrow b$ $p_1 \leftrightarrow p_2$	ltu a-1	1 .. 128, $2^{64}-127 .. 2^{64}$
gtu	$a > b$	ltu $a \leftrightarrow b$	ltu a-1 $p_1 \leftrightarrow p_2$	1 .. 128, $2^{64}-127 .. 2^{64}$
geu	$a \geq b$	ltu $p_1 \leftrightarrow p_2$	ltu $p_1 \leftrightarrow p_2$	0 .. 127, $2^{64}-128 .. 2^{64}-1$

The parallel compare types can be used only with a restricted set of relations and operands. They can be used with equal and not-equal comparisons between two registers or between a register and an immediate, or they can be used with inequality comparisons between a register and GR 0. Unsigned relations are not provided, since they are not of much use when one of the operands is zero. For the parallel inequality comparisons, hardware only directly implements the ones where the first operand (GR r_2) is GR 0. Comparisons where the second operand is GR 0 are pseudo-ops for which the assembler switches the register specifiers and uses the opposite relation.

Table 6-12. 64-bit Comparison Relations for Parallel Compares

<i>crel</i>	Compare Relation (<i>a rel b</i>)	Register Form is a Pseudo-op of	Immediate Range
eq	$a == b$		-128 .. 127
ne	$a != b$		-128 .. 127
lt	$0 < b$		no immediate forms
lt	$a < 0$	gt $a \leftrightarrow b$	
le	$0 \leq b$		
le	$a \leq 0$ signed	ge $a \leftrightarrow b$	
gt	$0 > b$		
gt	$a > 0$	lt $a \leftrightarrow b$	
ge	$0 \geq b$		
ge	$a \geq 0$	le $a \leftrightarrow b$	

```

Operation:  if (PR[qp]) {
            if (p1 == p2)
                illegal_operation_fault();

            tmp_nat = (register_form ? GR[r2].nat : 0) || GR[r3].nat;
            if (register_form)
                tmp_src = GR[r2];
            else if (imm8_form)
                tmp_src = sign_ext(imm8, 8);
            else // parallel_inequality_form
                tmp_src = 0;

            if (crel == 'eq') tmp_rel = tmp_src == GR[r3];
            else if (crel == 'ne') tmp_rel = tmp_src != GR[r3];
            else if (crel == 'lt') tmp_rel = lesser_signed(tmp_src, GR[r3]);
            else if (crel == 'le') tmp_rel = lesser_equal_signed(tmp_src, GR[r3]);
            else if (crel == 'gt') tmp_rel = greater_signed(tmp_src, GR[r3]);
            else if (crel == 'ge') tmp_rel = greater_equal_signed(tmp_src, GR[r3]);
            else if (crel == 'ltu') tmp_rel = lesser(tmp_src, GR[r3]);
            else if (crel == 'leu') tmp_rel = lesser_equal(tmp_src, GR[r3]);
            else if (crel == 'gtu') tmp_rel = greater(tmp_src, GR[r3]);
            else
                tmp_rel = greater_equal(tmp_src, GR[r3]); // 'geu'

            switch (ctype) {
                case 'and': // and-type compare
                    if (tmp_nat || !tmp_rel) {
                        PR[p1] = 0;
                        PR[p2] = 0;
                    }
                    break;
                case 'or': // or-type compare
                    if (!tmp_nat && tmp_rel) {
                        PR[p1] = 1;
                        PR[p2] = 1;
                    }
                    break;
                case 'or.andcm': // or.andcm-type compare
                    if (!tmp_nat && tmp_rel) {
                        PR[p1] = 1;
                        PR[p2] = 0;
                    }
                    break;
                case 'unc': // unc-type compare
                default: // normal compare
                    if (tmp_nat) {
                        PR[p1] = 0;
                        PR[p2] = 0;
                    } else {
                        PR[p1] = tmp_rel;
                        PR[p2] = !tmp_rel;
                    }
                    break;
            }
        } else {
            if (ctype == 'unc') {
                if (p1 == p2)
                    illegal_operation_fault();
                PR[p1] = 0;
                PR[p2] = 0;
            }
        }
    }

```

Compare Word

Format:	<i>(qp) cmp4.crel ctype p₁, p₂ = r₂, r₃</i>	register_form	A6
	<i>(qp) cmp4.crel ctype p₁, p₂ = imm₈, r₃</i>	imm8_form	A8
	<i>(qp) cmp4.crel ctype p₁, p₂ = r0, r₃</i>	parallel_inequality_form	A7
	<i>(qp) cmp4.crel ctype p₁, p₂ = r₃, r0</i>	pseudo-op	

Description: The least significant 32 bits from each of two source operands are compared for one of ten relations specified by *crel*. This produces a boolean result which is 1 if the comparison condition is true, and 0 otherwise. This result is written to the two predicate register destinations, *p₁* and *p₂*. The way the result is written to the destinations is determined by the compare type specified by *ctype*. See the Compare instruction and Table 6-10 on page 6-19.

In the *register_form* the first operand is GR *r₂*; in the *imm8_form* the first operand is taken from the sign extended *imm₈* encoding field; and in the *parallel_inequality_form* the first operand must be GR 0. The *parallel_inequality_form* is only used when the compare type is one of the parallel types, and the relation is an inequality (>, >=, <, <=). See the Compare instruction and Table 6-12 on page 6-20.

If the two predicate register destinations are the same (*p₁* and *p₂* specify the same predicate register), the instruction will take an Illegal Operation fault, if the qualifying predicate is set, or if the compare type is *unc*.

Of the ten relations, not all are directly implemented in hardware. Some are actually pseudo-ops. See the Compare instruction and Table 6-11 and Table 6-12 on page 6-20. The range for immediates is given below.

Table 6-13. Immediate Range for 32-bit Compares

<i>crel</i>	Compare Relation (<i>a rel b</i>)	Immediate Range
eq	<i>a == b</i>	-128 .. 127
ne	<i>a != b</i>	-128 .. 127
lt	<i>a < b</i>	-128 .. 127
le	<i>a <= b</i>	signed -127 .. 128
gt	<i>a > b</i>	
ge	<i>a >= b</i>	
ltu	<i>a < b</i>	unsigned 0 .. 127, 2 ³² -128 .. 2 ³² -1
leu	<i>a <= b</i>	
gtu	<i>a > b</i>	
geu	<i>a >= b</i>	

Operation:

```

if (PR[qp]) {
    if (p1 == p2)
        illegal_operation_fault();

    tmp_nat = (register_form ? GR[r2].nat : 0) || GR[r3].nat;

    if (register_form)
        tmp_src = GR[r2];
    else if (imm8_form)
        tmp_src = sign_ext(imm8, 8);
    else // parallel_inequality_form
        tmp_src = 0;

    if (crel == 'eq') tmp_rel = tmp_src{31:0} == GR[r3]{31:0};
    else if (crel == 'ne') tmp_rel = tmp_src{31:0} != GR[r3]{31:0};
    else if (crel == 'lt')
        tmp_rel = lesser_signed(sign_ext(tmp_src, 32), sign_ext(GR[r3], 32));
    else if (crel == 'le')
        tmp_rel = lesser_equal_signed(sign_ext(tmp_src, 32), sign_ext(GR[r3], 32));
    else if (crel == 'gt')

```



```

    tmp_rel = greater_signed(sign_ext(tmp_src, 32), sign_ext(GR[r3], 32));
else if (crel == 'ge')
    tmp_rel = greater_equal_signed(sign_ext(tmp_src, 32), sign_ext(GR[r3], 32));
else if (crel == 'ltu')
    tmp_rel = lesser(zero_ext(tmp_src, 32), zero_ext(GR[r3], 32));
else if (crel == 'leu')
    tmp_rel = lesser_equal(zero_ext(tmp_src, 32), zero_ext(GR[r3], 32));
else if (crel == 'gtu')
    tmp_rel = greater(zero_ext(tmp_src, 32), zero_ext(GR[r3], 32));
else // 'geu'
    tmp_rel = greater_equal(zero_ext(tmp_src, 32), zero_ext(GR[r3], 32));

switch (ctype) {
  case 'and': // and-type compare
    if (tmp_nat || !tmp_rel) {
      PR[p1] = 0;
      PR[p2] = 0;
    }
    break;
  case 'or': // or-type compare
    if (!tmp_nat && tmp_rel) {
      PR[p1] = 1;
      PR[p2] = 1;
    }
    break;
  case 'or.andcm': // or.andcm-type compare
    if (!tmp_nat && tmp_rel) {
      PR[p1] = 1;
      PR[p2] = 0;
    }
    break;
  case 'unc': // unc-type compare
  default: // normal compare
    if (tmp_nat) {
      PR[p1] = 0;
      PR[p2] = 0;
    } else {
      PR[p1] = tmp_rel;
      PR[p2] = !tmp_rel;
    }
    break;
}
} else {
  if (ctype == 'unc') {
    if (p1 == p2)
      illegal_operation_fault();
    PR[p1] = 0;
    PR[p2] = 0;
  }
}
}

```

Compare And Exchange

Format: $(qp) \text{ cmpxchg } sz.sem.lhint \ r_1 = [r_3], r_2, ar.ccv$ M16

Description: A value consisting of sz bytes is read from memory starting at the address specified by the value in GR r_3 . The value is zero extended and compared with the contents of the `cmpxchg` Compare Value application register (AR[CCV]). If the two are equal, then the least significant sz bytes of the value in GR r_2 are written to memory starting at the address specified by the value in GR r_3 . The zero-extended value read from memory is placed in GR r_1 and the NaT bit corresponding to GR r_1 is cleared.

The values of the sz completer are given in Table 6-14. The sem completer specifies the type of semaphore operation. These operations are described in Table 6-15.

Table 6-14. Memory Compare and Exchange Size

sz Completer	Bytes Accessed
1	1
2	2
4	4
8	8

Table 6-15. Compare and Exchange Semaphore Types

sem Completer	Ordering Semantics	Semaphore Operation
acq	Acquire	The memory read/write is made visible prior to all subsequent data memory accesses.
rel	Release	The memory read/write is made visible after all previous data memory accesses.

If the address specified by the value in GR r_3 is not naturally aligned to the size of the value being accessed in memory, an Unaligned Data Reference fault is taken independent of the state of the User Mask alignment checking bit, UM.ac (PSR.ac in the Processor Status Register).

The memory read and write are guaranteed to be atomic.

Both read and write access privileges for the referenced page are required. The write access privilege check is performed whether or not the memory write is performed.

The value of the $ldhint$ completer specifies the locality of the memory access. The values of the $ldhint$ completer are given in Table 6-28 on page 6-102. Locality hints do not affect program functionality and may be ignored by the implementation. See “Memory Hierarchy Control and Consistency” on page 4-16 for details.

```
Operation:  if (PR[qp]) {
            check_target_register(r1, SEMAPHORE);

            if (GR[r3].nat || GR[r2].nat)
                register_nat_consumption_fault(SEMAPHORE);

            paddr = tlb_translate(GR[r3], sz, SEMAPHORE, PSR.cpl, &mattr, &tmp_unused);

            if (!ma_supports_semaphores(mattr))
                unsupported_data_reference_fault(SEMAPHORE, GR[r3]);

            if (sem == 'acq') {
                val = mem_xchg_cond(AR[CCV], GR[r2], paddr, sz, UM.be, mattr, ACQUIRE,
                                   ldhint);
            } else { // 'rel'
                val = mem_xchg_cond(AR[CCV], GR[r2], paddr, sz, UM.be, mattr, RELEASE,
                                   ldhint);
            }
            val = zero_ext(val, sz * 8);

            if (AR[CCV] == val)
                alat_inval_multiple_entries(paddr, sz);

            GR[r1] = val;
            GR[r1].nat = 0;
        }
```

Compute Zero Index

Format:	<i>(qp)</i> <i>czx1.l</i> $r_1 = r_3$	one_byte_form, left_form	I29
	<i>(qp)</i> <i>czx1.r</i> $r_1 = r_3$	one_byte_form, right_form	I29
	<i>(qp)</i> <i>czx2.l</i> $r_1 = r_3$	two_byte_form, left_form	I29
	<i>(qp)</i> <i>czx2.r</i> $r_1 = r_3$	two_byte_form, right_form	I29

Description: GR r_3 is scanned for a zero element. The element is either an 8-bit aligned byte (*one_byte_form*) or a 16-bit aligned pair of bytes (*two_byte_form*). The index of the first zero element is placed in GR r_1 . If there are no zero elements in GR r_3 , a default value is placed in GR r_1 . Table 6-16 gives the possible result values. In the *left_form*, the source is scanned from most significant element to least significant element, and in the *right_form* it is scanned from least significant element to most significant element.

Table 6-16. Result Ranges for *czx*

Size	Element Width	Range of Result if Zero Element Found	Default Result if No Zero Element Found
1	8 bit	0-7	8
2	16 bit	0-3	4

```

Operation:  if (PR[qp]) {
                check_target_register( $r_1$ );

                if (one_byte_form) {
                    if (left_form) { // scan from most significant down
                        if ((GR[ $r_3$ ] & 0xff00000000000000) == 0) GR[ $r_1$ ] = 0;
                        else if ((GR[ $r_3$ ] & 0x00ff000000000000) == 0) GR[ $r_1$ ] = 1;
                        else if ((GR[ $r_3$ ] & 0x0000ff0000000000) == 0) GR[ $r_1$ ] = 2;
                        else if ((GR[ $r_3$ ] & 0x000000ff00000000) == 0) GR[ $r_1$ ] = 3;
                        else if ((GR[ $r_3$ ] & 0x00000000ff000000) == 0) GR[ $r_1$ ] = 4;
                        else if ((GR[ $r_3$ ] & 0x0000000000ff0000) == 0) GR[ $r_1$ ] = 5;
                        else if ((GR[ $r_3$ ] & 0x000000000000ff00) == 0) GR[ $r_1$ ] = 6;
                        else if ((GR[ $r_3$ ] & 0x00000000000000ff) == 0) GR[ $r_1$ ] = 7;
                        else GR[ $r_1$ ] = 8;
                    } else { // right_form scan from least significant up
                        if ((GR[ $r_3$ ] & 0x00000000000000ff) == 0) GR[ $r_1$ ] = 0;
                        else if ((GR[ $r_3$ ] & 0x000000000000ff00) == 0) GR[ $r_1$ ] = 1;
                        else if ((GR[ $r_3$ ] & 0x0000000000ff0000) == 0) GR[ $r_1$ ] = 2;
                        else if ((GR[ $r_3$ ] & 0x00000000ff000000) == 0) GR[ $r_1$ ] = 3;
                        else if ((GR[ $r_3$ ] & 0x000000ff00000000) == 0) GR[ $r_1$ ] = 4;
                        else if ((GR[ $r_3$ ] & 0x0000ff0000000000) == 0) GR[ $r_1$ ] = 5;
                        else if ((GR[ $r_3$ ] & 0x00ff000000000000) == 0) GR[ $r_1$ ] = 6;
                        else if ((GR[ $r_3$ ] & 0xff00000000000000) == 0) GR[ $r_1$ ] = 7;
                        else GR[ $r_1$ ] = 8;
                    }
                } else { // two_byte_form
                    if (left_form) { // scan from most significant down
                        if ((GR[ $r_3$ ] & 0xffff000000000000) == 0) GR[ $r_1$ ] = 0;
                        else if ((GR[ $r_3$ ] & 0x0000ffff00000000) == 0) GR[ $r_1$ ] = 1;
                        else if ((GR[ $r_3$ ] & 0x00000000ffff0000) == 0) GR[ $r_1$ ] = 2;
                        else if ((GR[ $r_3$ ] & 0x0000000000ffff) == 0) GR[ $r_1$ ] = 3;
                        else GR[ $r_1$ ] = 4;
                    } else { // right_form scan from least significant up
                        if ((GR[ $r_3$ ] & 0x000000000000ffff) == 0) GR[ $r_1$ ] = 0;
                        else if ((GR[ $r_3$ ] & 0x00000000ffff0000) == 0) GR[ $r_1$ ] = 1;
                        else if ((GR[ $r_3$ ] & 0x0000ffff00000000) == 0) GR[ $r_1$ ] = 2;
                        else if ((GR[ $r_3$ ] & 0xffff000000000000) == 0) GR[ $r_1$ ] = 3;
                        else GR[ $r_1$ ] = 4;
                    }
                }
                GR[ $r_1$ ].nat = GR[ $r_3$ ].nat;
            }
    
```

Deposit

Format:	(qp) dep $r_1 = r_2, r_3, pos_6, len_4$	merge_form, register_form	115
	(qp) dep $r_1 = imm_1, r_3, pos_6, len_6$	merge_form, imm_form	114
	(qp) dep.z $r_1 = r_2, pos_6, len_6$	zero_form, register_form	112
	(qp) dep.z $r_1 = imm_8, pos_6, len_6$	zero_form, imm_form	113

Description: In the merge_form, a right justified bit field taken from the first source operand is deposited into the value in GR r_3 at an arbitrary bit position and the result is placed in GR r_1 . In the register_form the first source operand is GR r_2 ; and in the imm_form it is the sign-extended value specified by imm_1 (either all ones or all zeroes). The deposited bit field begins at the bit position specified by the pos_6 immediate and extends to the left (towards the most significant bit) a number of bits specified by the len immediate. Note that len has a range of 1-16 in the register_form and 1-64 in the imm_form. The pos_6 immediate has a range of 0 to 63.

In the zero_form, a right justified bit field taken from either the value in GR r_2 (in the register_form) or the sign extended value in imm_8 (in the imm_form) is deposited into GR r_1 and all other bits in GR r_1 are cleared to zero. The deposited bit field begins at the bit position specified by the pos_6 immediate and extends to the left (towards the most significant bit) a number of bits specified by the len immediate. The len immediate has a range of 1-64 and the pos_6 immediate has a range of 0 to 63.

In the event that the deposited bit field extends beyond bit 63 of the target, i.e., $len + pos_6 > 64$, the most significant $len + pos_6 - 64$ bits of the deposited bit field are truncated. The len immediate is encoded as len minus 1 in the instruction.

The operation of `dep t = s, r, 36, 16` is illustrated in Figure 6-5.




Figure 6-5. Deposit Example

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    if (imm_form) {
        tmp_src = (merge_form ? sign_ext(imm1, 1) : sign_ext(imm8, 8));
        tmp_nat = merge_form ? GR[r3].nat : 0;
        tmp_len = len6 ;
    } else { // register_form
        tmp_src = GR[r2];
        tmp_nat = (merge_form ? GR[r3].nat : 0) || GR[r2].nat;
        tmp_len = merge_form ? len4 : len6 ;
    }
    if (pos6 + tmp_len > 64)
        tmp_len = 64 - pos6;

    if (merge_form)
        GR[r1] = GR[r3];
    else // zero_form
        GR[r1] = 0;

    GR[r1]{(pos6 + tmp_len - 1):pos6} = tmp_src{(tmp_len - 1):0};
    GR[r1].nat = tmp_nat;
}

```

Extract

Format: (qp) `extr r1 = r3, pos6, len6` signed_form I11
 (qp) `extr.u r1 = r3, pos6, len6` unsigned_form I11

Description: A field is extracted from GR r_3 , either zero extended or sign extended, and placed right-justified in GR r_1 . The field begins at the bit position given by the second operand and extends len_6 bits to the left. The bit position where the field begins is specified by the pos_6 immediate. The extracted field is sign extended in the `signed_form` or zero extended in the `unsigned_form`. The sign is taken from the most significant bit of the extracted field. If the specified field extends beyond the most significant bit of GR r_3 , the sign is taken from the most significant bit of GR r_3 . The immediate value len_6 can be any number in the range 1 to 64, and is encoded as len_6-1 in the instruction. The immediate value pos_6 can be any value in the range 0 to 63.

The operation of `extr t = r, 7, 50` is illustrated in Figure 6-6.




Figure 6-6. Extract Example

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    tmp_len = len6;

    if (pos6 + tmp_len > 64)
        tmp_len = 64 - pos6;

    if (unsigned_form)
        GR[r1] = zero_ext(shift_right_unsigned(GR[r3], pos6), tmp_len);
    else // signed_form
        GR[r1] = sign_ext(shift_right_unsigned(GR[r3], pos6), tmp_len);

    GR[r1].nat = GR[r3].nat;
}

```

Floating-Point Absolute Value

Format: (qp) fabs $f_1 = f_3$ pseudo-op of: (qp) fmerge.s $f_1 = f_0, f_3$

Description: The absolute value of the value in FR f_3 is computed and placed in FR f_1 .
If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation: See "Floating-Point Merge" on page 6-49.

Floating-Point Add

Format: $(qp) \text{ fadd}.pc.sf \ f_1 = f_3, f_2$ pseudo-op of: $(qp) \text{ fma}.pc.sf \ f_1 = f_3, f_1, f_2$

Description: FR f_3 and FR f_2 are added (computed to infinite precision), rounded to the precision indicated by pc (and possibly $FPSR.sf.pc$ and $FPSR.sf.wre$) using the rounding mode specified by $FPSR.sf.rc$, and placed in FR f_1 . If either FR f_3 or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

The mnemonic values for the opcode's pc are given in Table 6-17. The mnemonic values for sf are given in Table 6-18. For the encodings and interpretation of the status field's pc , wre , and rc , refer to Table 5-5 and Table 5-6 on page 5-5.

Table 6-17. Specified pc Mnemonic Values

pc Mnemonic	Precision Specified
.s	single
.d	double
<i>none</i>	dynamic (i.e., use pc value in status field)

Table 6-18. sf Mnemonic Values

sf Mnemonic	Status Field Accessed
.s0 or <i>none</i>	sf0
.s1	sf1
.s2	sf2
.s3	sf3

Operation: See "Floating-Point Multiply Add" on page 6-47.

Floating-Point Absolute Maximum

Format: $(qp) \text{ famax.sf } f_1 = f_2, f_3$ F8

Description: The operand with the larger absolute value is placed in FR f_1 . If the magnitude of FR f_2 equals the magnitude of FR f_3 , FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaN, FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as the `fcmp.lt` operation.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        fminmax_exception_fault_check(f2, f3, sf, &tmp_fp_env);
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_right = fp_reg_read(FR[f2]);
        tmp_left = fp_reg_read(FR[f3]);
        tmp_right.sign = FP_SIGN_POSITIVE;
        tmp_left.sign = FP_SIGN_POSITIVE;
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        FR[f1] = tmp_bool_res ? FR[f2] : FR[f3];

        fp_update_fpsr(sf, tmp_fp_env);
    }

    fp_update_psr(f1);
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Absolute Minimum

Format: (qp) $famin.sf f_1 = f_2, f_3$ F8

Description: The operand with the smaller absolute value is placed in FR f_1 . If the magnitude of FR f_2 equals the magnitude of FR f_3 , FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaN, FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as the `fcmp.lt` operation.

The mnemonic values for sf are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_isrcode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_isrcode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        fminmax_exception_fault_check(f2, f3, sf, &tmp_fp_env);
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_left = fp_reg_read(FR[f2]);
        tmp_right = fp_reg_read(FR[f3]);
        tmp_left.sign = FP_SIGN_POSITIVE;
        tmp_right.sign = FP_SIGN_POSITIVE;
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        FR[f1] = tmp_bool_res ? FR[f2] : FR[f3];

        fp_update_fpsr(sf, tmp_fp_env);
    }

    fp_update_psr(f1);
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Logical And

Format: (qp) fand $f_1 = f_2, f_3$ F9

Description: The bit-wise logical AND of the significand fields of FR f_2 and FR f_3 is computed. The resulting value is stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        FR[f1].significand = FR[f2].significand & FR[f3].significand;
        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;
    }
    fp_update_psr(f1);
}

```

FP Exceptions: None

Floating-Point And Complement

Format: (qp) fandcm $f_1 = f_2, f_3$ F9

Description: The bit-wise logical AND of the significand field of FR f_2 with the bit-wise complemented significand field of FR f_3 is computed. The resulting value is stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        FR[f1].significand = FR[f2].significand & ~FR[f3].significand;
        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;
    }
    fp_update_psr(f1);
}

```

FP Exceptions: None

Flush Cache

Format: $(qp) \text{ fc } r_3$

M28

Description: The cache line associated with the address specified by the value of GR r_3 is invalidated from all levels of the processor cache hierarchy. The invalidation is broadcast throughout the coherence domain. If, at any level of the cache hierarchy, the line is inconsistent with memory it is written to memory before invalidation.

The line size affected is at least 32-bytes (aligned on a 32-byte boundary). An implementation may flush a larger region.

This instruction follows data dependency rules; it is ordered with respect to preceding and following memory references to the same line. `fc` has data dependencies in the sense that any prior stores by this processor will be included in the data written back to memory. `fc` is an unordered operation, and is not affected by a memory fence (`mf`) instruction. It is ordered with respect to the `sync.i` instruction.

Operation:

```
if (PR[qp]) {
    itype = NON_ACCESS|FC|READ;
    if (GR[r3].nat)
        register_nat_consumption_fault(itype);
    tmp_paddr = tlb_translate_nonaccess(GR[r3], itype);
    mem_flush(tmp_paddr);
}
```

Floating-Point Check Flags

Format: (qp) fchkf.sf target₂₅

F14

Description: The flags in `FPSR.sf.flags` are compared with `FPSR.s0.flags` and `FPSR.traps`. If any flags set in `FPSR.sf.flags` correspond to `FPSR.traps` which are enabled, or if any flags set in `FPSR.sf.flags` are not set in `FPSR.s0.flags`, then a branch to `target25` is taken.

The `target25` operand, specifies a label to branch to. This is encoded in the instruction as a signed immediate displacement (`imm21`) between the target bundle and the bundle containing this instruction (`imm21 = target25 - IP >> 4`).

The mnemonic values for `sf` are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    switch (sf) {
        case 's0':
            tmp_flags = AR[FPSR].sf0.flags;
            break;
        case 's1':
            tmp_flags = AR[FPSR].sf1.flags;
            break;
        case 's2':
            tmp_flags = AR[FPSR].sf2.flags;
            break;
        case 's3':
            tmp_flags = AR[FPSR].sf3.flags;
            break;
    }
    if ((tmp_flags & ~AR[FPSR].traps) || (tmp_flags & ~AR[FPSR].sf0.flags)) {
        if (check_branch_implemented(FCHKF)) {
            taken_branch = 1;
            IP = IP + sign_ext((imm21 << 4), 25);
            if ((PSR.it && unimplemented_virtual_address(IP))
                || (!PSR.it && unimplemented_physical_address(IP)))
                unimplemented_instruction_address_trap(0, IP);
            if (PSR.tb)
                taken_branch_trap();
        } else
            speculation_fault(FCHKF, zero_ext(imm21, 21));
    }
}

```

FP Exceptions: None

Floating-Point Class

Format: $(qp) \text{ fclass.fcrel.fctype } p_1, p_2 = f_2, \text{fclass}_9$ F5

Description: The contents of FR f_2 are classified according to the fclass_9 completer as shown in Table 6-20. This produces a boolean result based on whether the contents of FR f_2 agrees with the floating-point number format specified by fclass_9 , as specified by the fcrel completer. This result is written to the two predicate register destinations, p_1 and p_2 . The result written to the destinations is determined by the compare type specified by fctype .

The allowed types are Normal (or *none*) and *unc*. See Table 6-21 on page 6-40. The assembly syntax allows the specification of membership or non-membership and the assembler swaps the target predicates to achieve the desired effect.

Table 6-19. Floating-point Class Relations

fcrel	Test Relation
m	FR f_2 agrees with the pattern specified by fclass_9 (is a member)
nm	FR f_2 does not agree with the pattern specified by fclass_9 (is not a member)

A number agrees with the pattern specified by fclass_9 if:

- the number is NaTVal and $\text{fclass}_9 \{8\}$ is 1, or
- the number is a quiet NaN and $\text{fclass}_9 \{7\}$ is 1, or
- the number is a signaling NaN and $\text{fclass}_9 \{6\}$ is 1, or
- the sign of the number agrees with the sign specified by one of the two low-order bits of fclass_9 , and the type of the number (disregarding the sign) agrees with the number-type specified by the next 4 bits of fclass_9 , as shown in Table 6-20.

Note: a fclass_9 of 0x1FF is equivalent to testing for any supported operand.

The class names used in Table 6-20 are defined in Table 5-2 on page 5-2.

Table 6-20. Floating-point Classes

fclass_9	Class	Mnemonic
Either these cases can be tested for		
0x0100	NaTVal	@nat
0x080	Quiet NaN	@qnan
0x040	Signaling NaN	@snan
or the OR of the following two cases		
0x001	Positive	@pos
0x002	Negative	@neg
AND'ed with OR of the following 4 cases		
0x004	Zero	@zero
0x008	Unnormalized	@unorm
0x010	Normalized	@norm
0x020	Infinity	@inf

```

Operation:  if (PR[qp]) {
            if (p1 == p2)
                illegal_operation_fault();

            if (tmp_israncode = fp_reg_disabled(f2, 0, 0, 0))
                disabled_fp_register_fault(tmp_israncode, 0);

            tmp_rel = ((fclass9{0} && !FR[f2].sign || fclass9{1} && FR[f2].sign)
                      && ((fclass9{2} && fp_is_zero(FR[f2])) ||
                          (fclass9{3} && fp_is_unorm(FR[f2])) ||
                          (fclass9{4} && fp_is_normal(FR[f2])) ||
                          (fclass9{5} && fp_is_inf(FR[f2]))
                      )
                      || (fclass9{6} && fp_is_snan(FR[f2]))
                      || (fclass9{7} && fp_is_qnan(FR[f2]))
                      || (fclass9{8} && fp_is_natval(FR[f2]));

            tmp_nat = fp_is_natval(FR[f2]) && (!fclass9{8});

            if (tmp_nat) {
                PR[p1] = 0;
                PR[p2] = 0;
            } else {
                PR[p1] = tmp_rel;
                PR[p2] = !tmp_rel;
            }
        } else {
            if (fctype == 'unc') {
                if (p1 == p2)
                    illegal_operation_fault();
                PR[p1] = 0;
                PR[p2] = 0;
            }
        }
    }

```

FP Exceptions: None

Floating-Point Clear Flags

Format: (qp) fclrf.sf

F13

Description: The status field's 6-bit flags field is reset to zero.
The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

Operation:

```
if (PR[qp]) {  
    fp_set_sf_flags(sf, 0);  
}
```

FP Exceptions: None

Floating-Point Compare

Format: $(qp) \text{ fcmp.frel.fctype.sf } p_1, p_2 = f_2, f_3$

F4

Description: The two source operands are compared for one of twelve relations specified by *frel*. This produces a boolean result which is 1 if the comparison condition is true, and 0 otherwise. This result is written to the two predicate register destinations, p_1 and p_2 . The way the result is written to the destinations is determined by the compare type specified by *fctype*. The allowed types are Normal (or *none*) and *unc*.

Table 6-21. Floating-point Comparison Types

<i>fctype</i>	PR[qp]==0		PR[qp]==1					
			result==0, No Source NaNVals		result==1, No Source NaNVals		One or More Source NaNVals	
	PR[p ₁]	PR[p ₂]	PR[p ₁]	PR[p ₂]	PR[p ₁]	PR[p ₂]	PR[p ₁]	PR[p ₂]
<i>none</i>			0	1	1	0	0	0
<i>unc</i>	0	0	0	1	1	0	0	0

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

The relations are defined for each of the comparison types in Table 6-22. Of the twelve relations, not all are directly implemented in hardware. Some are actually pseudo-ops. For these, the assembler simply switches the source operand specifiers and/or switches the predicate target specifiers and uses an implemented relation.

Table 6-22. Floating-point Comparison Relations

<i>frel</i>	<i>frel</i> Completer Unabbreviated	Relation	Pseudo-op of	Quiet NaN as Operand Signals Invalid
eq	equal	$f_2 == f_3$		No
lt	less than	$f_2 < f_3$		Yes
le	less than or equal	$f_2 \leq f_3$		Yes
gt	greater than	$f_2 > f_3$	lt $f_2 \leftrightarrow f_3$	Yes
ge	greater than or equal	$f_2 \geq f_3$	le $f_2 \leftrightarrow f_3$	Yes
unord	unordered	$f_2 ? f_3$		No
neq	not equal	$!(f_2 == f_3)$	eq $p_1 \leftrightarrow p_2$	No
nlt	not less than	$!(f_2 < f_3)$	lt $p_1 \leftrightarrow p_2$	Yes
nle	not less than or equal	$!(f_2 \leq f_3)$	le $p_1 \leftrightarrow p_2$	Yes
ngt	not greater than	$!(f_2 > f_3)$	lt $f_2 \leftrightarrow f_3$ $p_1 \leftrightarrow p_2$	Yes
nge	not greater than or equal	$!(f_2 \geq f_3)$	le $f_2 \leftrightarrow f_3$ $p_1 \leftrightarrow p_2$	Yes
ord	ordered	$!(f_2 ? f_3)$	unord $p_1 \leftrightarrow p_2$	No

```

Operation:  if (PR[qp]) {
            if (p1 == p2)
                illegal_operation_fault();

            if (tmp_israncode = fp_reg_disabled(f2, f3, 0, 0))
                disabled_fp_register_fault(tmp_israncode, 0);

            if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
                PR[p1] = 0;
                PR[p2] = 0;
            } else {
                fcmp_exception_fault_check(f2, f3, frel, sf, &tmp_fp_env);
                if (fp_raise_fault(tmp_fp_env))
                    fp_exception_fault(fp_decode_fault(tmp_fp_env));

                tmp_fr2 = fp_reg_read(FR[f2]);
                tmp_fr3 = fp_reg_read(FR[f3]);

                if (frel == 'eq') tmp_rel = fp_equal(tmp_fr2, tmp_fr3);
                else if (frel == 'lt') tmp_rel = fp_less_than(tmp_fr2, tmp_fr3);
                else if (frel == 'le') tmp_rel = fp_lesser_or_equal(tmp_fr2, tmp_fr3);
                else if (frel == 'gt') tmp_rel = fp_less_than(tmp_fr3, tmp_fr2);
                else if (frel == 'ge') tmp_rel = fp_lesser_or_equal(tmp_fr3, tmp_fr2);
                else if (frel == 'unord') tmp_rel = fp_unordered(tmp_fr2, tmp_fr3);
                else if (frel == 'neq') tmp_rel = !fp_equal(tmp_fr2, tmp_fr3);
                else if (frel == 'nlt') tmp_rel = !fp_less_than(tmp_fr2, tmp_fr3);
                else if (frel == 'nle') tmp_rel = !fp_lesser_or_equal(tmp_fr2, tmp_fr3);
                else if (frel == 'ngt') tmp_rel = !fp_less_than(tmp_fr3, tmp_fr2);
                else if (frel == 'nge') tmp_rel = !fp_lesser_or_equal(tmp_fr3, tmp_fr2);
                else tmp_rel = !fp_unordered(tmp_fr2, tmp_fr3); // 'ord'

                PR[p1] = tmp_rel;
                PR[p2] = !tmp_rel;

                fp_update_fpsr(sf, tmp_fp_env);
            }
        } else {
            if (fctype == 'unc') {
                if (p1 == p2)
                    illegal_operation_fault();
                PR[p1] = 0;
                PR[p2] = 0;
            }
        }
    }

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Convert Signed Integer to Floating-point

Format: (qp) fcvt.xf $f_1 = f_2$

F11

Description: The 64-bit significand of FR f_2 is treated as a signed integer and its register file precision floating-point representation is placed in FR f_1 .

If FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation is always exact and is unaffected by the rounding mode.

Operation:

```

if (PR[qp]) {
    fp_check_target_register( $f_1$ );
    if (tmp_israncode = fp_reg_disabled( $f_1$ ,  $f_2$ , 0, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[ $f_2$ ])) {
        FR[ $f_1$ ] = NATVAL;
    } else {
        tmp_res = FR[ $f_2$ ];
        if (tmp_res.significand{63}) {
            tmp_res.significand = (~tmp_res.significand) + 1;
            tmp_res.sign = 1;
        } else
            tmp_res.sign = 0;

        tmp_res.exponent = FP_INTEGER_EXP;
        tmp_res = fp_normalize(tmp_res);

        FR[ $f_1$ ].significand = tmp_res.significand;
        FR[ $f_1$ ].exponent = tmp_res.exponent;
        FR[ $f_1$ ].sign = tmp_res.sign;
    }
    fp_update_psr( $f_1$ );
}

```

FP Exceptions: None

Convert Unsigned Integer to Floating-point

Format: (qp) fcvt.xuf.pc.sf $f_1 = f_3$ (unsigned form) pseudo-op of: (qp) fma.pc.sf $f_1 = f_3, f_1, f_0$

Description: FR f_3 is multiplied with FR 1, rounded to the precision indicated by pc (and possibly $FPSR.sf.pc$ and $FPSR.sf.wre$) using the rounding mode specified by $FPSR.sf.rc$, and placed in FR f_1 .

Note: Multiplying FR f_3 with FR 1 (a 1.0) normalizes the canonical representation of an integer in the floating-point register file producing a normal floating-point value.

If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

The mnemonic values for the opcode's pc are given in Table 6-17 on page 6-30. The mnemonic values for sf are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field's pc , wre , and rc , refer to Table 5-5 and Table 5-6 on page 5-5.

Operation: See "Floating-Point Multiply Add" on page 6-47

Fetch And Add Immediate

Format: (*qp*) *fetchadd4.sem.ldhint* *r₁* = [*r₃*], *inc₃* four_byte_form M17
 (*qp*) *fetchadd8.sem.ldhint* *r₁* = [*r₃*], *inc₃* eight_byte_form M17

Description: A value consisting of four or eight bytes is read from memory starting at the address specified by the value in GR *r₃*. The value is zero extended and added to the sign-extended immediate value specified by *inc₃*. The values that may be specified by *inc₃* are: -16, -8, -4, -1, 1, 4, 8, 16. The least significant four or eight bytes of the sum are then written to memory starting at the address specified by the value in GR *r₃*. The zero-extended value read from memory is placed in GR *r₁* and the NaT bit corresponding to GR *r₁* is cleared.

The *sem* completer specifies the type of semaphore operation. These operations are described in Table 6-23.

Table 6-23. Fetch and Add Semaphore Types

<i>sem</i> Completer	Ordering Semantics	Semaphore Operation
acq	Acquire	The memory read/write is made visible prior to all subsequent data memory accesses.
rel	Release	The memory read/write is made visible after all previous data memory accesses.

The memory read and write are guaranteed to be atomic.

If the address specified by the value in GR *r₃* is not naturally aligned to the size of the value being accessed in memory, an Unaligned Data Reference fault is taken independent of the state of the User Mask alignment checking bit, UM.ac (PSR.ac in the Processor Status Register).

Both read and write access privileges for the referenced page are required. The write access privilege check is performed whether or not the memory write is performed.

The value of the *ldhint* completer specifies the locality of the memory access. The values of the *ldhint* completer are given in Table 6-28 on page 6-102. Locality hints do not affect program functionality and may be ignored by the implementation.

Operation:

```

if (PR[qp]) {
    check_target_register(r1, SEMAPHORE);

    if (GR[r3].nat)
        register_nat_consumption_fault(SEMAPHORE);

    size = four_byte_form ? 4 : 8;

    paddr = tlb_translate(GR[r3], size, SEMAPHORE, PSR.cpl, &mattr, &tmp_unused);
    if (!ma_supports_fetchadd(mattr))
        unsupported_data_reference_fault(SEMAPHORE, GR[r3]);

    if (sem == 'acq')
        val = mem_xchg_add(inc3, paddr, size, UM.be, mattr, ACQUIRE, ldhint);
    else // 'rel'
        val = mem_xchg_add(inc3, paddr, size, UM.be, mattr, RELEASE, ldhint);

    alat_inval_multiple_entries(paddr, size);

    GR[r1] = zero_ext(val, size * 8);
    GR[r1].nat = 0;
}

```

Flush Register Stack

Format: flushrs M25

Description: All stacked general registers in the dirty partition of the register stack are written to the backing store before execution continues. The dirty partition contains registers from previous procedure frames that have not yet been saved to the backing store.

After this instruction completes execution AR[BSPSTORE] is equal to AR[BSP].

This instruction must be the first instruction in an instruction group. Otherwise, the results are undefined. This instruction cannot be predicated.

Operation:

```
while (AR[BSPSTORE] != AR[BSP]) {  
    rse_store(MANDATORY);           // increments AR[BSPSTORE]  
    deliver_unmasked_pending_external_interrupt();  
}
```


Floating-Point Multiply Add

Format: $(qp) \text{ fma.pc.sf } f_1 = f_3, f_4, f_2$ F1

Description: The product of FR f_3 and FR f_4 is computed to infinite precision and then FR f_2 is added to this product, again in infinite precision. The resulting value is then rounded to the precision indicated by pc (and possibly FPSR.sf.pc and FPSR.sf.wre) using the rounding mode specified by FPSR.sf.rc . The rounded result is placed in FR f_1 .

If any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

If f_2 is f_0 , an IEEE multiply operation is performed instead of a multiply and add. See “Floating-Point Multiply” on page 6-54.

The mnemonic values for the opcode’s pc are given in Table 6-17 on page 6-30. The mnemonic values for sf are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field’s pc , wre , and rc , refer to Table 5-5 and Table 5-6 on page 5-5.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, f4))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3]) || fp_is_natval(FR[f4])) {
        FR[f1] = NATVAL;
        fp_update_psr(f1);
    } else {
        tmp_default_result = fma_exception_fault_check(f2, f3, f4,
                                                    pc, sf, &tmp_fp_env);

        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        if (fp_is_nan_or_inf(tmp_default_result)) {
            FR[f1] = tmp_default_result;
        } else {
            tmp_res = fp_mul(fp_reg_read(FR[f3]), fp_reg_read(FR[f4]));
            if (f2 != 0)
                tmp_res = fp_add(tmp_res, fp_reg_read(FR[f2]), tmp_fp_env);
            FR[f1] = fp_ieee_round(tmp_res, &tmp_fp_env);
        }

        fp_update_fpsr(sf, tmp_fp_env);
        fp_update_psr(f1);
        if (fp_raise_traps(tmp_fp_env))
            fp_exception_trap(fp_decode_trap(tmp_fp_env));
    }
}

```

FP Exceptions: Invalid Operation (V)	Overflow (O)
Denormal/Unnormal Operand (D)	Inexact (I)
Software Assist (SWA) fault	Software Assist (SWA) trap
Underflow (U)	

Floating-Point Maximum

Format: $(qp) \text{ fmax}.sf \ f_1 = f_2, f_3$

F8

Description: The operand with the larger value is placed in FR f_1 . If FR f_2 equals FR f_3 , FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaN, FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as the fcmp.lt operation.

The mnemonic values for sf are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_isrkode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_isrkode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        fminmax_exception_fault_check(f2, f3, sf, &tmp_fp_env);
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_bool_res = fp_less_than(fp_reg_read(FR[f3]), fp_reg_read(FR[f2]));
        FR[f1] = (tmp_bool_res ? FR[f2] : FR[f3]);

        fp_update_fpsr(sf, tmp_fp_env);
    }
    fp_update_psr(f1);
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Merge

Format:

<i>(qp)</i> fmerge.ns $f_1 = f_2, f_3$	neg_sign_form	F9
<i>(qp)</i> fmerge.s $f_1 = f_2, f_3$	sign_form	F9
<i>(qp)</i> fmerge.se $f_1 = f_2, f_3$	sign_exp_form	F9

Description: Sign, exponent and significand fields are extracted from FR f_2 and FR f_3 , combined, and the result is placed in FR f_1 .

For the neg_sign_form, the sign of FR f_2 is negated and concatenated with the exponent and the significand of FR f_3 . This form can be used to negate a floating-point number by using the same register for FR f_2 and FR f_3 .

For the sign_form, the sign of FR f_2 is concatenated with the exponent and the significand of FR f_3 .

For the sign_exp_form, the sign and exponent of FR f_2 is concatenated with the significand of FR f_3 .

For all forms, if either FR f_2 or FR f_3 is a NaNVal, FR f_1 is set to NaNVal instead of the computed result.




Figure 6-7. Floating-point Merge Negative Sign Operation

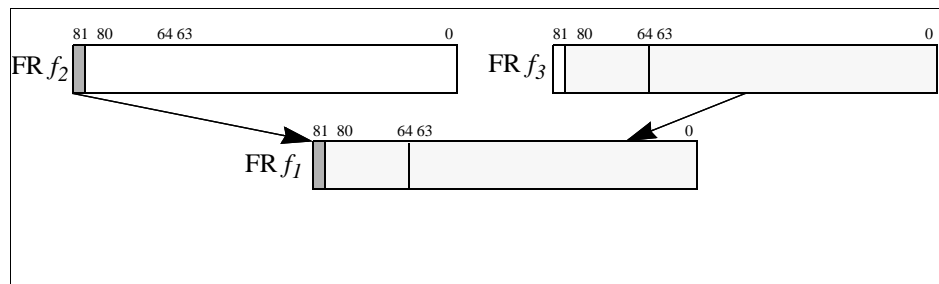


Figure 6-8. Floating-point Merge Sign Operation




Figure 6-9. Floating-point Merge Sign and Exponent Operation

```

Operation:  if (PR[qp]) {
            fp_check_target_register(f1);
            if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
                disabled_fp_register_fault(tmp_israncode, 0);

            if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
                FR[f1] = NATVAL;
            } else {
                FR[f1].significand = FR[f3].significand;
                if (neg_sign_form) {
                    FR[f1].exponent = FR[f3].exponent;
                    FR[f1].sign = !FR[f2].sign;
                } else if (sign_form) {
                    FR[f1].exponent = FR[f3].exponent;
                    FR[f1].sign = FR[f2].sign;
                } else {
                    FR[f1].exponent = FR[f2].exponent;
                    FR[f1].sign = FR[f2].sign;
                }
            }

            fp_update_psr(f1);
        }

```

FP Exceptions: None

Floating-Point Minimum

Format: $(qp) \text{ fmin.sf } f_1 = f_2, f_3$

F8

Description: The operand with the smaller value is placed in FR f_1 . If FR f_2 equals FR f_3 , FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaN, FR f_1 gets FR f_3 .

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as the `fcmp.lt` operation.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        fminmax_exception_fault_check(f2, f3, sf, &tmp_fp_env);
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_bool_res = fp_less_than(fp_reg_read(FR[f2]), fp_reg_read(FR[f3]));
        FR[f1] = tmp_bool_res ? FR[f2] : FR[f3];

        fp_update_fpsr(sf, tmp_fp_env);
    }
    fp_update_psr(f1);
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Parallel Mix

Format:	(qp) $fmix.l\ f_1 = f_2, f_3$	<code>mix_l_form</code>	F9
	(qp) $fmix.r\ f_1 = f_2, f_3$	<code>mix_r_form</code>	F9
	(qp) $fmix.lr\ f_1 = f_2, f_3$	<code>mix_lr_form</code>	F9

Description: For the `mix_l_form` (`mix_r_form`), the left (right) single precision value in FR f_2 is concatenated with the left (right) single precision value in FR f_3 . For the `mix_lr_form`, the left single precision value in FR f_2 is concatenated with the right single precision value in FR f_3 .

For all forms, the exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

For all forms, if either FR f_2 or FR f_3 is a NaNVal, FR f_1 is set to NaNVal instead of the computed result.




Figure 6-10. Floating-point Mix Left




Figure 6-11. Floating-point Mix Right




Figure 6-12. Floating-point Mix Left-Right

```

Operation:    if (PR[qp]) {
                  fp_check_target_register(f1);
                  if (tmp_isrkode = fp_reg_disabled(f1, f2, f3, 0))
                      disabled_fp_register_fault(tmp_isrkode, 0);

                  if (fp_is_natval(FR[f2] || fp_is_natval(FR[f3])) {
                      FR[f1] = NATVAL;
                  } else {
                      if (mix_l_form) {
                          tmp_res_hi = FR[f2].significand{63:32};
                          tmp_res_lo = FR[f3].significand{63:32};
                      } else if (mix_r_form) {
                          tmp_res_hi = FR[f2].significand{31:0};
                          tmp_res_lo = FR[f3].significand{31:0};
                      } else {
                          tmp_res_hi = FR[f2].significand{63:32};
                          tmp_res_lo = FR[f3].significand{31:0};
                      } // mix_lr_form
                      FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                      FR[f1].exponent = FP_INTEGER_EXP;
                      FR[f1].sign = FP_SIGN_POSITIVE;
                  }

                  fp_update_psr(f1);
                }

```

FP Exceptions: None

Floating-Point Multiply

Format: $(qp) \text{ fmpy}.pc.sf \ f_1 = f_3, f_4$ pseudo-op of: $(qp) \text{ fma}.pc.sf \ f_1 = f_3, f_4, f_0$

Description: The product FR f_3 and FR f_4 is computed to infinite precision. The resulting value is then rounded to the precision indicated by pc (and possibly $FPSR.sf.pc$ and $FPSR.sf.wre$) using the rounding mode specified by $FPSR.sf.rc$. The rounded result is placed in FR f_1 .

If either FR f_3 or FR f_4 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

The mnemonic values for the opcode's pc are given in Table 6-17 on page 6-30. The mnemonic values for sf are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field's pc , wre , and rc , refer to Table 5-5 and Table 5-6 on page 5-5.

Operation: See "Floating-Point Multiply Add" on page 6-47.

Floating-Point Multiply Subtract

Format: (qp) fms.pc.sf $f_1 = f_3, f_4, f_2$ F1

Description: The product of FR f_3 and FR f_4 is computed to infinite precision and then FR f_2 is subtracted from this product, again in infinite precision. The resulting value is then rounded to the precision indicated by *pc* (and possibly *FPSR.sf.pc* and *FPSR.sf.wre*) using the rounding mode specified by *FPSR.sf.rc*. The rounded result is placed in FR f_1 .

If any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, a NaTVal is placed in FR f_1 instead of the computed result.

If f_2 is f0, an IEEE multiply operation is performed instead of a multiply and subtract.

The mnemonic values for the opcode's *pc* are given in Table 6-17 on page 6-30. The mnemonic values for *sf* are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field's *pc*, *wre*, and *rc*, refer to Table 5-5 and Table 5-6 on page 5-5.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, f4))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3]) || fp_is_natval(FR[f4])) {
        FR[f1] = NATVAL;
        fp_update_psr(f1);
    } else {
        tmp_default_result = fms_fnma_exception_fault_check(f2, f3, f4,
                                                            pc, sf, &tmp_fp_env);

        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        if (fp_is_nan_or_inf(tmp_default_result)) {
            FR[f1] = tmp_default_result;
        } else {
            tmp_res = fp_mul(fp_reg_read(FR[f3]), fp_reg_read(FR[f4]));
            tmp_fr2 = fp_reg_read(FR[f2]);
            tmp_fr2.sign = !tmp_fr2.sign;
            if (f2 != 0)
                tmp_res = fp_add(tmp_res, tmp_fr2, tmp_fp_env);
            FR[f1] = fp_ieee_round(tmp_res, &tmp_fp_env);
        }

        fp_update_fpsr(sf, tmp_fp_env);
        fp_update_psr(f1);
        if (fp_raise_traps(tmp_fp_env))
            fp_exception_trap(fp_decode_trap(tmp_fp_env));
    }
}

```

FP Exceptions:	Invalid Operation (V)	Overflow (O)
	Denormal/Unnormal Operand (D)	Inexact (I)
	Software Assist (SWA) fault	Software Assist (SWA) trap
	Underflow (U)	

Floating-Point Negate

Format: (qp) fneg $f_1 = f_3$ pseudo-op of: (qp) fmerge.ns $f_1 = f_3, f_3$

Description: The value in FR f_3 is negated and placed in FR f_1 .
If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation: See "Floating-Point Merge" on page 6-49.

Floating-Point Negate Absolute Value

Format: (qp) fnegabs $f_1 = f_3$ pseudo-op of: (qp) fmerge.ns $f_1 = f_0, f_3$

Description: The absolute value of the value in FR f_3 is computed, negated, and placed in FR f_1 .
If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation: See "Floating-Point Merge" on page 6-49.

Floating-Point Negative Multiply Add

Format: (qp) fnma.pc.sf $f_1 = f_3, f_4, f_2$ F1

Description: The product of FR f_3 and FR f_4 is computed to infinite precision, negated, and then FR f_2 is added to this product, again in infinite precision. The resulting value is then rounded to the precision indicated by pc (and possibly FPSR.sf.pc and FPSR.sf.wre) using the rounding mode specified by FPSR.sf.rc. The rounded result is placed in FR f_1 .

If any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

If f_2 is f0, an IEEE multiply operation is performed, followed by negation of the product.

The mnemonic values for the opcode's pc are given in Table 6-17 on page 6-30. The mnemonic values for sf are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field's pc , wre , and rc , refer to Table 5-5 and Table 5-6 on page 5-5.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, f4))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3]) || fp_is_natval(FR[f4])) {
        FR[f1] = NATVAL;
        fp_update_psr(f1);
    } else {
        tmp_default_result = fms_fnma_exception_fault_check(f2, f3, f4,
                                                            pc, sf, &tmp_fp_env);

        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        if (fp_is_nan_or_inf(tmp_default_result)) {
            FR[f1] = tmp_default_result;
        } else {
            tmp_res = fp_mul(fp_reg_read(FR[f3]), fp_reg_read(FR[f4]));
            tmp_res.sign = !tmp_res.sign;
            if (f2 != 0)
                tmp_res = fp_add(tmp_res, fp_reg_read(FR[f2]), tmp_fp_env);
            FR[f1] = fp_ieee_round(tmp_res, &tmp_fp_env);
        }

        fp_update_fpsr(sf, tmp_fp_env);
        fp_update_psr(f1);
        if (fp_raise_traps(tmp_fp_env))
            fp_exception_trap(fp_decode_trap(tmp_fp_env));
    }
}

```

FP Exceptions: Invalid Operation (V)	Overflow (O)
Denormal/Unnormal Operand (D)	Inexact (I)
Software Assist (SWA) fault	Software Assist (SWA) trap
Underflow (U)	

Floating-Point Negative Multiply

Format: (qp) fnmpy.pc.sf $f_1 = f_3, f_4$ pseudo-op of: (qp) fnma.pc.sf $f_1 = f_3, f_4, f_0$

Description: The product FR f_3 and FR f_4 is computed to infinite precision and then negated. The resulting value is then rounded to the precision indicated by pc (and possibly FPSR.sf.pc and FPSR.sf.wre) using the rounding mode specified by FPSR.sf.rc. The rounded result is placed in FR f_1 .

If either FR f_3 or FR f_4 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

The mnemonic values for the opcode's pc are given in Table 6-17 on page 6-30. The mnemonic values for sf are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field's pc , wre , and rc , refer to Table 5-5 and Table 5-6 on page 5-5.

Operation: See "Floating-Point Negative Multiply Add" on page 6-58.

Floating-Point Normalize

Format: $(qp) \text{ fnorm.pc.sf } f_1 = f_3$ pseudo-op of: $(qp) \text{ fma.pc.sf } f_1 = f_3, f_1, f_0$

Description: FR f_3 is normalized and rounded to the precision indicated by pc (and possibly FPSR.sf.pc and FPSR.sf.wre) using the rounding mode specified by FPSR.sf.rc , and placed in FR f_1 .

If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

The mnemonic values for the opcode's pc are given in Table 6-17 on page 6-30. The mnemonic values for sf are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field's pc , wre , and rc , refer to Table 5-5 and Table 5-6 on page 5-5.

Operation: See "Floating-Point Multiply Add" on page 6-47.

Floating-Point Logical Or

Format: (qp) for $f_1 = f_2, f_3$ F9

Description: The bit-wise logical OR of the significand fields of FR f_2 and FR f_3 is computed. The resulting value is stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        FR[f1].significand = FR[f2].significand | FR[f3].significand;
        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;
    }

    fp_update_psr(f1);
}

```

FP Exceptions: None

Floating-Point Parallel Absolute Value

Format: (qp) fpabs $f_1 = f_3$ pseudo-op of: (qp) fpmerge.s $f_1 = f_0, f_3$

Description: The absolute values of the pair of single precision values in the significand field of FR f_3 are computed and stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation: See "Floating-Point Parallel Merge" on page 6-73.

Floating-Point Pack

Format: (*qp*) fpack $f_1 = f_2, f_3$ pack_form F9

Description: The register format numbers in FR f_2 and FR f_3 are converted to single precision memory format. These two single precision numbers are concatenated and stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.




Figure 6-13. Floating-point Pack

```

Operation:  if (PR[qp]) {
                fp_check_target_register( $f_1$ );
                if (tmp_israncode = fp_reg_disabled( $f_1, f_2, f_3, 0$ ))
                    disabled_fp_register_fault(tmp_israncode, 0);

                if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ])) {
                    FR[ $f_1$ ] = NATVAL;
                } else {
                    tmp_res_hi = fp_single(FR[ $f_2$ ]);
                    tmp_res_lo = fp_single(FR[ $f_3$ ]);

                    FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                    FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
                    FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;
                }
                fp_update_psr( $f_1$ );
            }
  
```

FP Exceptions: None

Floating-Point Parallel Absolute Maximum

Format: (qp) fpamax.sf $f_1 = f_2, f_3$ F8

Description: The paired single precision values in the significands of FR f_2 and FR f_3 are compared. The operands with the larger absolute value are returned in the significand field of FR f_1 .

If the magnitude of high (low) FR f_3 is less than the magnitude of high (low) FR f_2 , high (low) FR f_1 gets high (low) FR f_2 . Otherwise high (low) FR f_1 gets high (low) FR f_3 .

If high (low) FR f_2 or high (low) FR f_3 is a NaN, and neither FR f_2 or FR f_3 is a NaTVal, high (low) FR f_1 gets high (low) FR f_3 .

The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as for the fpcmp.lt operation.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register( $f_1$ );
    if (tmp_israncode = fp_reg_disabled( $f_1, f_2, f_3, 0$ ))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ])) {
        FR[ $f_1$ ] = NATVAL;
    } else {
        fpminmax_exception_fault_check( $f_2, f_3, sf, &tmp_fp_env$ );
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_fr2 = tmp_right = fp_reg_read_hi( $f_2$ );
        tmp_fr3 = tmp_left = fp_reg_read_hi( $f_3$ );
        tmp_right.sign = FP_SIGN_POSITIVE;
        tmp_left.sign = FP_SIGN_POSITIVE;
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_hi = fp_single(tmp_bool_res ? tmp_fr2: tmp_fr3);

        tmp_fr2 = tmp_right = fp_reg_read_lo( $f_2$ );
        tmp_fr3 = tmp_left = fp_reg_read_lo( $f_3$ );
        tmp_right.sign = FP_SIGN_POSITIVE;
        tmp_left.sign = FP_SIGN_POSITIVE;
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_lo = fp_single(tmp_bool_res ? tmp_fr2: tmp_fr3);

        FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
        FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
        FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;

        fp_update_fpsr(sf, tmp_fp_env);
    }
    fp_update_psr( $f_1$ );
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Parallel Absolute Minimum

Format: (qp) fpamin.sf $f_1 = f_2, f_3$ F8

Description: The paired single precision values in the significands of FR f_2 or FR f_3 are compared. The operands with the smaller absolute value is returned in the significand of FR f_1 .

If the magnitude of high (low) FR f_2 is less than the magnitude of high (low) FR f_3 , high (low) FR f_1 gets high (low) FR f_2 . Otherwise high (low) FR f_1 gets high (low) FR f_3 .

If high (low) FR f_2 or high (low) FR f_3 is a NaN, and neither FR f_2 or FR f_3 is a NaTVal, high (low) FR f_1 gets high (low) FR f_3 .

The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as for the fpcmp.lt operation.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register( $f_1$ );
    if (tmp_israncode = fp_reg_disabled( $f_1, f_2, f_3, 0$ ))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ])) {
        FR[ $f_1$ ] = NATVAL;
    } else {
        fpminmax_exception_fault_check( $f_2, f_3, sf, &tmp_fp_env$ );
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_fr2 = tmp_left = fp_reg_read_hi( $f_2$ );
        tmp_fr3 = tmp_right = fp_reg_read_hi( $f_3$ );
        tmp_left.sign = FP_SIGN_POSITIVE;
        tmp_right.sign = FP_SIGN_POSITIVE;
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_hi = fp_single(tmp_bool_res ? tmp_fr2 : tmp_fr3);

        tmp_fr2 = tmp_left = fp_reg_read_lo( $f_2$ );
        tmp_fr3 = tmp_right = fp_reg_read_lo( $f_3$ );
        tmp_left.sign = FP_SIGN_POSITIVE;
        tmp_right.sign = FP_SIGN_POSITIVE;
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_lo = fp_single(tmp_bool_res ? tmp_fr2 : tmp_fr3);

        FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
        FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
        FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;

        fp_update_fpsr(sf, tmp_fp_env);
    }
    fp_update_psr( $f_1$ );
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Parallel Compare

Format: $(qp) \text{ fpcmp.frel.sf } f1 = f_2, f_3$

F8

Description: The two pairs of single precision source operands in the significand fields of FR f_2 and FR f_3 are compared for one of twelve relations specified by *frel*. This produces a boolean result which is a mask of 32 1's if the comparison condition is true, and a mask of 32 0's otherwise. This result is written to a pair of 32-bit integers in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

Table 6-24. Floating-point Parallel Comparison Results

PR[qp]==0	PR[qp]==1		
	result==false, No Source NaNVals	result==true, No Source NaNVals	One or More Source NaNVal's
unchanged	0...0	1...1	NaNVal

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

The relations are defined for each of the comparison types in Table 6-24. Of the twelve relations, not all are directly implemented in hardware. Some are actually pseudo-ops. For these, the assembler simply switches the source operand specifiers and/or switches the predicate type specifiers and uses an implemented relation.

If either FR f_2 or FR f_3 is a NaNVal, FR f_1 is set to NaNVal instead of the computed result.

Table 6-25. Floating-point Parallel Comparison Relations

<i>frel</i>	<i>frel</i> Completer Unabbreviated	Relation	Pseudo-op of	Quiet NaN as Operand Signals Invalid
eq	equal	$f_2 == f_3$		No
lt	less than	$f_2 < f_3$		Yes
le	less than or equal	$f_2 <= f_3$		Yes
gt	greater than	$f_2 > f_3$	lt $f_2 \leftrightarrow f_3$	Yes
ge	greater than or equal	$f_2 >= f_3$	le $f_2 \leftrightarrow f_3$	Yes
unord	unordered	$f_2 ? f_3$		No
neq	not equal	$!(f_2 == f_3)$		No
nlt	not less than	$!(f_2 < f_3)$		Yes
nle	not less than or equal	$!(f_2 <= f_3)$		Yes
ngt	not greater than	$!(f_2 > f_3)$	nlt $f_2 \leftrightarrow f_3$	Yes
nge	not greater than or equal	$!(f_2 >= f_3)$	nle $f_2 \leftrightarrow f_3$	Yes
ord	ordered	$!(f_2 ? f_3)$		No

```

Operation:  if (PR[qp]) {
            fp_check_target_register( $f_1$ );
            if (tmp_israncode = fp_reg_disabled( $f_1$ ,  $f_2$ ,  $f_3$ , 0))
                disabled_fp_register_fault(tmp_israncode, 0);

            if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ])) {
                FR[ $f_1$ ] = NATVAL;
            } else {
                fpcmp_exception_fault_check( $f_2$ ,  $f_3$ ,  $frel$ ,  $sf$ , &tmp_fp_env);

                if (fp_raise_fault(tmp_fp_env))
                    fp_exception_fault(fp_decode_fault(tmp_fp_env));

                tmp_fr2 = fp_reg_read_hi( $f_2$ );
                tmp_fr3 = fp_reg_read_hi( $f_3$ );

                if      ( $frel$  == 'eq')  tmp_rel = fp_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'lt')  tmp_rel = fp_less_than(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'le')  tmp_rel = fp_lesser_or_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'gt')  tmp_rel = fp_less_than(tmp_fr3, tmp_fr2);
                else if ( $frel$  == 'ge')  tmp_rel = fp_lesser_or_equal(tmp_fr3, tmp_fr2);
                else if ( $frel$  == 'unord') tmp_rel = fp_unordered(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'neq') tmp_rel = !fp_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'nlt') tmp_rel = !fp_less_than(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'nle') tmp_rel = !fp_lesser_or_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'ngt') tmp_rel = !fp_less_than(tmp_fr3, tmp_fr2);
                else if ( $frel$  == 'nge') tmp_rel = !fp_lesser_or_equal(tmp_fr3, tmp_fr2);
                else                      tmp_rel = !fp_unordered(tmp_fr2, tmp_fr3); // 'ord'

                tmp_res_hi = (tmp_rel ? 0xFFFFFFFF : 0x00000000);

                tmp_fr2 = fp_reg_read_lo( $f_2$ );
                tmp_fr3 = fp_reg_read_lo( $f_3$ );

                if      ( $frel$  == 'eq')  tmp_rel = fp_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'lt')  tmp_rel = fp_less_than(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'le')  tmp_rel = fp_lesser_or_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'gt')  tmp_rel = fp_less_than(tmp_fr3, tmp_fr2);
                else if ( $frel$  == 'ge')  tmp_rel = fp_lesser_or_equal(tmp_fr3, tmp_fr2);
                else if ( $frel$  == 'unord') tmp_rel = fp_unordered(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'neq') tmp_rel = !fp_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'nlt') tmp_rel = !fp_less_than(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'nle') tmp_rel = !fp_lesser_or_equal(tmp_fr2, tmp_fr3);
                else if ( $frel$  == 'ngt') tmp_rel = !fp_less_than(tmp_fr3, tmp_fr2);
                else if ( $frel$  == 'nge') tmp_rel = !fp_lesser_or_equal(tmp_fr3, tmp_fr2);
                else                      tmp_rel = !fp_unordered(tmp_fr2, tmp_fr3); // 'ord'

                tmp_res_lo = (tmp_rel ? 0xFFFFFFFF : 0x00000000);

                FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
                FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;

                fp_update_fpsr( $sf$ , tmp_fp_env);
            }
            fp_update_psr( $f_1$ );
        }

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Convert Parallel Floating-Point to Integer

Format:	$(qp) \text{ fpcvt.fx.sf } f_1 = f_2$	signed_form	F10
	$(qp) \text{ fpcvt.fx.trunc.sf } f_1 = f_2$	signed_form, trunc_form	F10
	$(qp) \text{ fpcvt.fxu.sf } f_1 = f_2$	unsigned_form	F10
	$(qp) \text{ fpcvt.fxu.trunc.sf } f_1 = f_2$	unsigned_form, trunc_form	F10

Description: The pair of single precision values in the significand field of FR f_2 is converted to a pair of 32-bit signed integers (signed_form) or unsigned integers (unsigned_form) using either the rounding mode specified in the FPSR.sf.rc, or using Round-to-Zero if the trunc_form of the instruction is used. The result is written as a pair of 32-bit integers into the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0). If the result of the conversion doesn't fit in a 32-bit integer the 32-bit integer indefinite value 0x80000000 is used as the result if the IEEE Invalid Operation Floating-Point Exception fault is disabled.

If FR f_2 is a NatVal, FR f_1 is set to NatVal instead of the computed result.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

```

Operation:    if (PR[qp]) {
              fp_check_target_register(f1);
              if (tmp_israncode = fp_reg_disabled(f1, f2, 0, 0))
                  disabled_fp_register_fault(tmp_israncode, 0);

              if (fp_is_natval(FR[f2])) {
                  FR[f1] = NATVAL;
                  fp_update_psr(f1);
              } else {
                  tmp_default_result_pair = fpcvt_exception_fault_check(f2, sf,
                                                                 signed_form, trunc_form, &tmp_fp_env);

                  if (fp_raise_fault(tmp_fp_env))
                      fp_exception_fault(fp_decode_fault(tmp_fp_env));

                  if (fp_is_nan(tmp_default_result_pair.hi)) {
                      tmp_res_hi = INTEGER_INDEFINITE_32_BIT;
                  } else {
                      tmp_res = fp_ieee_rnd_to_int_sp(fp_reg_read_hi(f2), HIGH, &tmp_fp_env);
                      if (tmp_res.exponent)
                          tmp_res.significand = fp_U64_rsh(
                              tmp_res.significand, (FP_INTEGER_EXP - tmp_res.exponent));
                      if (signed_form && tmp_res.sign)
                          tmp_res.significand = (~tmp_res.significand) + 1;

                      tmp_res_hi = tmp_res.significand{31:0};
                  }

                  if (fp_is_nan(tmp_default_result_pair.lo)) {
                      tmp_res_lo = INTEGER_INDEFINITE_32_BIT;
                  } else {
                      tmp_res = fp_ieee_rnd_to_int_sp(fp_reg_read_lo(f2), LOW, &tmp_fp_env);
                      if (tmp_res.exponent)
                          tmp_res.significand = fp_U64_rsh(
                              tmp_res.significand, (FP_INTEGER_EXP - tmp_res.exponent));
                      if (signed_form && tmp_res.sign)
                          tmp_res.significand = (~tmp_res.significand) + 1;

                      tmp_res_lo = tmp_res.significand{31:0};
                  }

                  FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                  FR[f1].exponent = FP_INTEGER_EXP;
                  FR[f1].sign = FP_SIGN_POSITIVE;

                  fp_update_fpsr(sf, tmp_fp_env);
                  fp_update_psr(f1);
                  if (fp_raise_traps(tmp_fp_env))
                      fp_exception_trap(fp_decode_trap(tmp_fp_env));
              }
          }

```

FP Exceptions: Invalid Operation (V) Inexact (I)
 Denormal/Unnormal Operand (D)
 Software Assist (SWA) Fault

Floating-Point Parallel Multiply Add

Format: (qp) fpma.*sf* $f_1 = f_3, f_4, f_2$ F1

Description: The pair of products of the pairs of single precision values in the significand fields of FR f_3 and FR f_4 are computed to infinite precision and then the pair of single precision values in the significand field of FR f_2 is added to these products, again in infinite precision. The resulting values are then rounded to single precision using the rounding mode specified by FPSR.*sf.rc*. The pair of rounded results are stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed results.

Note: If f_2 is f0 in the fpma instruction, just the IEEE multiply operation is performed. (See “Floating-Point Parallel Multiply” on page 6-76.) FR f_1 , as an operand, is not a packed pair of 1.0 values, it is just the register file format’s 1.0 value.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

The encodings and interpretation for the status field’s *rc* are given in Table 5-6 on page 5-5.


```

Operation:   if (PR[qp]) {
              fp_check_target_register( $f_1$ );
              if (tmp_israncode = fp_reg_disabled( $f_1$ ,  $f_2$ ,  $f_3$ ,  $f_4$ ))
                  disabled_fp_register_fault(tmp_israncode, 0);

              if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ]) || fp_is_natval(FR[ $f_4$ ])) {
                  FR[ $f_1$ ] = NATVAL;
                  fp_update_psr( $f_1$ );
              } else {
                  tmp_default_result_pair = fpma_exception_fault_check( $f_2$ ,
                                                                            $f_3$ ,  $f_4$ ,  $sf$ , &tmp_fp_env);

                  if (fp_raise_fault(tmp_fp_env))
                      fp_exception_fault(fp_decode_fault(tmp_fp_env));

                  if (fp_is_nan_or_inf(tmp_default_result_pair.hi)) {
                      tmp_res_hi = fp_single(tmp_default_result_pair.hi);
                  } else {
                      tmp_res = fp_mul(fp_reg_read_hi( $f_3$ ), fp_reg_read_hi( $f_4$ ));
                      if ( $f_2$  != 0)
                          tmp_res = fp_add(tmp_res, fp_reg_read_hi( $f_2$ ), tmp_fp_env);
                      tmp_res_hi = fp_ieee_round_sp(tmp_res, HIGH, &tmp_fp_env);
                  }

                  if (fp_is_nan_or_inf(tmp_default_result_pair.lo)) {
                      tmp_res_lo = fp_single(tmp_default_result_pair.lo);
                  } else {
                      tmp_res = fp_mul(fp_reg_read_lo( $f_3$ ), fp_reg_read_lo( $f_4$ ));
                      if ( $f_2$  != 0)
                          tmp_res = fp_add(tmp_res, fp_reg_read_lo( $f_2$ ), tmp_fp_env);
                      tmp_res_lo = fp_ieee_round_sp(tmp_res, LOW, &tmp_fp_env);
                  }

                  FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                  FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
                  FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;

                  fp_update_fpsr( $sf$ , tmp_fp_env);
                  fp_update_psr( $f_1$ );
                  if (fp_raise_traps(tmp_fp_env))
                      fp_exception_trap(fp_decode_trap(tmp_fp_env));
              }
          }

```

FP Exceptions: Invalid Operation (V)	Underflow (U)
Denormal/Unnormal Operand (D)	Overflow (O)
Software Assist (SWA) Fault	Inexact (I)
	Software Assist (SWA) trap

Floating-Point Parallel Maximum

Format: (qp) fpmax.sf $f_1 = f_2, f_3$ F8

Description: The paired single precision values in the significands of FR f_2 or FR f_3 are compared. The operands with the larger value is returned in the significand of FR f_1 .

If the value of high (low) FR f_3 is less than the value of high (low) FR f_2 , high (low) FR f_1 gets high (low) FR f_2 . Otherwise high (low) FR f_1 gets high (low) FR f_3 .

If high (low) FR f_2 or high (low) FR f_3 is a NaN, high (low) FR f_1 gets high (low) FR f_3 .

The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as for the fpcmp.lt operation.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register( $f_1$ );
    if (tmp_israncode = fp_reg_disabled( $f_1, f_2, f_3, 0$ ))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ])) {
        FR[ $f_1$ ] = NATVAL;
    } else {
        fpminmax_exception_fault_check( $f_2, f_3, sf, &tmp_fp_env$ );
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_fr2 = tmp_right = fp_reg_read_hi( $f_2$ );
        tmp_fr3 = tmp_left = fp_reg_read_hi( $f_3$ );
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_hi = fp_single(tmp_bool_res ? tmp_fr2 : tmp_fr3);

        tmp_fr2 = tmp_right = fp_reg_read_lo( $f_2$ );
        tmp_fr3 = tmp_left = fp_reg_read_lo( $f_3$ );
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_lo = fp_single(tmp_bool_res ? tmp_fr2 : tmp_fr3);

        FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
        FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
        FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;

        fp_update_fpsr(sf, tmp_fp_env);
    }
    fp_update_psr( $f_1$ );
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Parallel Merge

Format:	(qp) fpmerge.ns $f_1 = f_2, f_3$	neg_sign_form	F9
	(qp) fpmerge.s $f_1 = f_2, f_3$	sign_form	F9
	(qp) fpmerge.se $f_1 = f_2, f_3$	sign_exp_form	F9

Description: For the `neg_sign_form`, the signs of the pair of single precision values in the significand field of FR f_2 are negated and concatenated with the exponents and the significands of the pair of single precision values in the significand field of FR f_3 and stored in the significand field of FR f_1 . This form can be used to negate a pair of single precision floating-point numbers by using the same register for f_2 and f_3 .

For the `sign_form`, the signs of the pair of single precision values in the significand field of FR f_2 are concatenated with the exponents and the significands of the pair of single precision values in the significand field of FR f_3 and stored in FR f_1 .

For the `sign_exp_form`, the signs and exponents of the pair of single precision values in the significand field of FR f_2 are concatenated with the pair of single precision significands in the significand field of FR f_3 and stored in the significand field of FR f_1 .

For all forms, the exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

For all forms, if either FR f_2 or FR f_3 is a NaNVal, FR f_1 is set to NaNVal instead of the computed result.




Figure 6-14. Floating-point Merge Negative Sign Operation




Figure 6-15. Floating-point Merge Sign Operation




Figure 6-16. Floating-point Merge Sign and Exponent Operation

```

Operation:  if (PR[qp]) {
            fp_check_target_register(f1);
            if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
                disabled_fp_register_fault(tmp_israncode, 0);

            if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
                FR[f1] = NATVAL;
            } else {
                if (neg_sign_form) {
                    tmp_res_hi = (!FR[f2].significand{63} << 31)
                                | (FR[f3].significand{62:32});
                    tmp_res_lo = (!FR[f2].significand{31} << 31)
                                | (FR[f3].significand{30:0});
                } else if (sign_form) {
                    tmp_res_hi = (FR[f2].significand{63} << 31)
                                | (FR[f3].significand{62:32});
                    tmp_res_lo = (FR[f2].significand{31} << 31)
                                | (FR[f3].significand{30:0});
                } else {
                    tmp_res_hi = (FR[f2].significand{63:55} << 23)
                                | (FR[f3].significand{54:32});
                    tmp_res_lo = (FR[f2].significand{31:23} << 23)
                                | (FR[f3].significand{22:0});
                }

                FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                FR[f1].exponent = FP_INTEGER_EXP;
                FR[f1].sign = FP_SIGN_POSITIVE;
            }

            fp_update_psr(f1);
        }

```

FP Exceptions: None

Floating-Point Parallel Minimum

Format: (qp) `fpmin.sf f1 = f2, f3` F8

Description: The paired single precision values in the significands of FR f_2 or FR f_3 are compared. The operands with the smaller value is returned in significand of FR f_1 .

If the value of high (low) FR f_2 is less than the value of high (low) FR f_3 , high (low) FR f_1 gets high (low) FR f_2 . Otherwise high (low) FR f_1 gets high (low) FR f_3 .

If high (low) FR f_2 or high (low) FR f_3 is a NaN, high (low) FR f_1 gets high (low) FR f_3 .

The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

This operation does not propagate NaNs the same way as other floating-point arithmetic operations. The Invalid Operation is signaled in the same manner as for the `fpcmp.lt` operation.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        fpminmax_exception_fault_check(f2, f3, sf, &tmp_fp_env);
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        tmp_fr2 = tmp_left = fp_reg_read_hi(f2);
        tmp_fr3 = tmp_right = fp_reg_read_hi(f3);
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_hi = fp_single(tmp_bool_res ? tmp_fr2 : tmp_fr3);

        tmp_fr2 = tmp_left = fp_reg_read_lo(f2);
        tmp_fr3 = tmp_right = fp_reg_read_lo(f3);
        tmp_bool_res = fp_less_than(tmp_left, tmp_right);
        tmp_res_lo = fp_single(tmp_bool_res ? tmp_fr2 : tmp_fr3);

        FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;

        fp_update_fpsr(sf, tmp_fp_env);
    }
    fp_update_psr(f1);
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Parallel Multiply

Format: (qp) fpm_{py}.*sf* $f_1 = f_3, f_4$ pseudo-op of: (qp) fpma.*sf* $f_1 = f_3, f_4, f_0$

Description: The pair of products of the pairs of single precision values in the significand fields of FR f_3 and FR f_4 are computed to infinite precision. The resulting values are then rounded to single precision using the rounding mode specified by FPSR.*sf.rc*. The pair of rounded results are stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_3 , or FR f_4 is a NaTVal, FR f_1 is set to NaTVal instead of the computed results.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

The encodings and interpretation for the status field's *rc* are given in Table 5-6 on page 5-5.

Operation: See "Floating-Point Parallel Multiply Add" on page 6-70.

Floating-Point Parallel Multiply Subtract

Format: (qp) fpms.*sf* $f_1 = f_3, f_4, f_2$

F1

Description: The pair of products of the pairs of single precision values in the significand fields of FR f_3 and FR f_4 are computed to infinite precision and then the pair of single precision values in the significand field of FR f_2 is subtracted from these products, again in infinite precision. The resulting values are then rounded to single precision using the rounding mode specified by FPSR.*sf.rc*. The pair of rounded results are stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed results.

Note: If f_2 is f0 in the fpms instruction, just the IEEE multiply operation is performed.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

The encodings and interpretation for the status field's *rc* are given in Table 5-6 on page 5-5.

```

Operation:   if (PR[qp]) {
                fp_check_target_register( $f_1$ );
                if (tmp_israncode = fp_reg_disabled( $f_1$ ,  $f_2$ ,  $f_3$ ,  $f_4$ ))
                    disabled_fp_register_fault(tmp_israncode, 0);

                if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ]) || fp_is_natval(FR[ $f_4$ ])) {
                    FR[ $f_1$ ] = NATVAL;
                    fp_update_psr( $f_1$ );
                } else {
                    tmp_default_result_pair = fpms_fpnma_exception_fault_check( $f_2$ ,  $f_3$ ,  $f_4$ ,
                                                                              sf, &tmp_fp_env);

                    if (fp_raise_fault(tmp_fp_env))
                        fp_exception_fault(fp_decode_fault(tmp_fp_env));

                    if (fp_is_nan_or_inf(tmp_default_result_pair.hi)) {
                        tmp_res_hi = fp_single(tmp_default_result_pair.hi);
                    } else {
                        tmp_res = fp_mul(fp_reg_read_hi( $f_3$ ), fp_reg_read_hi( $f_4$ ));
                        if ( $f_2$  != 0) {
                            tmp_sub = fp_reg_read_hi( $f_2$ );
                            tmp_sub.sign = !tmp_sub.sign;
                            tmp_res = fp_add(tmp_res, tmp_sub, tmp_fp_env);
                        }
                        tmp_res_hi = fp_ieee_round_sp(tmp_res, HIGH, &tmp_fp_env);
                    }

                    if (fp_is_nan_or_inf(tmp_default_result_pair.lo)) {
                        tmp_res_lo = fp_single(tmp_default_result_pair.lo);
                    } else {
                        tmp_res = fp_mul(fp_reg_read_lo( $f_3$ ), fp_reg_read_lo( $f_4$ ));
                        if ( $f_2$  != 0) {
                            tmp_sub = fp_reg_read_lo( $f_2$ );
                            tmp_sub.sign = !tmp_sub.sign;
                            tmp_res = fp_add(tmp_res, tmp_sub, tmp_fp_env);
                        }
                        tmp_res_lo = fp_ieee_round_sp(tmp_res, LOW, &tmp_fp_env);
                    }

                    FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                    FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
                    FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;

                    fp_update_fpsr(sf, tmp_fp_env);
                    fp_update_psr( $f_1$ );
                    if (fp_raise_traps(tmp_fp_env))
                        fp_exception_trap(fp_decode_trap(tmp_fp_env));
                }
            }

```

FP Exceptions: Invalid Operation (V)	Underflow (U)
Denormal/Unnormal Operand (D)	Overflow (O)
Software Assist (SWA) Fault	Inexact (I)
	Software Assist (SWA) trap

Floating-Point Parallel Negate

Format: (qp) fpneg $f_1 = f_3$ pseudo-op of: (qp) fpmerge.ns $f_1 = f_3, f_3$

Description: The pair of single precision values in the significand field of FR f_3 are negated and stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation: See "Floating-Point Parallel Merge" on page 6-73.

Floating-Point Parallel Negate Absolute Value

Format: (qp) fpnegabs $f_1 = f_3$ pseudo-op of: (qp) fpmerge.ns $f_1 = f_0, f_3$

Description: The absolute values of the pair of single precision values in the significand field of FR f_3 are computed, negated and stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation: See "Floating-Point Parallel Merge" on page 6-73.

Floating-Point Parallel Negative Multiply Add

Format: (qp) fpnma.*sf* $f_1 = f_3, f_4, f_2$ F1

Description: The pair of products of the pairs of single precision values in the significand fields of FR f_3 and FR f_4 are computed to infinite precision, negated, and then the pair of single precision values in the significand field of FR f_2 are added to these (negated) products, again in infinite precision. The resulting values are then rounded to single precision using the rounding mode specified by FPSR.*sf.rc*. The pair of rounded results are stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Note: If f_2 is f0 in the fpnma instruction, just the IEEE multiply operation (with the product being negated before rounding) is performed.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

The encodings and interpretation for the status field's *rc* are given in Table 5-6 on page 5-5.

```

Operation:   if (PR[qp]) {
                fp_check_target_register(f1);
                if (tmp_israncode = fp_reg_disabled(f1, f2, f3, f4))
                    disabled_fp_register_fault(tmp_israncode, 0);

                if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3]) || fp_is_natval(FR[f4])) {
                    FR[f1] = NATVAL;
                    fp_update_psr(f1);
                } else {
                    tmp_default_result_pair = fpms_fpnma_exception_fault_check(f2, f3, f4,
                                                                              sf, &tmp_fp_env);

                    if (fp_raise_fault(tmp_fp_env))
                        fp_exception_fault(fp_decode_fault(tmp_fp_env));

                    if (fp_is_nan_or_inf(tmp_default_result_pair.hi)) {
                        tmp_res_hi = fp_single(tmp_default_result_pair.hi);
                    } else {
                        tmp_res = fp_mul(fp_reg_read_hi(f3), fp_reg_read_hi(f4));
                        tmp_res.sign = !tmp_res.sign;
                        if (f2 != 0)
                            tmp_res = fp_add(tmp_res, fp_reg_read_hi(f2), tmp_fp_env);
                        tmp_res_hi = fp_ieee_round_sp(tmp_res, HIGH, &tmp_fp_env);
                    }

                    if (fp_is_nan_or_inf(tmp_default_result_pair.lo)) {
                        tmp_res_lo = fp_single(tmp_default_result_pair.lo);
                    } else {
                        tmp_res = fp_mul(fp_reg_read_lo(f3), fp_reg_read_lo(f4));
                        tmp_res.sign = !tmp_res.sign;
                        if (f2 != 0)
                            tmp_res = fp_add(tmp_res, fp_reg_read_lo(f2), tmp_fp_env);
                        tmp_res_lo = fp_ieee_round_sp(tmp_res, LOW, &tmp_fp_env);
                    }

                    FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                    FR[f1].exponent = FP_INTEGER_EXP;
                    FR[f1].sign = FP_SIGN_POSITIVE;

                    fp_update_fpsr(sf, tmp_fp_env);
                    fp_update_psr(f1);
                    if (fp_raise_traps(tmp_fp_env))
                        fp_exception_trap(fp_decode_trap(tmp_fp_env));
                }
            }

```

FP Exceptions: Invalid Operation (V) Underflow (U)
Denormal/Unnormal Operand (D) Overflow (O)
Software Assist (SWA) fault Inexact (I)
 Software Assist (SWA) trap

Floating-Point Parallel Negative Multiply

Format: (qp) fpmmpy.*sf* $f_1 = f_3, f_4$ pseudo-op of: (qp) fpmma.*sf* $f_1 = f_3, f_4, f_0$

Description: The pair of products of the pairs of single precision values in the significand fields of FR f_3 and FR f_4 are computed to infinite precision and then negated. The resulting values are then rounded to single precision using the rounding mode specified by FPSR.*sf.rc*. The pair of rounded results are stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either FR f_3 or FR f_4 is a NaTVal, FR f_1 is set to NaTVal instead of the computed results.

The mnemonic values for *sf* are given in Table 6-18 on page 6-30.

The encodings and interpretation for the status field's *rc* are given in Table 5-6 on page 5-5.

Operation: See "Floating-Point Parallel Negative Multiply Add" on page 6-81.

Floating-Point Parallel Reciprocal Approximation

Format: $(qp) \text{ fprcpa.sf } f_1, p_2 = f_2, f_3$

F6

Description: If PR qp is 0, PR p_2 is cleared and FR f_1 remains unchanged.

If PR qp is 1, the following will occur:

- Each half of the significand of FR f_1 is either set to an approximation (with a relative error $< 2^{-8.886}$) of the reciprocal of the corresponding half of FR f_3 , or set to the IEEE-754 mandated response for the quotient $\text{FR } f_2/\text{FR } f_3$ of the corresponding half — if that half of FR f_2 or of FR f_3 is in the set $\{-\text{Infinity}, -0, +0, +\text{Infinity}, \text{NaN}\}$.
- If either half of FR f_1 is set to the IEEE-754 mandated quotient, or is set to an approximation of the reciprocal which may cause the Newton-Raphson iterations to fail to produce the correct IEEE-754 divide result, then PR p_2 is set to 0, otherwise it is set to 1.
For correct IEEE divide results, when PR p_2 is cleared, user software is expected to compute the quotient $(\text{FR } f_2/\text{FR } f_3)$ for each half (using the non-parallel `frcpa` instruction), and merge the results into FR f_1 , keeping PR p_2 cleared.
- The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).
- If either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result, and PR p_2 is cleared.

The mnemonic values for sf are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
        PR[p2] = 0;
    } else {
        tmp_default_result_pair = fprcpa_exception_fault_check(f2, f3, sf,
                                                                &tmp_fp_env, &limits_check);

        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        if (fp_is_nan_or_inf(tmp_default_result_pair.hi) || limits_check.hi_fr3) {
            tmp_res_hi = fp_single(tmp_default_result_pair.hi);
            tmp_pred_hi = 0;
        } else {
            num = fp_normalize(fp_reg_read_hi(f2));
            den = fp_normalize(fp_reg_read_hi(f3));
            if (fp_is_inf(num) && fp_is_finite(den)) {
                tmp_res = FP_INFINITY;
                tmp_res.sign = num.sign ^ den.sign;
                tmp_pred_hi = 0;
            } else if (fp_is_finite(num) && fp_is_inf(den)) {
                tmp_res = FP_ZERO;
                tmp_res.sign = num.sign ^ den.sign;
                tmp_pred_hi = 0;
            } else if (fp_is_zero(num) && fp_is_finite(den)) {
                tmp_res = FP_ZERO;
                tmp_res.sign = num.sign ^ den.sign;
                tmp_pred_hi = 0;
            } else {
                tmp_res = fp_ieee_recip(den);
                if (limits_check.hi_fr2_or_quot)
                    tmp_pred_hi = 0;
                else
                    tmp_pred_hi = 1;
            }
        }
    }
}

```

```

        tmp_res_hi = fp_single(tmp_res);
    }
    if (fp_is_nan_or_inf(tmp_default_result_pair.lo) || limits_check.lo_fr3) {
        tmp_res_lo = fp_single(tmp_default_result_pair.lo);
        tmp_pred_lo = 0;
    } else {
        num = fp_normalize(fp_reg_read_lo(f2));
        den = fp_normalize(fp_reg_read_lo(f3));
        if (fp_is_inf(num) && fp_is_finite(den)) {
            tmp_res = FP_INFINITY;
            tmp_res.sign = num.sign ^ den.sign;
            tmp_pred_lo = 0;
        } else if (fp_is_finite(num) && fp_is_inf(den)) {
            tmp_res = FP_ZERO;
            tmp_res.sign = num.sign ^ den.sign;
            tmp_pred_lo = 0;
        } else if (fp_is_zero(num) && fp_is_finite(den)) {
            tmp_res = FP_ZERO;
            tmp_res.sign = num.sign ^ den.sign;
            tmp_pred_lo = 0;
        } else {
            tmp_res = fp_ieee_recip(den);
            if (limits_check.lo_fr2_or_quot)
                tmp_pred_lo = 0;
            else
                tmp_pred_lo = 1;
        }
        tmp_res_lo = fp_single(tmp_res);
    }

    FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
    FR[f1].exponent = FP_INTEGER_EXP;
    FR[f1].sign = FP_SIGN_POSITIVE;
    PR[p2] = tmp_pred_hi && tmp_pred_lo;

    fp_update_fpsr(sf, tmp_fp_env);
}
fp_update_psr(f1);
} else {
    PR[p2] = 0;
}

```

FP Exceptions: Invalid Operation (V)
 Zero Divide (Z)
 Denormal/Unnormal Operand (D)
 Software Assist (SWA) fault

Floating-Point Parallel Reciprocal Square Root Approximation

Format: $(qp) \text{ fprsqrrta.sf } f_1, p_2 = f_3$

F7

Description: If PR qp is 0, PR p_2 is cleared and FR f_1 remains unchanged.

If PR qp is 1, the following will occur:

- Each half of the significand of FR f_1 is either set to an approximation (with a relative error $< 2^{-8.831}$) of the reciprocal square root of the corresponding half of FR f_3 , or set to the IEEE-754 compliant response for the reciprocal square root of the corresponding half of FR f_3 — if that half of FR f_3 is in the set $\{-\text{Infinity}, -\text{Finite}, -0, +0, +\text{Infinity}, \text{NaN}\}$.
- If either half of FR f_1 is set to the IEEE-754 mandated reciprocal square root, or is set to an approximation of the reciprocal square root which may cause the Newton-Raphson iterations to fail to produce the correct IEEE-754 square root result, then PR p_2 is set to 0, otherwise it is set to 1.
For correct IEEE square root results, when PR p_2 is cleared, user software is expected to compute the square root for each half (using the non-parallel `frsqrrta` instruction), and merge the results in FR f_1 , keeping PR p_2 cleared.
- The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).
- If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result, and PR p_2 is cleared.

The mnemonic values for sf are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f3, 0, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
        PR[p2] = 0;
    } else {
        tmp_default_result_pair = fprsqrrta_exception_fault_check(f3, sf,
                                                                    &tmp_fp_env, &limits_check);

        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        if (fp_is_nan(tmp_default_result_pair.hi)) {
            tmp_res_hi = fp_single(tmp_default_result_pair.hi);
            tmp_pred_hi = 0;
        } else {
            tmp_fr3 = fp_normalize(fp_reg_read_hi(f3));
            if (fp_is_zero(tmp_fr3)) {
                tmp_res = FP_INFINITY;
                tmp_res.sign = tmp_fr3.sign;
                tmp_pred_hi = 0;
            } else if (fp_is_pos_inf(tmp_fr3)) {
                tmp_res = FP_ZERO;
                tmp_pred_hi = 0;
            } else {
                tmp_res = fp_ieee_recip_sqrt(tmp_fr3);
                if (limits_check.hi)
                    tmp_pred_hi = 0;
                else
                    tmp_pred_hi = 1;
            }
            tmp_res_hi = fp_single(tmp_res);
        }

        if (fp_is_nan(tmp_default_result_pair.lo)) {
            tmp_res_lo = fp_single(tmp_default_result_pair.lo);
            tmp_pred_lo = 0;
        } else {
            tmp_fr3 = fp_normalize(fp_reg_read_lo(f3));

```



```

        if (fp_is_zero(tmp_fr3)) {
            tmp_res = FP_INFINITY;
            tmp_res.sign = tmp_fr3.sign;
            tmp_pred_lo = 0;
        } else if (fp_is_pos_inf(tmp_fr3)) {
            tmp_res = FP_ZERO;
            tmp_pred_lo = 0;
        } else {
            tmp_res = fp_ieee_recip_sqrt(tmp_fr3);
            if (limits_check.lo)
                tmp_pred_lo = 0;
            else
                tmp_pred_lo = 1;
        }
        tmp_res_lo = fp_single(tmp_res);
    }

    FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
    FR[f1].exponent = FP_INTEGER_EXP;
    FR[f1].sign = FP_SIGN_POSITIVE;
    PR[p2] = tmp_pred_hi && tmp_pred_lo;

    fp_update_fpsr(sf, tmp_fp_env);
}
fp_update_psr(f1);
} else {
    PR[p2] = 0;
}

```

FP Exceptions: Invalid Operation (V)
 Denormal/Unnormal Operand (D)
 Software Assist (SWA) fault

Floating-Point Reciprocal Approximation

Format: (qp) frcpa.sf $f_1, p_2 = f_2, f_3$

F6

Description: If PR qp is 0, PR p₂ is cleared and FR f₁ remains unchanged.

If PR qp is 1, the following will occur:

- FR f₁ is either set to an approximation (with a relative error $< 2^{-8.886}$) of the reciprocal of FR f₃, or to the IEEE-754 mandated quotient of FR f₂/FR f₃ — if either FR f₂ or FR f₃ is in the set {-Infinity, -0, Pseudo-zero, +0, +Infinity, NaN, Unsupported}.
- If FR f₁ is set to the approximation of the reciprocal of FR f₃, then PR p₂ is set to 1; otherwise, it is set to 0.
- If FR f₂ and FR f₃ are such that the approximation of FR f₃'s reciprocal may cause the Newton-Raphson iterations to fail to produce the correct IEEE-754 result of FR f₂/FR f₃, then a Floating-point Exception fault for Software Assist occurs.
System software is expected to compute the IEEE-754 quotient (FR f₂/FR f₃), return the result in FR f₁, and set PR p₂ to 0.
- If either FR f₂ or FR f₃ is a NaTVal, FR f₁ is set to NaTVal instead of the computed result, and PR p₂ is cleared.

The mnemonic values for sf are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
        PR[p2] = 0;
    } else {
        tmp_default_result = frcpa_exception_fault_check(f2, f3, sf, &tmp_fp_env);
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        if (fp_is_nan_or_inf(tmp_default_result)) {
            FR[f1] = tmp_default_result;
            PR[p2] = 0;
        } else {
            num = fp_normalize(fp_reg_read(FR[f2]));
            den = fp_normalize(fp_reg_read(FR[f3]));
            if (fp_is_inf(num) && fp_is_finite(den)) {
                FR[f1] = FP_INFINITY;
                FR[f1].sign = num.sign ^ den.sign;
                PR[p2] = 0;
            } else if (fp_is_finite(num) && fp_is_inf(den)) {
                FR[f1] = FP_ZERO;
                FR[f1].sign = num.sign ^ den.sign;
                PR[p2] = 0;
            } else if (fp_is_zero(num) && fp_is_finite(den)) {
                FR[f1] = FP_ZERO;
                FR[f1].sign = num.sign ^ den.sign;
                PR[p2] = 0;
            } else {
                FR[f1] = fp_ieee_recip(den);
                PR[p2] = 1;
            }
        }
    }
    fp_update_fpsr(sf, tmp_fp_env);
}
fp_update_psr(f1);
} else {
    PR[p2] = 0;
}

```

```

// fp_ieee_recip()

fp_ieee_recip(den)
{
    const EM_uint_t RECIP_TABLE[256] = {
        0x3fc, 0x3f4, 0x3ec, 0x3e4, 0x3dd, 0x3d5, 0x3cd, 0x3c6,
        0x3be, 0x3b7, 0x3af, 0x3a8, 0x3a1, 0x399, 0x392, 0x38b,
        0x384, 0x37d, 0x376, 0x36f, 0x368, 0x361, 0x35b, 0x354,
        0x34d, 0x346, 0x340, 0x339, 0x333, 0x32c, 0x326, 0x320,
        0x319, 0x313, 0x30d, 0x307, 0x300, 0x2fa, 0x2f4, 0x2ee,
        0x2e8, 0x2e2, 0x2dc, 0x2d7, 0x2d1, 0x2cb, 0x2c5, 0x2bf,
        0x2ba, 0x2b4, 0x2af, 0x2a9, 0x2a3, 0x29e, 0x299, 0x293,
        0x28e, 0x288, 0x283, 0x27e, 0x279, 0x273, 0x26e, 0x269,
        0x264, 0x25f, 0x25a, 0x255, 0x250, 0x24b, 0x246, 0x241,
        0x23c, 0x237, 0x232, 0x22e, 0x229, 0x224, 0x21f, 0x21b,
        0x216, 0x211, 0x20d, 0x208, 0x204, 0x1ff, 0x1fb, 0x1f6,
        0x1f2, 0x1ed, 0x1e9, 0x1e5, 0x1e0, 0x1dc, 0x1d8, 0x1d4,
        0x1cf, 0x1cb, 0x1c7, 0x1c3, 0x1bf, 0x1bb, 0x1b6, 0x1b2,
        0x1ae, 0x1aa, 0x1a6, 0x1a2, 0x19e, 0x19a, 0x197, 0x193,
        0x18f, 0x18b, 0x187, 0x183, 0x17f, 0x17c, 0x178, 0x174,
        0x171, 0x16d, 0x169, 0x166, 0x162, 0x15e, 0x15b, 0x157,
        0x154, 0x150, 0x14d, 0x149, 0x146, 0x142, 0x13f, 0x13b,
        0x138, 0x134, 0x131, 0x12e, 0x12a, 0x127, 0x124, 0x120,
        0x11d, 0x11a, 0x117, 0x113, 0x110, 0x10d, 0x10a, 0x107,
        0x103, 0x100, 0x0fd, 0x0fa, 0x0f7, 0x0f4, 0x0f1, 0x0ee,
        0x0eb, 0x0e8, 0x0e5, 0x0e2, 0x0df, 0x0dc, 0x0d9, 0x0d6,
        0x0d3, 0x0d0, 0x0cd, 0x0ca, 0x0c8, 0x0c5, 0x0c2, 0x0bf,
        0x0bc, 0x0b9, 0x0b7, 0x0b4, 0x0b1, 0x0ae, 0x0ac, 0x0a9,
        0x0a6, 0x0a4, 0x0a1, 0x09e, 0x09c, 0x099, 0x096, 0x094,
        0x091, 0x08e, 0x08c, 0x089, 0x087, 0x084, 0x082, 0x07f,
        0x07c, 0x07a, 0x077, 0x075, 0x073, 0x070, 0x06e, 0x06b,
        0x069, 0x066, 0x064, 0x061, 0x05f, 0x05d, 0x05a, 0x058,
        0x056, 0x053, 0x051, 0x04f, 0x04c, 0x04a, 0x048, 0x045,
        0x043, 0x041, 0x03f, 0x03c, 0x03a, 0x038, 0x036, 0x033,
        0x031, 0x02f, 0x02d, 0x02b, 0x029, 0x026, 0x024, 0x022,
        0x020, 0x01e, 0x01c, 0x01a, 0x018, 0x015, 0x013, 0x011,
        0x00f, 0x00d, 0x00b, 0x009, 0x007, 0x005, 0x003, 0x001,
    };

    tmp_index = den.significand{62:55};
    tmp_res.significand = (1 << 63) | (RECIP_TABLE[tmp_index] << 53);
    tmp_res.exponent = FP_REG_EXP_ONES - 2 - den.exponent;
    tmp_res.sign = den.sign;
    return (tmp_res);
}

```

FP Exceptions: Invalid Operation (V)
Zero Divide (Z)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Reciprocal Square Root Approximation

Format: $(qp) \text{ frsqrta.sf } f_1, p_2 = f_3$

F7

Description: If PR qp is 0, PR p_2 is cleared and FR f_1 remains unchanged.

If PR qp is 1, the following will occur:

- FR f_1 is either set to an approximation (with a relative error $< 2^{-8.831}$) of the reciprocal square root of FR f_3 , or set to the IEEE-754 mandated square root of FR f_3 — if FR f_3 is in the set $\{-\text{Infinity}, -\text{Finite}, -0, \text{Pseudo-zero}, +0, +\text{Infinity}, \text{NaN}, \text{Unsupported}\}$.
- If FR f_1 is set to an approximation of the reciprocal square root of FR f_3 , then PR p_2 is set to 1; otherwise, it is set to 0.
- If FR f_3 is such the approximation of its reciprocal square root may cause the Newton-Raphson iterations to fail to produce the correct IEEE-754 square root result, then a Floating-point Exception fault for Software Assist occurs.
System software is expected to compute the IEEE-754 square root, return the result in FR f_1 , and set PR p_2 to 0.
- If FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result, and PR p_2 is cleared.

The mnemonic values for sf are given in Table 6-18 on page 6-30.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_isr = fp_reg_disabled(f1, f3, 0, 0))
        disabled_fp_register_fault(tmp_isr, 0);

    if (fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
        PR[p2] = 0;
    } else {
        tmp_default_result = frsqrta_exception_fault_check(f3, sf, &tmp_fp_env);
        if (fp_raise_fault(tmp_fp_env))
            fp_exception_fault(fp_decode_fault(tmp_fp_env));

        if (fp_is_nan(tmp_default_result)) {
            FR[f1] = tmp_default_result;
            PR[p2] = 0;
        } else {
            tmp_fr3 = fp_normalize(fp_reg_read(FR[f3]));
            if (fp_is_zero(tmp_fr3)) {
                FR[f1] = tmp_fr3;
                PR[p2] = 0;
            } else if (fp_is_pos_inf(tmp_fr3)) {
                FR[f1] = tmp_fr3;
                PR[p2] = 0;
            } else {
                FR[f1] = fp_ieee_recip_sqrt(tmp_fr3);
                PR[p2] = 1;
            }
        }
        fp_update_fpsr(sf, tmp_fp_env);
    }
    fp_update_psr(f1);
} else {
    PR[p2] = 0;
}

// fp_ieee_recip_sqrt()
fp_ieee_recip_sqrt(root)
{
    const EM_uint_t RECIP_SQRT_TABLE[256] = {
        0x1a5, 0x1a0, 0x19a, 0x195, 0x18f, 0x18a, 0x185, 0x180,
        0x17a, 0x175, 0x170, 0x16b, 0x166, 0x161, 0x15d, 0x158,

```

```

    0x153, 0x14e, 0x14a, 0x145, 0x140, 0x13c, 0x138, 0x133,
    0x12f, 0x12a, 0x126, 0x122, 0x11e, 0x11a, 0x115, 0x111,
    0x10d, 0x109, 0x105, 0x101, 0x0fd, 0x0fa, 0x0f6, 0x0f2,
    0x0ee, 0x0ea, 0x0e7, 0x0e3, 0x0df, 0x0dc, 0x0d8, 0x0d5,
    0x0d1, 0x0ce, 0x0ca, 0x0c7, 0x0c3, 0x0c0, 0x0bd, 0x0b9,
    0x0b6, 0x0b3, 0x0b0, 0x0ad, 0x0a9, 0x0a6, 0x0a3, 0x0a0,
    0x09d, 0x09a, 0x097, 0x094, 0x091, 0x08e, 0x08b, 0x088,
    0x085, 0x082, 0x07f, 0x07d, 0x07a, 0x077, 0x074, 0x071,
    0x06f, 0x06c, 0x069, 0x067, 0x064, 0x061, 0x05f, 0x05c,
    0x05a, 0x057, 0x054, 0x052, 0x04f, 0x04d, 0x04a, 0x048,
    0x045, 0x043, 0x041, 0x03e, 0x03c, 0x03a, 0x037, 0x035,
    0x033, 0x030, 0x02e, 0x02c, 0x029, 0x027, 0x025, 0x023,
    0x020, 0x01e, 0x01c, 0x01a, 0x018, 0x016, 0x014, 0x011,
    0x00f, 0x00d, 0x00b, 0x009, 0x007, 0x005, 0x003, 0x001,
    0x3fc, 0x3f4, 0x3ec, 0x3e5, 0x3dd, 0x3d5, 0x3ce, 0x3c7,
    0x3bf, 0x3b8, 0x3b1, 0x3aa, 0x3a3, 0x39c, 0x395, 0x38e,
    0x388, 0x381, 0x37a, 0x374, 0x36d, 0x367, 0x361, 0x35a,
    0x354, 0x34e, 0x348, 0x342, 0x33c, 0x336, 0x330, 0x32b,
    0x325, 0x31f, 0x31a, 0x314, 0x30f, 0x309, 0x304, 0x2fe,
    0x2f9, 0x2f4, 0x2ee, 0x2e9, 0x2e4, 0x2df, 0x2da, 0x2d5,
    0x2d0, 0x2cb, 0x2c6, 0x2c1, 0x2bd, 0x2b8, 0x2b3, 0x2ae,
    0x2aa, 0x2a5, 0x2a1, 0x29c, 0x298, 0x293, 0x28f, 0x28a,
    0x286, 0x282, 0x27d, 0x279, 0x275, 0x271, 0x26d, 0x268,
    0x264, 0x260, 0x25c, 0x258, 0x254, 0x250, 0x24c, 0x249,
    0x245, 0x241, 0x23d, 0x239, 0x235, 0x232, 0x22e, 0x22a,
    0x227, 0x223, 0x220, 0x21c, 0x218, 0x215, 0x211, 0x20e,
    0x20a, 0x207, 0x204, 0x200, 0x1fd, 0x1f9, 0x1f6, 0x1f3,
    0x1f0, 0x1ec, 0x1e9, 0x1e6, 0x1e3, 0x1df, 0x1dc, 0x1d9,
    0x1d6, 0x1d3, 0x1d0, 0x1cd, 0x1ca, 0x1c7, 0x1c4, 0x1c1,
    0x1be, 0x1bb, 0x1b8, 0x1b5, 0x1b2, 0x1af, 0x1ac, 0x1aa,
};

tmp_index = (root.exponent{0} << 7) | root.significand{62:56};
tmp_res.significand = (1 << 63) | (RECIP_SQRT_TABLE[tmp_index] << 53);
tmp_res.exponent = FP_REG_EXP_HALF - ((root.exponent - FP_REG_BIAS) >> 1);
tmp_res.sign = FP_SIGN_POSITIVE;
return (tmp_res);
}

```

FP Exceptions: Invalid Operation (V)
Denormal/Unnormal Operand (D)
Software Assist (SWA) fault

Floating-Point Select

Format: (qp) fselect $f_1 = f_3, f_4, f_2$ F3

Description: The significand field of FR f_3 is logically AND-ed with the significand field of FR f_2 and the significand field of FR f_4 is logically AND-ed with the one's complement of the significand field of FR f_2 . The two results are logically OR-ed together. The result is placed in the significand field of FR f_1 .

The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E). The sign bit field of FR f_1 is set to positive (0).

If any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, f4))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3]) || fp_is_natval(FR[f4])) {
        FR[f1] = NATVAL;
    } else {
        FR[f1].significand = (FR[f3].significand & FR[f2].significand)
            | (FR[f4].significand & ~FR[f2].significand);
        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;
    }

    fp_update_psr(f1);
}

```

FP Exceptions: None

Floating-Point Set Controls

Format: (qp) fsetc.sf amask7, omask7 F12

Description: The status field's control bits are initialized to the value obtained by logically AND-ing the sf0.controls and amask7 immediate field and logically OR-ing the omask7 immediate field.

The mnemonic values for sf are given in Table 6-18 on page 6-30.

Operation:

```
if (PR[qp]) {
    tmp_controls = (AR[FPSR].sf0.controls & amask7) | omask7;
    if (is_reserved_field(FSETC, sf, tmp_controls))
        reserved_register_field_fault();
    fp_set_sf_controls(sf, tmp_controls);
}
```

FP Exceptions: None

Floating-Point Subtract

Format: (qp) fsub.*pc.sf* $f_1 = f_3, f_2$ pseudo-op of: (qp) fms.*pc.sf* $f_1 = f_3, f_1, f_2$

Description: FR f_2 is subtracted from FR f_3 (computed to infinite precision), rounded to the precision indicated by *pc* (and possibly FPSR.*sf.pc* and FPSR.*sf.wre*) using the rounding mode specified by FPSR.*sf.rc*, and placed in FR f_1 .

If either FR f_3 or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

The mnemonic values for the opcode's *pc* are given in Table 6-17 on page 6-30. The mnemonic values for *sf* are given in Table 6-18 on page 6-30. For the encodings and interpretation of the status field's *pc*, *wre*, and *rc*, refer to Table 5-5 and Table 5-6 on page 5-5.

Operation: See "Floating-Point Multiply Subtract" on page 6-55.

Floating-Point Swap

Format:

<i>(qp)</i> fswap $f_1 = f_2, f_3$	swap_form	F9
<i>(qp)</i> fswap.nl $f_1 = f_2, f_3$	swap_nl_form	F9
<i>(qp)</i> fswap.nr $f_1 = f_2, f_3$	swap_nr_form	F9

Description: For the swap_form, the left single precision value in FR f_2 is concatenated with the right single precision value in FR f_3 . The concatenated pair is then swapped.

For the swap_nl_form, the left single precision value in FR f_2 is concatenated with the right single precision value in FR f_3 . The concatenated pair is then swapped, and the left single precision value is negated.

For the swap_nr_form, the left single precision value in FR f_2 is concatenated with the right single precision value in FR f_3 . The concatenated pair is then swapped, and the right single precision value is negated.

For all forms, the exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

For all forms, if either FR f_2 or FR f_3 is a NaNVal, FR f_1 is set to NaNVal instead of the computed result.




Figure 6-17. Floating-point Swap




Figure 6-18. Floating-point Swap Negate Left or Right

```

Operation:  if (PR[qp]) {
            fp_check_target_register( $f_1$ );
            if (tmp_israncode = fp_reg_disabled( $f_1$ ,  $f_2$ ,  $f_3$ , 0))
                disabled_fp_register_fault(tmp_israncode, 0);

            if (fp_is_natval(FR[ $f_2$ ]) || fp_is_natval(FR[ $f_3$ ])) {
                FR[ $f_1$ ] = NATVAL;
            } else {
                if (swap_form) {
                    tmp_res_hi = FR[ $f_3$ ].significand{31:0};
                    tmp_res_lo = FR[ $f_2$ ].significand{63:32};
                } else if (swap_nl_form) {
                    tmp_res_hi = (!FR[ $f_3$ ].significand{31} << 31)
                        | (FR[ $f_3$ ].significand{30:0});
                    tmp_res_lo = FR[ $f_2$ ].significand{63:32};
                } else { // swap_nr_form
                    tmp_res_hi = FR[ $f_3$ ].significand{31:0};
                    tmp_res_lo = (!FR[ $f_2$ ].significand{63} << 31)
                        | (FR[ $f_2$ ].significand{62:32});
                }

                FR[ $f_1$ ].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
                FR[ $f_1$ ].exponent = FP_INTEGER_EXP;
                FR[ $f_1$ ].sign = FP_SIGN_POSITIVE;
            }

            fp_update_psr( $f_1$ );
        }

```

FP Exceptions: None

Floating-Point Sign Extend

Format: (qp) fsxt.l $f_1 = f_2, f_3$ sxt_l_form F9
 (qp) fsxt.r $f_1 = f_2, f_3$ sxt_r_form F9

Description: For the sxt_l_form (sxt_r_form), the sign of the left (right) single precision value in FR f_2 is extended to 32-bits and is concatenated with the left (right) single precision value in FR f_3 .

For all forms, the exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

For all forms, if either FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.




Figure 6-19. Floating-point Sign Extend Left




Figure 6-20. Floating-point Sign Extend Right

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_isrkode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_isrkode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        if (sxt_l_form) {
            tmp_res_hi = (FR[f2].significand{63} ? 0xFFFFFFFF : 0x00000000);
            tmp_res_lo = FR[f3].significand{63:32};
        } else { // sxt_r_form
            tmp_res_hi = (FR[f2].significand{31} ? 0xFFFFFFFF : 0x00000000);
            tmp_res_lo = FR[f3].significand{31:0};
        }

        FR[f1].significand = fp_concatenate(tmp_res_hi, tmp_res_lo);
        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;
    }

    fp_update_psr(f1);
}

```

FP Exceptions: None

Floating-Point Exclusive Or

Format: (qp) fxor $f_1 = f_2, f_3$ F9

Description: The bit-wise logical exclusive-OR of the significand fields of FR f_2 and FR f_3 is computed. The resulting value is stored in the significand field of FR f_1 . The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

If either of FR f_2 or FR f_3 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3])) {
        FR[f1] = NATVAL;
    } else {
        FR[f1].significand = FR[f2].significand ^ FR[f3].significand;
        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;
    }

    fp_update_psr(f1);
}

```

FP Exceptions: None

Get Floating-Point Value or Exponent or Significand

Format:	(<i>qp</i>) getf.s $r_1 = f_2$	single_form	M19
	(<i>qp</i>) getf.d $r_1 = f_2$	double_form	M19
	(<i>qp</i>) getf.exp $r_1 = f_2$	exponent_form	M19
	(<i>qp</i>) getf.sig $r_1 = f_2$	significand_form	M19

Description: In the single and double forms, the value in FR f_2 is converted into a single precision (single_form) or double precision (double_form) memory representation and placed in GR r_1 . In the single_form, the most-significant 32 bits of GR r_1 are set to 0.

In the exponent_form, the exponent field of FR f_2 is copied to bits 16:0 of GR r_1 and the sign bit of the value in FR f_2 is copied to bit 17 of GR r_1 . The most-significant 46-bits of GR r_1 are set to zero.




Figure 6-21. Function of getf.exp

In the significand_form, the significand field of the value in FR f_2 is copied to GR r_1




Figure 6-22. Function of getf.sig

For all forms, if FR f_2 contains a NaTVal, then the NaT bit corresponding to GR r_1 is set to 1.

Operation:

```

if (PR[qp]) {
    check_target_register(r1);
    if (tmp_israncode = fp_reg_disabled(f2, 0, 0, 0))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (single_form) {
        GR[r1]{31:0} = fp_fr_to_mem_format(FR[f2], 4, 0);
        GR[r1]{63:32} = 0;
    } else if (double_form) {
        GR[r1] = fp_fr_to_mem_format(FR[f2], 8, 0);
    } else if (exponent_form) {
        GR[r1]{63:18} = 0;
        GR[r1]{16:0} = FR[f2].exponent;
        GR[r1]{17} = FR[f2].sign;
    } else // significand_form
        GR[r1] = FR[f2].significand;
    if (fp_is_natval(FR[f2]))
        GR[r1].nat = 1;
    else
        GR[r1].nat = 0;
}

```

Invalidate ALAT

Format:	(<i>qp</i>) invala	complete_form	M24
	(<i>qp</i>) invala.e <i>r_l</i>	gr_form, entry_form	M26
	(<i>qp</i>) invala.e <i>f_l</i>	fr_form, entry_form	M27

Description: The selected entry or entries in the ALAT are invalidated.

In the complete_form, all ALAT entries are invalidated. In the entry_form, the ALAT is queried using the general register specifier *r_l* (gr_form), or the floating-point register specifier *f_l* (fr_form), and if any ALAT entry matches, it is invalidated.

Operation:

```
if (PR[qp]) {
    if (complete_form)
        alat_inval();
    else { // entry_form
        if (gr_form)
            alat_inval_single_entry(GENERAL, rl);
        else // fr_form
            alat_inval_single_entry(FLOAT, fl);
    }
}
```

Load

Format:	(qp) ldsz.ldtype.ldhint $r_1 = [r_3]$	no_base_update_form	M1
	(qp) ldsz.ldtype.ldhint $r_1 = [r_3], r_2$	reg_base_update_form	M2
	(qp) ldsz.ldtype.ldhint $r_1 = [r_3], imm_9$	imm_base_update_form	M3
	(qp) ld8.fill.ldhint $r_1 = [r_3]$	fill_form, no_base_update_form	M1
	(qp) ld8.fill.ldhint $r_1 = [r_3], r_2$	fill_form, reg_base_update_form	M2
	(qp) ld8.fill.ldhint $r_1 = [r_3], imm_9$	fill_form, imm_base_update_form	M3

Description: A value consisting of *sz* bytes is read from memory starting at the address specified by the value in GR r_3 . The value is then zero extended and placed in GR r_1 . The values of the *sz* completer are given in Table 6-26. The NaT bit corresponding to GR r_1 is cleared, except as described below for speculative loads. The *ldtype* completer specifies special load operations, which are described in Table 6-27.

For the fill_form, an 8-byte value is loaded, and a bit in the UNAT application register is copied into the target register NaT bit. This instruction is used for reloading a spilled register/NaT pair. See “Control Speculation” on page 4-10 for details.

In the base update forms, the value in GR r_3 is added to either a signed immediate value (*imm₉*) or a value from GR r_2 , and the result is placed back in GR r_3 . This base register update is done after the load, and does not affect the load address. In the reg_base_update_form, if the NaT bit corresponding to GR r_2 is set, then the NaT bit corresponding to GR r_3 is set and no fault is raised.

Table 6-26. sz Completers

sz Completer	Bytes Accessed
1	1 byte
2	2 bytes
4	4 bytes
8	8 bytes

Table 6-27. Load Types

<i>ldtype</i> Completer	Interpretation	Special Load Operation
<i>none</i>	Normal load	
s	Speculative load	Certain exceptions may be deferred rather than generating a fault. Deferral causes the target register’s NaT bit to be set. The NaT bit is later used to detect deferral.
a	Advanced load	An entry is added to the ALAT. This allows later instructions to check for colliding stores. If the referenced data page has a non-speculative attribute, the target register and NaT bit is cleared, and the processor ensures that no ALAT entry exists for the target register. The absence of an ALAT entry is later used to detect deferral or collision.
sa	Speculative Advanced load	An entry is added to the ALAT, and certain exceptions may be deferred. Deferral causes the target register’s NaT bit to be set, and the processor ensures that no ALAT entry exists for the target register. The absence of an ALAT entry is later used to detect deferral or collision.
c.nc	Check load - no clear	The ALAT is searched for a matching entry. If found, no load is done and the target register is unchanged. Regardless of ALAT hit or miss, base register updates are performed, if specified. An implementation may optionally cause the ALAT lookup to fail independent of whether an ALAT entry matches. If not found, a load is performed, and an entry is added to the ALAT (unless the referenced data page has a non-speculative attribute, in which case no ALAT entry is allocated).

Table 6-27. Load Types (Continued)

<i>ldtype</i> Completer	Interpretation	Special Load Operation
c.clr	Check load - clear	The ALAT is searched for a matching entry. If found, the entry is removed, no load is done and the target register is unchanged. Regardless of ALAT hit or miss, base register updates are performed, if specified. An implementation may optionally cause the ALAT lookup to fail independent of whether an ALAT entry matches. If not found, a clear check load behaves like a normal load.
c.clr.acq	Ordered check load – clear	This type behaves the same as the unordered clear form, except that the ALAT lookup (and resulting load, if no ALAT entry is found) is performed with acquire semantics.
acq	Ordered load	An ordered load is performed with acquire semantics.
bias	Biased load	A hint is provided to the implementation to acquire exclusive ownership of the accessed cache line.

For more details on ordered, biased, speculative, advanced and check loads see “Control Speculation” on page 4-10 and “Data Speculation” on page 4-12. For more details on ordered loads see “Memory Access Ordering” on page 4-18. See “Memory Hierarchy Control and Consistency” on page 4-16 for details on biased loads.

For the non-speculative load types, if NaT bit associated with GR r_3 is 1, a Register NaT Consumption fault is taken. For speculative and speculative advanced loads, no fault is raised, and the exception is deferred. For the base-update calculation, if the NaT bit associated with GR r_2 is 1, the NaT bit associated with GR r_3 is set to 1 and no fault is raised.

The value of the *ldhint* completer specifies the locality of the memory access. The values of the *ldhint* completer are given in Table 6-28. A prefetch hint is implied in the base update forms. The address specified by the value in GR r_3 after the base update acts as a hint to prefetch the indicated cache line. This prefetch uses the locality hints specified by *ldhint*. Prefetch and locality hints do not affect program functionality and may be ignored by the implementation. See “Memory Hierarchy Control and Consistency” on page 4-16 for details.

Table 6-28. Load Hints

<i>ldhint</i> Completer	Interpretation
<i>none</i>	Temporal locality, level 1
nt1	No temporal locality, level 1
nta	No temporal locality, all levels

In the no_base_update form, the value in GR r_3 is not modified and no prefetch hint is implied.

For the base update forms, specifying the same register address in r_1 and r_3 will cause an Illegal Operation fault.


```

Operation:  if (PR[qp]) {
            size = fill_form ? 8 : sz;

            speculative = (ldtype == 's' || ldtype == 'sa');
            advanced = (ldtype == 'a' || ldtype == 'sa');
            check_clear = (ldtype == 'c.clr' || ldtype == 'c.clr.acq');
            check_no_clear = (ldtype == 'c.nc');
            check = check_clear || check_no_clear;
            acquire = (ldtype == 'acq' || ldtype == 'c.clr.acq');
            bias = (ldtype == 'bias') ? BIAS : 0 ;

            itype = READ;
            if (speculative) itype |= SPEC ;
            if (advanced) itype |= ADVANCE ;

            if ((reg_base_update_form || imm_base_update_form) && (r1 == r3))
                illegal_operation_fault();
            check_target_register(r1, itype);
            if (reg_base_update_form || imm_base_update_form)
                check_target_register(r3);

            if (reg_base_update_form) {
                tmp_r2 = GR[r2];
                tmp_r2nat = GR[r2].nat;
            }

            if (!speculative && GR[r3].nat) // fault on NaT address
                register_nat_consumption_fault(itype);
            defer = speculative && (GR[r3].nat || PSR.ed); // defer exception if spec

            if (check && alat_cmp(GENERAL, r1)) { // no load on ld.c & ALAT hit
                if (check_clear) // remove entry on ld.c.clr or ld.c.clr.acq
                    alat_inval_single_entry(GENERAL, r1);
            } else {
                if (!defer) {
                    paddr = tlb_translate(GR[r3], size, itype, PSR.cpl, &mattr,
                                         &defer);

                    if (!defer) {
                        otype = acquire ? ACQUIRE : UNORDERED;
                        val = mem_read(paddr, size, UM.be, mattr, otype, bias | ldhint);
                    }
                }
                if (check_clear || advanced) // remove any old ALAT entry
                    alat_inval_single_entry(GENERAL, r1);
                if (defer) {
                    if (speculative) {
                        GR[r1] = natd_gr_read(paddr, size, UM.be, mattr, otype,
                                              bias | ldhint);

                        GR[r1].nat = 1;
                    } else {
                        GR[r1] = 0; // ld.a to sequential memory
                        GR[r1].nat = 0;
                    }
                } else { // execute load normally
                    if (fill_form) // fill NaT on ld8.fill
                        bit_pos = GR[r3]{8:3};
                    GR[r1] = val;
                    GR[r1].nat = AR[UNAT]{bit_pos};
                } else { // clear NaT on other types
                    GR[r1] = zero_ext(val, size * 8);
                    GR[r1].nat = 0;
                }
                if ((check_no_clear || advanced) && ma_is_speculative(mattr))
                    // add entry to ALAT
                    alat_write(GENERAL, r1, paddr, size);
            }
        }
    }

```

```
if (imm_base_update_form) { // update base register
    GR[r3] = GR[r3] + sign_ext(imm9, 9);
    GR[r3].nat = GR[r3].nat;
} else if (reg_base_update_form) {
    GR[r3] = GR[r3] + tmp_r2;
    GR[r3].nat = GR[r3].nat || tmp_r2nat;
}

if ((reg_base_update_form || imm_base_update_form) && !GR[r3].nat)
    mem_implicit_prefetch(GR[r3], bias | ldhint);
}
```

Floating-Point Load

Format:	(qp) <i>ldfsz.fldtype.ldhint</i> $f_1 = [r_3]$	no_base_update_form	M6
	(qp) <i>ldfsz.fldtype.ldhint</i> $f_1 = [r_3], r_2$	reg_base_update_form	M7
	(qp) <i>ldfsz.fldtype.ldhint</i> $f_1 = [r_3], imm_9$	imm_base_update_form	M8
	(qp) <i>ldf8.fldtype.ldhint</i> $f_1 = [r_3]$	integer_form, no_base_update_form	M6
	(qp) <i>ldf8.fldtype.ldhint</i> $f_1 = [r_3], r_2$	integer_form, reg_base_update_form	M7
	(qp) <i>ldf8.fldtype.ldhint</i> $f_1 = [r_3], imm_9$	integer_form, imm_base_update_form	M8
	(qp) <i>ldf.fill.ldhint</i> $f_1 = [r_3]$	fill_form, no_base_update_form	M6
	(qp) <i>ldf.fill.ldhint</i> $f_1 = [r_3], r_2$	fill_form, reg_base_update_form	M7
	(qp) <i>ldf.fill.ldhint</i> $f_1 = [r_3], imm_9$	fill_form, imm_base_update_form	M8

Description: A value consisting of *fsz* bytes is read from memory starting at the address specified by the value in GR r_3 . The value is then converted into the floating-point register format and placed in FR f_1 . See “Data Types and Formats” on page 5-1 for details on conversion to floating-point register format. The values of the *fsz* completer are given in Table 6-29. The *fldtype* completer specifies special load operations, which are described in Table 6-30.

For the *integer_form*, an 8-byte value is loaded and placed in the significand field of FR f_1 without conversion. The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).

For the *fill_form*, a 16-byte value is loaded, and the appropriate fields are placed in FR f_1 without conversion. This instruction is used for reloading a spilled register. See “Control Speculation” on page 4-10 for details.

In the base update forms, the value in GR r_3 is added to either a signed immediate value (*imm₉*) or a value from GR r_2 , and the result is placed back in GR r_3 . This base register update is done after the load, and does not affect the load address. In the *reg_base_update_form*, if the NaT bit corresponding to GR r_2 is set, then the NaT bit corresponding to GR r_3 is set and no fault is raised.

Table 6-29. *fsz* Completers

<i>fsz</i> Completer	Bytes Accessed	Memory Format
s	4 bytes	Single precision
d	8 bytes	Double precision
e	10 bytes	Extended precision

Table 6-30. FP Load Types

<i>fldtype</i> Completer	Interpretation	Special Load Operation
<i>none</i>	Normal load	
s	Speculative load	Certain exceptions may be deferred rather than generating a fault. Deferral causes NaTVal to be placed in the target register. The NaTVal value is later used to detect deferral.
a	Advanced load	An entry is added to the ALAT. This allows later instructions to check for colliding stores. If the referenced data page has a non-speculative attribute, no ALAT entry is added to the ALAT and the target register is set as follows: for the <i>integer_form</i> , the exponent is set to 0x1003E and the sign and significand are set to zero; for all other forms, the sign, exponent and significand are set to zero. The absence of an ALAT entry is later used to detect deferral or collision.
sa	Speculative Advanced load	An entry is added to the ALAT, and certain exceptions may be deferred. Deferral causes NaTVal to be placed in the target register, and the processor ensures that no ALAT entry exists for the target register. The absence of an ALAT entry is later used to detect deferral or collision.

Table 6-30. FP Load Types (Continued)

<i>fldtype</i> Completer	Interpretation	Special Load Operation
c.nc	Check load - no clear	The ALAT is searched for a matching entry. If found, no load is done and the target register is unchanged. Regardless of ALAT hit or miss, base register updates are performed, if specified. An implementation may optionally cause the ALAT lookup to fail independent of whether an ALAT entry matches. If not found, a load is performed, and an entry is added to the ALAT (unless the referenced data page has a non-speculative attribute, in which case no ALAT entry is allocated).
c.clr	Check load – clear	The ALAT is searched for a matching entry. If found, the entry is removed, no load is done and the target register is unchanged. Regardless of ALAT hit or miss, base register updates are performed, if specified. An implementation may optionally cause the ALAT lookup to fail independent of whether an ALAT entry matches. If not found, a clear check load behaves like a normal load.

For more details on speculative, advanced and check loads see “Control Speculation” on page 4-10 and “Data Speculation” on page 4-12.

For the non-speculative load types, if NaT bit associated with GR r_3 is 1, a Register NaT Consumption fault is taken. For speculative and speculative advanced loads, no fault is raised, and the exception is deferred. For the base-update calculation, if the NaT bit associated with GR r_2 is 1, the NaT bit associated with GR r_3 is set to 1 and no fault is raised.

The value of the *ldhint* modifier specifies the locality of the memory access. The mnemonic values of *ldhint* are given in Table 6-28 on page 6-102. A prefetch hint is implied in the base update forms. The address specified by the value in GR r_3 after the base update acts as a hint to prefetch the indicated cache line. This prefetch uses the locality hints specified by *ldhint*. Prefetch and locality hints do not affect program functionality and may be ignored by the implementation. See “Memory Hierarchy Control and Consistency” on page 4-16 for details.

In the *no_base_update* form, the value in GR r_3 is not modified and no prefetch hint is implied.

The PSR.mfl and PSR.mfh bits are updated to reflect the modification of FR f_1 .

```

Operation:  if (PR[qp]) {
            size = (fill_form ? 16 : (integer_form ? 8 : fsz));
            speculative = (fldtype == 's' || fldtype == 'sa');
            advanced = (fldtype == 'a' || fldtype == 'sa');
            check_clear = (fldtype == 'c.clr' );
            check_no_clear = (fldtype == 'c.nc');
            check = check_clear || check_no_clear;

            itype = READ;
            if (speculative) itype |= SPEC ;
            if (advanced) itype |= ADVANCE ;

            if (reg_base_update_form || imm_base_update_form)
                check_target_register(r3);
            fp_check_target_register(f1);
            if (tmp_isrkode = fp_reg_disabled(f1, 0, 0, 0))
                disabled_fp_register_fault(tmp_isrkode, itype);

            if (!speculative && GR[r3].nat) // fault on NaT address
                register_nat_consumption_fault(itype);

            defer = speculative && (GR[r3].nat || PSR.ed); // defer exception if spec

            if (check && alat_cmp(FLOAT, f1)) { // no load on ldf.c & ALAT hit
                if (check_clear) // remove entry on ldf.c.clr
                    alat_inval_single_entry(FLOAT, f1);
            } else {
                if (!defer) {
                    paddr = tlb_translate(GR[r3], size, itype, PSR.cpl, &mattr,
                                         &defer);
                    if (!defer)
                        val = mem_read(paddr, size, UM.be, mattr, UNORDERED, ldhint);
                }
                if (check_clear || advanced) // remove any old ALAT entry
                    alat_inval_single_entry(FLOAT, f1);
                if (speculative && defer) {
                    FR[f1] = NATVAL;
                } else if (advanced && !speculative && defer) {
                    FR[f1] = (integer_form ? FP_INT_ZERO : FP_ZERO);
                } else { // execute load normally
                    FR[f1] = fp_mem_to_fr_format(val, size, integer_form);

                    if ((check_no_clear || advanced) && ma_is_speculative(mattr))
                        // add entry to ALAT
                        alat_write(FLOAT, f1, paddr, size);
                }
            }
        }

        if (imm_base_update_form) { // update base register
            GR[r3] = GR[r3] + sign_ext(imm9, 9);
            GR[r3].nat = GR[r3].nat;
        } else if (reg_base_update_form) {
            GR[r3] = GR[r3] + GR[r2];
            GR[r3].nat = GR[r3].nat || GR[r2].nat;
        }

        if ((reg_base_update_form || imm_base_update_form) && !GR[r3].nat)
            mem_implicit_prefetch(GR[r3], ldhint);

        fp_update_psr(f1);
    }

```

Floating-Point Load Pair

Format:	<i>(qp)</i> <i>ldfps.fldtype.ldhint</i> $f_1, f_2 = [r_3]$	single_form, no_base_update_form	M11
	<i>(qp)</i> <i>ldfps.fldtype.ldhint</i> $f_1, f_2 = [r_3], 8$	single_form, base_update_form	M12
	<i>(qp)</i> <i>ldfpd.fldtype.ldhint</i> $f_1, f_2 = [r_3]$	double_form, no_base_update_form	M11
	<i>(qp)</i> <i>ldfpd.fldtype.ldhint</i> $f_1, f_2 = [r_3], 16$	double_form, base_update_form	M12
	<i>(qp)</i> <i>ldfp8.fldtype.ldhint</i> $f_1, f_2 = [r_3]$	integer_form, no_base_update_form	M11
	<i>(qp)</i> <i>ldfp8.fldtype.ldhint</i> $f_1, f_2 = [r_3], 16$	integer_form, base_update_form	M12

Description: Eight (single_form) or sixteen (double_form/integer_form) bytes are read from memory starting at the address specified by the value in GR r_3 . The value read is treated as a contiguous pair of floating-point numbers for the single_form/double_form and as integer/Parallel FP data for the integer_form. Each number is converted into the floating-point register format. The value at the lowest address is placed in FR f_1 , and the value at the highest address is placed in FR f_2 . See “Data Types and Formats” on page 5-1 for details on conversion to floating-point register format. The *fldtype* completer specifies special load operations, which are described in Table 6-30 on page 6-105.

For more details on speculative, advanced and check loads see “Control Speculation” on page 4-10 and “Data Speculation” on page 4-12.

For the non-speculative load types, if NaT bit associated with GR r_3 is 1, a Register NaT Consumption fault is taken. For speculative and speculative advanced loads, no fault is raised, and the exception is deferred.

In the base_update_form, the value in GR r_3 is added to an implied immediate value (equal to double the data size) and the result is placed back in GR r_3 . This base register update is done after the load, and does not affect the load address.

The value of the *ldhint* modifier specifies the locality of the memory access. The mnemonic values of *ldhint* are given in Table 6-28 on page 6-102. A prefetch hint is implied in the base update form. The address specified by the value in GR r_3 after the base update acts as a hint to prefetch the indicated cache line. This prefetch uses the locality hints specified by *ldhint*. Prefetch and locality hints do not affect program functionality and may be ignored by the implementation. See “Memory Hierarchy Control and Consistency” on page 4-16 for details.

In the no_base_update form, the value in GR r_3 is not modified and no prefetch hint is implied.

The PSR.mfl and PSR.mfh bits are updated to reflect the modification of FR f_1 and FR f_2 .

There is a restriction on the choice of target registers. Register specifiers f_1 and f_2 must specify one odd-numbered physical FR and one even-numbered physical FR. Specifying two odd or two even registers will cause an Illegal Operation fault to be raised. The restriction is on physical register numbers after register rotation. This means that if f_1 and f_2 both specify static registers or both specify rotating registers, then f_1 and f_2 must be odd/even or even/odd. If f_1 and f_2 specify one static and one rotating register, the restriction depends on CFM.rrb.fr. If CFM.rrb.fr is even, the restriction is the same; f_1 and f_2 must be odd/even or even/odd. If CFM.rrb.fr is odd, then f_1 and f_2 must be even/even or odd/odd. Specifying one static and one rotating register should only be done when CFM.rrb.fr will have a predictable value (such as 0).

Operation:

```

if (PR[qp]) {
    size = single_form ? 8 : 16;

    speculative = (fldtype == 's' || fldtype == 'sa');
    advanced = (fldtype == 'a' || fldtype == 'sa');
    check_clear = (fldtype == 'c.clr');
    check_no_clear = (fldtype == 'c.nc');
    check = check_clear || check_no_clear;

    itype = READ;
    if (speculative) itype |= SPEC;
    if (advanced) itype |= ADVANCE;

    if (fp_reg_bank_conflict(f1, f2))
        illegal_operation_fault();

```

```

if (base_update_form)
    check_target_register(r3);

fp_check_target_register(f1);
fp_check_target_register(f2);
if (tmp_israncode = fp_reg_disabled(f1, f2, 0, 0))
    disabled_fp_register_fault(tmp_israncode, itype);

if (!speculative && GR[r3].nat) // fault on NaT address
    register_nat_consumption_fault(itype);

defer = speculative && (GR[r3].nat || PSR.ed); // defer exception if spec

if (check && alat_cmp(FLOAT, f1)) { // no load on ldfp.c & ALAT hit
    if (check_clear) // remove entry on ldfp.c.clr
        alat_inval_single_entry(FLOAT, f1);
} else {
    if (!defer) {
        paddr = tlb_translate(GR[r3], size, itype, PSR.cpl, &mattr,
                               &defer);
        if (!defer)
            val = mem_read(paddr, size, UM.be, mattr, UNORDERED, ldhint);
    }
    if (check_clear || advanced) // remove any old ALAT entry
        alat_inval_single_entry(FLOAT, f1);
    if (speculative && defer) {
        FR[f1] = NATVAL;
        FR[f2] = NATVAL;
    } else if (advanced && !speculative && defer) {
        FR[f1] = (integer_form ? FP_INT_ZERO : FP_ZERO);
        FR[f2] = (integer_form ? FP_INT_ZERO : FP_ZERO);
    } else { // execute load normally
        if (UM.be) {
            FR[f1] = fp_mem_to_fr_format(val u>> (size/2*8), size/2,
                                         integer_form);
            FR[f2] = fp_mem_to_fr_format(val, size/2, integer_form);
        } else {
            FR[f1] = fp_mem_to_fr_format(val, size/2, integer_form);
            FR[f2] = fp_mem_to_fr_format(val u>> (size/2*8), size/2,
                                         integer_form);
        }
    }

    if ((check_no_clear || advanced) && ma_is_speculative(mattr))
        // add entry to ALAT
        alat_write(FLOAT, f1, paddr, size);
}
}

if (base_update_form) { // update base register
    GR[r3] = GR[r3] + size;
    GR[r3].nat = GR[r3].nat;
    if (!GR[r3].nat)
        mem_implicit_prefetch(GR[r3], ldhint);
}

fp_update_psr(f1);
fp_update_psr(f2);
}

```

Line Prefetch

Format:	<i>(qp)</i> lfetch. <i>lftype.lfhint</i> [<i>r</i> ₃]	no_base_update_form	M13
	<i>(qp)</i> lfetch. <i>lftype.lfhint</i> [<i>r</i> ₃], <i>r</i> ₂	reg_base_update_form	M14
	<i>(qp)</i> lfetch. <i>lftype.lfhint</i> [<i>r</i> ₃], <i>imm</i> ₉	imm_base_update_form	M15
	<i>(qp)</i> lfetch. <i>lftype.excl.lfhint</i> [<i>r</i> ₃]	no_base_update_form, exclusive_form	M13
	<i>(qp)</i> lfetch. <i>lftype.excl.lfhint</i> [<i>r</i> ₃], <i>r</i> ₂	reg_base_update_form, exclusive_form	M14
	<i>(qp)</i> lfetch. <i>lftype.excl.lfhint</i> [<i>r</i> ₃], <i>imm</i> ₉	imm_base_update_form, exclusive_form	M15

Description: The line containing the address specified by the value in GR *r*₃ is moved to the highest level of the data memory hierarchy. The value of the *lfhint* modifier specifies the locality of the memory access. The mnemonic values of *lfhint* are given in Table 6-32.

The behavior of the memory read is also determined by the memory attribute associated with the accessed page. Line size is implementation dependent but must be a power of two greater than or equal to 32 bytes. In the exclusive form, the cache line is allowed to be marked in an exclusive state. This qualifier is used when the program expects soon to modify a location in that line. If the memory attribute for the page containing the line is not cacheable, then no reference is made.

The completer, *lftype*, specifies whether or not the instruction raises faults normally associated with a regular load. Table 6-31 defines these two options.

Table 6-31. *lftype* Mnemonic Values

<i>lftype</i> Mnemonic	Interpretation
<i>none</i>	Ignore faults
<i>fault</i>	Raise faults

In the base update forms, after being used to address memory, the value in GR *r*₃ is incremented by either the sign extended value in *imm*₉ (in the *imm_base_update_form*) or the value in GR *r*₂ (in the *reg_base_update_form*). In the *reg_base_update_form*, if the NaT bit corresponding to GR *r*₂ is set, then the NaT bit corresponding to GR *r*₃ is set – no fault is raised.

In the *reg_base_update_form* and the *imm_base_update_form*, if the NaT bit corresponding to GR *r*₃ is clear, then the address specified by the value in GR *r*₃ after the post-increment acts as a hint to implicitly prefetch the indicated cache line. This implicit prefetch uses the locality hints specified by *lfhint*. The implicit prefetch does not affect program functionality, does not raise any faults, and may be ignored by the implementation.

In the *no_base_update_form*, the value in GR *r*₃ is not modified and no implicit prefetch hint is implied.

If the NaT bit corresponding to GR *r*₃ is set then the state of memory is not affected. In the *reg_base_update_form* and *imm_base_update_form*, the post increment of GR *r*₃ is performed and prefetch is hinted as described above.

Table 6-32. *lfhint* Mnemonic Values

<i>lfhint</i> Mnemonic	Interpretation
<i>none</i>	Temporal locality, level 1
<i>nt1</i>	No temporal locality, level 1
<i>nt2</i>	No temporal locality, level 2
<i>nta</i>	No temporal locality, all levels


```

Operation:  if (PR[qp]) {
            itype = READ|NON_ACCESS;
            itype |= (lftype == 'fault') ? LFETCH_FAULT : LFETCH;

            if (reg_base_update_form || imm_base_update_form)
                check_target_register(r3);

            if (lftype == 'fault') {                // faulting form
                if (GR[r3].nat && !PSR.ed)          // fault on NaT address
                    register_nat_consumption_fault(itype);
            }

            if (exclusive_form)
                excl_hint = EXCLUSIVE;
            else
                excl_hint = 0;

            if (!GR[r3].nat && !PSR.ed) { // faulting form already faulted if r3 is nat'ed
                paddr = tlb_translate(GR[r3], 1, itype, PSR.cpl, &mattr, &defer);
                if (!defer)
                    mem_promote(paddr, mattr, lfhint | excl_hint);
            }

            if (imm_base_update_form) {
                GR[r3] = GR[r3] + sign_ext(imm9, 9);
                GR[r3].nat = GR[r3].nat;
            } else if (reg_base_update_form) {
                GR[r3] = GR[r3] + GR[r2];
                GR[r3].nat = GR[r2].nat || GR[r3].nat;
            }

            if ((reg_base_update_form || imm_base_update_form) && !GR[r3].nat)
                mem_implicit_prefetch(GR[r3], lfhint | excl_hint);
        }

```

Memory Fence

Format: *(qp)* mf ordering_form M24
(qp) mf.a acceptance_form M24

Description: This instruction forces ordering between prior and subsequent memory accesses. The `ordering_form` ensures all prior data memory accesses are made visible prior to any subsequent data memory accesses being made visible. It does not ensure prior data memory references have been accepted by the external platform, nor that prior data memory references are visible.

The `acceptance_form` prevents any subsequent data memory accesses by the processor from initiating transactions to the external platform until:

- all prior loads have returned data, and
- all prior stores have been accepted by the external platform.

The definition of “acceptance” is platform dependent. The `acceptance_form` is typically used to ensure the processor has “waited” until a memory-mapped IO transaction has been “accepted”, before initiating additional external transactions. The `acceptance_form` does not ensure ordering.

Operation:

```
if (PR[qp]){
  if (acceptance_form)
    acceptance_fence();
  else
    ordering_fence();
}
```

Mix

Format:	<i>(qp)</i> mix1.l $r_1 = r_2, r_3$	one_byte_form, left_form	I2
	<i>(qp)</i> mix2.l $r_1 = r_2, r_3$	two_byte_form, left_form	I2
	<i>(qp)</i> mix4.l $r_1 = r_2, r_3$	four_byte_form, left_form	I2
	<i>(qp)</i> mix1.r $r_1 = r_2, r_3$	one_byte_form, right_form	I2
	<i>(qp)</i> mix2.r $r_1 = r_2, r_3$	two_byte_form, right_form	I2
	<i>(qp)</i> mix4.r $r_1 = r_2, r_3$	four_byte_form, right_form	I2

Description: The data elements of GR r_2 and r_3 are mixed as shown in Figure 6-23, and the result placed in GR r_1 . The data elements in the source registers are grouped in pairs, and one element from each pair is selected for the result. In the left_form, the result is formed from the leftmost elements from each of the pairs. In the right_form, the result is formed from the rightmost elements. Elements are selected alternately from the two source registers.




Figure 6-23. Mix Example

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) { // one-byte elements
                x[0] = GR[r2]{7:0};      y[0] = GR[r3]{7:0};
                x[1] = GR[r2]{15:8};     y[1] = GR[r3]{15:8};
                x[2] = GR[r2]{23:16};   y[2] = GR[r3]{23:16};
                x[3] = GR[r2]{31:24};   y[3] = GR[r3]{31:24};
                x[4] = GR[r2]{39:32};   y[4] = GR[r3]{39:32};
                x[5] = GR[r2]{47:40};   y[5] = GR[r3]{47:40};
                x[6] = GR[r2]{55:48};   y[6] = GR[r3]{55:48};
                x[7] = GR[r2]{63:56};   y[7] = GR[r3]{63:56};

                if (left_form)
                    GR[r1] = concatenate8(x[7], y[7], x[5], y[5],
                                             x[3], y[3], x[1], y[1]);
                else
                    GR[r1] = concatenate8(x[6], y[6], x[4], y[4],
                                             x[2], y[2], x[0], y[0]);
            } else if (two_byte_form) { // two-byte elements
                x[0] = GR[r2]{15:0};      y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};     y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};     y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};     y[3] = GR[r3]{63:48};

                if (left_form)
                    GR[r1] = concatenate4(x[3], y[3], x[1], y[1]);
                else
                    GR[r1] = concatenate4(x[2], y[2], x[0], y[0]);
            } else { // four-byte elements
                x[0] = GR[r2]{31:0};      y[0] = GR[r3]{31:0};
                x[1] = GR[r2]{63:32};     y[1] = GR[r3]{63:32};

                if (left_form)
                    GR[r1] = concatenate2(x[1], y[1]);
                else
                    GR[r1] = concatenate2(x[0], y[0]);
            }
            GR[r1].nat = GR[r2].nat || GR[r3].nat;
        }

```

Move Application Register

Format:	<i>(qp)</i> mov $r_1 = ar_3$	pseudo-op	
	<i>(qp)</i> mov $ar_3 = r_2$	pseudo-op	
	<i>(qp)</i> mov $ar_3 = imm_8$	pseudo-op	
	<i>(qp)</i> mov.i $r_1 = ar_3$	i_form, from_form	I28
	<i>(qp)</i> mov.i $ar_3 = r_2$	i_form, register_form, to_form	I26
	<i>(qp)</i> mov.i $ar_3 = imm_8$	i_form, immediate_form, to_form	I27
	<i>(qp)</i> mov.m $r_1 = ar_3$	m_form, from_form	M31
	<i>(qp)</i> mov.m $ar_3 = r_2$	m_form, register_form, to_form	M29
	<i>(qp)</i> mov.m $ar_3 = imm_8$	m_form, immediate_form, to_form	M30

Description: The source operand is copied to the destination register.

In the from_form, the application register specified by ar_3 is copied into GR r_1 and the corresponding NaT bit is cleared.

In the to_form, the value in GR r_2 (in the register_form), or the sign extended value in imm_8 (in the immediate_form), is placed in AR ar_3 . In the register_form if the NaT bit corresponding to GR r_2 is set, then a Register NaT Consumption fault is raised.

Only a subset of the application registers can be accessed by each execution unit (M or I). Table 3-3 on page 3-5 in indicates which application registers may be accessed from which execution unit type. An access to an application register from the wrong unit type causes an Illegal Operation fault.

This instruction has multiple forms with the pseudo operation eliminating the need for specifying the execution unit. Accesses of the ARs are always implicitly serialized. While implicitly serialized, read-after-write and write-after-write dependencies must be avoided (e.g., setting CCV, followed by `cmpxchg` in the same instruction group, or simultaneous writes to the UNAT register by `ld.fill` and `mov` to UNAT).

```

Operation:  if (PR[qp]) {
            tmp_type = (i_form ? AR_I_TYPE : AR_M_TYPE);
            if (is_reserved_reg(tmp_type, ar3))
                illegal_operation_fault();

            if (from_form) {
                check_target_register(r1);
                if ((ar3 == BSPSTORE) || (ar3 == RNAT)) && (AR[RSC].mode != 0)
                    illegal_operation_fault();

                if (ar3 == ITC && PSR.si && PSR.cpl != 0)
                    privileged_register_fault();

                GR[r1] = (is_ignored_reg(ar3)) ? 0 : AR[ar3];
                GR[r1].nat = 0;
            } else { // to_form
                tmp_val = (register_form) ? GR[r2] : sign_ext(imm8, 8);

                if (ar3 == BSP)
                    illegal_operation_fault();

                if ((ar3 == BSPSTORE) || (ar3 == RNAT)) && (AR[RSC].mode != 0)
                    illegal_operation_fault();

                if (register_form && GR[r2].nat)
                    register_nat_consumption_fault(0);

                if (is_reserved_field(AR_TYPE, ar3, tmp_val))
                    reserved_register_field_fault();

                if ((is_kernel_reg(ar3) || ar3 == ITC) && (PSR.cpl != 0))
                    privileged_register_fault();

                if (!is_ignored_reg(ar3)) {
                    tmp_val = ignored_field_mask(AR_TYPE, ar3, tmp_val);
                    // check for illegal promotion
                    if (ar3 == RSC && tmp_val{3:2} u< PSR.cpl)
                        tmp_val{3:2} = PSR.cpl;
                    AR[ar3] = tmp_val;

                    if (ar3 == BSPSTORE) {
                        AR[BSP] = rse_update_internal_stack_pointers(tmp_val);
                        AR[RNAT] = undefined();
                    }
                }
            }
        }
    }

```

Move Branch Register

Format:	<i>(qp)</i> mov $r_1 = b_2$	from_form	I22
	<i>(qp)</i> mov $b_1 = r_2$	to_form	I21
	<i>(qp)</i> mov.ret $b_1 = r_2$	return_form, to_form	I21

Description: The source operand is copied to the destination register.

In the *from_form*, the branch register specified by b_2 is copied into GR r_1 . The NaT bit corresponding to GR r_1 is cleared.

In the *to_form*, the value in GR r_2 is copied into BR b_1 . If the NaT bit corresponding to GR r_2 is 1, then a Register NaT Consumption fault is taken.

Operation:

```

if (PR[qp]) {
  if (from_form) {
    check_target_register( $r_1$ );
    GR[ $r_1$ ] = BR[ $b_2$ ];
    GR[ $r_1$ ].nat = 0;
  } else { // to_form
    if (GR[ $r_2$ ].nat)
      register_nat_consumption_fault(0);
    BR[ $b_1$ ] = GR[ $r_2$ ];
  }
}

```


Move Floating-Point Register

Format: $(qp) \text{ mov } f_1 = f_3$

pseudo-op of: $(qp) \text{ fmerge.s } f_1 = f_3, f_3$

Description: The value of FR f_3 is copied to FR f_1 .

Operation: See "Floating-Point Merge" on page 6-49.

Move General Register

Format: $(qp) \text{ mov } r_1 = r_3$

pseudo-op of: $(qp) \text{ adds } r_1 = 0, r_3$

Description: The value of GR r_3 is copied to GR r_1 .

Operation: See "Add" on page 6-3.

Move Immediate

Format: $(qp) \text{ mov } r_1 = imm_{22}$ pseudo-op of: $(qp) \text{ addl } r_1 = imm_{22}, r_0$

Description: The immediate value, imm_{22} , is sign extended to 64 bits and placed in GR r_1 .

Operation: See "Add" on page 6-3.

Move Indirect Register

Format: $(qp) \text{ mov } r_1 = \text{ireg}[r_3]$ from_form M43

Description: The source operand is copied to the destination register.

For move from indirect register, GR r_3 is read and the value used as an index into the register file specified by *ireg* (see Table 6-33 below). The indexed register is read and its value is copied into GR r_1 .

Table 6-33. Indirect Register File Mnemonics

<i>ireg</i>	Register File
cpuid	Processor Identification Register
pmd	Performance Monitor Data Register

Bits {7:0} of GR r_3 are used as the index. The remainder of the bits are ignored.

Apart from the PMD register file, access of a non-existent register results in a Reserved Register/Field fault. All accesses to the implementation-dependent portion of the PMD register file result in implementation dependent behavior but do not fault.

Operation:

```

if (PR[qp]) {
    tmp_index = GR[r3]{7:0};

    if (from_form) {
        check_target_register(r1);

        if (GR[r3].nat)
            register_nat_consumption_fault(0);

        if (is_reserved_reg(ireg, tmp_index))
            reserved_register_field_fault();

        if (ireg == PMD_TYPE) {
            GR[r1] = pmd_read(tmp_index);
        } else
            switch (ireg) {
                case CPUID_TYPE: GR[r1] = CPUID[tmp_index]; break;
            }
        GR[r1].nat = 0;
    }
}

```

Move Instruction Pointer

Format: (*qp*) mov *r_I* = ip

I25

Description: The Instruction Pointer (IP) for the bundle containing this instruction is copied into GR *r_I*.

Operation: if (PR[*qp*]) {
 check_target_register(*r_I*);

 GR[*r_I*] = IP;
 GR[*r_I*].nat = 0;
 }

Move Predicates

Format:	(qp) mov r ₁ = pr	from_form	I25
	(qp) mov pr = r ₂ , mask ₁₇	to_form	I23
	(qp) mov pr.rot = imm ₄₄	to_rotate_form	I24

Description: The source operand is copied to the destination register.

For moving the predicates to a GR, PR *i* is copied to bit position *i* within GR *r₁*.

For moving to the predicates, the source can either be a general register, or an immediate value. In the *to_form*, the source operand is GR *r₂* and only those predicates specified by the immediate value *mask₁₇* are written. The value *mask₁₇* is encoded in the instruction in an *imm₁₆* field such that: *imm₁₆* = *mask₁₇* >> 1. Predicate register 0 is always one. The *mask₁₇* value is sign extended. The most significant bit of *mask₁₇*, therefore, is the mask bit for all of the rotating predicates. If there is a deferred exception for GR *r₂* (the NaT bit is 1), a Register NaT Consumption fault is taken.

In the *to_rotate_form*, only the 48 rotating predicates can be written. The source operand is taken from the *imm₄₄* operand (which is encoded in the instruction in an *imm₂₈* field, such that: *imm₂₈* = *imm₄₄* >> 16). The low 16-bits correspond to the static predicates. The immediate is sign extended to set the top 21 predicates. Bit position *i* in the source operand is copied to PR *i*.

This instruction operates as if the predicate rotation base in the Current Frame Marker (CFM.rrb.pr) were zero.

Operation:

```

if (PR[qp]) {
    if (from_form) {
        check_target_register(r1);
        GR[r1] = 1; // PR[0] is always 1
        for (i = 1; i <= 63; i++) {
            GR[r1]{i} = PR[pr_phys_to_virt(i)];
        }
        GR[r1].nat = 0;
    } else if (to_form) {
        if (GR[r2].nat)
            register_nat_consumption_fault(0);
        tmp_src = sign_ext(mask17, 17);
        for (i = 1; i <= 63; i++) {
            if (tmp_src{i})
                PR[pr_phys_to_virt(i)] = GR[r2]{i};
        }
    } else { // to_rotate_form
        tmp_src = sign_ext(imm44, 44);
        for (i = 16; i <= 63; i++) {
            PR[pr_phys_to_virt(i)] = tmp_src{i};
        }
    }
}

```


Move Long Immediate

Format: (qp) movl $r_1 = imm_{64}$ X2

Description: The immediate value imm_{64} is copied to GR r_1 . The L slot of the bundle contains 41 bits of imm_{64} .

Operation:

```
if (PR[qp]) {
    check_target_register( $r_1$ );

    GR[ $r_1$ ] =  $imm_{64}$ ;
    GR[ $r_1$ ].nat = 0;
}
```


Mux

Format: (qp) mux1 $r_1 = r_2, mbtype_4$ one_byte_form 13
 (qp) mux2 $r_1 = r_2, mhbyte_8$ two_byte_form 14

Description: A permutation is performed on the packed elements in a single source register, GR r_2 , and the result is placed in GR r_1 . For 8-bit elements, only some of all possible permutations can be specified. The five possible permutations are given in Table 6-34 and shown in Figure 6-24.

Table 6-34. Mux Permutations for 8-bit Elements

$mbtype_4$	Function
@rev	Reverse the order of the bytes
@mix	Perform a Mix operation on the two halves of GR r_2
@shuf	Perform a Shuffle operation on the two halves of GR r_2
@alt	Perform an Alternate operation on the two halves of GR r_2
@brcst	Perform a Broadcast operation on the least significant byte of GR r_2

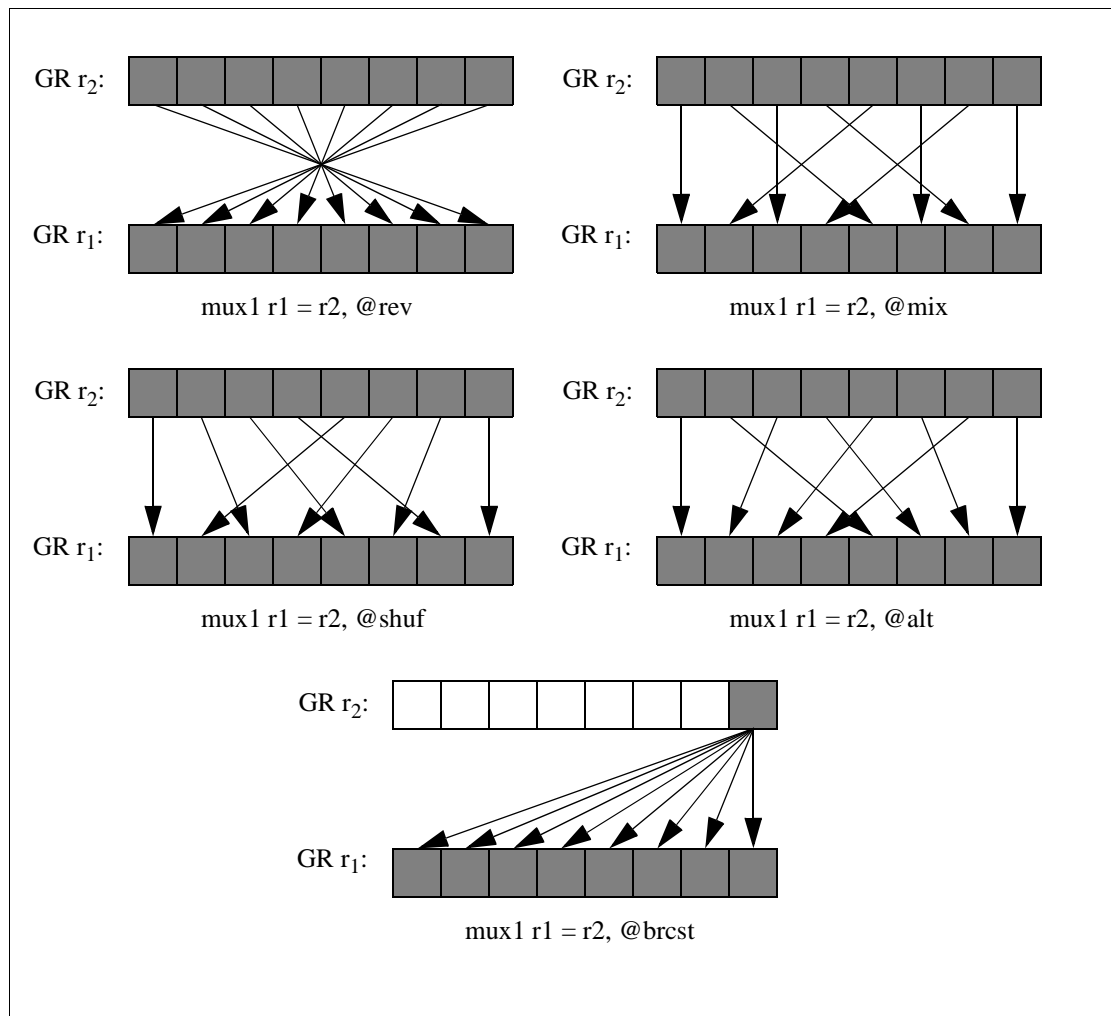


Figure 6-24. Mux1 Operation (8-bit elements)

For 16-bit elements, all possible permutations, with and without repetitions can be specified. They are expressed with an 8-bit $mhbyte_8$ field, which encodes the indices of the four 16-bit data elements. The indexed 16-bit elements of GR r_2 are copied to corresponding 16-bit positions in the target register GR r_1 . The indices are encoded in little-endian order. (The 8 bits of $mhbyte_8[7:0]$ are grouped in pairs of bits and named $mhbyte_8[3]$, $mhbyte_8[2]$, $mhbyte_8[1]$, $mhbyte_8[0]$ in the operation section).

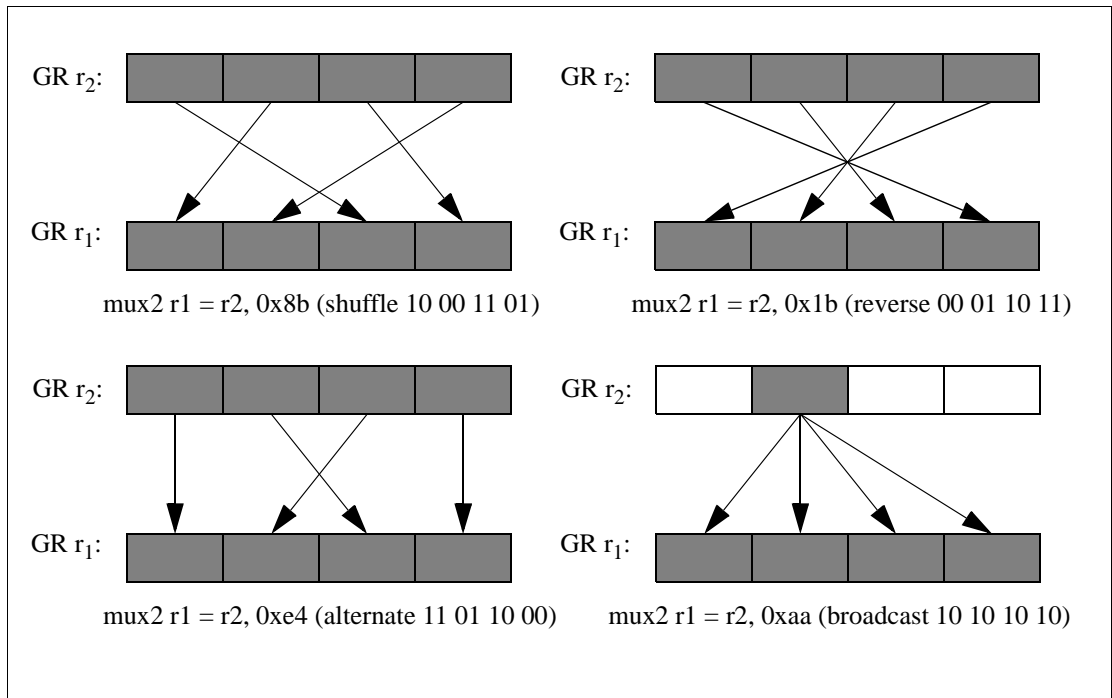


Figure 6-25. Mux2 Examples (16-bit elements)

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) {
                x[0] = GR[r2]{7:0};
                x[1] = GR[r2]{15:8};
                x[2] = GR[r2]{23:16};
                x[3] = GR[r2]{31:24};
                x[4] = GR[r2]{39:32};
                x[5] = GR[r2]{47:40};
                x[6] = GR[r2]{55:48};
                x[7] = GR[r2]{63:56};

                switch (mbtype) {
                    case '@rev':
                        GR[r1] = concatenate8(x[0], x[1], x[2], x[3],
                                                x[4], x[5], x[6], x[7]);
                        break;

                    case '@mix':
                        GR[r1] = concatenate8(x[7], x[3], x[5], x[1],
                                                x[6], x[2], x[4], x[0]);
                        break;

                    case '@shuf':
                        GR[r1] = concatenate8(x[7], x[3], x[6], x[2],
                                                x[5], x[1], x[4], x[0]);
                        break;

                    case '@alt':
                        GR[r1] = concatenate8(x[7], x[5], x[3], x[1],
                                                x[6], x[4], x[2], x[0]);
                        break;

                    case '@brcst':
                        GR[r1] = concatenate8(x[0], x[0], x[0], x[0],
                                                x[0], x[0], x[0], x[0]);
                        break;
                }
            } else { // two_byte_form
                x[0] = GR[r2]{15:0};
                x[1] = GR[r2]{31:16};
                x[2] = GR[r2]{47:32};
                x[3] = GR[r2]{63:48};

                res[0] = x[mhbyte8{1:0}];
                res[1] = x[mhbyte8{3:2}];
                res[2] = x[mhbyte8{5:4}];
                res[3] = x[mhbyte8{7:6}];

                GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
            }
            GR[r1].nat = GR[r2].nat;
        }
    }

```

No Operation

Format:	<i>(qp)</i> nop <i>imm</i> ₂₁	pseudo-op	
	<i>(qp)</i> nop.i <i>imm</i> ₂₁	i_unit_form	I19
	<i>(qp)</i> nop.b <i>imm</i> ₂₁	b_unit_form	B9
	<i>(qp)</i> nop.m <i>imm</i> ₂₁	m_unit_form	M37
	<i>(qp)</i> nop.f <i>imm</i> ₂₁	f_unit_form	F15
	<i>(qp)</i> nop.x <i>imm</i> ₆₂	x_unit_form	X1

Description: No operation is done.

The immediate, *imm*₂₁ or *imm*₆₂, can be used by software as a marker in program code. It is ignored by hardware.

For the *x_unit_form*, the L slot of the bundle contains the upper 41 bits of *imm*₆₂.

This instruction has five forms, each of which can be executed only on a particular execution unit type. The pseudo-op can be used if the unit type to execute on is unimportant.

Operation:

```

if (PR[qp]) {
    ; // no operation
}

```


Pack

Format: (qp) pack2.sss $r_1 = r_2, r_3$ two_byte_form, signed_saturation_form I2
 (qp) pack2.uss $r_1 = r_2, r_3$ two_byte_form, unsigned_saturation_form I2
 (qp) pack4.sss $r_1 = r_2, r_3$ four_byte_form, signed_saturation_form I2

Description: 32-bit or 16-bit elements from GR r_2 and GR r_3 are converted into 16-bit or 8-bit elements respectively, and the results are placed GR r_1 . The source elements are treated as signed values. If a source element cannot be represented in the result element, then saturation clipping is performed. The saturation can either be signed or unsigned. If an element is larger than the upper limit value, the result is the upper limit value. If it is smaller than the lower limit value, the result is the lower limit value. The saturation limits are given in Table 6-35.

Table 6-35. Pack Saturation Limits

Size	Source Element Width	Result Element Width	Saturation	Upper Limit	Lower Limit
2	16 bit	8 bit	signed	0x7f	0x80
2	16 bit	8 bit	unsigned	0xff	0x00
4	32 bit	16 bit	signed	0x7fff	0x8000




Figure 6-26. Pack Operation

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (two_byte_form) {
                if (signed_saturation_form) {
                    max = sign_ext(0x7f, 8);
                    min = sign_ext(0x80, 8);
                } else {
                    max = 0xff;
                    min = 0x00;
                }
                temp[0] = sign_ext(GR[r2]{15:0}, 16);
                temp[1] = sign_ext(GR[r2]{31:16}, 16);
                temp[2] = sign_ext(GR[r2]{47:32}, 16);
                temp[3] = sign_ext(GR[r2]{63:48}, 16);
                temp[4] = sign_ext(GR[r3]{15:0}, 16);
                temp[5] = sign_ext(GR[r3]{31:16}, 16);
                temp[6] = sign_ext(GR[r3]{47:32}, 16);
                temp[7] = sign_ext(GR[r3]{63:48}, 16);

                for (i = 0; i < 8; i++) {
                    if (temp[i] > max)
                        temp[i] = max;

                    if (temp[i] < min)
                        temp[i] = min;
                }

                GR[r1] = concatenate8(temp[7], temp[6], temp[5], temp[4],
                                       temp[3], temp[2], temp[1], temp[0]);
            } else {
                max = sign_ext(0x7fff, 16);
                min = sign_ext(0x8000, 16);
                temp[0] = sign_ext(GR[r2]{31:0}, 32);
                temp[1] = sign_ext(GR[r2]{63:32}, 32);
                temp[2] = sign_ext(GR[r3]{31:0}, 32);
                temp[3] = sign_ext(GR[r3]{63:32}, 32);

                for (i = 0; i < 4; i++) {
                    if (temp[i] > max)
                        temp[i] = max;

                    if (temp[i] < min)
                        temp[i] = min;
                }

                GR[r1] = concatenate4(temp[3], temp[2], temp[1], temp[0]);
            }
            GR[r1].nat = GR[r2].nat || GR[r3].nat;
        }

```

Parallel Add

Format:	(qp) padd1 $r_1 = r_2, r_3$	one_byte_form, modulo_form	A9
	(qp) padd1.sss $r_1 = r_2, r_3$	one_byte_form, sss_saturation_form	A9
	(qp) padd1.uus $r_1 = r_2, r_3$	one_byte_form, uus_saturation_form	A9
	(qp) padd1.uuu $r_1 = r_2, r_3$	one_byte_form, uuu_saturation_form	A9
	(qp) padd2 $r_1 = r_2, r_3$	two_byte_form, modulo_form	A9
	(qp) padd2.sss $r_1 = r_2, r_3$	two_byte_form, sss_saturation_form	A9
	(qp) padd2.uus $r_1 = r_2, r_3$	two_byte_form, uus_saturation_form	A9
	(qp) padd2.uuu $r_1 = r_2, r_3$	two_byte_form, uuu_saturation_form	A9
	(qp) padd4 $r_1 = r_2, r_3$	four_byte_form, modulo_form	A9

Description: The sets of elements from the two source operands are added, and the results placed in GR r_1 .

If a sum of two elements cannot be represented in the result element and a saturation completer is specified, then saturation clipping is performed. The saturation can either be signed or unsigned, as given in Table 6-36. If the sum of two elements is larger than the upper limit value, the result is the upper limit value. If it is smaller than the lower limit value, the result is the lower limit value. The saturation limits are given in Table 6-37.

Table 6-36. Parallel Add Saturation Completers

Completer	Result r_1 Treated as	Source r_2 Treated as	Source r_3 Treated as
sss	signed	signed	signed
uus	unsigned	unsigned	signed
uuu	unsigned	unsigned	unsigned

Table 6-37. Parallel Add Saturation Limits

Size	Element Width	Result r_1 Signed		Result r_1 Unsigned	
		Upper Limit	Lower Limit	Upper Limit	Lower Limit
1	8 bit	0x7f	0x80	0xff	0x00
2	16 bit	0x7fff	0x8000	0xffff	0x0000




Figure 6-27. Parallel Add Examples


```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) { // one-byte elements
                x[0] = GR[r2]{7:0};    y[0] = GR[r3]{7:0};
                x[1] = GR[r2]{15:8};   y[1] = GR[r3]{15:8};
                x[2] = GR[r2]{23:16};  y[2] = GR[r3]{23:16};
                x[3] = GR[r2]{31:24};  y[3] = GR[r3]{31:24};
                x[4] = GR[r2]{39:32};  y[4] = GR[r3]{39:32};
                x[5] = GR[r2]{47:40};  y[5] = GR[r3]{47:40};
                x[6] = GR[r2]{55:48};  y[6] = GR[r3]{55:48};
                x[7] = GR[r2]{63:56};  y[7] = GR[r3]{63:56};

                if (sss_saturation_form) { // sss_saturation_form
                    max = sign_ext(0x7f, 8);
                    min = sign_ext(0x80, 8);

                    for (i = 0; i < 8; i++) {
                        temp[i] = sign_ext(x[i], 8) + sign_ext(y[i], 8);
                    }
                } else if (uus_saturation_form) { // uus_saturation_form
                    max = 0xff;
                    min = 0x00;

                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) + sign_ext(y[i], 8);
                    }
                } else if (uuu_saturation_form) { // uuu_saturation_form
                    max = 0xff;
                    min = 0x00;

                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) + zero_ext(y[i], 8);
                    }
                } else { // modulo_form
                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) + zero_ext(y[i], 8);
                    }
                }

                if (sss_saturation_form || uus_saturation_form || uuu_saturation_form) {
                    for (i = 0; i < 8; i++) {
                        if (temp[i] > max)
                            temp[i] = max;

                        if (temp[i] < min)
                            temp[i] = min;
                    }
                }

                GR[r1] = concatenate8(temp[7], temp[6], temp[5], temp[4],
                                     temp[3], temp[2], temp[1], temp[0]);
            } else if (two_byte_form) { // 2-byte elements
                x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};   y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};   y[3] = GR[r3]{63:48};

                if (sss_saturation_form) { // sss_saturation_form
                    max = sign_ext(0x7fff, 16);
                    min = sign_ext(0x8000, 16);

                    for (i = 0; i < 4; i++) {
                        temp[i] = sign_ext(x[i], 16) + sign_ext(y[i], 16);
                    }
                } else if (uus_saturation_form) { // uus_saturation_form
                    max = 0xffff;
                    min = 0x0000;
                }
            }
        }

```

```

        for (i = 0; i < 4; i++) {
            temp[i] = zero_ext(x[i], 16) + sign_ext(y[i], 16);
        }
    } else if (uuu_saturation_form) { // uuu_saturation_form
        max = 0xffff;
        min = 0x0000;

        for (i = 0; i < 4; i++) {
            temp[i] = zero_ext(x[i], 16) + zero_ext(y[i], 16);
        }
    } else { // modulo_form
        for (i = 0; i < 4; i++) {
            temp[i] = zero_ext(x[i], 16) + zero_ext(y[i], 16);
        }
    }

    if (sss_saturation_form || uus_saturation_form || uuu_saturation_form) {
        for (i = 0; i < 4; i++) {
            if (temp[i] > max)
                temp[i] = max;

            if (temp[i] < min)
                temp[i] = min;
        }
    }
    GR[r1] = concatenate4(temp[3], temp[2], temp[1], temp[0]);
} else { // four-byte elements
    x[0] = GR[r2]{31:0};    y[0] = GR[r3]{31:0};
    x[1] = GR[r2]{63:32};  y[1] = GR[r3]{63:32};

    for (i = 0; i < 2; i++) { // modulo_form
        temp[i] = zero_ext(x[i], 32) + zero_ext(y[i], 32);
    }

    GR[r1] = concatenate2(temp[1], temp[0]);
}
GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Parallel Average

Format:	(qp) pavg1 $r_1 = r_2, r_3$	normal_form, one_byte_form	A9
	(qp) pavg1.raz $r_1 = r_2, r_3$	raz_form, one_byte_form	A9
	(qp) pavg2 $r_1 = r_2, r_3$	normal_form, two_byte_form	A9
	(qp) pavg2.raz $r_1 = r_2, r_3$	raz_form, two_byte_form	A9

Description: The unsigned data elements of GR r_2 are added to the unsigned data elements of GR r_3 . The results of the add are then each independently shifted to the right by one bit position. The high-order bits of each element are filled with the carry bits of the sums. To prevent cumulative round-off errors, an averaging is performed. The unsigned results are placed in GR r_1 .

The averaging operation works as follows. In the normal_form, the low-order bit of each result is set to 1 if at least one of the two least significant bits of the corresponding sum is 1. In the raz_form, the average rounds away from zero by adding 1 to each of the sums.

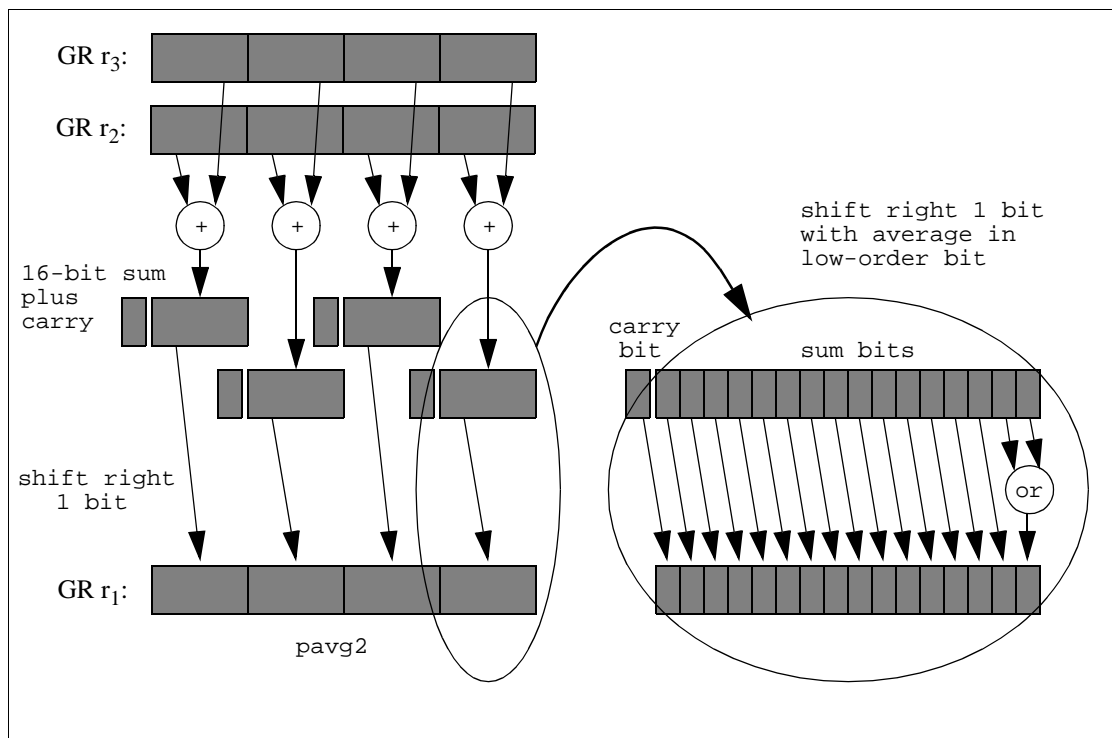


Figure 6-28. Parallel Average Example




Figure 6-29. Parallel Average with Round Away from Zero Example

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) { // one_byte_form
                x[0] = GR[r2]{7:0};    y[0] = GR[r3]{7:0};
                x[1] = GR[r2]{15:8};   y[1] = GR[r3]{15:8};
                x[2] = GR[r2]{23:16};  y[2] = GR[r3]{23:16};
                x[3] = GR[r2]{31:24};  y[3] = GR[r3]{31:24};
                x[4] = GR[r2]{39:32};  y[4] = GR[r3]{39:32};
                x[5] = GR[r2]{47:40};  y[5] = GR[r3]{47:40};
                x[6] = GR[r2]{55:48};  y[6] = GR[r3]{55:48};
                x[7] = GR[r2]{63:56};  y[7] = GR[r3]{63:56};

                if (raz_form) {
                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) + zero_ext(y[i], 8) + 1;
                        res[i] = shift_right_unsigned(temp[i], 1);
                    }
                } else { // normal form
                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) + zero_ext(y[i], 8);
                        res[i] = shift_right_unsigned(temp[i], 1) | (temp[i]{0});
                    }
                }
                GR[r1] = concatenate8(res[7], res[6], res[5], res[4],
                                       res[3], res[2], res[1], res[0]);
            } else { // two_byte_form
                x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};   y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};   y[3] = GR[r3]{63:48};

                if (raz_form) {
                    for (i = 0; i < 4; i++) {
                        temp[i] = zero_ext(x[i], 16) + zero_ext(y[i], 16) + 1;
                        res[i] = shift_right_unsigned(temp[i], 1);
                    }
                } else { // normal form
                    for (i = 0; i < 4; i++) {
                        temp[i] = zero_ext(x[i], 16) + zero_ext(y[i], 16);
                        res[i] = shift_right_unsigned(temp[i], 1) | (temp[i]{0});
                    }
                }
                GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
            }
            GR[r1].nat = GR[r2].nat || GR[r3].nat;
        }
    }

```

Parallel Average Subtract

Format: (qp) pavgsub1 $r_1 = r_2, r_3$ one_byte_form A9
 (qp) pavgsub2 $r_1 = r_2, r_3$ two_byte_form A9

Description: The unsigned data elements of GR r_3 are subtracted from the unsigned data elements of GR r_2 . The results of the subtraction are then each independently shifted to the right by one bit position. The high-order bits of each element are filled with the borrow bits of the subtraction (the complements of the ALU carries). To prevent cumulative round-off errors, an averaging is performed. The low-order bit of each result is set to 1 if at least one of the two least significant bits of the corresponding difference is 1. The signed results are placed in GR r_1 .




Figure 6-30. Parallel Average Subtract Example

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) { // one_byte_form
                x[0] = GR[r2]{7:0};      y[0] = GR[r3]{7:0};
                x[1] = GR[r2]{15:8};     y[1] = GR[r3]{15:8};
                x[2] = GR[r2]{23:16};    y[2] = GR[r3]{23:16};
                x[3] = GR[r2]{31:24};    y[3] = GR[r3]{31:24};
                x[4] = GR[r2]{39:32};    y[4] = GR[r3]{39:32};
                x[5] = GR[r2]{47:40};    y[5] = GR[r3]{47:40};
                x[6] = GR[r2]{55:48};    y[6] = GR[r3]{55:48};
                x[7] = GR[r2]{63:56};    y[7] = GR[r3]{63:56};

                for (i = 0; i < 8; i++) {
                    temp[i] = zero_ext(x[i], 8) - zero_ext(y[i], 8);
                    res[i] = (temp[i]{8:0} u>> 1) | (temp[i]{0});
                }
                GR[r1] = concatenate8(res[7], res[6], res[5], res[4],
                                       res[3], res[2], res[1], res[0]);
            } else { // two_byte_form
                x[0] = GR[r2]{15:0};      y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};     y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};     y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};     y[3] = GR[r3]{63:48};

                for (i = 0; i < 4; i++) {
                    temp[i] = zero_ext(x[i], 16) - zero_ext(y[i], 16);
                    res[i] = (temp[i]{16:0} u>> 1) | (temp[i]{0});
                }
                GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
            }
            GR[r1].nat = GR[r2].nat || GR[r3].nat;
        }

```

Parallel Compare

Format:

<i>(qp)</i> pcmp1.prel $r_1 = r_2, r_3$	one_byte_form	A9
<i>(qp)</i> pcmp2.prel $r_1 = r_2, r_3$	two_byte_form	A9
<i>(qp)</i> pcmp4.prel $r_1 = r_2, r_3$	four_byte_form	A9

Description: The two source operands are compared for one of the two relations shown in Table 6-38. If the comparison condition is true for corresponding data elements of GR r_2 and GR r_3 , then the corresponding data element in GR r_1 is set to all ones. If the comparison condition is false, the corresponding data element in GR r_1 is set to all zeros. For the '>' relation, both operands are interpreted as signed.

Table 6-38. Pcmp Relations

<i>prel</i>	Compare Relation (r_2 <i>prel</i> r_3)
eq	$r_2 == r_3$
gt	$r_2 > r_3$ (signed)

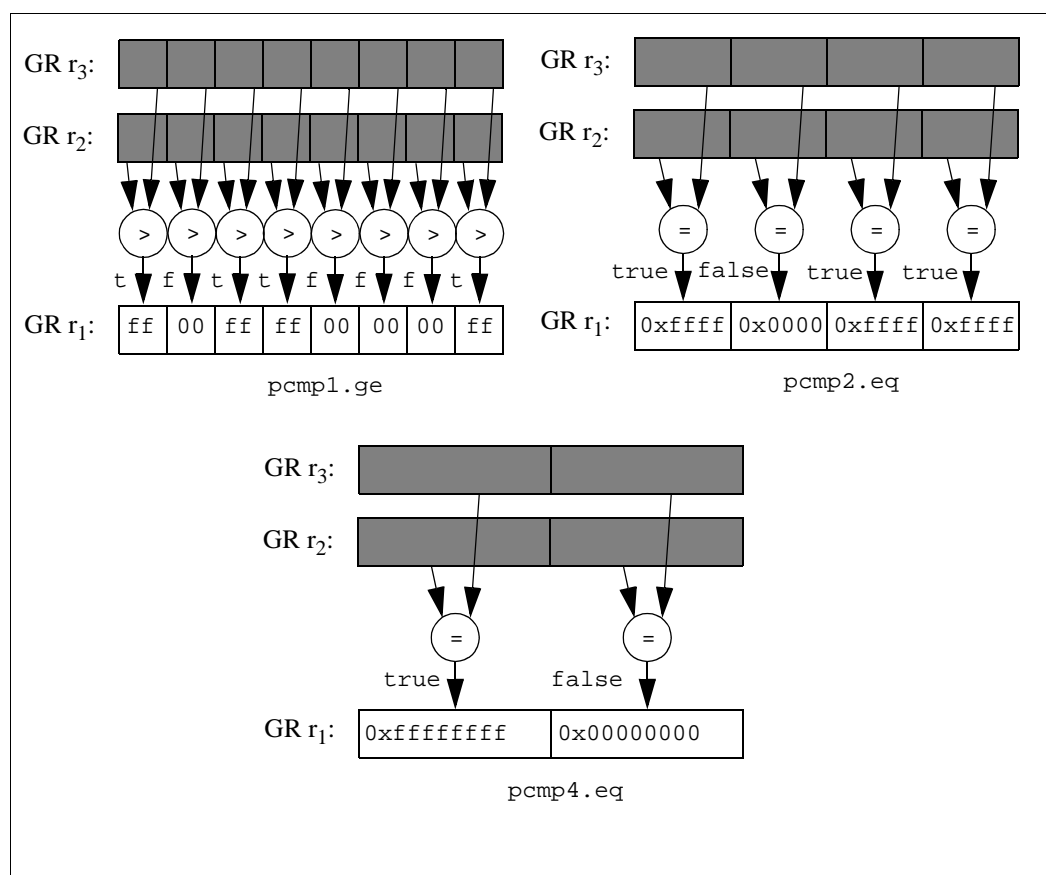


Figure 6-31. Parallel Compare Example


```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) { // one-byte elements
                x[0] = GR[r2]{7:0};    y[0] = GR[r3]{7:0};
                x[1] = GR[r2]{15:8};   y[1] = GR[r3]{15:8};
                x[2] = GR[r2]{23:16};  y[2] = GR[r3]{23:16};
                x[3] = GR[r2]{31:24};  y[3] = GR[r3]{31:24};
                x[4] = GR[r2]{39:32};  y[4] = GR[r3]{39:32};
                x[5] = GR[r2]{47:40};  y[5] = GR[r3]{47:40};
                x[6] = GR[r2]{55:48};  y[6] = GR[r3]{55:48};
                x[7] = GR[r2]{63:56};  y[7] = GR[r3]{63:56};
                for (i = 0; i < 8; i++) {
                    if (prel == 'eq')
                        tmp_rel = x[i] == y[i];
                    else
                        tmp_rel = greater_signed(sign_ext(x[i], 8), sign_ext(y[i], 8));

                    if (tmp_rel)
                        res[i] = 0xff;
                    else
                        res[i] = 0x00;
                }
                GR[r1] = concatenate8(res[7], res[6], res[5], res[4],
                                     res[3], res[2], res[1], res[0]);
            } else if (two_byte_form) { // two-byte elements
                x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};   y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};   y[3] = GR[r3]{63:48};
                for (i = 0; i < 4; i++) {
                    if (prel == 'eq')
                        tmp_rel = x[i] == y[i];
                    else
                        tmp_rel = greater_signed(sign_ext(x[i], 16), sign_ext(y[i], 16));

                    if (tmp_rel)
                        res[i] = 0xffff;
                    else
                        res[i] = 0x0000;
                }
                GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
            } else { // four-byte elements
                x[0] = GR[r2]{31:0};    y[0] = GR[r3]{31:0};
                x[1] = GR[r2]{63:32};   y[1] = GR[r3]{63:32};
                for (i = 0; i < 2; i++) {
                    if (prel == 'eq')
                        tmp_rel = x[i] == y[i];
                    else
                        tmp_rel = greater_signed(sign_ext(x[i], 32), sign_ext(y[i], 32));

                    if (tmp_rel)
                        res[i] = 0xffffffff;
                    else
                        res[i] = 0x00000000;
                }
                GR[r1] = concatenate2(res[1], res[0]);
            }
            GR[r1].nat = GR[r2].nat || GR[r3].nat;
        }
    
```

Parallel Maximum

Format: (qp) pmax1.u $r_1 = r_2, r_3$ one_byte_form I2
 (qp) pmax2 $r_1 = r_2, r_3$ two_byte_form I2

Description: The maximum of the two source operands is placed in the result register. In the one_byte_form, each unsigned 8-bit element of GR r_2 is compared with the corresponding unsigned 8-bit element of GR r_3 and the greater of the two is placed in the corresponding 8-bit element of GR r_1 . In the two_byte_form, each signed 16-bit element of GR r_2 is compared with the corresponding signed 16-bit element of GR r_3 and the greater of the two is placed in the corresponding 16-bit element of GR r_1 .

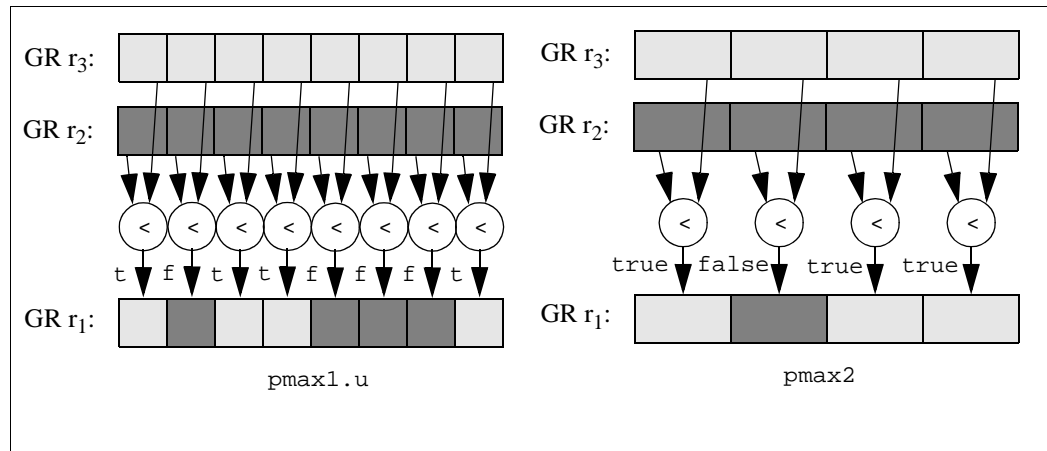


Figure 6-32. Parallel Maximum Example

Operation:

```

if (PR[qp]) {
  check_target_register(r1);

  if (one_byte_form) {
    // one-byte elements
    x[0] = GR[r2]{7:0};      y[0] = GR[r3]{7:0};
    x[1] = GR[r2]{15:8};    y[1] = GR[r3]{15:8};
    x[2] = GR[r2]{23:16};   y[2] = GR[r3]{23:16};
    x[3] = GR[r2]{31:24};   y[3] = GR[r3]{31:24};
    x[4] = GR[r2]{39:32};   y[4] = GR[r3]{39:32};
    x[5] = GR[r2]{47:40};   y[5] = GR[r3]{47:40};
    x[6] = GR[r2]{55:48};   y[6] = GR[r3]{55:48};
    x[7] = GR[r2]{63:56};   y[7] = GR[r3]{63:56};
    for (i = 0; i < 8; i++) {
      res[i] = (zero_ext(x[i],8) < zero_ext(y[i],8)) ? y[i] : x[i];
    }
    GR[r1] = concatenate8(res[7], res[6], res[5], res[4],
                          res[3], res[2], res[1], res[0]);
  } else {
    // two-byte elements
    x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
    x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
    x[2] = GR[r2]{47:32};   y[2] = GR[r3]{47:32};
    x[3] = GR[r2]{63:48};   y[3] = GR[r3]{63:48};
    for (i = 0; i < 4; i++) {
      res[i] = (sign_ext(x[i],16) < sign_ext(y[i],16)) ? y[i] : x[i];
    }
    GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
  }
  GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Parallel Minimum

Format: (qp) pmin1.u $r_1 = r_2, r_3$ one_byte_form I2
 (qp) pmin2 $r_1 = r_2, r_3$ two_byte_form I2

Description: The minimum of the two source operands is placed in the result register. In the one_byte_form, each unsigned 8-bit element of GR r_2 is compared with the corresponding unsigned 8-bit element of GR r_3 and the smaller of the two is placed in the corresponding 8-bit element of GR r_1 . In the two_byte_form, each signed 16-bit element of GR r_2 is compared with the corresponding signed 16-bit element of GR r_3 and the smaller of the two is placed in the corresponding 16-bit element of GR r_1 .

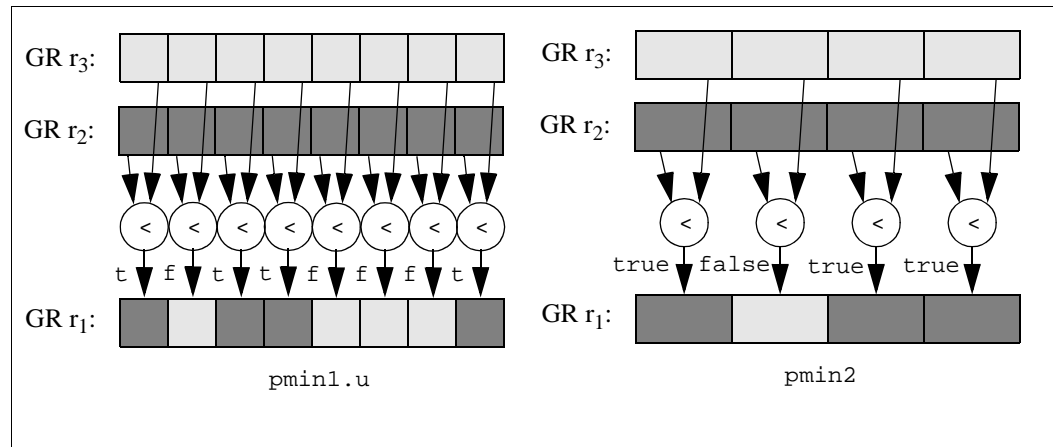


Figure 6-33. Parallel Minimum Example

Operation:

```

if (PR[qp]) {
  check_target_register(r1);

  if (one_byte_form) {
    // one-byte elements
    x[0] = GR[r2]{7:0};      y[0] = GR[r3]{7:0};
    x[1] = GR[r2]{15:8};    y[1] = GR[r3]{15:8};
    x[2] = GR[r2]{23:16};   y[2] = GR[r3]{23:16};
    x[3] = GR[r2]{31:24};   y[3] = GR[r3]{31:24};
    x[4] = GR[r2]{39:32};   y[4] = GR[r3]{39:32};
    x[5] = GR[r2]{47:40};   y[5] = GR[r3]{47:40};
    x[6] = GR[r2]{55:48};   y[6] = GR[r3]{55:48};
    x[7] = GR[r2]{63:56};   y[7] = GR[r3]{63:56};
    for (i = 0; i < 8; i++) {
      res[i] = (zero_ext(x[i],8) < zero_ext(y[i],8)) ? x[i] : y[i];
    }
    GR[r1] = concatenate8(res[7], res[6], res[5], res[4],
                          res[3], res[2], res[1], res[0]);
  } else {
    // two-byte elements
    x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
    x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
    x[2] = GR[r2]{47:32};   y[2] = GR[r3]{47:32};
    x[3] = GR[r2]{63:48};   y[3] = GR[r3]{63:48};
    for (i = 0; i < 4; i++) {
      res[i] = (sign_ext(x[i],16) < sign_ext(y[i],16)) ? x[i] : y[i];
    }
    GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
  }
  GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Parallel Multiply

Format: (qp) pmpy2.r $r_1 = r_2, r_3$ right_form I2
 (qp) pmpy2.l $r_1 = r_2, r_3$ left_form I2

Description: Two signed 16-bit data elements of GR r_2 are multiplied by the corresponding two signed 16-bit data elements of GR r_3 as shown in Figure 6-34. The two 32-bit results are placed in GR r_1 .

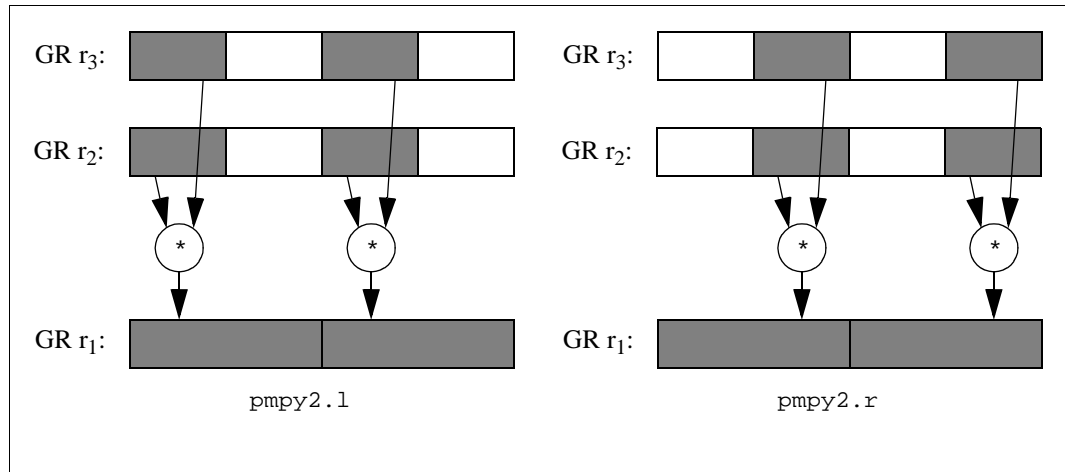


Figure 6-34. Parallel Multiply Operation

```

Operation:  if (PR[qp]) {
                check_target_register(r1);

                if (right_form) {
                    GR[r1]{31:0} = sign_ext(GR[r2]{15:0}, 16) * sign_ext(GR[r3]{15:0}, 16);
                    GR[r1]{63:32} = sign_ext(GR[r2]{47:32}, 16) * sign_ext(GR[r3]{47:32}, 16);
                } else {
                    // left_form
                    GR[r1]{31:0} = sign_ext(GR[r2]{31:16}, 16) * sign_ext(GR[r3]{31:16}, 16);
                    GR[r1]{63:32} = sign_ext(GR[r2]{63:48}, 16) * sign_ext(GR[r3]{63:48}, 16);
                }

                GR[r1].nat = GR[r2].nat || GR[r3].nat;
            }
  
```

Parallel Multiply and Shift Right

Format: (qp) pmpyshr2 $r_1 = r_2, r_3, count_2$ signed_form I1
 (qp) pmpyshr2.u $r_1 = r_2, r_3, count_2$ unsigned_form I1

Description: The four 16-bit data elements of GR r_2 are multiplied by the corresponding four 16-bit data elements of GR r_3 as shown in Figure 6-35. This multiplication can either be signed (pmpyshr2), or unsigned (pmpyshr2.u). Each product is then shifted to the right $count_2$ bits, and the least-significant 16-bits of each shifted product form 4 16-bit results, which are placed in GR r_1 . A $count_2$ of 0 gives the 16 low bits of the results, a $count_2$ of 16 gives the 16 high bits of the results. The allowed values for $count_2$ are given in Table 6-39.

Table 6-39. PMPYSHR Shift Options

$count_2$	Selected Bit Field from each 32-bit Product
0	15:0
7	22:7
15	30:15
16	31:16

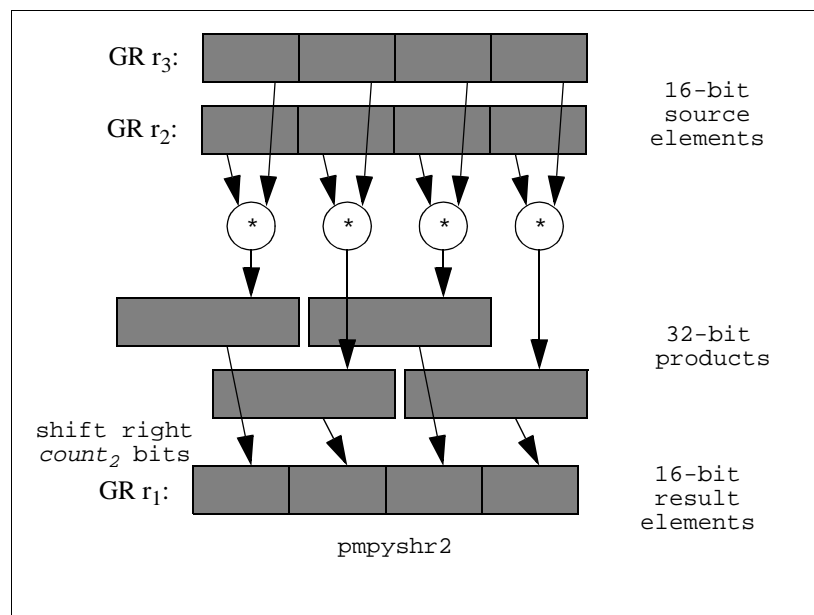


Figure 6-35. Parallel Multiply and Shift Right Operation

```

Operation:  if (PR[qp]) {
                check_target_register(r1);
                x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};   y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};   y[3] = GR[r3]{63:48};
                for (i = 0; i < 4; i++) {
                    if (unsigned_form) // unsigned multiplication
                        temp[i] = zero_ext(x[i], 16) * zero_ext(y[i], 16);
                    else // signed multiplication
                        temp[i] = sign_ext(x[i], 16) * sign_ext(y[i], 16);

                    res[i] = temp[i]{(count2 + 15):count2};
                }

                GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
                GR[r1].nat = GR[r2].nat || GR[r3].nat;
            }
  
```

Population Count

Format: (qp) popcnt $r_1 = r_3$

19

Description: The number of bits in GR r_3 having the value 1 is counted, and the resulting sum is placed in GR r_1 .

Operation:

```
if (PR[qp]) {
    check_target_register(r1);

    res = 0;
    // Count up all the one bits
    for (i = 0; i < 64; i++) {
        res += GR[r3]{i};
    }

    GR[r1] = res;
    GR[r1].nat = GR[r3].nat;
}
```

Parallel Sum of Absolute Difference

Format: (qp) psadl $r_1 = r_2, r_3$

12

Description: The unsigned 8-bit elements of GR r_2 are subtracted from the unsigned 8-bit elements of GR r_3 . The absolute value of each difference is accumulated across the elements and placed in GR r_1 .

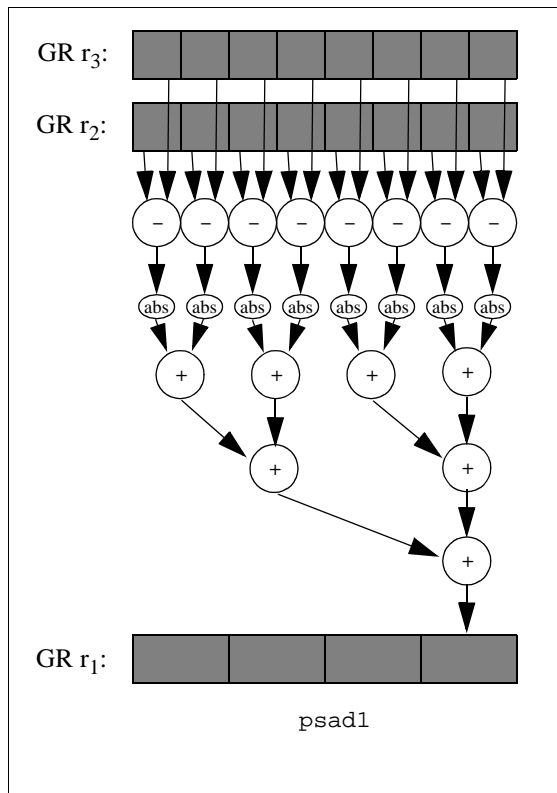


Figure 6-36. Parallel Sum of Absolute Difference Example

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            x[0] = GR[r2]{7:0};      y[0] = GR[r3]{7:0};
            x[1] = GR[r2]{15:8};    y[1] = GR[r3]{15:8};
            x[2] = GR[r2]{23:16};   y[2] = GR[r3]{23:16};
            x[3] = GR[r2]{31:24};   y[3] = GR[r3]{31:24};
            x[4] = GR[r2]{39:32};   y[4] = GR[r3]{39:32};
            x[5] = GR[r2]{47:40};   y[5] = GR[r3]{47:40};
            x[6] = GR[r2]{55:48};   y[6] = GR[r3]{55:48};
            x[7] = GR[r2]{63:56};   y[7] = GR[r3]{63:56};

            GR[r1] = 0;
            for (i = 0; i < 8; i++) {
                temp[i] = zero_ext(x[i], 8) - zero_ext(y[i], 8);
                if (temp[i] < 0)
                    temp[i] = -temp[i];
                GR[r1] += temp[i];
            }

            GR[r1].nat = GR[r2].nat || GR[r3].nat;
        }

```

Parallel Shift Left

Format:	(qp) pshl2 $r_1 = r_2, r_3$	two_byte_form, variable_form	I7
	(qp) pshl2 $r_1 = r_2, count_5$	two_byte_form, fixed_form	I8
	(qp) pshl4 $r_1 = r_2, r_3$	four_byte_form, variable_form	I7
	(qp) pshl4 $r_1 = r_2, count_5$	four_byte_form, fixed_form	I8

Description: The data elements of GR r_2 are each independently shifted to the left by the scalar shift count in GR r_3 , or in the immediate field $count_5$. The low-order bits of each element are filled with zeros. The shift count is interpreted as unsigned. Shift counts greater than 15 (for 16-bit quantities) or 31 (for 32-bit quantities) yield all zero results. The results are placed in GR r_1 .

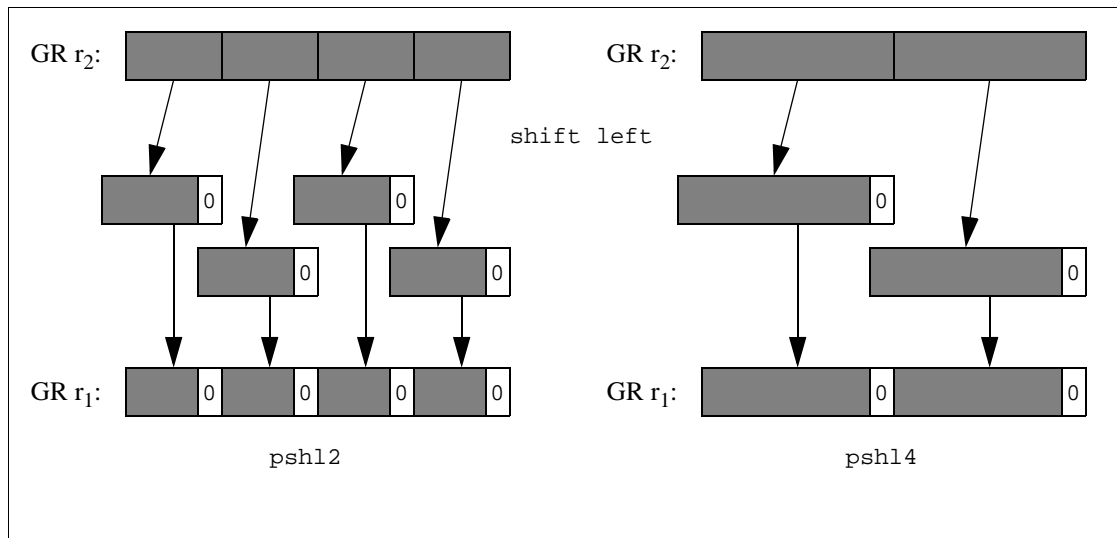


Figure 6-37. Parallel Shift Left Example

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    shift_count = (variable_form ? GR[r3] : count5);
    tmp_nat = (variable_form ? GR[r3].nat : 0);

    if (two_byte_form) { // two_byte_form
        if (shift_count > 16)
            shift_count = 16;
        GR[r1]{15:0} = GR[r2]{15:0} << shift_count;
        GR[r1]{31:16} = GR[r2]{31:16} << shift_count;
        GR[r1]{47:32} = GR[r2]{47:32} << shift_count;
        GR[r1]{63:48} = GR[r2]{63:48} << shift_count;
    } else { // four_byte_form
        if (shift_count > 32)
            shift_count = 32;
        GR[r1]{31:0} = GR[r2]{31:0} << shift_count;
        GR[r1]{63:32} = GR[r2]{63:32} << shift_count;
    }

    GR[r1].nat = GR[r2].nat || tmp_nat;
}

```


Parallel Shift Left and Add

Format: (qp) pshladd2 $r_1 = r_2, count_2, r_3$

A10

Description: The four signed 16-bit data elements of GR r_2 are each independently shifted to the left by $count_2$ bits (shifting zeros into the low-order bits), and added to the four signed 16-bit data elements of GR r_3 . Both the left shift and the add operations are saturating: if the result of either the shift or the add is not representable as a signed 16-bit value, the final result is saturated. The four signed 16-bit results are placed in GR r_1 . The first operand can be shifted by 1, 2 or 3 bits.

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
    x[1] = GR[r2]{31:16};  y[1] = GR[r3]{31:16};
    x[2] = GR[r2]{47:32};  y[2] = GR[r3]{47:32};
    x[3] = GR[r2]{63:48};  y[3] = GR[r3]{63:48};

    max = sign_ext(0x7fff, 16);
    min = sign_ext(0x8000, 16);

    for (i = 0; i < 4; i++) {
        temp[i] = sign_ext(x[i], 16) << count2;

        if (temp[i] > max)
            res[i] = max;
        else if (temp[i] < min)
            res[i] = min;
        else {
            res[i] = temp[i] + sign_ext(y[i], 16);
            if (res[i] > max)
                res[i] = max;
            if (res[i] < min)
                res[i] = min;
        }
    }

    GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
    GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Parallel Shift Right

Format:	(qp) pshr2 $r_1 = r_3, r_2$	signed_form, two_byte_form, variable_form	15
	(qp) pshr2 $r_1 = r_3, count_5$	signed_form, two_byte_form, fixed_form	16
	(qp) pshr2.u $r_1 = r_3, r_2$	unsigned_form, two_byte_form, variable_form	15
	(qp) pshr2.u $r_1 = r_3, count_5$	unsigned_form, two_byte_form, fixed_form	16
	(qp) pshr4 $r_1 = r_3, r_2$	signed_form, four_byte_form, variable_form	15
	(qp) pshr4 $r_1 = r_3, count_5$	signed_form, four_byte_form, fixed_form	16
	(qp) pshr4.u $r_1 = r_3, r_2$	unsigned_form, four_byte_form, variable_form	15
	(qp) pshr4.u $r_1 = r_3, count_5$	unsigned_form, four_byte_form, fixed_form	16

Description: The data elements of GR r_3 are each independently shifted to the right by the scalar shift count in GR r_2 , or in the immediate field $count_5$. The high-order bits of each element are filled with either the initial value of the sign bits of the data elements in GR r_3 (arithmetic shift) or zeros (logical shift). The shift count is interpreted as unsigned. Shift counts greater than 15 (for 16-bit quantities) or 31 (for 32-bit quantities) yield all zero or all one results depending on the initial values of the sign bits of the data elements in GR r_3 and whether a signed or unsigned shift is done. The results are placed in GR r_1 .

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    shift_count = (variable_form ? GR[r2] : count5);
    tmp_nat = (variable_form ? GR[r2].nat : 0);

    if (two_byte_form) { // two_byte_form
        if (shift_count > 16)
            shift_count = 16;
        if (unsigned_form) { // unsigned shift
            GR[r1]{15:0} = shift_right_unsigned(zero_ext(GR[r3]{15:0}, 16),
            shift_count);
            GR[r1]{31:16} = shift_right_unsigned(zero_ext(GR[r3]{31:16}, 16),
            shift_count);
            GR[r1]{47:32} = shift_right_unsigned(zero_ext(GR[r3]{47:32}, 16),
            shift_count);
            GR[r1]{63:48} = shift_right_unsigned(zero_ext(GR[r3]{63:48}, 16),
            shift_count);
        } else { // signed shift
            GR[r1]{15:0} = shift_right_signed(sign_ext(GR[r3]{15:0}, 16),
            shift_count);
            GR[r1]{31:16} = shift_right_signed(sign_ext(GR[r3]{31:16}, 16),
            shift_count);
            GR[r1]{47:32} = shift_right_signed(sign_ext(GR[r3]{47:32}, 16),
            shift_count);
            GR[r1]{63:48} = shift_right_signed(sign_ext(GR[r3]{63:48}, 16),
            shift_count);
        }
    } else { // four_byte_form
        if (shift_count > 32)
            shift_count = 32;
        if (unsigned_form) { // unsigned shift
            GR[r1]{31:0} = shift_right_unsigned(zero_ext(GR[r3]{31:0}, 32),
            shift_count);
            GR[r1]{63:32} = shift_right_unsigned(zero_ext(GR[r3]{63:32}, 32),
            shift_count);
        } else { // signed shift
            GR[r1]{31:0} = shift_right_signed(sign_ext(GR[r3]{31:0}, 32),
            shift_count);
            GR[r1]{63:32} = shift_right_signed(sign_ext(GR[r3]{63:32}, 32),
            shift_count);
        }
    }

    GR[r1].nat = GR[r3].nat || tmp_nat;
}

```

Parallel Shift Right and Add

Format: (qp) pshradd2 $r_1 = r_2, count_2, r_3$

A10

Description: The four signed 16-bit data elements of GR r_2 are each independently shifted to the right by $count_2$ bits, and added to the four signed 16-bit data elements of GR r_3 . The right shift operation fills the high-order bits of each element with the initial value of the sign bits of the data elements in GR r_2 . The add operation is performed with signed saturation. The four signed 16-bit results of the add are placed in GR r_1 . The first operand can be shifted by 1, 2 or 3 bits.

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
    x[1] = GR[r2]{31:16};  y[1] = GR[r3]{31:16};
    x[2] = GR[r2]{47:32};  y[2] = GR[r3]{47:32};
    x[3] = GR[r2]{63:48};  y[3] = GR[r3]{63:48};

    max = sign_ext(0x7fff, 16);
    min = sign_ext(0x8000, 16);

    for (i = 0; i < 4; i++) {
        temp[i] = shift_right_signed(sign_ext(x[i], 16), count2);

        res[i] = temp[i] + sign_ext(y[i], 16);
        if (res[i] > max)
            res[i] = max;
        if (res[i] < min)
            res[i] = min;
    }

    GR[r1] = concatenate4(res[3], res[2], res[1], res[0]);
    GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Parallel Subtract

Format:	(qp) psub1 $r_1 = r_2, r_3$	one_byte_form, modulo_form	A9
	(qp) psub1.sss $r_1 = r_2, r_3$	one_byte_form, sss_saturation_form	A9
	(qp) psub1.uus $r_1 = r_2, r_3$	one_byte_form, uus_saturation_form	A9
	(qp) psub1.uuu $r_1 = r_2, r_3$	one_byte_form, uuu_saturation_form	A9
	(qp) psub2 $r_1 = r_2, r_3$	two_byte_form, modulo_form	A9
	(qp) psub2.sss $r_1 = r_2, r_3$	two_byte_form, sss_saturation_form	A9
	(qp) psub2.uus $r_1 = r_2, r_3$	two_byte_form, uus_saturation_form	A9
	(qp) psub2.uuu $r_1 = r_2, r_3$	two_byte_form, uuu_saturation_form	A9
	(qp) psub4 $r_1 = r_2, r_3$	four_byte_form, modulo_form	A9

Description: The sets of elements from the two source operands are subtracted, and the results placed in GR r_1 .

If the difference between two elements cannot be represented in the result element and a saturation completer is specified, then saturation clipping is performed. The saturation can either be signed or unsigned, as given in Table 6-40. If the difference of two elements is larger than the upper limit value, the result is the upper limit value. If it is smaller than the lower limit value, the result is the lower limit value. The saturation limits are given in Table 6-41.

Table 6-40. Parallel Subtract Saturation Completers

Completer	Result r_1 Treated as	Source r_2 Treated as	Source r_3 Treated as
sss	signed	signed	signed
uus	unsigned	unsigned	signed
uuu	unsigned	unsigned	unsigned

Table 6-41. Parallel Subtract Saturation Limits

Size	Element Width	Result r_1 Signed		Result r_1 Unsigned	
		Upper Limit	Lower Limit	Upper Limit	Lower Limit
1	8 bit	0x7f	0x80	0xff	0x00
2	16 bit	0x7fff	0x8000	0xffff	0x0000

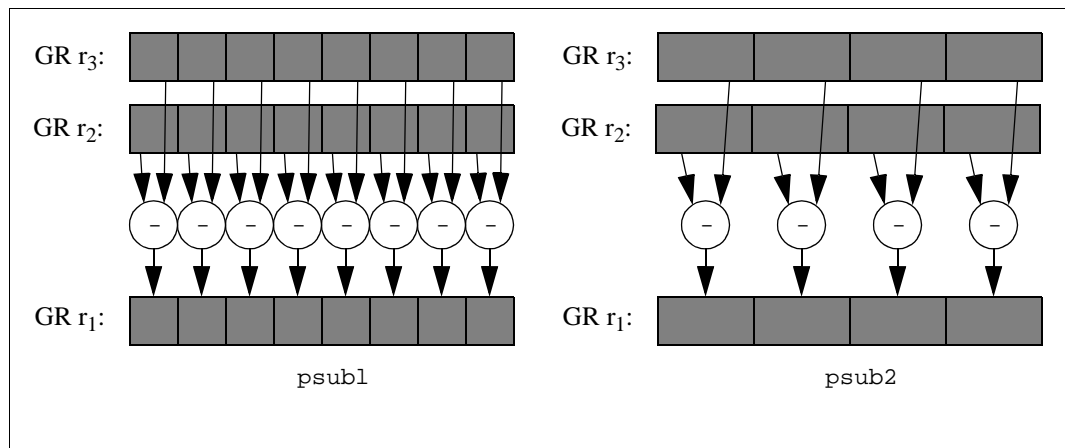


Figure 6-38. Parallel Subtract Example

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) { // one-byte elements
                x[0] = GR[r2]{7:0};    y[0] = GR[r3]{7:0};
                x[1] = GR[r2]{15:8};   y[1] = GR[r3]{15:8};
                x[2] = GR[r2]{23:16};  y[2] = GR[r3]{23:16};
                x[3] = GR[r2]{31:24};  y[3] = GR[r3]{31:24};
                x[4] = GR[r2]{39:32};  y[4] = GR[r3]{39:32};
                x[5] = GR[r2]{47:40};  y[5] = GR[r3]{47:40};
                x[6] = GR[r2]{55:48};  y[6] = GR[r3]{55:48};
                x[7] = GR[r2]{63:56};  y[7] = GR[r3]{63:56};

                if (sss_saturation_form) { // sss_saturation_form
                    max = sign_ext(0x7f, 8);
                    min = sign_ext(0x80, 8);
                    for (i = 0; i < 8; i++) {
                        temp[i] = sign_ext(x[i], 8) - sign_ext(y[i], 8);
                    }
                } else if (uus_saturation_form) { // uus_saturation_form
                    max = 0xff;
                    min = 0x00;
                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) - sign_ext(y[i], 8);
                    }
                } else if (uuu_saturation_form) { // uuu_saturation_form
                    max = 0xff;
                    min = 0x00;
                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) - zero_ext(y[i], 8);
                    }
                } else { // modulo_form
                    for (i = 0; i < 8; i++) {
                        temp[i] = zero_ext(x[i], 8) - zero_ext(y[i], 8);
                    }
                }

                if (sss_saturation_form || uus_saturation_form || uuu_saturation_form) {
                    for (i = 0; i < 8; i++) {
                        if (temp[i] > max)
                            temp[i] = max;
                        if (temp[i] < min)
                            temp[i] = min;
                    }
                }

                GR[r1] = concatenate8(temp[7], temp[6], temp[5], temp[4],
                                       temp[3], temp[2], temp[1], temp[0]);
            } else if (two_byte_form) { // two-byte elements
                x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};  y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};  y[3] = GR[r3]{63:48};

                if (sss_saturation_form) { // sss_saturation_form
                    max = sign_ext(0x7fff, 16);
                    min = sign_ext(0x8000, 16);
                    for (i = 0; i < 4; i++) {
                        temp[i] = sign_ext(x[i], 16) - sign_ext(y[i], 16);
                    }
                } else if (uus_saturation_form) { // uus_saturation_form
                    max = 0xffff;
                    min = 0x0000;
                    for (i = 0; i < 4; i++) {
                        temp[i] = zero_ext(x[i], 16) - sign_ext(y[i], 16);
                    }
                } else if (uuu_saturation_form) { // uuu_saturation_form
                    max = 0xffff;
                }
            }
        }

```

```

        min = 0x0000;
        for (i = 0; i < 4; i++) {
            temp[i] = zero_ext(x[i], 16) - zero_ext(y[i], 16);
        }
    } else {
        // modulo_form
        for (i = 0; i < 4; i++) {
            temp[i] = zero_ext(x[i], 16) - zero_ext(y[i], 16);
        }
    }

    if (sss_saturation_form || uus_saturation_form || uuu_saturation_form) {
        for (i = 0; i < 4; i++) {
            if (temp[i] > max)
                temp[i] = max;
            if (temp[i] < min)
                temp[i] = min;
        }
    }

    GR[r1] = concatenate4(temp[3], temp[2], temp[1], temp[0]);
} else {
    // four-byte elements
    x[0] = GR[r2]{31:0};    y[0] = GR[r3]{31:0};
    x[1] = GR[r2]{63:32};  y[1] = GR[r3]{63:32};

    for (i = 0; i < 2; i++) {
        temp[i] = zero_ext(x[i], 32) - zero_ext(y[i], 32);
    }

    GR[r1] = concatenate2(temp[1], temp[0]);
}

GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Reset User Mask

Format: (qp) rum *imm*₂₄

M44

Description: The complement of the *imm*₂₄ operand is ANDed with the user mask (PSR{5:0}) and the result is placed in the user mask.

PSR.up is only cleared if the secure performance monitor bit (PSR.sp) is zero. Otherwise PSR.up is not modified.

Operation:

```
if (PR[qp]) {
    if (is_reserved_field(PSR_TYPE, PSR_UM, imm24))
        reserved_register_field_fault();

    if (imm24{1})    PSR{1} = 0;
    if (imm24{2} && PSR.sp == 0)    //non-secure perf monitor
        PSR{2} = 0;
    if (imm24{3})    PSR{3} = 0;
    if (imm24{4})    PSR{4} = 0;
    if (imm24{5})    PSR{5} = 0;
}
```

Set Floating-Point Value, Exponent, or Significand

Format:	(qp) setf.s $f_1 = r_2$	single_form	M18
	(qp) setf.d $f_1 = r_2$	double_form	M18
	(qp) setf.exp $f_1 = r_2$	exponent_form	M18
	(qp) setf.sig $f_1 = r_2$	significand_form	M18

Description: In the single and double forms, GR r_2 is treated as a single precision (in the single_form) or double precision (in the double_form) memory representation, converted into floating-point register format, and placed in FR f_1 .

In the exponent_form, bits 16:0 of GR r_2 are copied to the exponent field of FR f_1 and bit 17 of GR r_2 is copied to the sign bit of FR f_1 . The significand field of FR f_1 is set to one (0x800...000).




Figure 6-39. Function of setf.exp

In the significand_form, the value in GR r_2 is copied to the significand field of FR f_1 .

The exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0).




Figure 6-40. Function of setf.sig

For all forms, if the NaT bit corresponding to r_2 is equal to 1, FR f_1 is set to NaTVal instead of the computed result.


```

Operation:  if (PR[qp]) {
            fp_check_target_register(f1);
            if (tmp_israncode = fp_reg_disabled(f1, 0, 0, 0))
                disabled_fp_register_fault(tmp_israncode, 0);

            if (!GR[r2].nat) {
                if (single_form)
                    FR[f1] = fp_mem_to_fr_format(GR[r2], 4, 0);
                else if (double_form)
                    FR[f1] = fp_mem_to_fr_format(GR[r2], 8, 0);
                else if (significand_form) {
                    FR[f1].significand = GR[r2];
                    FR[f1].exponent = FP_INTEGER_EXP;
                    FR[f1].sign = 0;
                } else {
                    FR[f1].significand = 0x8000000000000000;           // exponent_form
                    FR[f1].exp = GR[r2]{16:0};
                    FR[f1].sign = GR[r2]{17};
                }
            } else
                FR[f1] = NATVAL;

            fp_update_psr(f1);
        }

```

Shift Left

Format: (qp) shl $r_1 = r_2, r_3$ I7
 (qp) shl $r_1 = r_2, count_6$ pseudo-op of: (qp) dep.z $r_1 = r_2, count_6, 64 - count_6$

Description: The value in GR r_2 is shifted to the left, with the vacated bit positions filled with zeroes, and placed in GR r_1 . The number of bit positions to shift is specified by the value in GR r_3 or by an immediate value $count_6$. The shift count is interpreted as an unsigned number. If the value in GR r_3 is greater than 63, then the result is all zeroes.

For the immediate form, See “Deposit” on page 6-27.

Operation:

```

if (PR[qp]) {
    check_target_register( $r_1$ );

    count = GR[ $r_3$ ];
    GR[ $r_1$ ] = (count > 63) ? 0 : GR[ $r_2$ ] << count;

    GR[ $r_1$ ].nat = GR[ $r_2$ ].nat || GR[ $r_3$ ].nat;
}

```

Shift Left and Add

Format: (qp) shladd $r_1 = r_2, count_2, r_3$ A2

Description: The first source operand is shifted to the left by $count_2$ bits and then added to the second source operand and the result placed in GR r_1 . The first operand can be shifted by 1, 2, 3, or 4 bits.

Operation:

```
if (PR[qp]) {
    check_target_register(r1);

    GR[r1] = (GR[r2] << count2) + GR[r3];
    GR[r1].nat = GR[r2].nat || GR[r3].nat;
}
```

Shift Left and Add Pointer

Format: $(qp) \text{ shladdp4 } r_1 = r_2, \text{count}_2, r_3$

A2

Description: The first source operand is shifted to the left by count_2 bits and then is added to the second source operand. The upper 32 bits of the result are forced to zero, and then bits {31:30} of GR r_3 are copied to bits {62:61} of the result. This result is placed in GR r_1 . The first operand can be shifted by 1, 2, 3, or 4 bits.

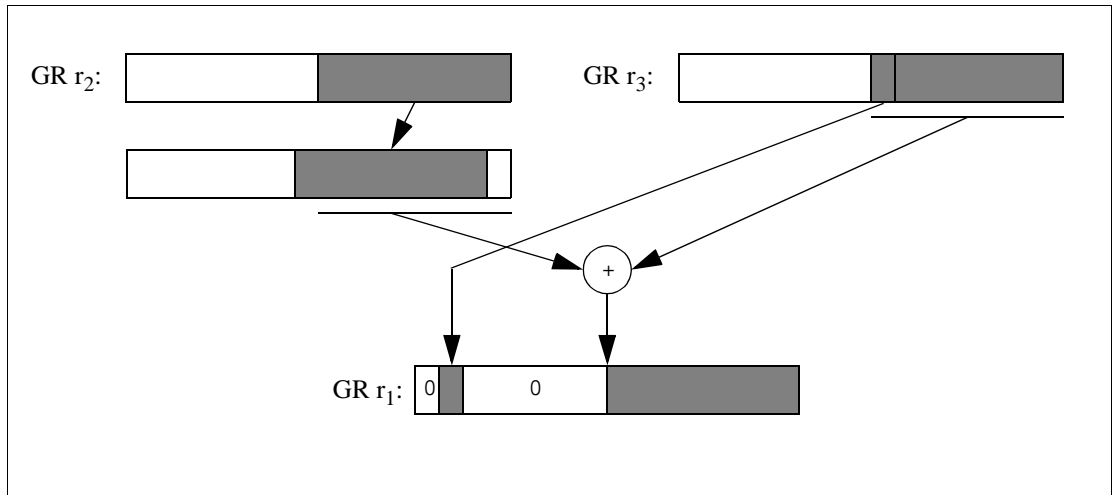


Figure 6-41. Shift Left and Add Pointer

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    tmp_res = (GR[r2] << count2) + GR[r3];
    tmp_res = zero_ext(tmp_res{31:0}, 32);
    tmp_res{62:61} = GR[r3]{31:30};
    GR[r1] = tmp_res;
    GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Shift Right

Format:	<i>(qp)</i> shr $r_1 = r_3, r_2$	signed_form	I5
	<i>(qp)</i> shr.u $r_1 = r_3, r_2$	unsigned_form	I5
	<i>(qp)</i> shr $r_1 = r_3, count_6$	pseudo-op of: <i>(qp)</i> extr $r_1 = r_3, count_6, 64-count_6$	
	<i>(qp)</i> shr.u $r_1 = r_3, count_6$	pseudo-op of: <i>(qp)</i> extr.u $r_1 = r_3, count_6, 64-count_6$	

Description: The value in GR r_3 is shifted to the right and placed in GR r_1 . In the signed_form the vacated bit positions are filled with bit 63 of GR r_3 ; in the unsigned_form the vacated bit positions are filled with zeroes. The number of bit positions to shift is specified by the value in GR r_2 or by an immediate value $count_6$. The shift count is interpreted as an unsigned number. If the value in GR r_2 is greater than 63, then the result is all zeroes (for the unsigned_form, or if bit 63 of GR r_3 was 0) or all ones (for the signed_form if bit 63 of GR r_3 was 1).

If the .u completer is specified, the shift is unsigned (logical), otherwise it is signed (arithmetic).

For the immediate forms, See “Extract” on page 6-28.

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    if (signed_form) {
        count = (GR[r2] > 63) ? 63 : GR[r2];
        GR[r1] = shift_right_signed(GR[r3], count);
    } else {
        count = GR[r2];
        GR[r1] = (count > 63) ? 0 : shift_right_unsigned(GR[r3], count);
    }

    GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Shift Right Pair

Format: (qp) shrp $r_1 = r_2, r_3, count_6$

110

Description: The two source operands, GR r_2 and GR r_3 , are concatenated to form a 128-bit value and shifted to the right $count_6$ bits. The least-significant 64 bits of the result are placed in GR r_1 .

The immediate value $count_6$ can be any number in the range 0 to 63.




Figure 6-42. Shift Right Pair

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    temp1 = shift_right_unsigned(GR[r3], count6);
    temp2 = GR[r2] << (64 - count6);
    GR[r1] = zero_ext(temp1, 64 - count6) | temp2;
    GR[r1].nat = GR[r2].nat || GR[r3].nat;
}

```

Serialize

Format: (qp) srlz.i M24

Description: Instruction serialization (`srlz.i`) ensures:

- prior modifications to processor register resources that affect fetching of subsequent instruction groups are observed,
- prior modifications to processor register resources that affect subsequent execution or data memory accesses are observed,
- prior memory synchronization (`sync.i`) operations have taken effect on the local processor instruction cache,
- subsequent instruction group fetches are re-initiated after `srlz.i` completes.

The `srlz.i` instruction must be in an instruction group after the instruction group containing the operation that is to be serialized. Operations dependent on the serialization must be in an instruction group after the instruction group containing the `srlz.i`.

Operation:

```
if (PR[qp]) {  
    instruction_serialize();  
}
```

Store

Format:	<i>(qp)</i> <i>stsz.sttype.sthint</i> [r_3] = r_2	normal_form, no_base_update_form	M4
	<i>(qp)</i> <i>stsz.sttype.sthint</i> [r_3] = r_2 , <i>imm₉</i>	normal_form, imm_base_update_form	M5
	<i>(qp)</i> <i>st8.spill.sthint</i> [r_3] = r_2	spill_form, no_base_update_form	M4
	<i>(qp)</i> <i>st8.spill.sthint</i> [r_3] = r_2 , <i>imm₉</i>	spill_form, imm_base_update_form	M5

Description: A value consisting of the least significant *sz* bytes of the value in GR r_2 is written to memory starting at the address specified by the value in GR r_3 . The values of the *sz* completer are given in Table 6-26 on page 6-101. The *sttype* completer specifies special store operations, which are described in Table 6-42. If the NaT bit corresponding to GR r_3 is 1 (or in the normal_form, if the NaT bit corresponding to GR r_2 is 1), a Register NaT Consumption fault is taken.

In the spill_form, an 8-byte value is stored, and the NaT bit corresponding to GR r_2 is copied to a bit in the UNAT application register. This instruction is used for spilling a register/NaT pair. See “Control Speculation” on page 4-10 for details.

In the imm_base_update form, the value in GR r_3 is added to a signed immediate value (*imm₉*) and the result is placed back in GR r_3 . This base register update is done after the store, and does not affect the store address, nor the value stored (for the case where r_2 and r_3 specify the same register).

Table 6-42. Store Types

<i>sttype</i> Completer	Interpretation	Special Store Operation
<i>none</i>	Normal store	
rel	Ordered store	An ordered store is performed with release semantics.

For more details on ordered stores see “Memory Access Ordering” on page 4-18 .

The ALAT is queried using the physical memory address and the access size, and all overlapping entries are invalidated.

The value of the *sthint* completer specifies the locality of the memory access. The values of the *sthint* completer are given in Table 6-43. See “Memory Hierarchy Control and Consistency” on page 4-16.

Table 6-43. Store Hints

<i>sthint</i> Completer	Interpretation
<i>none</i>	Temporal locality, level 1
nta	Non-temporal locality, all levels


```

Operation:  if (PR[gp]) {
            size = spill_form ? 8 : sz;
            otype = (sttype == 'rel') ? RELEASE : UNORDERED;

            if (imm_base_update_form)
                check_target_register(r3);
            if (GR[r3].nat || (normal_form && GR[r2].nat))
                register_nat_consumption_fault(WRITE);

            paddr = tlb_translate(GR[r3], size, WRITE, PSR.cpl, &mattr,
                                &tmp_unused);
            if (spill_form && GR[r2].nat)
                natd_gr_write(GR[r2], paddr, size, UM.be, mattr, otype, sthint);
            else
                mem_write(GR[r2], paddr, size, UM.be, mattr, otype, sthint);

            if (spill_form) {
                bit_pos = GR[r3]{8:3};
                AR[UNAT]{bit_pos} = GR[r2].nat;
            }

            alat_inval_multiple_entries(paddr, size);

            if (imm_base_update_form) {
                GR[r3] = GR[r3] + sign_ext(imm9, 9);
                GR[r3].nat = 0;
            }
        }

```

Floating-Point Store

Format:	<i>(qp)</i> stf <i>fsz.sthint</i> [<i>r</i> ₃] = <i>f</i> ₂	normal_form, no_base_update_form	M9
	<i>(qp)</i> stf <i>fsz.sthint</i> [<i>r</i> ₃] = <i>f</i> ₂ , <i>imm</i> ₉	normal_form, imm_base_update_form	M10
	<i>(qp)</i> stf8. <i>sthint</i> [<i>r</i> ₃] = <i>f</i> ₂	integer_form, no_base_update_form	M9
	<i>(qp)</i> stf8. <i>sthint</i> [<i>r</i> ₃] = <i>f</i> ₂ , <i>imm</i> ₉	integer_form, imm_base_update_form	M10
	<i>(qp)</i> stf.spill. <i>sthint</i> [<i>r</i> ₃] = <i>f</i> ₂	spill_form, no_base_update_form	M9
	<i>(qp)</i> stf.spill. <i>sthint</i> [<i>r</i> ₃] = <i>f</i> ₂ , <i>imm</i> ₉	spill_form, imm_base_update_form	M10

Description: A value, consisting of *fsz* bytes, is generated from the value in FR *f*₂ and written to memory starting at the address specified by the value in GR *r*₃. In the normal_form, the value in FR *f*₂ is converted to the memory format and then stored. In the integer_form, the significand of FR *f*₂ is stored. The values of the *fsz* completer are given in Table 6-29 on page 6-105. In the normal_form or the integer_form, if the NaT bit corresponding to GR *r*₃ is 1 or if FR *f*₂ contains NaTVal, a Register NaT Consumption fault is taken. See “Data Types and Formats” on page 5-1 for details on conversion from floating-point register format.

In the spill_form, a 16-byte value from FR *f*₂ is stored without conversion. This instruction is used for spilling a register. See “Control Speculation” on page 4-10 for details.

In the imm_base_update form, the value in GR *r*₃ is added to a signed immediate value (*imm*₉) and the result is placed back in GR *r*₃. This base register update is done after the store, and does not affect the store address.

The ALAT is queried using the physical memory address and the access size, and all overlapping entries are invalidated.

The value of the *sthint* completer specifies the locality of the memory access. The values of the *sthint* completer are given in Table 6-43 on page 6-166. See “Memory Hierarchy Control and Consistency” on page 4-16.

Operation:

```

if (PR[qp]) {
    if (imm_base_update_form)
        check_target_register(r3);
    if (tmp_israncode = fp_reg_disabled(f2, 0, 0, 0))
        disabled_fp_register_fault(tmp_israncode, WRITE);

    if (GR[r3].nat || (!spill_form && (FR[f2] == NATVAL)))
        register_nat_consumption_fault(WRITE);

    size = spill_form ? 16 : (integer_form ? 8 : fsz);

    paddr = tlb_translate(GR[r3], size, WRITE, PSR.cpl, &matr, &tmp_unused);
    val = fp_fr_to_mem_format(FR[f2], size, integer_form);
    mem_write(val, paddr, size, UM.be, matr, UNORDERED, sthint);

    alat_inval_multiple_entries(paddr, size);

    if (imm_base_update_form) {
        GR[r3] = GR[r3] + sign_ext(imm9, 9);
        GR[r3].nat = 0;
    }
}

```

Subtract

Format:	<i>(qp)</i> sub $r_1 = r_2, r_3$	register_form	A1
	<i>(qp)</i> sub $r_1 = r_2, r_3, 1$	minus1_form, register_form	A1
	<i>(qp)</i> sub $r_1 = imm_8, r_3$	imm8_form	A3

Description: The second source operand (and an optional constant 1) are subtracted from the first operand and the result placed in GR r_1 . In the register form the first operand is GR r_2 ; in the immediate form the first operand is taken from the sign extended imm_8 encoding field.

The minus1_form is available only in the register_form (although the equivalent effect can be achieved by adjusting the immediate).

Operation:

```

if (PR[qp]) {
    check_target_register(r1);

    tmp_src = (register_form ? GR[r2] : sign_ext(imm8, 8));
    tmp_nat = (register_form ? GR[r2].nat : 0);

    if (minus1_form)
        GR[r1] = tmp_src - GR[r3] - 1;
    else
        GR[r1] = tmp_src - GR[r3];

    GR[r1].nat = tmp_nat || GR[r3].nat;
}

```

Set User Mask

Format: (qp) sum imm_{24}

M44

Description: The imm_{24} operand is ORed with the user mask (PSR{5:0}) and the result is placed in the user mask.

PSR.up can only be set if the secure performance monitor bit (PSR.sp) is zero. Otherwise PSR.up is not modified.

Operation:

```
if (PR[qp]) {
    if (is_reserved_field(PSR_TYPE, PSR_UM, imm24))
        reserved_register_field_fault();

    if (imm24{1})    PSR{1} = 1;
    if (imm24{2} && PSR.sp == 0)    //non-secure perf monitor
        PSR{2} = 1;
    if (imm24{3})    PSR{3} = 1;
    if (imm24{4})    PSR{4} = 1;
    if (imm24{5})    PSR{5} = 1;
}
```

Sign Extend

Format: (qp) sxt x sz $r_1 = r_3$

I29

Description: The value in GR r_3 is sign extended from the bit position specified by x sz and the result is placed in GR r_1 . The mnemonic values for x sz are given in Table 6-44.

Table 6-44. x sz Mnemonic Values

x sz Mnemonic	Bit Position
1	7
2	15
4	31

Operation:

```
if (PR[qp]) {
    check_target_register( $r_1$ );

    GR[ $r_1$ ] = sign_ext(GR[ $r_3$ ],  $x$ sz * 8);
    GR[ $r_1$ ].nat = GR[ $r_3$ ].nat;
}
```

Memory Synchronization

Format: (*qp*) sync.i

M24

Description: `sync.i` ensures that when previously initiated Flush Cache (`fc`) operations issued by the local processor become visible to local data memory references, prior Flush Cache operations are also observed by the local processor instruction fetch stream. `sync.i` also ensures that at the time previously initiated Flush Cache (`fc`) operations are observed on a remote processor by data memory references they are also observed by instruction memory references on the remote processor. `sync.i` is ordered with respect to all cache flush operations as observed by another processor. A `sync.i` and a previous `fc` must be in separate instruction groups. If semantically required, the programmer must explicitly insert ordered data references (acquire, release or fence type) to appropriately constrain `sync.i` (and hence `fc`) visibility to the data stream on other processors.

`sync.i` is used to maintain an ordering relationship between instruction and data caches on local and remote processors. An instruction serialize operation be used to ensure synchronization initiated by `sync.i` on the local processor has been observed by a given point in program execution.

An example of self-modifying code (local processor):

```

    st [L1] = data    //store into local instruction stream
    fc L1            //flush stale datum from instruction/data cache
    ;;              //require instruction boundary between fc and sync.i
    sync.i           //ensure local and remote data/inst caches are synchronized
    ;;
    srlz.i           //ensure sync has been observed by the local processor,
    ;;              //ensure subsequent instructions observe modified memory
L1: target          //instruction modified

```

Operation:

```

if (PR[qp]) {
    instruction_synchronize();
}

```

Test Bit

Format: (qp) tbit.trel.ctype $p_1, p_2 = r_3, pos_6$

I16

Description: The bit specified by the pos_6 immediate is selected from GR r_3 . The selected bit forms a single bit result either complemented or not depending on the *trel* completer. This result is written to the two predicate register destinations p_1 and p_2 . The way the result is written to the destinations is determined by the compare type specified by *ctype*. See the Compare instruction and Table 6-10 on page 6-19.

The *trel* completer values *.nz* and *.z* indicate non-zero and zero sense of the test. For normal and *unc* types, only the *.z* value is directly implemented in hardware; the *.nz* value is actually a pseudo-op. For it, the assembler simply switches the predicate target specifiers and uses the implemented relation. For the parallel types, both relations are implemented in hardware.

Table 6-45. Test Bit Relations for Normal and *unc* tbits

<i>trel</i>	Test Relation	Pseudo-op of
<i>nz</i>	selected bit == 1	<i>z</i> $p_1 \leftrightarrow p_2$
<i>z</i>	selected bit == 0	

Table 6-46. Test Bit Relations for Parallel tbits

<i>trel</i>	Test Relation
<i>nz</i>	selected bit == 1
<i>z</i>	selected bit == 0

If the two predicate register destinations are the same (p_1 and p_2 specify the same predicate register), the instruction will take an Illegal Operation fault, if the qualifying predicate is set, or if the compare type is *unc*.

```

Operation:  if (PR[qp]) {
            if (p1 == p2)
                illegal_operation_fault();

            if (trel == 'nz') // 'nz' - test for 1
                tmp_rel = GR[r3]{pos6};
            else // 'z' - test for 0
                tmp_rel = !GR[r3]{pos6};

            switch (ctype) {
            case 'and': // and-type compare
                if (GR[r3].nat || !tmp_rel) {
                    PR[p1] = 0;
                    PR[p2] = 0;
                }
                break;
            case 'or': // or-type compare
                if (!GR[r3].nat && tmp_rel) {
                    PR[p1] = 1;
                    PR[p2] = 1;
                }
                break;
            case 'or.andcm': // or.andcm-type compare
                if (!GR[r3].nat && tmp_rel) {
                    PR[p1] = 1;
                    PR[p2] = 0;
                }
                break;
            case 'unc': // unc-type compare
            default: // normal compare
                if (GR[r3].nat) {
                    PR[p1] = 0;
                    PR[p2] = 0;
                } else {
                    PR[p1] = tmp_rel;
                    PR[p2] = !tmp_rel;
                }
                break;
            }
        } else {
            if (ctype == 'unc') {
                if (p1 == p2)
                    illegal_operation_fault();
                PR[p1] = 0;
                PR[p2] = 0;
            }
        }
    }

```


Test NaT

Format: $(qp) \text{ tнат.trel.ctype } p_1, p_2 = r_3$

117

Description: The NaT bit from GR r_3 forms a single bit result, either complemented or not depending on the *trel* completer. This result is written to the two predicate register destinations, p_1 and p_2 . The way the result is written to the destinations is determined by the compare type specified by *ctype*. See the Compare instruction and Table 6-10 on page 6-19.

The *trel* completer values *.nz* and *.z* indicate non-zero and zero sense of the test. For normal and *unc* types, only the *.z* value is directly implemented in hardware; the *.nz* value is actually a pseudo-op. For it, the assembler simply switches the predicate target specifiers and uses the implemented relation. For the parallel types, both relations are implemented in hardware.

Table 6-47. Test NaT Relations for Normal and *unc* tнатs

<i>trel</i>	Test Relation	Pseudo-op of
<i>nz</i>	selected bit == 1	<i>z</i> $p_1 \leftrightarrow p_2$
<i>z</i>	selected bit == 0	

Table 6-48. Test NaT Relations for Parallel tнатs

<i>trel</i>	Test Relation
<i>nz</i>	selected bit == 1
<i>z</i>	selected bit == 0

If the two predicate register destinations are the same (p_1 and p_2 specify the same predicate register), the instruction will take an Illegal Operation fault, if the qualifying predicate is set, or if the compare type is *unc*.

```

Operation:  if (PR[qp]) {
            if (p1 == p2)
                illegal_operation_fault();

            if (trel == 'nz') // 'nz' - test for 1
                tmp_rel = GR[r3].nat;
            else // 'z' - test for 0
                tmp_rel = !GR[r3].nat;

            switch (ctype) {
                case 'and': // and-type compare
                    if (!tmp_rel) {
                        PR[p1] = 0;
                        PR[p2] = 0;
                    }
                    break;
                case 'or': // or-type compare
                    if (tmp_rel) {
                        PR[p1] = 1;
                        PR[p2] = 1;
                    }
                    break;
                case 'or.andcm': // or.andcm-type compare
                    if (tmp_rel) {
                        PR[p1] = 1;
                        PR[p2] = 0;
                    }
                    break;
                case 'unc': // unc-type compare
                default: // normal compare
                    PR[p1] = tmp_rel;
                    PR[p2] = !tmp_rel;
                    break;
            }
        } else {
            if (ctype == 'unc') {
                if (p1 == p2)
                    illegal_operation_fault();
                PR[p1] = 0;
                PR[p2] = 0;
            }
        }
    }

```

Unpack

Format:	<i>(qp)</i> unpack1.h $r_1 = r_2, r_3$	one_byte_form, high_form	I2
	<i>(qp)</i> unpack2.h $r_1 = r_2, r_3$	two_byte_form, high_form	I2
	<i>(qp)</i> unpack4.h $r_1 = r_2, r_3$	four_byte_form, high_form	I2
	<i>(qp)</i> unpack1.l $r_1 = r_2, r_3$	one_byte_form, low_form	I2
	<i>(qp)</i> unpack2.l $r_1 = r_2, r_3$	two_byte_form, low_form	I2
	<i>(qp)</i> unpack4.l $r_1 = r_2, r_3$	four_byte_form, low_form	I2

Description: The data elements of GR r_2 and r_3 are unpacked, and the result placed in GR r_1 . In the high_form, the most significant elements of each source register are selected, while in the low_form the least significant elements of each source register are selected. Elements are selected alternately from the source registers.




Figure 6-43. Unpack Operation

```

Operation:  if (PR[qp]) {
            check_target_register(r1);

            if (one_byte_form) { // one-byte elements
                x[0] = GR[r2]{7:0};    y[0] = GR[r3]{7:0};
                x[1] = GR[r2]{15:8};   y[1] = GR[r3]{15:8};
                x[2] = GR[r2]{23:16};  y[2] = GR[r3]{23:16};
                x[3] = GR[r2]{31:24};  y[3] = GR[r3]{31:24};
                x[4] = GR[r2]{39:32};  y[4] = GR[r3]{39:32};
                x[5] = GR[r2]{47:40};  y[5] = GR[r3]{47:40};
                x[6] = GR[r2]{55:48};  y[6] = GR[r3]{55:48};
                x[7] = GR[r2]{63:56};  y[7] = GR[r3]{63:56};

                if (high_form)
                    GR[r1] = concatenate8( x[7], y[7], x[6], y[6],
                                             x[5], y[5], x[4], y[4]);
                else
                    GR[r1] = concatenate8( x[3], y[3], x[2], y[2],
                                             x[1], y[1], x[0], y[0]);
            } else if (two_byte_form) { // two-byte elements
                x[0] = GR[r2]{15:0};    y[0] = GR[r3]{15:0};
                x[1] = GR[r2]{31:16};   y[1] = GR[r3]{31:16};
                x[2] = GR[r2]{47:32};   y[2] = GR[r3]{47:32};
                x[3] = GR[r2]{63:48};   y[3] = GR[r3]{63:48};

                if (high_form)
                    GR[r1] = concatenate4(x[3], y[3], x[2], y[2]);
                else
                    GR[r1] = concatenate4(x[1], y[1], x[0], y[0]);
            } else { // four-byte elements
                x[0] = GR[r2]{31:0};    y[0] = GR[r3]{31:0};
                x[1] = GR[r2]{63:32};   y[1] = GR[r3]{63:32};

                if (high_form)
                    GR[r1] = concatenate2(x[1], y[1]);
                else
                    GR[r1] = concatenate2(x[0], y[0]);
            }
            GR[r1].nat = GR[r2].nat || GR[r3].nat;
        }

```

Exchange

Format: $(qp) \text{ xchgsz.lhint } r_1 = [r_3], r_2$ M16

Description: A value consisting of sz bytes is read from memory starting at the address specified by the value in GR r_3 . The least significant sz bytes of the value in GR r_2 are written to memory starting at the address specified by the value in GR r_3 . The value read from memory is then zero extended and placed in GR r_1 and the NaT bit corresponding to GR r_1 is cleared. The values of the sz completer are given in Table 6-49.

If the address specified by the value in GR r_3 is not naturally aligned to the size of the value being accessed in memory, an Unaligned Data Reference fault is taken independent of the state of the User Mask alignment checking bit, UM.ac (PSR.ac in the Processor Status Register).

Both read and write access privileges for the referenced page are required.

Table 6-49. Memory Exchange Size

<i>sz</i> Completer	Bytes Accessed
1	1 byte
2	2 bytes
4	4 bytes
8	8 bytes

The exchange is performed with acquire semantics, i.e., the memory read/write is made visible prior to all subsequent data memory accesses.

The memory read and write are guaranteed to be atomic.

The value of the *ldhint* completer specifies the locality of the memory access. The values of the *ldhint* completer are given in Table 6-28 on page 6-102. Locality hints do not affect program functionality and may be ignored by the implementation. See “Memory Hierarchy Control and Consistency” on page 4-16 for details.

Operation:

```

if (PR[qp]) {
    check_target_register(r1, SEMAPHORE);

    if (GR[r3].nat || GR[r2].nat)
        register_nat_consumption_fault(SEMAPHORE);

    paddr = tlb_translate(GR[r3], sz, SEMAPHORE, PSR.cpl, &mattr, &tmp_unused);

    if (!ma_supports_semaphores(mattr))
        unsupported_data_reference_fault(SEMAPHORE, GR[r3]);

    val = mem_xchg(GR[r2], paddr, sz, UM.be, mattr, ACQUIRE, ldhint);

    alat_inval_multiple_entries(paddr, sz);

    GR[r1] = zero_ext(val, sz * 8);
    GR[r1].nat = 0;
}

```

Fixed-Point Multiply Add

Format:	<i>(qp)</i> xma.l $f_1 = f_3, f_4, f_2$	low_form	F2
	<i>(qp)</i> xma.lu $f_1 = f_3, f_4, f_2$	pseudo-op of: <i>(qp)</i> xma.l $f_1 = f_3, f_4, f_2$	
	<i>(qp)</i> xma.h $f_1 = f_3, f_4, f_2$	high_form	F2
	<i>(qp)</i> xma.hu $f_1 = f_3, f_4, f_2$	high_unsigned_form	F2

Description: Two source operands (FR f_3 and FR f_4) are treated as either signed or unsigned integers and multiplied. The third source operand (FR f_2) is zero extended and added to the product. The upper or lower 64 bits of the resultant sum are selected and placed in FR f_1 .

In the high_unsigned_form, the significant fields of FR f_3 and FR f_4 are treated as unsigned integers and multiplied to produce a full 128-bit unsigned result. The significant field of FR f_2 is zero extended and added to the product. The most significant 64-bits of the resultant sum are placed in the significant field of FR f_1 .

In the high_form, the significant fields of FR f_3 and FR f_4 are treated as signed integers and multiplied to produce a full 128-bit signed result. The significant field of FR f_2 is zero extended and added to the product. The most significant 64-bits of the resultant sum are placed in the significant field of FR f_1 .

In the other forms, the significant fields of FR f_3 and FR f_4 are treated as signed integers and multiplied to produce a full 128-bit signed result. The significant field of FR f_2 is zero extended and added to the product. The least significant 64-bits of the resultant sum are placed in the significant field of FR f_1 .

In all forms, the exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0). Note: f1 as an operand is not an integer 1; it is just the register file format's 1.0 value.

In all forms, if any of FR f_3 , FR f_4 , or FR f_2 is a NaTVal, FR f_1 is set to NaTVal instead of the computed result.

Operation:

```

if (PR[qp]) {
    fp_check_target_register(f1);
    if (tmp_israncode = fp_reg_disabled(f1, f2, f3, f4))
        disabled_fp_register_fault(tmp_israncode, 0);

    if (fp_is_natval(FR[f2]) || fp_is_natval(FR[f3]) || fp_is_natval(FR[f4])) {
        FR[f1] = NATVAL;
    } else {
        if (low_form || high_form)
            tmp_res_128 =
                fp_I64_x_I64_to_I128(FR[f3].significand, FR[f4].significand);
        else // high_unsigned_form
            tmp_res_128 =
                fp_U64_x_U64_to_U128(FR[f3].significand, FR[f4].significand);

        tmp_res_128 =
            fp_U128_add(tmp_res_128, fp_U64_to_U128(FR[f2].significand));

        if (high_form || high_unsigned_form)
            FR[f1].significand = tmp_res_128.hi;
        else // low_form
            FR[f1].significand = tmp_res_128.lo;

        FR[f1].exponent = FP_INTEGER_EXP;
        FR[f1].sign = FP_SIGN_POSITIVE;
    }

    fp_update_psr(f1);
}

```

Fixed-Point Multiply

Format:

(qp) xmpy.l $f_1 = f_3, f_4$	pseudo-op of: (qp) xma.l $f_1 = f_3, f_4, f_0$
(qp) xmpy.lu $f_1 = f_3, f_4$	pseudo-op of: (qp) xma.l $f_1 = f_3, f_4, f_0$
(qp) xmpy.h $f_1 = f_3, f_4$	pseudo-op of: (qp) xma.h $f_1 = f_3, f_4, f_0$
(qp) xmpy.hu $f_1 = f_3, f_4$	pseudo-op of: (qp) xma.hu $f_1 = f_3, f_4, f_0$

Description: Two source operands (FR f_3 and FR f_4) are treated as either signed or unsigned integers and multiplied. The upper or lower 64 bits of the resultant product are selected and placed in FR f_1 .

In the high_unsigned_form, the significand fields of FR f_3 and FR f_4 are treated as unsigned integers and multiplied to produce a full 128-bit unsigned result. The most significant 64-bits of the resultant product are placed in the significand field of FR f_1 .

In the high_form, the significand fields of FR f_3 and FR f_4 are treated as signed integers and multiplied to produce a full 128-bit signed result. The most significant 64-bits of the resultant product are placed in the significand field of FR f_1 .

In the other forms, the significand fields of FR f_3 and FR f_4 are treated as signed integers and multiplied to produce a full 128-bit signed result. The least significant 64-bits of the resultant product are placed in the significand field of FR f_1 .

In all forms, the exponent field of FR f_1 is set to the biased exponent for 2.0^{63} (0x1003E) and the sign field of FR f_1 is set to positive (0). Note: f1 as an operand is not an integer 1; it is just the register file format's 1.0 value.

Operation: See "Fixed-Point Multiply Add" on page 6-181.

Zero Extend

Format: $(qp) \text{ zxt}xsz \ r_1 = r_3$

129

Description: The value in GR r_3 is zero extended above the bit position specified by xsz and the result is placed in GR r_1 . The mnemonic values for xsz are given in Table 6-44 on page 6-171.

Operation:

```
if (PR[qp]) {
    check_target_register( $r_1$ );

    GR[ $r_1$ ] = zero_ext(GR[ $r_3$ ],  $xsz * 8$ );
    GR[ $r_1$ ].nat = GR[ $r_3$ ].nat;
}
```

A Instruction Sequencing Considerations

Instruction execution consists of four phases:

1. Read the instruction from memory (*fetch*)
2. Read architectural state, if necessary (*read*)
3. Perform the specified operation (*execute*)
4. Update architectural state, if necessary (*update*).

An **instruction group** is a sequence of instructions starting at a given bundle address and slot number and including all instructions at sequentially increasing slot numbers and bundle addresses up to the first stop or taken branch. For the instructions in an instruction group to have well-defined behavior, they must meet the ordering and dependency requirements described below.

If the instructions in instruction groups meet the resource-dependency requirements, then the behavior of a program will be as though each individual instruction is sequenced through these phases in the order listed above. The order of a phase of a given instruction relative to any phase of a previous instruction is prescribed by the instruction sequencing rules below.

- There is no a priori relationship between the *fetch* of an instruction and the *read*, *execute*, or *update* of any dynamically previous instruction. The `sync.i` and `sr1z.i` instructions can be used to enforce a sequential relationship between the *fetch* of all succeeding instructions and the *update* of all previous instructions.
- Between instruction groups, every instruction in a given instruction group will behave as though its read occurred after the update of all the instructions from the previous instruction group. All instructions are assumed to have unit latency. Instructions on opposing sides of a stop are architecturally considered to be separated by at least one unit of latency.

Some system state updates require more stringent requirements than those described here.

- Within an instruction group, every instruction will behave as though its read of the memory and ALAT state occurred after the update of the memory and ALAT state of all prior instructions in that instruction group.
- Within an instruction group, every instruction will behave as though its read of the register state occurred before the update of the register state by any instruction (prior or later) in that instruction group, except as noted in the dependency restrictions section below.

The ordering rules above form the context for register dependency restrictions, memory dependency restrictions and the order of exception reporting. These dependency restrictions apply only between instructions whose resource reads and writes are not dynamically disabled by predication.

- Register dependencies: Within an instruction group, read-after-write (RAW) and write-after-write (WAW) register dependencies are not allowed (except as noted in “RAW Ordering Exceptions” on page A-2 and “WAW Ordering Exceptions” on page A-3). Write-after-read (WAR) register dependencies are allowed (except as noted in “WAR Ordering Exceptions” on page A-3).

These dependency restrictions apply to both explicit register accesses (from the instruction’s operands) and implicit register accesses (such as application and control registers implicitly accessed by certain instructions). Predicate register PR0 is excluded from these register dependency restrictions, since writes to PR0 are ignored and reads always return 1 (one).

- Memory dependencies: Within an instruction group, RAW, WAW, and WAR memory dependencies and ALAT dependencies are allowed. A load will observe the results of the most recent store to the same memory address. In the event that multiple stores to the same address are present in the same instruction group, memory will contain the result of the latest store after execution of the instruction group. A store following a load to the same address will not

affect the data loaded by the load. Advanced loads, check loads, advanced load checks, stores, and memory semaphore instructions implicitly access the ALAT. RAW, WAW, and WAR ALAT dependencies are allowed within an instruction group and behave as described for memory dependencies.

The net effect of the dependency restrictions stated above is that a processor may execute all (or any subset) of the instructions within a legal instruction group concurrently or serially with the end result being identical. If these dependency restrictions are not met, the behavior of the program is undefined.

The instruction sequencing resulting from the rules stated above is termed sequential execution.

The ordering rules and the dependency restrictions allow the processor to dynamically re-order instructions, execute instructions with non-unit latency, or even concurrently execute instructions on opposing sides of a stop or taken branch, provided that correct sequencing is enforced and the appearance of sequential execution is presented to the programmer.

IP is a special resource in that reads and writes of IP behave as though the instruction stream was being executed serially, rather than in parallel. RAW dependencies on IP are allowed, and the reader gets the IP of the bundle in which it is contained. So, each bundle being executed in parallel logically reads IP, increments it and writes it back. WAW is also allowed.

Ignored ARs are not exceptional for dependency checking purposes. RAW and WAW dependencies to ignored ARs are not allowed.

A.1 RAW Ordering Exceptions

There are four exceptions to the rule prohibiting RAW register dependencies within an instruction group. These exceptions are the `alloc` instruction, check load instructions, instructions that affect branching, and the `ld8.fill` and `st8.spill` instructions.

- The `alloc` instruction implicitly writes the Current Frame Marker (CFM) which is implicitly read by all instructions accessing the stacked subset of the general register file. Instructions that access the stacked subset of the general register file may appear in the same instruction group as `alloc` and will see the stack frame specified by the `alloc`. Note that some instructions have RAW or WAW dependences on resources other than CFM affected by `alloc` and are thus not allowed in the same instruction group after an `alloc`: `flushrs`, `move from AR[BSPSTORE]`, `move from AR[RNAT]`, `br.cexit`, `br.ctop`, `br.wexit`, `br.wtop`, `br.call`, `br.ia`, `br.ret`, `clrrrb`. Note that `alloc` is required to be the first instruction in an instruction group.
- A check load instruction may or may not perform a load since it is dependent upon its corresponding advanced load. If the check load misses the ALAT it will execute a load from memory. A check load and a subsequent instruction that reads the target of the check load may exist in the same instruction group. The dependent instruction will get the new value loaded by the check load.
- A branch may read branch registers and may implicitly read predicate registers, the LC, EC, and PFS application registers, as well as CFM. Except for LC, EC and predicate registers, writes to any of these registers by a non-branch instruction will be visible to a subsequent branch in the same instruction group. Writes to predicate registers by any non-floating-point instruction will be visible to a subsequent branch in the same instruction group. RAW register dependencies within the same instruction group are not allowed for LC and EC. Dynamic RAW dependencies where the predicate writer is a floating-point instruction and the reader is a branch are also not allowed within the same instruction group. Branches `br.cond`, `br.call`, `br.ret` and `br.ia` work like other instructions for the purposes of register dependency; i.e., if their qualifying predicate is 0, they are not considered readers or writers of other resources. Branches `br.cloop`, `br.cexit`, `br.ctop`, `br.wexit`, and `br.wtop` are exceptional in that they are always readers or writers of their resources, regardless of the value of their qualifying predicate.
- The `ld8.fill` and `st8.spill` instructions implicitly access the User NaT Collection application register (UNAT). For these instructions the restriction on dynamic RAW register dependencies with respect to UNAT applies at the bit level. These instructions may appear in the same instruction group provided they do not access the same bit of UNAT. RAW UNAT dependencies between `ld8.fill` or `st8.spill` instructions and `mov ar=` or `mov =ar` instructions accessing UNAT must not occur within the same instruction group.

For the purposes of resource dependencies, CFM is treated as a single resource.

A.2 WAW Ordering Exceptions

There are three exceptions to the rule prohibiting WAW register dependencies within an instruction group. The exceptions are compare-type instructions, floating-point instructions, and the `st8.spill` instruction.

- The set of compare-type instructions includes: `cmp`, `cmp4`, `tbit`, `tnat`, `fcmp`, `frsqrta`, `frcpa`, and `fclass`. Compare-type instructions in the same instruction group may target the same predicate register provided:
 - The compare-type instructions are either all AND-type compares or all OR-type compares (AND-type compares correspond to “.and” and “.andcm” completers; OR-type compares correspond to “.or” and “.orcm” completers), or
 - The compare-type instructions all target PR0. All WAW dependencies for PR0 are allowed; the compares can be of any types and can be of differing types.

All other WAW dependencies within an instruction group are disallowed, including dynamic WAW register dependencies with move to PR instructions that access the same predicate registers as another writer. Note that the move to PR instruction only writes those PRs indicated by its mask, but the move from PR instruction always reads all the predicate registers.

- Floating-point instructions implicitly write the Floating-Point Status Register (FPSR) and the Processor Status Register (PSR). Multiple floating-point instructions may appear in the same instruction group since the restriction on WAW register dependencies with respect to the FPSR and PSR do not apply. The state of FPSR and PSR after executing the instruction group will be the logical OR of all writes.
- The `st8.spill` instruction implicitly writes the UNAT register. For this instruction the restriction on WAW register dependencies with respect to UNAT applies at the bit level. Multiple `st8.spill` instructions may appear in the same instruction group provided they do not write the same bit of UNAT. WAW register dependencies between `st8.spill` instructions and `mov ar=` instructions targeting UNAT must not occur within the same instruction group.

WAW dependencies to ignored ARs are not allowed.

A.3 WAR Ordering Exceptions

WAR dependence between the reading of PR63 by a branch instruction and the subsequent writing of PR63 by a loop closing branch (`br.ctop`, `br.cexit`, `br.wtop`, or `br.wexit`) in the same instruction group is not allowed. Otherwise, WAR dependencies are allowed.

B

IA-64 Pseudo-Code Functions

This appendix contains a table of pseudo-code functions used in Chapter 6, "IA-64 Instruction Reference".

Table B-1. Pseudo-Code Functions

Function	Operation
<code>xxx_fault(parameters ...)</code>	There are several fault functions. Each fault function accepts parameters specific to the fault, e.g., exception code values, virtual addresses, etc. If the fault is deferred for speculative load exceptions the fault function will return with a deferral indication. Otherwise, fault routines do not return and terminate the instruction sequence.
<code>xxx_trap(parameters ...)</code>	There are several trap functions. Each trap function accepts parameters specific to the trap, e.g., trap code values, virtual addresses, etc. Trap routines do not return.
<code>acceptance_fence()</code>	Ensures prior data memory references to uncached ordered-sequential memory pages are “accepted”, before subsequent data memory references are performed by the processor.
<code>alat_cmp(rtype, raddr)</code>	Returns a one if the implementation finds an ALAT entry which matches the register type specified by <code>rtype</code> and the register address specified by <code>raddr</code> , else returns zero. This function is implementation specific. Note that an implementation may optionally choose to return zero (indicating no match) even if a matching entry exists in the ALAT. This provides implementation flexibility in designing fast ALAT lookup circuits.
<code>alat_frame_update(delta_bof, delta_sof)</code>	Notifies the ALAT of a change in the bottom of frame and/or size of frame. This allows management of the ALAT’s tag bits or other management functions it might need.
<code>alat_inval()</code>	Invalidate all entries in the ALAT.
<code>alat_inval_multiple_entries(paddr, size)</code>	The ALAT is queried using the physical memory address specified by <code>paddr</code> and the access size specified by <code>size</code> . All matching ALAT entries are invalidated. No value is returned.
<code>alat_inval_single_entry(rtype, rega)</code>	The ALAT is queried using the register type specified by <code>rtype</code> and the register address specified by <code>rega</code> . At most one matching ALAT entry is invalidated. No value is returned.
<code>alat_write(rtype, raddr, paddr, size)</code>	Allocates a new ALAT entry using the register type specified by <code>rtype</code> , the register address specified by <code>raddr</code> , the physical memory address specified by <code>paddr</code> , and the access size specified by <code>size</code> . No value is returned. This function guarantees that only one ALAT entry exists for a given <code>raddr</code> . If a <code>ld.c.nc</code> , <code>ldf.c.nc</code> , or <code>ldfp.c.nc</code> instruction’s <code>raddr</code> matches an existing ALAT entry’s register tag, but the instruction’s <code>size</code> and/or <code>paddr</code> are different than that of the existing entry’s; then this function may either preserve the existing entry, or invalidate it and write a new entry with the instruction’s specified <code>size</code> and <code>paddr</code> .
<code>check_target_register(r1)</code>	If <code>r1</code> targets an out-of-frame stacked register (as defined by CFM), an illegal operation fault is delivered, and this function does not return.
<code>check_target_register_sof(r1, new-sof)</code>	If <code>r1</code> targets an out-of-frame stacked register (as defined by the <code>newsof</code> parameter), an illegal operation fault is delivered, and this function does not return.

Table B-1. Pseudo-Code Functions (Continued)

Function	Operation
concatenate2(x1, x2)	Concatenates the lower 32 bits of the 2 arguments, and returns the 64-bit result.
concatenate4(x1, x2, x3, x4)	Concatenates the lower 16 bits of the 4 arguments, and returns the 64-bit result.
concatenate8(x1, x2, x3, x4, x5, x6, x7, x8)	Concatenates the lower 8 bits of the 8 arguments, and returns the 64-bit result.
fadd(fp_dp, fr2)	Adds a floating-point register value to the infinitely precise product and return the infinitely precise sum, ready for rounding.
fcmp_exception_fault_check(fr2, fr3, fr1, sf, *tmp_fp_env)	Checks for all floating-point faulting conditions for the fcmp instruction.
fcvt_fx_exception_fault_check(fr2, trunc, sf *tmp_fp_env)	Checks for all floating-point faulting conditions for the fcvt.fx and fcvt.fx.trunc instructions. It propagates NaNs, and NaTVals.
fcvt_fxu_exception_fault_check(fr2, trunc, sf, *tmp_fp_env)	Checks for all floating-point faulting conditions for the fcvt.fxu and fcvt.fxu.trunc instructions. It propagates NaNs, and NaTVals.
fma_exception_fault_check(fr2, fr3, fr4, pc, sf, *tmp_fp_env)	Checks for all floating-point faulting conditions for the fma instruction. It propagates NaNs, NaTVals, and special IEEE results.
fminmax_exception_fault_check(fr2, fr3, sf, *tmp_fp_env)	Checks for all floating-point faulting conditions for the famax, famin, fmax, and fmin instructions.
fms_fnma_exception_fault_check(fr2, fr3, fr4, pc, sf, *tmp_fp_env)	Checks for all floating-point faulting conditions for the fms and fnma instructions. It propagates NaNs, NaTVals, and special IEEE results.
fmul(fr3, fr4)	Performs an infinitely precise multiply of two floating-point register values.
followed_by_stop()	Returns TRUE if the current instruction is followed by a stop; otherwise, returns FALSE.
fp_check_target_register(f1)	If the specified floating-point register identifier is 0 or 1, this function causes an illegal operation fault.
fp_decode_fault(tmp_fp_env)	Returns floating-point exception fault code values for ISR.code.
fp_decode_traps(tmp_fp_env)	Returns floating-point trap code values for ISR.code.
fp_is_nan_or_inf(freg)	Returns true if the floating-point exception_fault_check functions returned a IEEE fault disabled default result or a propagated NaN.
fp_equal(fr1, fr2)	IEEE standard equality relationship test.
fp_ieee_recip(num, den)	Returns the true quotient for special sets of operands, or an approximation to the reciprocal of the divisor to be used in the software divide algorithm.
fp_ieee_recip_sqrt(root)	Returns the true square root result for special operands, or an approximation to the reciprocal square root to be used in the software square root algorithm.
fp_is_nan(freg)	Returns true when floating register contains a NaN.
fp_is_natval(freg)	Returns true when floating register contains a NaTVal
fp_is_normal(freg)	Returns true when floating register contains a normal number.
fp_is_pos_inf(freg)	Returns true when floating register contains a positive infinity.
fp_is_qnan(freg)	Returns true when floating register contains a quiet NaN.
fp_is_snan(freg)	Returns true when floating register contains a signalling NaN.
fp_is_unorm(freg)	Returns true when floating register contains an unnormalized number.
fp_is_unsupported(freg)	Returns true when floating register contains an unsupported format.
fp_less_than(fr1, fr2)	IEEE standard less-than relationship test.
fp_lesser_or_equal(fr1, fr2)	IEEE standard less-than or equal-to relationship test
fp_normalize(fr1)	Normalizes an unnormalized fp value. This function flushes to zero any unnormal values which can not be represented in the register file
fp_raise_fault(tmp_fp_env)	Checks the local instruction state for any faulting conditions which require an interruption to be raised.

Table B-1. Pseudo-Code Functions (Continued)

Function	Operation
<code>fp_raise_traps(tmp_fp_env)</code>	Checks the local instruction state for any trapping conditions which require an interruption to be raised.
<code>fp_reg_bank_conflict(f1, f2)</code>	Returns true if the two specified FRs are in the same bank.
<code>fp_reg_disabled(f1, f2, f3, f4)</code>	Check for possible disabled floating-point register faults.
<code>fp_reg_read(freg)</code>	Reads the FR and gives canonical double-extended denormals (and pseudo-denormals) their true mathematical exponent. Other classes of operands are unaltered.
<code>fp_unordered(fr1, fr2)</code>	IEEE standard unordered relationship
<code>fp_fr_to_mem_format(freg, size)</code>	Converts a floating-point value in register format to floating-point memory format. It assumes that the floating-point value in the register has been previously rounded to the correct precision which corresponds with the <code>size</code> parameter.
<code>frcpa_exception_fault_check(fr2, fr3, sf, *tmp_fp_env)</code>	Checks for all floating-point faulting conditions for the <code>frcpa</code> instruction. It propagates NaNs, NaTVals, and special IEEE results.
<code>frsqrrta_exception_fault_check(fr3, sf, *tmp_fp_env)</code>	Checks for all floating-point faulting conditions for the <code>frsqrrta</code> instruction. It propagates NaNs, NaTVals, and special IEEE results
<code>ignored_field_mask(regclass, reg, value)</code>	Boolean function that returns value with bits cleared to 0 corresponding to ignored bits for the specified register and register type.
<code>instruction_serialize()</code>	Ensures all prior register updates with side-effects are observed before subsequent instruction and data memory references are performed. Also ensures prior SYNC.i operations have been observed by the instruction cache.
<code>instruction_synchronize</code>	Synchronizes the instruction and data stream for Flush Cache operations. This function ensures that when prior FC operations are observed by the local data cache they are observed by the local instruction cache, and when prior FC operations are observed by another processor's data cache they are observed within the same processor's instruction cache.
<code>is_finite(freg)</code>	Returns true when floating register contains a finite number.
<code>is_ignored_reg(regnum)</code>	Boolean function that returns true if <code>regnum</code> is an ignored application register, otherwise false.
<code>is_inf(freg)</code>	Returns true when floating register contains an infinite number.
<code>is_kernel_reg(ar_addr)</code>	Returns a one if <code>ar_addr</code> is the address of a kernel register application register
<code>is_reserved_field(regclass, arg2, arg3)</code>	Returns true if the specified data would write a one in a reserved field.
<code>is_reserved_reg(regclass, regnum)</code>	Returns true if register <code>regnum</code> is reserved in the <code>regclass</code> register file.
<code>mem_flush(paddr)</code>	The line addressed by the physical address <code>paddr</code> is invalidated in all levels of the memory hierarchy above memory and written back to memory if it is inconsistent with memory.
<code>mem_implicit_prefetch(vaddr, hint)</code>	Moves the line addressed by <code>vaddr</code> to the location of the memory hierarchy specified by <code>hint</code> . This function is implementation dependent and can be ignored.
<code>mem_promote(paddr, mtype, hint)</code>	Moves the line addressed by <code>paddr</code> to the highest level of the memory hierarchy conditioned by the access hints specified by <code>hint</code> . Implementation dependent and can be ignored.
<code>mem_read(paddr, size, border, mattr, otype, hint)</code>	Returns the <code>size</code> bytes starting at the physical memory location specified by <code>paddr</code> with byte order specified by <code>border</code> , memory attributes specified by <code>mattr</code> , and access hint specified by <code>hint</code> . <code>otype</code> specifies the memory ordering attribute of this access, and must be UNORDERED or ACQUIRE.
<code>fp_mem_to_fr_format(mem, size)</code>	Converts a floating-point value in memory format to floating-point register format.

Table B-1. Pseudo-Code Functions (Continued)

Function	Operation
<code>mem_write(value, paddr, size, border, mattr, otype, hint)</code>	Writes the least significant <code>size</code> bytes of <code>value</code> into memory starting at the physical memory address specified by <code>paddr</code> with byte order specified by <code>border</code> , memory attributes specified by <code>mattr</code> , and access hint specified by <code>hint</code> . <code>otype</code> specifies the memory ordering attribute of this access, and must be <code>UNORDERED</code> or <code>RELEASE</code> . No value is returned.
<code>mem_xchg(data, paddr, size, byte_order, mattr, otype, hint)</code>	Returns <code>size</code> bytes from memory starting at the physical address specified by <code>paddr</code> . The read is conditioned by the locality hint specified by <code>hint</code> . After the read, the least significant <code>size</code> bytes of <code>data</code> are written to <code>size</code> bytes in memory starting at the physical address specified by <code>paddr</code> . The read and write are performed atomically. Both the read and the write are conditioned by the memory attribute specified by <code>mattr</code> and the byte ordering in memory is specified by <code>byte_order</code> . <code>otype</code> specifies the memory ordering attribute of this access, and must be <code>ACQUIRE</code> .
<code>mem_xchg_add(add_val, paddr, size, byte_order, mattr, otype, hint)</code>	Returns <code>size</code> bytes from memory starting at the physical address specified by <code>paddr</code> . The read is conditioned by the locality hint specified by <code>hint</code> . The least significant <code>size</code> bytes of the sum of the value read from memory and <code>add_val</code> is then written to <code>size</code> bytes in memory starting at the physical address specified by <code>paddr</code> . The read and write are performed atomically. Both the read and the write are conditioned by the memory attribute specified by <code>mattr</code> and the byte ordering in memory is specified by <code>byte_order</code> . <code>otype</code> specifies the memory ordering attribute of this access, and has the value <code>ACQUIRE</code> or <code>RELEASE</code> .
<code>mem_xchg_cond(cmp_val, data, paddr, size, byte_order, mattr, otype, hint)</code>	Returns <code>size</code> bytes from memory starting at the physical address specified by <code>paddr</code> . The read is conditioned by the locality hint specified by <code>hint</code> . If the value read from memory is equal to <code>cmp_val</code> , then the least significant <code>size</code> bytes of <code>data</code> are written to <code>size</code> bytes in memory starting at the physical address specified by <code>paddr</code> . If the write is performed, the read and write are performed atomically. Both the read and the write are conditioned by the memory attribute specified by <code>mattr</code> and the byte ordering in memory is specified by <code>byte_order</code> . <code>otype</code> specifies the memory ordering attribute of this access, and has the value <code>ACQUIRE</code> or <code>RELEASE</code> .
<code>ordering_fence()</code>	Ensures prior data memory references are made visible before future data memory references are made visible by the processor.
<code>pr_phys_to_virt(phys_id)</code>	Returns the virtual register id of the predicate from the physical register id, <code>phys_id</code> of the predicate.
<code>rotate_regs()</code>	Decrements the Register Rename Base registers, effectively rotating the register files. <code>CFM.rrb.gr</code> is decremented only if <code>CFM.sor</code> is non-zero.
<code>rse_enable_current_frame_load()</code>	If the RSE load pointer (<code>RSE.BSPload</code>) is greater than <code>AR[BSP]</code> , the <code>RSE.CFLE</code> bit is set to indicate that mandatory RSE loads are allowed to restore registers in the current frame (in no other case does the RSE spill or fill registers in the current frame). This function does not perform mandatory RSE loads. This procedure does not cause any interruptions.
<code>rse_invalidate_non_current_regs()</code>	All registers outside the current frame are invalidated.
<code>rse_new_frame(current_frame_size, new_frame_size)</code>	A new frame is defined without changing any register renaming. The new frame size is completely defined by the <code>new_frame_size</code> parameter (successive calls are not cumulative). If <code>new_frame_size</code> is larger than <code>current_frame_size</code> and the number of registers in the invalid and clean partitions is less than the size of frame growth then mandatory RSE stores are issued until enough registers are available. The resulting sequence of RSE stores may be interrupted. Mandatory RSE stores may cause interruptions; see <code>rse_store</code> for a list.

Table B-1. Pseudo-Code Functions (Continued)

Function	Operation
rse_preserve_frame(preserved_frame_size)	The number of registers specified by <code>preserved_frame_size</code> are marked to be preserved by the RSE. Register renaming causes the <code>preserved_frame_size</code> registers after <code>GR[32]</code> to be renamed to <code>GR[32]</code> . <code>AR[BSP]</code> is updated to contain the backing store address where the new <code>GR[32]</code> will be stored.
rse_store(type)	Saves a register or NaT collection to the backing store (<code>store_address = AR[BSPSTORE]</code>). If <code>store_address{8:3}</code> is equal to <code>0x3f</code> then the NaT collection <code>AR[RNAT]</code> is stored. If <code>store_address{8:3}</code> is not equal to <code>0x3f</code> then the register <code>RSE.StoreReg</code> is stored and the NaT bit from that register is deposited in <code>AR[RNAT]{store_address{8:3}}</code> . If the store is successful <code>AR[BSPSTORE]</code> is incremented by 8. If the store is successful and a register was stored <code>RSE.StoreReg</code> is incremented by 1 (possibly wrapping in the stacked registers). This store moves a register from the dirty partition to the clean partition. The privilege level of the store is obtained from <code>AR[RSC].pl</code> . The byte order of the store is obtained from <code>AR[RSC].be</code> . For mandatory RSE stores, <code>type</code> is <code>MANDATORY</code> . RSE stores do not invalidate ALAT entries.
rse_update_internal_stack_pointers(new_store_pointer)	Given a new value for <code>AR[BSPSTORE]</code> (<code>new_store_pointer</code>) this function computes the new value for <code>AR[BSP]</code> . This value is equal to <code>new_store_pointer</code> plus the number of dirty registers plus the number of intervening NaT collections. This means that the size of the dirty partition is the same before and after a write to <code>AR[BSPSTORE]</code> . All clean registers are moved to the invalid partition.
sign_ext(value, pos)	Returns a 64 bit number with bits <code>pos-1</code> through 0 taken from <code>value</code> and bit <code>pos-1</code> of <code>value</code> replicated in bit positions <code>pos</code> through 63. If <code>pos</code> is greater than or equal to 64, <code>value</code> is returned.
tlb_translate(vaddr, size, type, cpl, *attr, *defer)	Returns the translated data physical address for the specified virtual memory address (<code>vaddr</code>) when translation enabled; otherwise, returns <code>vaddr</code> . <code>size</code> specifies the size of the access, <code>type</code> specifies the type of access (e.g., read, write, advance, spec). <code>cpl</code> specifies the privilege level for access checking purposes. <code>*attr</code> returns the mapped physical memory attribute. If any fault conditions are detected and deferred, <code>tlb_translate</code> returns with <code>*defer</code> set. If a fault is generated but the fault is not deferred, <code>tlb_translate</code> does not return.
tlb_translate_nonaccess(vaddr, type)	Returns the translated data physical address for the specified virtual memory address (<code>vaddr</code>). <code>type</code> specifies the type of access (e.g., FC). If a fault is generated, <code>tlb_translate_nonaccess</code> does not return.
unimplemented_physical_address(paddr)	Return <code>TRUE</code> if the presented physical address is unimplemented on this processor model; <code>FALSE</code> otherwise. This function is model-specific.
impl_undefined_natd_gr_read(paddr, size, be, mattr, otype, ldhint)	defines register return data for a speculative load to a NaTed address. This function may return data from another address space.
unimplemented_virtual_address(vaddr)	Return <code>TRUE</code> if the presented virtual address is unimplemented on this processor model; <code>FALSE</code> otherwise. This function is model-specific.
fp_update_fpsr(sf, tmp_fp_env)	Copies a floating-point instruction's local state into the global FPSR.
zero_ext(value, pos)	Returns a 64 bit unsigned number with bits <code>pos-1</code> through 0 taken from <code>value</code> and zeroes in bit positions <code>pos</code> through 63. If <code>pos</code> is greater than or equal to 64, <code>value</code> is returned.

C

IA-64 Instruction Formats

Each IA-64 instruction is categorized into one of six types; each instruction type may be executed on one or more execution unit types. Table C-1 lists the instruction types and the execution unit type on which they are executed:

Table C-1. Relationship Between Instruction Type and Execution Unit Type

Instruction Type	Description	Execution Unit Type
A	Integer ALU	I-unit or M-unit
I	Non-ALU integer	I-unit
M	Memory	M-unit
F	Floating-point	F-unit
B	Branch	B-unit
L+X	Extended	I-unit

Three instructions are grouped together into 128-bit sized and aligned containers called **bundles**. Each bundle contains three 41-bit **instruction slots** and a 5-bit template field. The format of a bundle is depicted in Figure C-1.




Figure C-1. Bundle Format

The template field specifies two properties: stops within the current bundle, and the mapping of instruction slots to execution unit types. Not all combinations of these two properties are allowed - Table C-2 indicates the defined combinations. The three rightmost columns correspond to the three instruction slots in a bundle; listed within each column is the execution unit type controlled by that instruction slot for each encoding of the template field. A double line to the right of an instruction slot indicates that a stop occurs at that point within the current bundle. See “Instruction Encoding Overview” on page 3-11 for the definition of a stop. Within a bundle, execution order proceeds from slot 0 to slot 2. Unused template values (appearing as empty rows in Table C-2) are reserved and cause an Illegal Operation fault.

Extended instructions, used for long immediate integer instructions, occupy two instruction slots.

Table C-2. Template Field Encoding and Instruction Slot Mapping

Template	Slot 0	Slot 1	Slot 2
00	M-unit	I-unit	I-unit
01	M-unit	I-unit	I-unit
02	M-unit	I-unit	I-unit
03	M-unit	I-unit	I-unit
04	M-unit	L-unit	X-unit
05	M-unit	L-unit	X-unit
06			
07			
08	M-unit	M-unit	I-unit
09	M-unit	M-unit	I-unit
0A	M-unit	M-unit	I-unit
0B	M-unit	M-unit	I-unit
0C	M-unit	F-unit	I-unit
0D	M-unit	F-unit	I-unit
0E	M-unit	M-unit	F-unit
0F	M-unit	M-unit	F-unit
10	M-unit	I-unit	B-unit
11	M-unit	I-unit	B-unit
12	M-unit	B-unit	B-unit
13	M-unit	B-unit	B-unit
14			
15			
16	B-unit	B-unit	B-unit
17	B-unit	B-unit	B-unit
18	M-unit	M-unit	B-unit
19	M-unit	M-unit	B-unit
1A			
1B			
1C	M-unit	F-unit	B-unit
1D	M-unit	F-unit	B-unit
1E			
1F			

C.1 Format Summary

All instructions in the instruction set are 41 bits in length. The leftmost 4 bits (40:37) of each instruction are the major opcode. Table C-3 shows the major opcode assignments for each of the 5 instruction types — ALU (A), Integer (I), Memory (M), Floating-point (F), and Branch (B). Bundle template bits are used to distinguish among the 4 columns, so the same major op values can be reused in each column.

Unused major ops (appearing as blank entries in Table C-3) behave in one of three ways:

- Ignored major ops (white entries in Table C-3) execute as `nop` instructions.
- Reserved major ops (light gray in the grayscale version of Table C-3, brown in the color version) cause an Illegal Operation fault.
- Reserved if PR[qp] is 1 major ops (dark gray in the grayscale version of Table C-3, purple in the color version) cause an Illegal Operation fault if the predicate register specified by the qp field of the instruction (bits 5:0) is 1 and execute as a `nop` instruction if 0.

Table C-3. Major Opcode Assignments

Major Op (bits 40:37)	Instruction Type				
	I/A	M/A	F	B	L+X
0	Misc ⁰	Mem Mgmt ⁰	FP Misc ⁰	Misc/Indirect Branch ⁰	Misc ⁰
1		Mem Mgmt ¹	FP Misc ¹	Indirect Call ¹	
2				Nop ²	
3					
4	Deposit ⁴	Int Ld +Reg/getf ⁴	FP Compare ⁴	IP-relative Branch ⁴	
5	Shift/Test Bit ⁵	Int Ld/St +Imm ⁵	FP Class ⁵	IP-rel Call ⁵	
6		FP Ld/St +Reg/setf ⁶			movl ⁶
7	MM Mpy/Shift ⁷	FP Ld/St +Imm ⁷			
8	ALU/MM ALU ⁸	ALU/MM ALU ⁸	fma ⁸		
9	Add Imm ₂₂ ⁹	Add Imm ₂₂ ⁹	fma ⁹		
A			fms ^A		
B			fms ^B		
C	Compare ^C	Compare ^C	fnma ^C		
D	Compare ^D	Compare ^D	fnma ^D		
E	Compare ^E	Compare ^E	fselect/xma ^E		
F					

Table C-4 on page C-4 summarizes all the instruction formats. The instruction fields are color-coded for ease of identification, as described in Table C-5 on page C-6.

The instruction field names, used throughout this chapter, are described in Table C-6 on page C-7. The set of special notations (such as whether an instruction must be first in an instruction group) are listed in Table C-7 on page C-7. These notations appear in the “Instruction” column of the opcode tables.

Most instruction containing immediates encode those immediates in more than one instruction field. For example, the 14-bit immediate in the Add Imm₁₄ instruction (format A4) is formed from the imm_{7b}, imm_{6d}, and s fields. Table C-65 on page C-57 shows how the immediates are formed from the instruction fields for each instruction which has an immediate.

Table C-4. Instruction Format Summary

		40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0					
ALU	A1	8			x_{2a}	v_e		x_4		x_{2b}			r_3								r_2																						r_1		QP		
Shift L and Add	A2	8			x_{2a}	v_e		x_4		ct_{2d}			r_3								r_2																							r_1		QP	
ALU Imm ₈	A3	8	s		x_{2a}	v_e		x_4		x_{2b}			r_3								imm_{7b}																						r_1		QP		
Add Imm ₁₄	A4	8	s		x_{2a}	v_e		imm_{6d}					r_3								imm_{7b}																						r_1		QP		
Add Imm ₂₂	A5	9	s					imm_{9d}					imm_{5c}	r_3							imm_{7b}																						r_1		QP		
Compare	A6	C - E	t_b	x_2	t_a			p_2					r_3								r_2		c																		p_1		QP				
Compare to Zero	A7	C - E	t_b	x_2	t_a			p_2					r_3								0		c																		p_1		QP				
Compare Imm ₈	A8	C - E	s	x_2	t_a			p_2					r_3								imm_{7b}		c																		p_1		QP				
MM ALU	A9	8		z_a	x_{2a}	z_b		x_4		x_{2b}			r_3								r_2																					r_1		QP			
MM Shift and Add	A10	8		z_a	x_{2a}	z_b		x_4		ct_{2d}			r_3								r_2																					r_1		QP			
MM Multiply Shift	I1	7		z_a	x_{2a}	z_b	v_e	ct_{2d}		x_{2b}			r_3								r_2																						r_1		QP		
MM Mpy/Mix/Pack	I2	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}			r_3								r_2																						r_1		QP		
MM Mux1	I3	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}											mbt_{4c}																						r_2		r_1		QP
MM Mux2	I4	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}											mht_{8c}																						r_2		r_1		QP
Shift R Variable	I5	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}			r_3								r_2																						r_1		QP		
MM Shift R Fixed	I6	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}			r_3									$count_{5b}$																				r_1		QP			
Shift L Variable	I7	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}			r_3								r_2																						r_1		QP		
MM Shift L Fixed	I8	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}											$ccount_{5c}$																						r_2		r_1		QP
Popcount	I9	7		z_a	x_{2a}	z_b	v_e	x_{2c}		x_{2b}			r_3								0																						r_1		QP		
Shift Right Pair	I10	5			x_2	x		$count_{6d}$					r_3								r_2																					r_1		QP			
Extract	I11	5			x_2	x		len_{6d}					r_3								pos_{6b}		y																			r_1		QP			
Dep.Z	I12	5			x_2	x		len_{6d}		y		$cpos_{6c}$		r_2																												r_1		QP			
Dep.Z Imm ₈	I13	5	s		x_2	x		len_{6d}		y		$cpos_{6c}$		imm_{7b}																													r_1		QP		
Deposit Imm ₁	I14	5	s		x_2	x		len_{6d}					r_3								$cpos_{6b}$																						r_1		QP		
Deposit	I15	4			$cpos_{6d}$			len_{4d}					r_3								r_2																						r_1		QP		
Test Bit	I16	5		t_b	x_2	t_a		p_2					r_3								pos_{6b}		y	c																		p_1		QP			
Test NaT	I17	5		t_b	x_2	t_a		p_2					r_3										y	c																		p_1		QP			
Break/Nop	I19	0	i		x_3			x_6													imm_{20a}																								QP		
Int Spec Check	I20	0	s		x_3			imm_{13c}													r_2																						imm_{7a}		QP		
Move to BR	I21	0			x_3																r_2																						b_1		QP		
Move from BR	I22	0			x_3			x_6														b_2																					r_1		QP		
Move to Pred	I23	0	s		x_3			$mask_{8c}$													r_2																						$mask_{7a}$		QP		
Move to Pred Imm ₄₄	I24	0	s		x_3			imm_{27a}																																					QP		
Move from Pred/IP	I25	0			x_3			x_6																																			r_1		QP		
Move to AR	I26	0			x_3			x_6					ar_3								r_2																							QP			
Move to AR Imm ₈	I27	0	s		x_3			x_6					ar_3								imm_{7b}																							QP			
Move from AR	I28	0			x_3			x_6					ar_3																														r_1		QP		
Sxt/Zxt/Czx	I29	0			x_3			x_6					r_3																														r_1		QP		

Table C-4. Instruction Format Summary (Continued)

	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
IP-Relative Branch B1	4	s	d	wh	imm _{20b}																p		btype	qp																	
Counted Branch B2	4	s	d	wh	imm _{20b}																p		btype	0																	
IP-Relative Call B3	5	s	d	wh	imm _{20b}																p		b ₁	qp																	
Indirect Branch B4	0		d	wh	x ₆												b ₂	p		btype	qp																				
Indirect Call B5	1		d	wh													b ₂	p		b ₁	qp																				
Misc B8	0							x ₆												0																					
Break/Nop B9	0/2	i							x ₆						imm _{20a}						qp																				
FP Arithmetic F1	8 - D	x	sf	f ₄						f ₃						f ₂						f ₁						qp													
Fixed Multiply Add F2	E	x	x ₂	f ₄						f ₃						f ₂						f ₁						qp													
FP Select F3	E	x	f ₄						f ₃						f ₂						f ₁						qp														
FP Compare F4	4	r _b	sf	r _a	p ₂						f ₃						f ₂						t _a	p ₁						qp											
FP Class F5	5	fc ₂						p ₂						fclass _{7c}						f ₂						t _a	p ₁						qp								
FP Recip Approx F6	0 - 1	q	sf	x	p ₂						f ₃						f ₂						f ₁						qp												
FP Recip Sqrt App F7	0 - 1	q	sf	x	p ₂						f ₃												f ₁						qp												
FP Min/Max/Pcmp F8	0 - 1		sf	x	x ₆						f ₃						f ₂						f ₁						qp												
FP Merge/Logical F9	0 - 1	x						x ₆						f ₃						f ₂						f ₁						qp									
Convert FP to Fixed F10	0 - 1		sf	x	x ₆												f ₂						f ₁						qp												
Convert Fixed to FP F11	0	x						x ₆												f ₂						f ₁						qp									
FP Set Controls F12	0		sf	x	x ₆						omask _{7c}						amask _{7b}												qp												
FP Clear Flags F13	0		sf	x	x ₆																		qp																		
FP Check Flags F14	0	s	sf	x	x ₆						imm _{20a}												qp																		
Break/Nop F15	0	i	x						x ₆						imm _{20a}												qp														
Break/Nop X1	0	i	x ₃	x ₆						imm _{20a}												qp	imm ₄₁																		
Move Imm ₆₄ X2	6	i	imm _{9d}						imm _{5c}						i _c	v _c	imm _{7b}						r ₁						qp	imm ₄₁											

Table C-5. Instruction Field Color Key

Field & Color	
ALU Instruction	Opcode Extension
Integer Instruction	Opcode Hint Extension
Memory Instruction	Immediate
Branch Instruction	Indirect Source
Floating-point Instruction	Predicate Destination
Integer Source	Integer Destination
Memory Source	Memory Source & Destination
Shift Source	Shift Immediate
Special Register Source	Special Register Destination
Floating-point Source	Floating-point Destination
Branch Source	Branch Destination
Address Source	Reserved Instruction
Qualifying Predicate	Reserved Inst if PR[qp] is 1
Ignored Field/Instruction	

Table C-6. Instruction Field Names

Field Name	Description
ar ₃	application register source/target
b ₁ , b ₂	branch register source/target
btype	branch type opcode extension
c	complement compare relation opcode extension
ccount _{5c}	multimedia shift left complemented shift count immediate
count _{5b} , count _{6d}	multimedia shift right/shift right pair shift count immediate
cpos _{sx}	deposit complemented bit position immediate
ct _{2d}	multimedia multiply shift/shift and add shift count immediate
d	branch cache deallocation hint opcode extension
f _n	floating-point register source/target
fc ₂ , fclass _{7c}	floating-point class immediate
hint	memory reference hint opcode extension
i, i _{2b} , i _{2d} , imm _x	immediate of length 1, 2, or <i>x</i>
len _{4d} , len _{6d}	extract/deposit length immediate
m	memory reference post-modify opcode extension
mask _x	predicate immediate mask
mbt _{4c} , mht _{8c}	multimedia mux1/mux2 immediate
p	sequential prefetch hint opcode extension
p ₁ , p ₂	predicate register target
pos _{6b}	test bit/extract bit position immediate
q	floating-point reciprocal/reciprocal square-root opcode extension
qp	qualifying predicate register source
r _n	general register source/target
s	immediate sign bit
sf	floating-point status field opcode extension
sof, sol, sor	alloc size of frame, size of locals, size of rotating immediates
t _a , t _b	compare type opcode extension
v _x	reserved opcode extension field
wh	branch whether hint opcode extension
x, x _n	opcode extension of length 1 or <i>n</i>
y	extract/deposit/test bit/test NaT opcode extension
z _a , z _b	multimedia operand size opcode extension

Table C-7. Special Instruction Notations

Notation	Description
f	instruction must be the first in an instruction group
l	instruction must be the last in an instruction group
t	instruction is only allowed in instruction slot 2

The remaining sections of this chapter present the detailed encodings of all instructions. The “A-Unit Instruction encodings” are presented first, followed by the “I-Unit Instruction Encodings” on page C-16, “M-Unit Instruction Encodings” on page C-26, “B-Unit Instruction Encodings” on page C-45, “F-Unit Instruction Encodings” on page C-49, and “X-Unit Instruction Encodings” on page C-56.

Within each section, the instructions are grouped by function, and appear with their instruction format in the same order as in Table C-4, “Instruction Format Summary,” on page C-4. The opcode extension fields are briefly described and tables present the opcode extension assignments. Unused instruction encodings (appearing as blank entries in the opcode extensions tables) behave in one of three ways:

- Ignored instructions (white entries in the tables) execute as `nop` instructions.
- Reserved instructions (light gray in the grayscale version of the tables, brown in the color version) cause an Illegal Operation fault.
- Reserved if `PR[qp]` is 1 instructions (dark gray in the grayscale version of the tables, purple in the color version) cause an Illegal Operation fault if the predicate register specified by the `qp` field of the instruction (bits 5:0) is 1 and execute as a `nop` instruction if 0.

Constant 0 fields in instructions must be 0 or undefined operation results. The undefined operation may include checking that the constant field is 0 and causing an Illegal Operation fault if it is not. If an instruction having a constant 0 field also has a qualifying predicate (`qp` field), the fault or other undefined operation must not occur if `PR[qp]` is 0. For constant 0 fields in instruction bits 5:0 (normally used for `qp`), the fault or other undefined operation may or may not depend on the `PR` addressed by those bits.

Ignored (white space) fields in instructions should be coded as 0. Although ignored in this revision of the architecture, future architecture revisions may define these fields as hint extensions. These hint extensions will be defined such that the 0 value in each field corresponds to the default hint. It is expected that assemblers will automatically set these fields to zero by default.

C.2 A-Unit Instruction Encodings

C.2.1 Integer ALU

All integer ALU instructions are encoded within major opcode 8 using a 2-bit opcode extension field in bits 35:34 (x_{2a}) and most have a second 2-bit opcode extension field in bits 28:27 (x_{2b}), a 4-bit opcode extension field in bits 32:29 (x_4), and a 1-bit reserved opcode extension field in bit 33 (v_e). Table C-8 shows the 2-bit x_{2a} and 1-bit v_e assignments, Table C-9 shows the integer ALU 4-bit+2-bit assignments, and Table C-12 on page C-13 shows the multimedia ALU 1-bit+2-bit assignments (which also share major opcode 8).


Table C-8. Integer ALU 2-bit+1-bit Opcode Extensions

Opcode Bits 40:37	x_{2a} Bits 35:34	v_e Bit 33	
		0	1
8	0	Integer ALU 4-bit+2-bit Ext (Table C-9)	
	1	Multimedia ALU 1-bit+2-bit Ext (Table C-12)	
	2	<code>adds – imm₁₄ A4</code>	
	3	<code>addp4 – imm₁₄ A4</code>	

Table C-9. Integer ALU 4-bit+2-bit Opcode Extensions


Opcode Bits 40:37	x_{2a} Bits 35:34	v_e Bit 33	x_4 Bits 32:29	x_{2b} Bits 28:27			
				0	1	2	3
8	0	0	0	<code>add A1</code>	<code>add +1 A1</code>		
			1	<code>sub –1 A1</code>	<code>sub A1</code>		
			2	<code>addp4 A1</code>			
			3	<code>and A1</code>	<code>andcm A1</code>	<code>or A1</code>	<code>xor A1</code>
			4	<code>shladd A2</code>			
			5				
			6	<code>shladdp4 A2</code>			
			7				
			8				
			9		<code>sub – imm₈ A3</code>		
			A				
			B	<code>and – imm₈ A3</code>	<code>andcm – imm₈ A3</code>	<code>or – imm₈ A3</code>	<code>xor – imm₈ A3</code>
			C				
			D				
			E				
			F				

C.2.1.1 Integer ALU – Register-Register




Instruction	Operands	Opcode	Extension				
			x _{2a}	v _e	x ₄	x _{2b}	
add	$r_1 = r_2, r_3$ $r_1 = r_2, r_3, 1$	8	0	0	0	0	
sub	$r_1 = r_2, r_3$ $r_1 = r_2, r_3, 1$				1	1	
addp4					2	0	
and	$r_1 = r_2, r_3$				3	1	0
andcm							
or xor		3					

C.2.1.2 Shift Left and Add




Instruction	Operands	Opcode	Extension		
			x _{2a}	v _e	x ₄
shladd shladdp4	$r_1 = r_2, count_2, r_3$	8	0	0	4 6

C.2.1.3 Integer ALU – Immediate₈-Register




Instruction	Operands	Opcode	Extension						
			x _{2a}	v _e	x ₄	x _{2b}			
sub	$r_1 = imm_8, r_3$	8	0	0	9	1			
and					B	0	1		
andcm								2	
or									3
xor									

C.2.1.4 Add Immediate₁₄



Instruction	Operands	Opcode	Extension	
			x _{2a}	v _e
adds addp4	$r_1 = imm_{14}, r_3$	8	2 3	0

C.2.1.5 Add Immediate₂₂



Instruction	Operands	Opcode
addl	$r_1 = imm_{22}, r_3$	9

C.2.2 Integer Compare

The integer compare instructions are encoded within major opcodes C – E using a 2-bit opcode extension field (x_2) in bits 35:34 and three 1-bit opcode extension fields in bits 33 (t_a), 36 (t_b), and 12 (c), as shown in Table C-10. The integer compare immediate instructions are encoded within major opcodes C – E using a 2-bit opcode extension field (x_2) in bits 35:34 and two 1-bit opcode extension fields in bits 33 (t_a) and 12 (c), as shown in Table C-11.


Table C-10. Integer Compare Opcode Extensions

x_2 Bits 35:34	t_b Bit 36	t_a Bit 33	c Bit 12	Opcode Bits 40:37		
				C	D	E
0	0	0	0	cmp.lt A6	cmp.ltu A6	cmp.eq A6
			1	cmp.lt.unc A6	cmp.ltu.unc A6	cmp.eq.unc A6
		1	0	cmp.eq.and A6	cmp.eq.or A6	cmp.eq.or.andcm A6
			1	cmp.ne.and A6	cmp.ne.or A6	cmp.ne.or.andcm A6
	1	0	0	cmp.gt.and A7	cmp.gt.or A7	cmp.gt.or.andcm A7
			1	cmp.le.and A7	cmp.le.or A7	cmp.le.or.andcm A7
		1	0	cmp.ge.and A7	cmp.ge.or A7	cmp.ge.or.andcm A7
			1	cmp.lt.and A7	cmp.lt.or A7	cmp.lt.or.andcm A7
1	0	0	0	cmp4.lt A6	cmp4.ltu A6	cmp4.eq A6
			1	cmp4.lt.unc A6	cmp4.ltu.unc A6	cmp4.eq.unc A6
		1	0	cmp4.eq.and A6	cmp4.eq.or A6	cmp4.eq.or.andcm A6
			1	cmp4.ne.and A6	cmp4.ne.or A6	cmp4.ne.or.andcm A6
		1	0	cmp4.gt.and A7	cmp4.gt.or A7	cmp4.gt.or.andcm A7
			1	cmp4.le.and A7	cmp4.le.or A7	cmp4.le.or.andcm A7
	1	0	0	cmp4.ge.and A7	cmp4.ge.or A7	cmp4.ge.or.andcm A7
			1	cmp4.lt.and A7	cmp4.lt.or A7	cmp4.lt.or.andcm A7

Table C-11. Integer Compare Immediate Opcode Extensions


x_2 Bits 35:34	t_a Bit 33	c Bit 12	Opcode Bits 40:37		
			C	D	E
2	0	0	cmp.lt – imm ₈ A8	cmp.ltu – imm ₈ A8	cmp.eq – imm ₈ A8
		1	cmp.lt.unc – imm ₈ A8	cmp.ltu.unc – imm ₈ A8	cmp.eq.unc – imm ₈ A8
	1	0	cmp.eq.and – imm ₈ A8	cmp.eq.or – imm ₈ A8	cmp.eq.or.andcm – imm ₈ A8
		1	cmp.ne.and – imm ₈ A8	cmp.ne.or – imm ₈ A8	cmp.ne.or.andcm – imm ₈ A8
3	0	0	cmp4.lt – imm ₈ A8	cmp4.ltu – imm ₈ A8	cmp4.eq – imm ₈ A8
		1	cmp4.lt.unc – imm ₈ A8	cmp4.ltu.unc – imm ₈ A8	cmp4.eq.unc – imm ₈ A8
	1	0	cmp4.eq.and – imm ₈ A8	cmp4.eq.or – imm ₈ A8	cmp4.eq.or.andcm – imm ₈ A8
		1	cmp4.ne.and – imm ₈ A8	cmp4.ne.or – imm ₈ A8	cmp4.ne.or.andcm – imm ₈ A8

C.2.2.1 Integer Compare – Register-Register




Instruction	Operands	Opcode	Extension				
			x ₂	t _b	t _a	c	
cmp.lt	<i>p₁, p₂ = r₂, r₃</i>	C	0	0	0	0	
cmp.ltu		D					
cmp.eq		E					
cmp.lt.unc		C					
cmp.ltu.unc		D					
cmp.eq.unc		E					
cmp.eq.and		C			1	0	
cmp.eq.or		D					
cmp.eq.or.andcm		E					
cmp.ne.and		C					
cmp.ne.or		D					
cmp.ne.or.andcm		E					
cmp4.lt		1	C	1	0	0	0
cmp4.ltu			D				
cmp4.eq			E				
cmp4.lt.unc			C				
cmp4.ltu.unc	D						
cmp4.eq.unc	E						
cmp4.eq.and	C		1			0	
cmp4.eq.or	D						
cmp4.eq.or.andcm	E						
cmp4.ne.and	C						
cmp4.ne.or	D						
cmp4.ne.or.andcm	E						

C.2.2.2 Integer Compare to Zero – Register



Instruction	Operands	Opcode	Extension					
			x ₂	t _b	t _a	c		
cmp.gt.and cmp.gt.or cmp.gt.or.andcm	<i>p₁, p₂ = r0, r₃</i>	C D E	0	1	0	0		
cmp.le.and cmp.le.or cmp.le.or.andcm		C D E				1		
cmp.ge.and cmp.ge.or cmp.ge.or.andcm		C D E				1	0	
cmp.lt.and cmp.lt.or cmp.lt.or.andcm		C D E			1	1		
cmp4.gt.and cmp4.gt.or cmp4.gt.or.andcm		C D E			1	1	0	0
cmp4.le.and cmp4.le.or cmp4.le.or.andcm		C D E						1
cmp4.ge.and cmp4.ge.or cmp4.ge.or.andcm		C D E	1	0				
cmp4.lt.and cmp4.lt.or cmp4.lt.or.andcm		C D E	1	1				

C.2.2.3 Integer Compare – Immediate-Register



Instruction	Operands	Opcode	Extension		
			x ₂	t _a	c
cmp.lt	<i>p₁, p₂ = imm₈, r₃</i>	C	2	0	0
cmp.ltu		D			
cmp.eq		E			
cmp.lt.unc		C			1
cmp.ltu.unc		D			
cmp.eq.unc		E			
cmp.eq.and		C	3	1	0
cmp.eq.or		D			
cmp.eq.or.andcm		E			
cmp.ne.and		C			1
cmp.ne.or		D			
cmp.ne.or.andcm		E			
cmp4.lt		C	3	0	0
cmp4.ltu		D			
cmp4.eq		E			
cmp4.lt.unc		C			
cmp4.ltu.unc	D				
cmp4.eq.unc	E				
cmp4.eq.and	C	1			
cmp4.eq.or	D				
cmp4.eq.or.andcm	E				
cmp4.ne.and	C				1
cmp4.ne.or	D				
cmp4.ne.or.andcm	E				

C.2.3 Multimedia

All multimedia ALU instructions are encoded within major opcode 8 using two 1-bit opcode extension fields in bits 36 (*z_a*) and 33 (*z_b*) and a 2-bit opcode extension field in bits 35:34 (*x_{2a}*) as shown in Table C-12. The multimedia ALU instructions also have a 4-bit opcode extension field in bits 32:29 (*x₄*), and a 2-bit opcode extension field in bits 28:27 (*x_{2b}*) as shown in Table C-13 on page C-14.

Table C-12. Multimedia ALU 2-bit+1-bit Opcode Extensions

Opcode Bits 40:37	x _{2a} Bits 35:34	z _a Bit 36	z _b Bit 33	
8	1	0	0	Multimedia ALU Size 1 (Table C-13)
			1	Multimedia ALU Size 2 (Table C-14)
		1	0	Multimedia ALU Size 4 (Table C-15)
			1	

Table C-13. Multimedia ALU Size 1 4-bit+2-bit Opcode Extensions

Opcode Bits 40:37	x _{2a} Bits 35:34	z _a Bit 36	z _b Bit 33	x ₄ Bits 32:29	x _{2b} Bits 28:27			
					0	1	2	3
8	1	0	0	0	padd1 A9	padd1.sss A9	padd1.uuu A9	padd1.uus A9
				1	psub1 A9	psub1.sss A9	psub1.uuu A9	psub1.uus A9
				2			pavg1 A9	pavg1.raz A9
				3			pavgsub1 A9	
				4				
				5				
				6				
				7				
				8				
				9	pcmp1.eq A9	pcmp1.gt A9		
				A				
				B				
				C				
				D				
				E				
				F				


Table C-14. Multimedia ALU Size 2 4-bit+2-bit Opcode Extensions

Opcode Bits 40:37	x _{2a} Bits 35:34	z _a Bit 36	z _b Bit 33	x ₄ Bits 32:29	x _{2b} Bits 28:27			
					0	1	2	3
8	1	0	1	0	padd2 A9	padd2.sss A9	padd2.uuu A9	padd2.uus A9
				1	psub2 A9	psub2.sss A9	psub2.uuu A9	psub2.uus A9
				2			pavg2 A9	pavg2.raz A9
				3			pavgsub2 A9	
				4	pshladd2 A10			
				5				
				6	pshradd2 A10			
				7				
				8				
				9	pcmp2.eq A9	pcmp2.gt A9		
				A				
				B				
				C				
				D				
				E				
				F				

Table C-15. Multimedia ALU Size 4 4-bit+2-bit Opcode Extensions


Opcode Bits 40:37	x _{2a} Bits 35:34	z _a Bit 36	z _b Bit 33	x ₄ Bits 32:29	x _{2b} Bits 28:27			
					0	1	2	3
8	1	1	0	0	padd4 A9			
				1	psub4 A9			
				2				
				3				
				4				
				5				
				6				
				7				
				8				
				9	pcmp4.eq A9	pcmp4.gt A9		
				A				
				B				
				C				
				D				
				E				
				F				

C.2.3.1 Multimedia ALU



Instruction	Operands	Opcode	Extension						
			x _{2a}	z _a	z _b	x ₄	x _{2b}		
padd1	<i>r₁ = r₂, r₃</i>	8	1	0	0	0	0		
padd2				1	1				
padd4				1	0				
padd1.sss				0	0			1	
padd2.sss				0	1				
padd1.uuu				0	0			2	
padd2.uuu				0	1				
padd1.uus				0	0			3	
padd2.uus				0	1				
psub1				0	0			0	
psub2				1	1				
psub4				1	0				
psub1.sss				0	0			1	1
psub2.sss				0	1				
psub1.uuu				0	0			2	
psub2.uuu				0	1				
psub1.uus				0	0			3	
psub2.uus				0	1				
pavg1				0	0			2	2
pavg2				0	1				
pavg1.raz	0	0	3	3					
pavg2.raz	0	1							
pavgsub1	0	0	3	2					
pavgsub2	0	1							
pcmp1.eq	0	0	9	0					
pcmp2.eq	1	1							
pcmp4.eq	1	0							
pcmp1.gt	0	0	9	1					
pcmp2.gt	0	1							
pcmp4.gt	1	0							

C.2.3.2 Multimedia Shift and Add



Instruction	Operands	Opcode	Extension			
			x _{2a}	z _a	z _b	x ₄
pshladd2	$r_1 = r_2, count_2, r_3$	8	1	0	1	4
pshradd2			1	0	1	6

C.3 I-Unit Instruction Encodings

C.3.1 Multimedia and Variable Shifts

All multimedia multiply/shift/max/min/mix/mux/pack/unpack and variable shift instructions are encoded within major opcode 7 using two 1-bit opcode extension fields in bits 36 (z_a) and 33 (z_b) and a 1-bit reserved opcode extension in bit 32 (v_e) as shown in Table C-16. They also have a 2-bit opcode extension field in bits 35:34 (x_{2a}) and a 2-bit field in bits 29:28 (x_{2b}) and most have a 2-bit field in bits 31:30 (x_{2c}) as shown in Table C-17.

Table C-16. Multimedia and Variable Shift 1-bit Opcode Extensions

Opcode Bits 40:37	z _a Bit 36	z _b Bit 33	v _e Bit 32	
			0	1
			7	0
1	Multimedia Size 2 (Table C-18)			
1	0	Multimedia Size 4 (Table C-19)		
	1	Variable Shift (Table C-20)		

Table C-17. Multimedia Max/Min/Mix/Pack/Unpack Size 1 2-bit Opcode Extensions

Opcode Bits 40:37	z _a Bit 36	z _b Bit 33	v _e Bit 32	x _{2a} Bits 35:34	x _{2b} Bits 29:28	x _{2c} Bits 31:30					
						0	1	2	3		
						7	0	0	0	0	0
1											
2											
3											
1	0										
	1										
	2										
	3										
2	0			unpack1.h I2	mix1.r I2						
	1	pmin1.u I2		pmax1.u I2							
	2			unpack1.l I2	mix1.l I2						
	3				psad1 I2						
3	0										
	1										
	2				mux1 I3						
	3										

Table C-18. Multimedia Multiply/Shift/Max/Min/Mix/Pack/Unpack Size 2 2-bit Opcode Extensions

Opcode Bits 40:37	z _a Bit 36	z _b Bit 33	v _e Bit 32	x _{2a} Bits 35:34	x _{2b} Bits 29:28	x _{2c} Bits 31:30			
						0	1	2	3
7	0	1	0	0	0	pshr2.u – var I5	pshl2 – var I7		
					1	pmpyshr2.u I1			
					2	pshr2 – var I5			
					3	pmpyshr2 I1			
				1	0				
					1	pshr2.u – fixed I6		popcnt I9	
					2				
					3	pshr2 – fixed I6			
				2	0	pack2.uss I2	unpack2.h I2	mix2.r I2	
					1			pmpy2.r I2	
					2	pack2.sss I2	unpack2.l I2	mix2.l I2	
					3	pmin2 I2	pmax2 I2	pmpy2.l I2	
				3	0				
					1		pshl2 – fixed I8		
					2			mux2 I4	
					3				


Table C-19. Multimedia Shift/Mix/Pack/Unpack Size 4 2-bit Opcode Extensions

Opcode Bits 40:37	z _a Bit 36	z _b Bit 33	v _e Bit 32	x _{2a} Bits 35:34	x _{2b} Bits 29:28	x _{2c} Bits 31:30			
						0	1	2	3
7	1	0	0	0	0	pshr4.u – var I5	pshl4 – var I7		
					1				
					2	pshr4 – var I5			
					3				
				1	0				
					1	pshr4.u – fixed I6			
					2				
					3	pshr4 – fixed I6			
				2	0		unpack4.h I2	mix4.r I2	
					1				
					2	pack4.sss I2	unpack4.l I2	mix4.l I2	
					3				
				3	0				
					1		pshl4 – fixed I8		
					2				
					3				

Table C-20. Variable Shift 2-bit Opcode Extensions


Opcode Bits 40:37	z _a Bit 36	z _b Bit 33	v _e Bit 32	x _{2a} Bits 35:34	x _{2b} Bits 29:28	x _{2c} Bits 31:30			
						0	1	2	3
7	1	1	0	0	0	shr.u – var I5	shl – var I7		
						1			
						2	shr – var I5		
						3			
						1	0		
							1		
							2		
						2	3		
							0		
							1		
						3	2		
							3		
0									

C.3.1.1 Multimedia Multiply and Shift




Instruction	Operands	Opcode	Extension				
			z _a	z _b	v _e	x _{2a}	x _{2b}
pmpyshr2 pmpyshr2.u	r ₁ = r ₂ , r ₃ , count ₂	7	0	1	0	0	3 1

C.3.1.2 Multimedia Multiply/Mix/Pack/Unpack




Instruction	Operands	Opcode	Extension					
			z _a	z _b	v _e	x _{2a}	x _{2b}	x _{2c}
pmpy2.r pmpy2.l	r ₁ = r ₂ , r ₃	7	0	1	0	2	1 3	3
mix1.r mix2.r mix4.r			0 0 1	0 1 0			0	2
mix1.l mix2.l mix4.l			0 0 1	0 1 0			2	
pack2.uss pack2.sss pack4.sss			0 0 1	1 1 0			0	0
unpack1.h unpack2.h unpack4.h			0 0 1	0 1 0			0	1
unpack1.l unpack2.l unpack4.l			0 0 1	0 1 0			2	
pmin1.u pmax1.u			0 0	0 0			1	0 1
pmin2 pmax2			0 0	1 1			3 3	0 1
psad1			0	0			3	2

C.3.1.3 Multimedia Mux1




Instruction	Operands	Opcode	Extension					
			z _a	z _b	v _e	x _{2a}	x _{2b}	x _{2c}
mux1	$r_1 = r_2, mbtype_4$	7	0	0	0	3	2	2

C.3.1.4 Multimedia Mux2




Instruction	Operands	Opcode	Extension					
			z _a	z _b	v _e	x _{2a}	x _{2b}	x _{2c}
mux2	$r_1 = r_2, mhtype_8$	7	0	1	0	3	2	2

C.3.1.5 Shift Right – Variable




Instruction	Operands	Opcode	Extension					
			z _a	z _b	v _e	x _{2a}	x _{2b}	x _{2c}
pshr2	$r_1 = r_3, r_2$	7	0	1	0	0	2	0
pshr4			1	0				
shr			1	1				
pshr2.u			0	1				
pshr4.u			1	0				
shr.u	1	1						

C.3.1.6 Multimedia Shift Right – Fixed




Instruction	Operands	Opcode	Extension					
			z _a	z _b	v _e	x _{2a}	x _{2b}	x _{2c}
pshr2	$r_1 = r_3, count_5$	7	0	1	0	1	3	0
pshr4			1	0				
pshr2.u			0	1				
pshr4.u			1	0				

C.3.1.7 Shift Left – Variable




Instruction	Operands	Opcode	Extension					
			z _a	z _b	v _e	x _{2a}	x _{2b}	x _{2c}
pshl2	$r_1 = r_2, r_3$	7	0	1	0	0	0	1
pshl4			1	0				
shl			1	1				

C.3.1.8 Multimedia Shift Left – Fixed



C.3.1.9 Population Count



C.3.2 Integer Shifts

The integer shift, test bit, and test NaT instructions are encoded within major opcode 5 using a 2-bit opcode extension field in bits 35:34 (x_2) and a 1-bit opcode extension field in bit 33 (x). The extract and test bit instructions also have a 1-bit opcode extension field in bit 13 (y). Table C-21 shows the test bit, extract, and shift right pair assignments.

Table C-21. Integer Shift/Test Bit/Test NaT 2-bit Opcode Extensions


Opcode bits 40:37	x_2 bits 35:34	x bit 33	y bit 13	
			0	1
5	0	0	Test Bit (Table C-23)	Test NaT (Table C-23)
	1		extr.u I11	extr I11
	2		shrp I10	
	3			

Most deposit instructions also have a 1-bit opcode extension field in bit 26 (y). Table C-22 shows these assignments.


Table C-22. Deposit Opcode Extensions

Opcode bits 40:37	x_2 bits 35:34	x bit 33	y bit 26	
			0	1
5	0	1	Test Bit/Test NaT (Table C-23)	
	1		dep.z I12	dep.z – imm ₈ I13
	2		dep – imm ₁ I14	
	3			

C.3.2.1 Shift Right Pair




C.3.2.2 Extract




Instruction	Operands	Opcode	Extension		
			x ₂	x	y
extr.u	$r_1 = r_3, pos_6, len_6$	5	1	0	0
extr					1

C.3.2.3 Zero and Deposit




Instruction	Operands	Opcode	Extension		
			x ₂	x	y
dep.z	$r_1 = r_2, pos_6, len_6$	5	1	1	0

C.3.2.4 Zero and Deposit Immediate₈




Instruction	Operands	Opcode	Extension		
			x ₂	x	y
dep.z	$r_1 = imm_8, pos_6, len_6$	5	1	1	1

C.3.2.5 Deposit Immediate₁



Instruction	Operands	Opcode	Extension	
			x ₂	x
dep	$r_1 = imm_1, r_3, pos_6, len_6$	5	3	1

C.3.2.6 Deposit



Instruction	Operands	Opcode
dep	$r_1 = r_2, r_3, pos_6, len_4$	4


C.3.3 Test Bit

All test bit instructions are encoded within major opcode 5 using a 2-bit opcode extension field in bits 35:34 (x_2) plus four 1-bit opcode extension fields in bits 33 (t_a), 36 (t_b), 12 (c), and 19 (y). Table C-23 summarizes these assignments.

Table C-23. Test Bit Opcode Extensions


Opcode bits 40:37	x_2 bits 35:34	t_a bit 33	t_b bit 36	c bit 12	y bit 13	
					0	1
5	0	0	0	0	tbit.z I16	tnat.z I17
				1	tbit.z.unc I16	tnat.z.unc I17
			1	0	tbit.z.and I16	tnat.z.and I17
				1	tbit.nz.and I16	tnat.nz.and I17
		1	0	0	tbit.z.or I16	tnat.z.or I17
				1	tbit.nz.or I16	tnat.nz.or I17
			1	0	tbit.z.or.andcm I16	tnat.z.or.andcm I17
				1	tbit.nz.or.andcm I16	tnat.nz.or.andcm I17

C.3.3.1 Test Bit



Instruction	Operands	Opcode	Extension							
			x_2	t_a	t_b	y	c			
tbit.z	$p_1, p_2 = r_3, pos_6$	5	0	0	0	0	0			
tbit.z.unc					1		1			
tbit.z.and					0		0	0		
tbit.nz.and							1	1		
tbit.z.or				1	0	0	0	0		
tbit.nz.or							1	1		
tbit.z.or.andcm							0	1	0	0
tbit.nz.or.andcm									1	1

C.3.3.2 Test NaT



Instruction	Operands	Opcode	Extension							
			x_2	t_a	t_b	y	c			
tnat.z	$p_1, p_2 = r_3$	5	0	0	0	1	0			
tnat.z.unc					1		1			
tnat.z.and					0		0	0		
tnat.nz.and							1	1		
tnat.z.or				1	0	0	0	0		
tnat.nz.or							1	1		
tnat.z.or.andcm							0	1	0	0
tnat.nz.or.andcm									1	1

C.3.4 Miscellaneous I-Unit Instructions

The miscellaneous I-unit instructions are encoded in major opcode 0 using a 3-bit opcode extension field (x_3) in bits 35:33. Some also have a 6-bit opcode extension field (x_6) in bits 32:27. Table C-24 shows the 3-bit assignments and Table C-25 summarizes the 6-bit assignments.


Table C-24. Misc I-Unit 3-bit Opcode Extensions

Opcode Bits 40:37	x_3 Bits 35:33	
0	0	6-bit Ext (Table C-25)
	1	chk.s.i – int I20
	2	mov to pr.rot – imm ₄₄ I24
	3	mov to pr I23
	4	
	5	
	6	
	7	mov to b I22

Table C-25. Misc I-Unit 6-bit Opcode Extensions


Opcode Bits 40:37	x_3 Bits 35:33	x_6 Bits 32:31				
		Bits 30:27				
		0	1	2	3	
0	0	0	break.i I19	zxt1 I29		mov from ip I25
		1	nop.i I19	zxt2 I29		mov from b I22
		2		zxt4 I29		mov.i from ar I28
		3				mov from pr I25
		4		sxt1 I29		
		5		sxt2 I29		
		6		sxt4 I29		
		7				
		8		czx1.l I29		
		9		czx2.l I29		
		A	mov.i to ar – imm ₈ I27			mov.i to ar I26
		B				
		C		czx1.r I29		
		D		czx2.r I29		
		E				
		F				

C.3.4.1 Break/Nop (I-Unit)



Instruction	Operands	Opcode	Extension	
			x_3	x_6
break.i		0	0	00
nop.i	imm ₂₁	0	0	01

C.3.4.2 Integer Speculation Check (I-Unit)




Instruction	Operands	Opcode	Extension
			x_3
chk.s.i	$r_2, target_{25}$	0	1

C.3.5 GR/BR Moves


The GR/BR move instructions are encoded in major opcode 0. See “Miscellaneous I-Unit Instructions” on page C-23 for a summary of the opcode extensions. The mov to BR instruction uses a 1-bit opcode extension field (x) in bit 22 to distinguish the return form from the normal form.

C.3.5.1 Move to BR



Instruction	Operands	Opcode	Extension	
			x ₃	x
mov	$b_1 = r_2$	0	7	0
mov.ret				1

C.3.5.2 Move from BR




Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
mov	$r_1 = b_2$	0	0	31

C.3.6 GR/Predicate/IP Moves


The GR/Predicate/IP move instructions are encoded in major opcode 0. See “Miscellaneous I-Unit Instructions” on page C-23 for a summary of the opcode extensions.

C.3.6.1 Move to Predicates – Register




Instruction	Operands	Opcode	Extension
			x ₃
mov	$pr = r_2, mask_{17}$	0	3

C.3.6.2 Move to Predicates – Immediate₄₄



Instruction	Operands	Opcode	Extension
			x ₃
mov	$pr.rot = imm_{44}$	0	2

C.3.6.3 Move from Predicates/IP




Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
mov	$r_1 = ip$ $r_1 = pr$	0	0	30 33

C.3.7 GR/AR Moves (I-Unit)


The I-Unit GR/AR move instructions are encoded in major opcode 0. (some ARs are accessed using system/memory management instructions on the M-unit. See “GR/AR Moves (M-Unit)” on page C-42.) See “Miscellaneous I-Unit Instructions” on page C-23 for a summary of the I-Unit GR/AR opcode extensions.

C.3.7.1 Move to AR – Register (I-Unit)




Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
mov.i	$ar_3 = r_2$	0	0	2A

C.3.7.2 Move to AR – Immediate₈ (I-Unit)




Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
mov.i	$ar_3 = imm_8$	0	0	0A

C.3.7.3 Move from AR (I-Unit)



Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
mov.i	$r_1 = ar_3$	0	0	32

C.3.8 Sign/Zero Extend/Compute Zero Index



Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
zxt1 zxt2 zxt4	$r_1 = r_3$	0	0	10
sxt1				11
sxt2				12
sxt4				14
czx1.l				15
czx2.l				16
czx1.r				18
czx2.r				19
				1C
	1D			

C.4 M-Unit Instruction Encodings

C.4.1 Loads and Stores

All load and store instructions are encoded within major opcodes 4, 5, 6, and 7 using a 6-bit opcode extension field in bits 35:30 (x_6). Instructions in major opcode 4 (integer load/store, semaphores, and get FR) use two 1-bit opcode extension fields in bit 36 (m) and bit 27 (x) as shown in Table C-26. Instructions in major opcode 6 (floating-point load/store, load pair, and set FR) use two 1-bit opcode extension fields in bit 36 (m) and bit 27 (x) as shown in Table C-27.

Table C-26. Integer Load/Store/Semaphore/Get FR 1-bit Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	
4	0	0	Load/Store (Table C-28)
	0	1	Semaphore/get FR (Table C-31)
	1	0	Load +Reg (Table C-29)
	1	1	

Table C-27. Floating-point Load/Store/Load Pair/Set FR 1-bit Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	
6	0	0	FP Load/Store (Table C-32)
	0	1	FP Load Pair/set FR (Table C-35)
	1	0	FP Load +Reg (Table C-33)
	1	1	FP Load Pair +Imm (Table C-36)

The integer load/store opcode extensions are summarized in Table C-28 on page C-26, Table C-29 on page C-27, and Table C-30 on page C-27, and the semaphore and get FR opcode extensions in Table C-31 on page C-28. The floating-point load/store opcode extensions are summarized in Table C-32 on page C-28, Table C-33 on page C-29, and Table C-34 on page C-29, the floating-point load pair and set FR opcode extensions in Table C-35 on page C-30 and Table C-36 on page C-30.

Table C-28. Integer Load/Store Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	x_6				
			Bits 35:32	Bits 31:30			
				0	1	2	3
4	0	0	0	ld1 M1	ld2 M1	ld4 M1	ld8 M1
			1	ld1.s M1	ld2.s M1	ld4.s M1	ld8.s M1
			2	ld1.a M1	ld2.a M1	ld4.a M1	ld8.a M1
			3	ld1.sa M1	ld2.sa M1	ld4.sa M1	ld8.sa M1
			4	ld1.bias M1	ld2.bias M1	ld4.bias M1	ld8.bias M1
			5	ld1.acq M1	ld2.acq M1	ld4.acq M1	ld8.acq M1
			6				ld8.fill M1
			7				
			8	ld1.c.clr M1	ld2.c.clr M1	ld4.c.clr M1	ld8.c.clr M1
			9	ld1.c.nc M1	ld2.c.nc M1	ld4.c.nc M1	ld8.c.nc M1
			A	ld1.c.clr.acq M1	ld2.c.clr.acq M1	ld4.c.clr.acq M1	ld8.c.clr.acq M1
			B				
			C	st1 M4	st2 M4	st4 M4	st8 M4
			D	st1.rel M4	st2.rel M4	st4.rel M4	st8.rel M4
			E				st8.spill M4
			F				

Table C-29. Integer Load +Reg Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	x ₆				
			Bits 35:32	Bits 31:30			
				0	1	2	3
4	1	0	0	ld1 M2	ld2 M2	ld4 M2	ld8 M2
			1	ld1.s M2	ld2.s M2	ld4.s M2	ld8.s M2
			2	ld1.a M2	ld2.a M2	ld4.a M2	ld8.a M2
			3	ld1.sa M2	ld2.sa M2	ld4.sa M2	ld8.sa M2
			4	ld1.bias M2	ld2.bias M2	ld4.bias M2	ld8.bias M2
			5	ld1.acq M2	ld2.acq M2	ld4.acq M2	ld8.acq M2
			6				ld8.fill M2
			7				
			8	ld1.c.clr M2	ld2.c.clr M2	ld4.c.clr M2	ld8.c.clr M2
			9	ld1.c.nc M2	ld2.c.nc M2	ld4.c.nc M2	ld8.c.nc M2
			A	ld1.c.clr.acq M2	ld2.c.clr.acq M2	ld4.c.clr.acq M2	ld8.c.clr.acq M2
			B				
			C				
			D				
			E				
			F				

Table C-30. Integer Load/Store +Imm Opcode Extensions

Opcode Bits 40:37	Bits 35:32	x ₆			
		Bits 31:30			
		0	1	2	3
5	0	ld1 M3	ld2 M3	ld4 M3	ld8 M3
	1	ld1.s M3	ld2.s M3	ld4.s M3	ld8.s M3
	2	ld1.a M3	ld2.a M3	ld4.a M3	ld8.a M3
	3	ld1.sa M3	ld2.sa M3	ld4.sa M3	ld8.sa M3
	4	ld1.bias M3	ld2.bias M3	ld4.bias M3	ld8.bias M3
	5	ld1.acq M3	ld2.acq M3	ld4.acq M3	ld8.acq M3
	6				ld8.fill M3
	7				
	8	ld1.c.clr M3	ld2.c.clr M3	ld4.c.clr M3	ld8.c.clr M3
	9	ld1.c.nc M3	ld2.c.nc M3	ld4.c.nc M3	ld8.c.nc M3
	A	ld1.c.clr.acq M3	ld2.c.clr.acq M3	ld4.c.clr.acq M3	ld8.c.clr.acq M3
	B				
	C	st1 M5	st2 M5	st4 M5	st8 M5
	D	st1.rel M5	st2.rel M5	st4.rel M5	st8.rel M5
	E				st8.spill M5
	F				

Table C-31. Semaphore/Get FR Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	x ₆				
			Bits 35:32	Bits 31:30			
				0	1	2	3
4	0	1	0	cmpxchg1.acq M16	cmpxchg2.acq M16	cmpxchg4.acq M16	cmpxchg8.acq M16
			1	cmpxchg1.rel M16	cmpxchg2.rel M16	cmpxchg4.rel M16	cmpxchg8.rel M16
			2	xchg1 M16	xchg2 M16	xchg4 M16	xchg8 M16
			3				
			4			fetchadd4.acq M17	fetchadd8.acq M17
			5			fetchadd4.rel M17	fetchadd8.rel M17
			6				
			7	getf.sig M19	getf.exp M19	getf.s M19	getf.d M19
			8				
			9				
			A				
			B				
			C				
			D				
			E				
			F				

Table C-32. Floating-point Load/Store/Lfetch Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	x ₆				
			Bits 35:32	Bits 31:30			
				0	1	2	3
6	0	0	0	ldfe M6	ldf8 M6	ldfs M6	ldfd M6
			1	ldfe.s M6	ldf8.s M6	ldfs.s M6	ldfd.s M6
			2	ldfe.a M6	ldf8.a M6	ldfs.a M6	ldfd.a M6
			3	ldfe.sa M6	ldf8.sa M6	ldfs.sa M6	ldfd.sa M6
			4				
			5				
			6				ldf.fill M6
			7				
			8	ldfe.c.clr M6	ldf8.c.clr M6	ldfs.c.clr M6	ldfd.c.clr M6
			9	ldfe.c.nc M6	ldf8.c.nc M6	ldfs.c.nc M6	ldfd.c.nc M6
			A				
			B	lfetch M13	lfetch.excl M13	lfetch.fault M13	lfetch.fault.excl M13
			C	stfe M9	stf8 M9	stfs M9	stfd M9
			D				
			E				stf.spill M9
			F				

Table C-33. Floating-point Load/Lfetch +Reg Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	x ₆				
			Bits 35:32	Bits 31:30			
				0	1	2	3
6	1	0	0	ldfe M7	ldf8 M7	ldfs M7	ldfd M7
			1	ldfe.s M7	ldf8.s M7	ldfs.s M7	ldfd.s M7
			2	ldfe.a M7	ldf8.a M7	ldfs.a M7	ldfd.a M7
			3	ldfe.sa M7	ldf8.sa M7	ldfs.sa M7	ldfd.sa M7
			4				
			5				
			6				ldf.fill M7
			7				
			8	ldfe.c.clr M7	ldf8.c.clr M7	ldfs.c.clr M7	ldfd.c.clr M7
			9	ldfe.c.nc M7	ldf8.c.nc M7	ldfs.c.nc M7	ldfd.c.nc M7
			A				
			B	lfetch M14	lfetch.excl M14	lfetch.fault M14	lfetch.fault.excl M14
			C				
			D				
			E				
			F				

Table C-34. Floating-point Load/Store/Lfetch +Imm Opcode Extensions

Opcode Bits 40:37	Bits 35:32	x ₆			
		Bits 31:30			
		0	1	2	3
7	0	ldfe M8	ldf8 M8	ldfs M8	ldfd M8
	1	ldfe.s M8	ldf8.s M8	ldfs.s M8	ldfd.s M8
	2	ldfe.a M8	ldf8.a M8	ldfs.a M8	ldfd.a M8
	3	ldfe.sa M8	ldf8.sa M8	ldfs.sa M8	ldfd.sa M8
	4				
	5				
	6				ldf.fill M8
	7				
	8	ldfe.c.clr M8	ldf8.c.clr M8	ldfs.c.clr M8	ldfd.c.clr M8
	9	ldfe.c.nc M8	ldf8.c.nc M8	ldfs.c.nc M8	ldfd.c.nc M8
	A				
	B	lfetch M15	lfetch.excl M15	lfetch.fault M15	lfetch.fault.excl M15
	C	stfe M10	stf8 M10	stfs M10	stfd M10
	D				
	E				stf.spill M10
	F				

Table C-35. Floating-point Load Pair/Set FR Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	x ₆				
			Bits 35:32	Bits 31:30			
				0	1	2	3
6	0	1	0		ldfp8 M11	ldfps M11	ldfpd M11
			1		ldfp8.s M11	ldfps.s M11	ldfpd.s M11
			2		ldfp8.a M11	ldfps.a M11	ldfpd.a M11
			3		ldfp8.sa M11	ldfps.sa M11	ldfpd.sa M11
			4				
			5				
			6				
			7	setf.sig M18	setf.exp M18	setf.s M18	setf.d M18
			8		ldfp8.c.clr M11	ldfps.c.clr M11	ldfpd.c.clr M11
			9		ldfp8.c.nc M11	ldfps.c.nc M11	ldfpd.c.nc M11
			A				
			B				
			C				
			D				
			E				
F							

Table C-36. Floating-point Load Pair +Imm Opcode Extensions

Opcode Bits 40:37	m Bit 36	x Bit 27	x ₆				
			Bits 35:32	Bits 31:30			
				0	1	2	3
6	1	1	0		ldfp8 M12	ldfps M12	ldfpd M12
			1		ldfp8.s M12	ldfps.s M12	ldfpd.s M12
			2		ldfp8.a M12	ldfps.a M12	ldfpd.a M12
			3		ldfp8.sa M12	ldfps.sa M12	ldfpd.sa M12
			4				
			5				
			6				
			7				
			8		ldfp8.c.clr M12	ldfps.c.clr M12	ldfpd.c.clr M12
			9		ldfp8.c.nc M12	ldfps.c.nc M12	ldfpd.c.nc M12
			A				
			B				
			C				
			D				
			E				
F							

The load and store instructions all have a 2-bit opcode extension field in bits 29:28 (hint) which encodes locality hint information. Table C-37 and Table C-38 summarize these assignments.


Table C-37. Load Hint Completer

Hint Bits 29:28	ldhint
0	none
1	.ntl
2	
3	.nta

Table C-38. Store Hint Completer


Hint Bits 29:28	sthint
0	none
1	
2	
3	.nta

C.4.1.1 Integer Load




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
ld1.l _{dhint}	$r_1 = [r_3]$	4	0	0	00	See Table C-37 on page C-30
ld2.l _{dhint}					01	
ld4.l _{dhint}					02	
ld8.l _{dhint}					03	
ld1.s.l _{dhint}					04	
ld2.s.l _{dhint}					05	
ld4.s.l _{dhint}					06	
ld8.s.l _{dhint}					07	
ld1.a.l _{dhint}					08	
ld2.a.l _{dhint}					09	
ld4.a.l _{dhint}					0A	
ld8.a.l _{dhint}					0B	
ld1.sa.l _{dhint}					0C	
ld2.sa.l _{dhint}					0D	
ld4.sa.l _{dhint}					0E	
ld8.sa.l _{dhint}					0F	
ld1.bias.l _{dhint}					10	
ld2.bias.l _{dhint}					11	
ld4.bias.l _{dhint}					12	
ld8.bias.l _{dhint}					13	
ld1.acq.l _{dhint}					14	
ld2.acq.l _{dhint}					15	
ld4.acq.l _{dhint}					16	
ld8.acq.l _{dhint}					17	
ld8.fill.l _{dhint}					1B	
ld1.c.clr.l _{dhint}					20	
ld2.c.clr.l _{dhint}					21	
ld4.c.clr.l _{dhint}					22	
ld8.c.clr.l _{dhint}					23	
ld1.c.nc.l _{dhint}					24	
ld2.c.nc.l _{dhint}					25	
ld4.c.nc.l _{dhint}					26	
ld8.c.nc.l _{dhint}	27					
ld1.c.clr.acq.l _{dhint}	28					
ld2.c.clr.acq.l _{dhint}	29					
ld4.c.clr.acq.l _{dhint}	2A					
ld8.c.clr.acq.l _{dhint}	2B					

C.4.1.2 Integer Load – Increment by Register




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
ld1.l _{dhint}	$r_1 = [r_3], r_2$	4	1	0	00	See Table C-37 on page C-30
ld2.l _{dhint}					01	
ld4.l _{dhint}					02	
ld8.l _{dhint}					03	
ld1.s.l _{dhint}					04	
ld2.s.l _{dhint}					05	
ld4.s.l _{dhint}					06	
ld8.s.l _{dhint}					07	
ld1.a.l _{dhint}					08	
ld2.a.l _{dhint}					09	
ld4.a.l _{dhint}					0A	
ld8.a.l _{dhint}					0B	
ld1.sa.l _{dhint}					0C	
ld2.sa.l _{dhint}					0D	
ld4.sa.l _{dhint}					0E	
ld8.sa.l _{dhint}					0F	
ld1.bias.l _{dhint}					10	
ld2.bias.l _{dhint}					11	
ld4.bias.l _{dhint}					12	
ld8.bias.l _{dhint}					13	
ld1.acq.l _{dhint}					14	
ld2.acq.l _{dhint}					15	
ld4.acq.l _{dhint}					16	
ld8.acq.l _{dhint}					17	
ld8.fill.l _{dhint}					1B	
ld1.c.clr.l _{dhint}					20	
ld2.c.clr.l _{dhint}					21	
ld4.c.clr.l _{dhint}					22	
ld8.c.clr.l _{dhint}					23	
ld1.c.nc.l _{dhint}					24	
ld2.c.nc.l _{dhint}					25	
ld4.c.nc.l _{dhint}					26	
ld8.c.nc.l _{dhint}	27					
ld1.c.clr.acq.l _{dhint}	28					
ld2.c.clr.acq.l _{dhint}	29					
ld4.c.clr.acq.l _{dhint}	2A					
ld8.c.clr.acq.l _{dhint}	2B					

C.4.1.3 Integer Load – Increment by Immediate




Instruction	Operands	Opcode	Extension	
			x ₆	hint
ld1.l _{dhint}	$r_1 = [r_3], imm_9$	5	00	See Table C-37 on page C-30
ld2.l _{dhint}			01	
ld4.l _{dhint}			02	
ld8.l _{dhint}			03	
ld1.s.l _{dhint}			04	
ld2.s.l _{dhint}			05	
ld4.s.l _{dhint}			06	
ld8.s.l _{dhint}			07	
ld1.a.l _{dhint}			08	
ld2.a.l _{dhint}			09	
ld4.a.l _{dhint}			0A	
ld8.a.l _{dhint}			0B	
ld1.sa.l _{dhint}			0C	
ld2.sa.l _{dhint}			0D	
ld4.sa.l _{dhint}			0E	
ld8.sa.l _{dhint}			0F	
ld1.bias.l _{dhint}			10	
ld2.bias.l _{dhint}			11	
ld4.bias.l _{dhint}			12	
ld8.bias.l _{dhint}			13	
ld1.acq.l _{dhint}			14	
ld2.acq.l _{dhint}			15	
ld4.acq.l _{dhint}			16	
ld8.acq.l _{dhint}			17	
ld8.fill.l _{dhint}			1B	
ld1.c.clr.l _{dhint}			20	
ld2.c.clr.l _{dhint}			21	
ld4.c.clr.l _{dhint}			22	
ld8.c.clr.l _{dhint}			23	
ld1.c.nc.l _{dhint}			24	
ld2.c.nc.l _{dhint}			25	
ld4.c.nc.l _{dhint}			26	
ld8.c.nc.l _{dhint}	27			
ld1.c.clr.acq.l _{dhint}	28			
ld2.c.clr.acq.l _{dhint}	29			
ld4.c.clr.acq.l _{dhint}	2A			
ld8.c.clr.acq.l _{dhint}	2B			

C.4.1.4 Integer Store




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
st1.st _{hint}	$[r_3] = r_2$	4	0	0	30	See Table C-38 on page C-30
st2.st _{hint}					31	
st4.st _{hint}					32	
st8.st _{hint}					33	
st1.rel.st _{hint}					34	
st2.rel.st _{hint}					35	
st4.rel.st _{hint}					36	
st8.rel.st _{hint}					37	
st8.spill.st _{hint}					3B	

C.4.1.5 Integer Store – Increment by Immediate




Instruction	Operands	Opcode	Extension	
			x ₆	hint
st1. <i>sthint</i>	[r ₃] = r ₂ , imm ₉	5	30	See Table C-38 on page C-30
st2. <i>sthint</i>			31	
st4. <i>sthint</i>			32	
st8. <i>sthint</i>			33	
st1.rel. <i>sthint</i>			34	
st2.rel. <i>sthint</i>			35	
st4.rel. <i>sthint</i>			36	
st8.rel. <i>sthint</i>			37	
st8.spill. <i>sthint</i>			3B	

C.4.1.6 Floating-point Load




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
ldfs. <i>ldhint</i>	f ₁ = [r ₃]	6	0	0	02	See Table C-37 on page C-30
ldfd. <i>ldhint</i>					03	
ldf8. <i>ldhint</i>					01	
ldfe. <i>ldhint</i>					00	
ldfs.s. <i>ldhint</i>					06	
ldfd.s. <i>ldhint</i>					07	
ldf8.s. <i>ldhint</i>					05	
ldfe.s. <i>ldhint</i>					04	
ldfs.a. <i>ldhint</i>					0A	
ldfd.a. <i>ldhint</i>					0B	
ldf8.a. <i>ldhint</i>					09	
ldfe.a. <i>ldhint</i>					08	
ldfs.sa. <i>ldhint</i>					0E	
ldfd.sa. <i>ldhint</i>					0F	
ldf8.sa. <i>ldhint</i>					0D	
ldfe.sa. <i>ldhint</i>					0C	
ldf.fill. <i>ldhint</i>					1B	
ldfs.c.clr. <i>ldhint</i>					22	
ldfd.c.clr. <i>ldhint</i>					23	
ldf8.c.clr. <i>ldhint</i>					21	
ldfe.c.clr. <i>ldhint</i>					20	
ldfs.c.nc. <i>ldhint</i>					26	
ldfd.c.nc. <i>ldhint</i>					27	
ldf8.c.nc. <i>ldhint</i>					25	
ldfe.c.nc. <i>ldhint</i>	24					

C.4.1.7 Floating-point Load – Increment by Register




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
ldfs. <i>ldhint</i>	$f_1 = [r_3], r_2$	6	1	0	02	See Table C-37 on page C-30
ldfd. <i>ldhint</i>					03	
ldf8. <i>ldhint</i>					01	
ldfe. <i>ldhint</i>					00	
ldfs.s. <i>ldhint</i>					06	
ldfd.s. <i>ldhint</i>					07	
ldf8.s. <i>ldhint</i>					05	
ldfe.s. <i>ldhint</i>					04	
ldfs.a. <i>ldhint</i>					0A	
ldfd.a. <i>ldhint</i>					0B	
ldf8.a. <i>ldhint</i>					09	
ldfe.a. <i>ldhint</i>					08	
ldfs.sa. <i>ldhint</i>					0E	
ldfd.sa. <i>ldhint</i>					0F	
ldf8.sa. <i>ldhint</i>					0D	
ldfe.sa. <i>ldhint</i>					0C	
ldf.fill. <i>ldhint</i>					1B	
ldfs.c.clr. <i>ldhint</i>					22	
ldfd.c.clr. <i>ldhint</i>					23	
ldf8.c.clr. <i>ldhint</i>					21	
ldfe.c.clr. <i>ldhint</i>					20	
ldfs.c.nc. <i>ldhint</i>					26	
ldfd.c.nc. <i>ldhint</i>					27	
ldf8.c.nc. <i>ldhint</i>					25	
ldfe.c.nc. <i>ldhint</i>	24					

C.4.1.8 Floating-point Load – Increment by Immediate




Instruction	Operands	Opcode	Extension	
			x ₆	hint
ldfs.l _{dhint}	$f_1 = [r_3], imm_9$	7	02	See Table C-37 on page C-30
ldfd.l _{dhint}			03	
ldf8.l _{dhint}			01	
ldfe.l _{dhint}			00	
ldfs.s.l _{dhint}			06	
ldfd.s.l _{dhint}			07	
ldf8.s.l _{dhint}			05	
ldfe.s.l _{dhint}			04	
ldfs.a.l _{dhint}			0A	
ldfd.a.l _{dhint}			0B	
ldf8.a.l _{dhint}			09	
ldfe.a.l _{dhint}			08	
ldfs.sa.l _{dhint}			0E	
ldfd.sa.l _{dhint}			0F	
ldf8.sa.l _{dhint}			0D	
ldfe.sa.l _{dhint}			0C	
ldf.fill.l _{dhint}			1B	
ldfs.c.clr.l _{dhint}			22	
ldfd.c.clr.l _{dhint}			23	
ldf8.c.clr.l _{dhint}			21	
ldfe.c.clr.l _{dhint}			20	
ldfs.c.nc.l _{dhint}			26	
ldfd.c.nc.l _{dhint}			27	
ldf8.c.nc.l _{dhint}			25	
ldfe.c.nc.l _{dhint}	24			

C.4.1.9 Floating-point Store




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
stfs.s _{thint}	$[r_3] = f_2$	6	0	0	32	See Table C-38 on page C-30
stfd.s _{thint}					33	
stf8.s _{thint}					31	
stfe.s _{thint}					30	
stf.spill.s _{thint}					3B	

C.4.1.10 Floating-point Store – Increment by Immediate




Instruction	Operands	Opcode	Extension	
			x ₆	hint
stfs.s _{thint}	$[r_3] = f_2, imm_9$	7	32	See Table C-38 on page C-30
stfd.s _{thint}			33	
stf8.s _{thint}			31	
stfe.s _{thint}			30	
stf.spill.s _{thint}			3B	

C.4.1.11 Floating-point Load Pair



Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
ldfps.ldhint	$f_1, f_2 = [r_3]$	6	0	1	02	See Table C-37 on page C-30
ldfpd.ldhint					03	
ldfp8.ldhint					01	
ldfps.s.ldhint					06	
ldfpd.s.ldhint					07	
ldfp8.s.ldhint					05	
ldfps.a.ldhint					0A	
ldfpd.a.ldhint					0B	
ldfp8.a.ldhint					09	
ldfps.sa.ldhint					0E	
ldfpd.sa.ldhint					0F	
ldfp8.sa.ldhint					0D	
ldfps.c.clr.ldhint					22	
ldfpd.c.clr.ldhint					23	
ldfp8.c.clr.ldhint					21	
ldfps.c.nc.ldhint					26	
ldfpd.c.nc.ldhint	27					
ldfp8.c.nc.ldhint	25					

C.4.1.12 Floating-point Load Pair – Increment by Immediate



Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
ldfps.ldhint	$f_1, f_2 = [r_3], 8$	6	1	1	02	See Table C-37 on page C-30
ldfpd.ldhint					03	
ldfp8.ldhint	01					
ldfps.s.ldhint	06					
ldfpd.s.ldhint	07					
ldfp8.s.ldhint	05					
ldfps.a.ldhint	0A					
ldfpd.a.ldhint	0B					
ldfp8.a.ldhint	09					
ldfps.sa.ldhint	0E					
ldfpd.sa.ldhint	0F					
ldfp8.sa.ldhint	0D					
ldfps.c.clr.ldhint	22					
ldfpd.c.clr.ldhint	23					
ldfp8.c.clr.ldhint	21					
ldfps.c.nc.ldhint	26					
ldfpd.c.nc.ldhint	27					
ldfp8.c.nc.ldhint	25					

C.4.2 Line Prefetch


The line prefetch instructions are encoded in major opcodes 6 and 7 along with the floating-point load/store instructions. See “Loads and Stores” on page C-26 for a summary of the opcode extensions.

The line prefetch instructions all have a 2-bit opcode extension field in bits 29:28 (hint) which encodes locality hint information as shown in Table C-39.

Table C-39. Line Prefetch Hint Completer


Hint Bits 29:28	<i>lfhint</i>
0	<i>none</i>
1	<i>.nt1</i>
2	<i>.nt2</i>
3	<i>.nta</i>

C.4.2.1 Line Prefetch




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
<i>lfetch.lfhint</i> <i>lfetch.excl.lfhint</i> <i>lfetch.fault.lfhint</i> <i>lfetch.fault.excl.lfhint</i>	[<i>r</i> ₃]	6	0	0	2C 2D 2E 2F	See Table C-39 on page C-38

C.4.2.2 Line Prefetch – Increment by Register



Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
<i>lfetch.lfhint</i> <i>lfetch.excl.lfhint</i> <i>lfetch.fault.lfhint</i> <i>lfetch.fault.excl.lfhint</i>	[<i>r</i> ₃], <i>r</i> ₂	6	1	0	2C 2D 2E 2F	See Table C-39 on page C-38

C.4.2.3 Line Prefetch – Increment by Immediate




Instruction	Operands	Opcode	Extension	
			x ₆	hint
<i>lfetch.lfhint</i> <i>lfetch.excl.lfhint</i> <i>lfetch.fault.lfhint</i> <i>lfetch.fault.excl.lfhint</i>	[<i>r</i> ₃], <i>imm</i> ₉	7	2C 2D 2E 2F	See Table C-39 on page C-38

C.4.3 Semaphores


The semaphore instructions are encoded in major opcode 4 along with the integer load/store instructions. See “Loads and Stores” on page C-26 for a summary of the opcode extensions.

C.4.3.1 Exchange/Compare and Exchange



Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
cmpxchg1.acq.l ^{dhint}	$r_1 = [r_3], r_2, ar.ccv$	4	0	1	00	See Table C-37 on page C-30
cmpxchg2.acq.l ^{dhint}					01	
cmpxchg4.acq.l ^{dhint}					02	
cmpxchg8.acq.l ^{dhint}					03	
cmpxchg1.rel.l ^{dhint}					04	
cmpxchg2.rel.l ^{dhint}					05	
cmpxchg4.rel.l ^{dhint}					06	
cmpxchg8.rel.l ^{dhint}					07	
xchg1.l ^{dhint}	$r_1 = [r_3], r_2$	4	0	1	08	
xchg2.l ^{dhint}					09	
xchg4.l ^{dhint}					0A	
xchg8.l ^{dhint}					0B	

C.4.3.2 Fetch and Add – Immediate




Instruction	Operands	Opcode	Extension			
			m	x	x ₆	hint
fetchadd4.acq.l ^{dhint}	$r_1 = [r_3], inc_3$	4	0	1	12	See Table C-37 on page C-30
fetchadd8.acq.l ^{dhint}					13	
fetchadd4.rel.l ^{dhint}					16	
fetchadd8.rel.l ^{dhint}					17	

C.4.4 Set/Get FR


The set FR instructions are encoded in major opcode 6 along with the floating-point load/store instructions. The get FR instructions are encoded in major opcode 4 along with the integer load/store instructions. See “Loads and Stores” on page C-26 for a summary of the opcode extensions.

C.4.4.1 Set FR



Instruction	Operands	Opcode	Extension		
			m	x	x ₆
setf.sig	$f_1 = r_2$	6	0	1	1C
setf.exp					1D
setf.s					1E
setf.d					1F

C.4.4.2 Get FR




Instruction	Operands	Opcode	Extension		
			m	x	x ₆
getf.sig getf.exp getf.s getf.d	$r_1 = f_2$	4	0	1	1C 1D 1E 1F

C.4.5 Speculation and Advanced Load Checks


The speculation and advanced load check instructions are encoded in major opcodes 0 and 1 along with the system/memory management instructions. See “Memory Management” on page C-43 for a summary of the opcode extensions.

C.4.5.1 Integer Speculation Check (M-Unit)




Instruction	Operands	Opcode	Extension
			x ₃
chk.s.m	$r_2, target_{25}$	1	1

C.4.5.2 Floating-point Speculation Check




Instruction	Operands	Opcode	Extension
			x ₃
chk.s	$f_2, target_{25}$	1	3

C.4.5.3 Integer Advanced Load Check



Instruction	Operands	Opcode	Extension
			x ₃
chk.a.nc chk.a.clr	$r_1, target_{25}$	0	4 5

C.4.5.4 Floating-point Advanced Load Check




Instruction	Operands	Opcode	Extension
			x ₃
chk.a.nc chk.a.clr	$f_1, target_{25}$	0	6 7

C.4.6 Cache/Synchronization/RSE/ALAT


The cache/synchronization/RSE/ALAT instructions are encoded in major opcode 0 along with the memory management instructions. See “Memory Management” on page C-43 for a summary of the opcode extensions.

C.4.6.1 Sync/Fence/Serialize/ALAT Control




Instruction	Opcode	Extension		
		x ₃	x ₄	x ₂
invala	0	0	0	1
mf			2	2
mf.a			3	2
srlz.i			1	3
sync.i			3	3

C.4.6.2 RSE Control




Instruction	Opcode	Extension		
		x ₃	x ₄	x ₂
flushrs ^f	0	0	C	0

C.4.6.3 Integer ALAT Entry Invalidate




Instruction	Operands	Opcode	Extension		
			x ₃	x ₄	x ₂
invala.e	r _I	0	0	2	1

C.4.6.4 Floating-point ALAT Entry Invalidate



Instruction	Operands	Opcode	Extension		
			x ₃	x ₄	x ₂
invala.e	f _I	0	0	3	1

C.4.6.5 Flush Cache




Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
fc	r ₃	1	0	30


C.4.7 GR/AR Moves (M-Unit)

The M-Unit GR/AR move instructions are encoded in major opcode 0 along with the system/memory management instructions. (some ARs are accessed using system control instructions on the I-unit. See “GR/AR Moves (I-Unit)” on page C-25.) See “Memory Management” on page C-43 for a summary of the M-Unit GR/AR opcode extensions.


C.4.7.1 Move to AR – Register (M-Unit)



C.4.7.2 Move to AR – Immediate₈ (M-Unit)




C.4.7.3 Move from AR (M-Unit)



C.4.8 Miscellaneous M-Unit Instructions

The miscellaneous M-unit instructions are encoded in major opcode 0 along with the memory management instructions. See “Memory Management” on page C-43 for a summary of the opcode extensions.

C.4.8.1 Allocate Register Stack Frame




NOTE: The three immediates in the instruction encoding are formed from the operands as follows:

$$sof = i + l + o$$


$$sol = i + l$$

$$sor = r \gg 3$$

C.4.8.2 Move to PSR




C.4.8.3 Move from PSR



Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
mov	$r_1 = \text{psr.um}$	1	0	21

C.4.8.4 Break/Nop (M-Unit)



Instruction	Operands	Opcode	Extension		
			x ₃	x ₄	x ₂
break.m	imm_{21}	0	0	0	0
nop.m			0	1	0

C.4.9 Memory Management

All system/memory management instructions are encoded within major opcodes 0 and 1 using a 3-bit opcode extension field (x_3) in bits 35:33. Some instructions also have a 4-bit opcode extension field (x_4) in bits 30:27, or a 6-bit opcode extension field (x_6) in bits 32:27. Most of the instructions having a 4-bit opcode extension field also have a 2-bit extension field (x_2) in bits 32:31. Table C-40 shows the 3-bit assignments for opcode 0, Table C-41 summarizes the 4-bit+2-bit assignments for opcode 0, Table C-42 shows the 3-bit assignments for opcode 1, and Table C-43 summarizes the 6-bit assignments for opcode 1.

Table C-40. Opcode 0 Memory Management 3-bit Opcode Extensions

Opcode Bits 40:37	x ₃ Bits 35:33	
0	0	System/Memory Management 4-bit+2-bit Ext (Table C-41)
	1	
	2	
	3	
	4	chk.a.nc – int M22
	5	chk.a.clr – int M22
	6	chk.a.nc – fp M23
	7	chk.a.clr – fp M23

Table C-41. Opcode 0 Memory Management 4-bit+2-bit Opcode Extensions

Opcode Bits 40:37	x ₃ Bits 35:33	x ₄ Bits 30:27	x ₂ Bits 32:31			
			0	1	2	3
0	0	0	break.m M37	invala M24		
		1	nop.m M37			srlz.i M24
		2		invala.e – int M26	mf M24	
		3		invala.e – fp M27	mf.a M24	sync.i M24
		4	sum M44			
		5	rum M44			
		6				
		7				
		8			mov.m to ar – imm ₈ M30	
		9				
		A				
		B				
		C	flushrs M25			
		D				
		E				
		F				


Table C-42. Opcode 1 Memory Management 3-bit Opcode Extensions

Opcode Bits 40:37	x ₃ Bits 35:33	
1	0	Memory Management 6-bit Ext (Table C-43)
	1	chk.s.m – int M20
	2	
	3	chk.s – fp M21
	4	
	5	
	6	alloc M34
	7	

Table C-43. Opcode 1 Memory Management 6-bit Opcode Extensions


Opcode Bits 40:37	x ₃ Bits 35:33	Bits 30:27	x ₆ Bits 32:31			
			0	1	2	3
1	0	0				fc M28
		1			mov from psr.um M36	
		2			mov.m from ar M31	
		3				
		4				
		5		mov from pmd M43		
		6				
		7		mov from cpuid M43		
		8				
		9				mov to psr.um M35
		A				mov.m to ar M29
		B				
		C				
		D				
		E				
		F				

C.4.9.1 Move from Indirect Register



Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
mov	$r_I = \text{pmd}[r_3]$ $r_I = \text{cpuid}[r_3]$	1	0	15 17

C.4.9.2 Set/Reset User Mask



Instruction	Operands	Opcode	Extension	
			x ₃	x ₄
sum	imm_{24}	0	0	4
rum			0	5

C.5 B-Unit Instruction Encodings

The branch-unit includes branch and miscellaneous instructions.

C.5.1 Branches

Opcode 0 is used for indirect branch, opcode 1 for indirect call, opcode 4 for IP-relative branch, and opcode 5 for IP-relative call.

The IP-relative branch instructions encoded within major opcode 4 use a 3-bit opcode extension field in bits 8:6 (btype) to distinguish the branch types as shown in Table C-44.

Table C-44. IP-Relative Branch Types

Opcode Bits 40:37	btype bits 8:6	
4	0	br.cond B1
	1	
	2	br.wexit B1
	3	br.wtop B1
	4	
	5	br.cloop B2
	6	br.cexit B2
	7	br.ctop B2

The indirect branch, indirect return, and miscellaneous branch-unit instructions are encoded within major opcode 0 using a 6-bit opcode extension field in bits 32:27 (x₆). Table C-45 summarizes these assignments.

Table C-45. Indirect/Miscellaneous Branch Opcode Extensions

Opcode Bits 40:37	Bits 30:27	x ₆ Bits 32:31			
		0	1	2	3
		0	break.b B9		Indirect Branch (Table C-46)
1			Indirect Return (Table C-47)		
2					
3					
4	clrrb B8				
5	clrrb.pr B8				
6					
7					
8					
9					
A					
B					
C					
D					
E					
F					

The indirect branch instructions encoded within major opcodes 0 use a 3-bit opcode extension field in bits 8:6 (btype) to distinguish the branch types as shown in Table C-46.

Table C-46. Indirect Branch Types

Opcode Bits 40:37	x ₆ Bits 32:27	btype Bits 8:6	
0	20	0	br.cond B4
		1	br.ia B4
		2	
		3	
		4	
		5	
		6	
		7	

The indirect return branch instructions encoded within major opcodes 0 use a 3-bit opcode extension field in bits 8:6 (btype) to distinguish the branch types as shown in Table C-47.

Table C-47. Indirect Return Branch Types

Opcode Bits 40:37	x ₆ Bits 32:27	btype Bits 8:6	
0	21	0	
		1	
		2	
		3	
		4	br.ret B4
		5	
		6	
		7	

All of the branch instructions have a 1-bit opcode extension field, p, in bit 12 which provides a sequential prefetch hint. Table C-48 summarizes these assignments.

Table C-48. Sequential Prefetch Hint Completer

P Bit 12	ph
0	.few
1	.many

The IP-relative and indirect branch instructions all have a 2-bit opcode extension field in bits 34:33 (wh) which encodes branch prediction “whether” hint information as shown in Table C-49. Indirect call instructions have a 3-bit opcode extension field in bits 34:32 (wh) for “whether” hint information as shown in Table C-50.

Table C-49. Branch Whether Hint Completer

wh Bits 34:33	bwh
0	.sptk
1	.spnt
2	.dptk
3	.dpnt

Table C-50. Indirect Call Whether Hint Completer


wh Bits 34:32	bwh
0	
1	.sptk
2	
3	.spnt
4	
5	.dptk
6	
7	.dpnt

The branch instructions also have a 1-bit opcode extension field in bit 35 (d) which encodes a branch cache deallocation hint as shown in Table C-51.

Table C-51. Branch Cache Deallocation Hint Completer


d Bit 35	dh
0	none
1	.clr

C.5.1.1 IP-Relative Branch




Instruction	Operands	Opcode	Extension			
			btype	p	wh	d
br.cond.bwh.ph.dh	target ₂₅	4	0	See	See	See
br.wexit.bwh.ph.dh ^t			2	Table C-48 on page C-46	Table C-49 on page C-47	Table C-51 on page C-47
br.wtop.bwh.ph.dh ^t			3			

C.5.1.2 IP-Relative Counted Branch



Instruction	Operands	Opcode	Extension			
			btype	p	wh	d
br.cloop.bwh.ph.dh ^t	target ₂₅	4	5	See	See	See
br.cexit.bwh.ph.dh ^t			6	Table C-48 on page C-46	Table C-49 on page C-47	Table C-51 on page C-47
br.ctop.bwh.ph.dh ^t			7			

C.5.1.3 IP-Relative Call




Instruction	Operands	Opcode	Extension		
			p	wh	d
br.call.bwh.ph.dh	$b_1 = target_{25}$	5	See Table C-48 on page C-46	See Table C-49 on page C-47	See Table C-51 on page C-47

C.5.1.4 Indirect Branch



Instruction	Operands	Opcode	Extension				
			x ₆	btype	p	wh	d
br.cond.bwh.ph.dh				0	See Table C-48 on page C-46	See Table C-49 on page C-47	See Table C-51 on page C-47
br.ia.bwh.ph.dh	b_2	0	20	1	6	7	7
br.ret.bwh.ph.dh			21	4	6	7	7

C.5.1.5 Indirect Call



Instruction	Operands	Opcode	Extension		
			p	wh	d
br.call.bwh.ph.dh	$b_1 = b_2$	1	See Table C-48 on page C-46	See Table C-50 on page C-47	See Table C-51 on page C-47

C.5.2 Nop

The nop instruction is encoded in major opcode 2. The nop instruction in major opcode 2 uses a 6-bit opcode extension field in bits 32:27 (x_6). Table C-52 summarizes these assignments.


Table C-52. Indirect Predict/Nop Opcode Extensions

Opcode Bits 40:37	Bits 30:27	x_6 Bits 32:31			
		0	1	2	3
		2	0	nop.b B9	
	1				
	2				
	3				
	4				
	5				
	6				
	7				
	8				
	9				
	A				
	B				
	C				
	D				
	E				
	F				

C.5.3 Miscellaneous B-Unit Instructions


The miscellaneous branch-unit instructions include a number of instructions encoded within major opcode 0 using a 6-bit opcode extension field in bits 32:27 (x_6) as described in Table C-45 on page C-46.

C.5.3.1 Miscellaneous (B-Unit)



Instruction	Opcode	Extension
		x_6
clrrrb ¹	0	04
clrrb.pr ¹	0	05

C.5.3.2 Break/Nop (B-Unit)



Instruction	Operands	Opcode	Extension
			x_6
break.b	imm_{21}	0	00
nop.b		2	

C.6 F-Unit Instruction Encodings

The floating-point instructions are encoded in major opcodes 8 – E for floating-point and fixed-point arithmetic, opcode 4 for floating-point compare, opcode 5 for floating-point class, and opcodes 0 and 1 for miscellaneous floating-point instructions.

The miscellaneous and reciprocal approximation floating-point instructions are encoded within major opcodes 0 and 1 using a 1-bit opcode extension field (x) in bit 33 and either a second 1-bit extension field in bit 36 (q) or a 6-bit opcode extension field (x_6) in bits 32:27. Table C-53 shows the 1-bit x assignments, Table C-56 shows the additional 1-bit q assignments for the reciprocal approximation instructions; Table C-54 and Table C-55 summarize the 6-bit x_6 assignments.

Table C-53. Miscellaneous Floating-point 1-bit Opcode Extensions

Opcode Bits 40:37	x Bit 33	
0	0	6-bit Ext (Table C-54)
	1	Reciprocal Approximation (Table C-56)
1	0	6-bit Ext (Table C-55)
	1	Reciprocal Approximation (Table C-56)

Table C-54. Opcode 0 Miscellaneous Floating-point 6-bit Opcode Extensions

Opcode Bits 40:37	x Bit 33	x ₆				
		Bits 30:27	Bits 32:31			
			0	1	2	3
0	0	0	break.f F15	fmerge.s F9		
		1	nop.f F15	fmerge.ns F9		
		2		fmerge.se F9		
		3				
		4	fsetc F12	fmin F8		fswap F9
		5	fclrf F13	fmax F8		fswap.nl F9
		6		famin F8		fswap.nr F9
		7		famax F8		
		8	fchkf F14	fcvt.fx F10	fpack F9	
		9		fcvt.fxu F10		fmix.lr F9
		A		fcvt.fx.trunc F10		fmix.r F9
		B		fcvt.fxu.trunc F10		fmix.l F9
		C		fcvt.xf F11	fand F9	fsxt.r F9
		D			fandem F9	fsxt.l F9
		E			for F9	
		F			fxor F9	

Table C-55. Opcode 1 Miscellaneous Floating-point 6-bit Opcode Extensions

Opcode Bits 40:37	x Bit 33	x ₆				
		Bits 30:27	Bits 32:31			
			0	1	2	3
1	0	0		fpmerge.s F9		fpcmp.eq F8
		1		fpmerge.ns F9		fpcmp.lt F8
		2		fpmerge.se F9		fpcmp.le F8
		3				fpcmp.unord F8
		4		fpmin F8		fpcmp.neq F8
		5		fpmax F8		fpcmp.nlt F8
		6		fpamin F8		fpcmp.nle F8
		7		fpamax F8		fpcmp.ord F8
		8		fpcvt.fx F10		
		9		fpcvt.fxu F10		
		A		fpcvt.fx.trunc F10		
		B		fpcvt.fxu.trunc F10		
		C				
		D				
		E				
		F				

Table C-56. Reciprocal Approximation 1-bit Opcode Extensions

Opcode Bits 40:37	x Bit 33	q Bit 36	
0	1	0	frcpa F6
		1	frsqta F7
1		0	fprcpa F6
		1	fprsqta F7

Most floating-point instructions have a 2-bit opcode extension field in bits 35:34 (sf) which encodes the FPSR status field to be used. Table C-57 summarizes these assignments.

Table C-57. Floating-point Status Field Completer

sf Bits 35:34	sf
0	.s0
1	.s1
2	.s2
3	.s3

C.6.1 Arithmetic

The floating-point arithmetic instructions are encoded within major opcodes 8 – D using a 1-bit opcode extension field (x) in bit 36 and a 2-bit opcode extension field (sf) in bits 35:34. The opcode and x assignments are shown in Table C-58.

Table C-58. Floating-point Arithmetic 1-bit Opcode Extensions


x Bit 36	Opcode Bits 40:37					
	8	9	A	B	C	D
0	fma F1	fma.d F1	fms F1	fms.d F1	fnma F1	fnma.d F1
1	fma.s F1	fpma F1	fms.s F1	fpms F1	fnma.s F1	fpnma F1

The fixed-point arithmetic and parallel floating-point select instructions are encoded within major opcode E using a 1-bit opcode extension field (x) in bit 36. The fixed-point arithmetic instructions also have a 2-bit opcode extension field (x₂) in bits 35:34. These assignments are shown in Table C-59.

Table C-59. Fixed-point Multiply Add and Select Opcode Extensions


Opcode Bits 40:37	x Bit 36	x ₂ Bits 35:34			
		0	1	2	3
E	0	fselect F3			
	1	xma.l F2		xma.hu F2	xma.h F2

C.6.1.1 Floating-point Multiply Add




Instruction	Operands	Opcode	Extension	
			x	sf
fma.sf	$f_1 = f_3, f_4, f_2$	8	0	See Table C-57 on page C-51
fma.s.sf			1	
fma.d.sf		9	0	
fpma.sf			1	
fms.sf		A	0	
fms.s.sf			1	
fms.d.sf			B	
fpms.sf		1		
fnma.sf		C	0	
fnma.s.sf			1	
fnma.d.sf			D	
fpnma.sf		1		

C.6.1.2 Fixed-point Multiply Add



Instruction	Operands	Opcode	Extension	
			x	x ₂
xma.l	$f_1 = f_3, f_4, f_2$	E	1	0
xma.h			1	3
xma.hu			1	2

C.6.2 Parallel Floating-point Select



Instruction	Operands	Opcode	Extension
			x
fselect	$f_1 = f_3, f_4, f_2$	E	0

C.6.3 Compare and Classify

The predicate setting floating-point compare instructions are encoded within major opcode 4 using three 1-bit opcode extension fields in bits 33 (r_a), 36 (r_b), and 12 (t_a), and a 2-bit opcode extension field (sf) in bits 35:34. The opcode, r_a , r_b , and t_a assignments are shown in Table C-60. The sf assignments are shown in Table C-57 on page C-51.

The parallel floating-point compare instructions are described on page C-54.

Table C-60. Floating-point Compare Opcode Extensions


Opcode Bits 40:37	r_a Bit 33	r_b Bit 36	t_a Bit 12	
			0	1
4	0	0	fcmp.eq F4	fcmp.eq.unc F4
		1	fcmp.lt F4	fcmp.lt.unc F4
	1	0	fcmp.le F4	fcmp.le.unc F4
		1	fcmp.unord F4	fcmp.unord.unc F4

The floating-point class instructions are encoded within major opcode 5 using a 1-bit opcode extension field in bit 12 (t_a) as shown in Table C-61.

Table C-61. Floating-point Class 1-bit Opcode Extensions


Opcode Bits 40:37	t_a Bit 12	
5	0	fclass.m F5
	1	fclass.m.unc F5

C.6.3.1 Floating-point Compare



Instruction	Operands	Opcode	Extension			
			r_a	r_b	t_a	sf
fcmp.eq.sf	$p_1, p_2 = f_2, f_3$	4	0	0	0	See Table C-57 on page C-51
fcmp.lt.sf				1		
fcmp.le.sf			0			
fcmp.unord.sf			1			
fcmp.eq.unc.sf			0	1		
fcmp.lt.unc.sf			1			
fcmp.le.unc.sf			0			
fcmp.unord.unc.sf			1			

C.6.3.2 Floating-point Class




Instruction	Operands	Opcode	Extension
			ta
fclass.m	$p_1, p_2 = f_2, fclass_9$	5	0
fclass.m.unc			1

C.6.4 Approximation

C.6.4.1 Floating-point Reciprocal Approximation


There are two Reciprocal Approximation instructions. The first, in major op 0, encodes the full register variant. The second, in major op 1, encodes the parallel variant.



Instruction	Operands	Opcode	Extension		
			x	q	sf
frcpa.sf	$f_1, p_2 = f_2, f_3$	0	1	0	See Table C-57 on page C-51
fprcpa.sf		1	0		

C.6.4.2 Floating-point Reciprocal Square Root Approximation


There are two Reciprocal Square Root Approximation instructions. The first, in major op 0, encodes the full register variant. The second, in major op 1, encodes the parallel variant.



Instruction	Operands	Opcode	Extension		
			x	q	sf
frsqrta.sf	$f_1, p_2 = f_3$	0	1	1	See Table C-57 on page C-51
fprsqrta.sf		1			


C.6.5 Minimum/Maximum and Parallel Compare

There are 2 groups of Minimum/Maximum instructions. The first group, in major op 0, encodes the full register variants. The second group, in major op 1, encodes the parallel variants. The parallel compare instructions are all encoded in major op 1.



Instruction	Operands	Opcode	Extension			
			x	x ₆	sf	
fmin. <i>sf</i>	$f_1 = f_2, f_3$	0	0	14	See Table C-57 on page C-51	
fmax. <i>sf</i>				15		
famin. <i>sf</i>				16		
famax. <i>sf</i>				17		
fpmin. <i>sf</i>				1		14
fpmax. <i>sf</i>						15
fpamin. <i>sf</i>						16
fpamax. <i>sf</i>						17
fpcmp.eq. <i>sf</i>						30
fpcmp.lt. <i>sf</i>						31
fpcmp.le. <i>sf</i>						32
fpcmp.unord. <i>sf</i>						33
fpcmp.neq. <i>sf</i>						34
fpcmp.nlt. <i>sf</i>						35
fpcmp.nle. <i>sf</i>	36					
fpcmp.ord. <i>sf</i>	37					


C.6.6 Merge and Logical



Instruction	Operands	Opcode	Extension	
			x	x ₆
fmerge.s	$f_1 = f_2, f_3$	0	0	10
fmerge.ns				11
fmerge.se				12
fmix.lr				39
fmix.r				3A
fmix.l				3B
fsxt.r				3C
fsxt.l				3D
fpack				28
fswap				34
fswap.nl				35
fswap.nr				36
fand				2C
fandcm				2D
for	2E			
fxor	2F			
fpmerge.s	1	1	0	10
fpmerge.ns				11
fpmerge.se				12


C.6.7 Conversion

C.6.7.1 Convert Floating-point to Fixed-point



Instruction	Operands	Opcode	Extension		
			x	x ₆	sf
fcvt.fx.sf fcvt.fxu.sf fcvt.fx.trunc.sf fcvt.fxu.trunc.sf	$f_1 = f_2$	0	0	18 19 1A 1B	See Table C-57 on page C-51
fpcvt.fx.sf fpcvt.fxu.sf fpcvt.fx.trunc.sf fpcvt.fxu.trunc.sf		1		18 19 1A 1B	


C.6.7.2 Convert Fixed-point to Floating-point



Instruction	Operands	Opcode	Extension	
			x	x ₆
fcvt.xf	$f_1 = f_2$	0	0	1C


C.6.8 Status Field Manipulation

C.6.8.1 Floating-point Set Controls




Instruction	Operands	Opcode	Extension		
			x	x ₆	sf
fsetc.sf	$amask_7, omask_7$	0	0	04	See Table C-57 on page C-51

C.6.8.2 Floating-point Clear Flags



Instruction	Opcode	Extension		
		x	x ₆	sf
fclrf.sf	0	0	05	See Table C-57 on page C-51


C.6.8.3 Floating-point Check Flags



Instruction	Operands	Opcode	Extension		
			x	x ₆	sf
fchkf.sf	$target_{25}$	0	0	08	See Table C-57 on page C-51

C.6.9 Miscellaneous F-Unit Instructions

C.6.9.1 Break/Nop (F-Unit)



Instruction	Operands	Opcode	Extension	
			x	x ₆
break.f	<i>imm₂₁</i>	0	0	00
nop.f			0	01

C.7 X-Unit Instruction Encodings

The X-unit instructions occupy two instruction slots, L+X. The major opcode, opcode extensions and hints, qp, and small immediate fields occupy the X instruction slot. For movl, break.x, and nop.x, the imm₄₁ field occupies the L instruction slot.

C.7.1 Miscellaneous X-Unit Instructions

The miscellaneous X-unit instructions are encoded in major opcode 0 using a 3-bit opcode extension field (x₃) in bits 35:33 and a 6-bit opcode extension field (x₆) in bits 32:27. Table C-62 shows the 3-bit assignments and Table C-63 summarizes the 6-bit assignments. These instructions are executed by an I-unit.


Table C-62. Misc X-Unit 3-bit Opcode Extensions

Opcode Bits 40:37	x ₃ Bits 35:33	
0	0	6-bit Ext (Table C-63)
	1	
	2	
	3	
	4	
	5	
	6	
	7	

Table C-63. Misc X-Unit 6-bit Opcode Extensions

Opcode Bits 40:37	x ₃ Bits 35:33	x ₆ Bits 32:31				
		Bits 30:27	0	1	2	3
		0	0	break.x X1		
		1	nop.x X1			
		2				
		3				
		4				
		5				
		6				
		7				
		8				
		9				
		A				
		B				
		C				
		D				
		E				
		F				

C.7.1.1 Break/Nop (X-Unit)




Instruction	Operands	Opcode	Extension	
			x ₃	x ₆
break.x	<i>imm</i> ₆₂	0	0	00
nop.x				01

C.7.2 Move Long Immediate₆₄

The move long immediate instruction is encoded within major opcode 6 using a 1-bit reserved opcode extension in bit 20 (*v_c*) as shown in Table C-64. This instruction is executed by an I-unit.

Table C-64. Move Long 1-bit Opcode Extensions



Instruction	Operands	Opcode	Extension
			<i>v_c</i>
movl	<i>r₁</i> = <i>imm</i> ₆₄	6	0

C.8 Immediate Formation

The following table shows, for each instruction format that has one or more immediates, how those immediates are formed. In each equation, the symbol to the left of the equals is the assembly language name for the immediate. The symbols to the right are the field names in the instruction encoding.

Table C-65. Immediate Formation

Instruction Format	Immediate Formation
A2	$count_2 = ct_{2d} + 1$
A3 A8 I27 M30	$imm_8 = sign_ext(s \ll 7 \mid imm_{7b}, 8)$
A4	$imm_{14} = sign_ext(s \ll 13 \mid imm_{6d} \ll 7 \mid imm_{7b}, 14)$
A5	$imm_{22} = sign_ext(s \ll 21 \mid imm_{5c} \ll 16 \mid imm_{9d} \ll 7 \mid imm_{7b}, 22)$
A10	$count_2 = (ct_{2d} > 2) ? reservedQP^a : ct_{2d} + 1$
I1	$count_2 = (ct_{2d} == 0) ? 0 : (ct_{2d} == 1) ? 7 : (ct_{2d} == 2) ? 15 : 16$
I3	$motype_4 = (m_{bt_{4c}} == 0) ? @brst : (m_{bt_{4c}} == 8) ? @mix : (m_{bt_{4c}} == 9) ? @shuf : (m_{bt_{4c}} == 0xA) ? @alt : (m_{bt_{4c}} == 0xB) ? @rev : reservedQP^a$
I4	$m_{htype}_8 = m_{ht}_{8c}$
I6	$count_5 = count_{5b}$
I8	$count_5 = 31 - ccount_{5c}$
I10	$count_6 = count_{6d}$
I11	$len_6 = len_{6d} + 1$ $pos_6 = pos_{6b}$
I12	$len_6 = len_{6d} + 1$ $pos_6 = 63 - cpos_{6c}$

Table C-65. Immediate Formation (Continued)

Instruction Format	Immediate Formation
I13	$len_6 = len_{6d} + 1$ $pos_6 = 63 - cpos_{6c}$ $imm_8 = sign_ext(s \ll 7 imm_{7b}, 8)$
I14	$len_6 = len_{6d} + 1$ $pos_6 = 63 - cpos_{6b}$ $imm_1 = sign_ext(s, 1)$
I15	$len_4 = len_{4d} + 1$ $pos_6 = 63 - cpos_{6d}$
I16	$pos_6 = pos_{6b}$
I19 M37	$imm_{21} = i \ll 20 imm_{20a}$
I23	$mask_{17} = sign_ext(s \ll 16 mask_{8c} \ll 8 mask_{7a} \ll 1, 17)$
I24	$imm_{44} = sign_ext(s \ll 43 imm_{27a} \ll 16, 44)$
M3 M8 M15	$imm_9 = sign_ext(s \ll 8 i \ll 7 imm_{7b}, 9)$
M5 M10	$imm_9 = sign_ext(s \ll 8 i \ll 7 imm_{7a}, 9)$
M17	$inc_3 = sign_ext(((s) ? -1 : 1) * ((i_{2b} == 3) ? 1 : 1 \ll (4 - i_{2b})), 6)$
I20 M20 M21	$target_{25} = IP + (sign_ext(s \ll 20 imm_{13c} \ll 7 imm_{7a}, 21) \ll 4)$
M22 M23	$target_{25} = IP + (sign_ext(s \ll 20 imm_{20b}, 21) \ll 4)$
M34	$il = sol$ $o = sof - sol$ $r = sor \ll 3$
M44	$imm_{24} = i \ll 23 i_{2d} \ll 21 imm_{21a}$
B1 B2 B3	$target_{25} = IP + (sign_ext(s \ll 20 imm_{20b}, 21) \ll 4)$
B9	$imm_{21} = i \ll 20 imm_{20a}$
F5	$fclass_9 = fclass_{7c} \ll 2 fc_2$
F12	$amask_7 = amask_{7b}$ $omask_7 = omask_{7c}$
F14	$target_{25} = IP + (sign_ext(s \ll 20 imm_{20a}, 21) \ll 4)$
F15	$imm_{21} = i \ll 20 imm_{20a}$
X1	$imm_{62} = imm_{41} \ll 21 i \ll 20 imm_{20a}$
X2	$imm_{64} = i \ll 63 imm_{41} \ll 22 i_c \ll 21 imm_{5c} \ll 16 imm_{9d} \ll 7 imm_{7b}$

a. This encoding causes an Illegal Operation fault if the value of the qualifying predicate is 1.